On the evaluation complexity of constrained nonlinear least-squares and general constrained nonlinear optimization using second-order methods

Cartis, Coralia; Gould, Nicholas I M; Toint, Philippe L.

Published in:
SIAM Journal on Numerical Analysis

DOI:
10.1137/130915546

Publication date:
2015

Document Version
Publisher's PDF, also known as Version of record

SIAM J. Numer. Anal.
Vol. 53, No. 2, pp. 836–851

ON THE EVALUATION COMPLEXITY OF CONSTRAINED NONLINEAR LEAST-SQUARES AND GENERAL CONSTRAINED NONLINEAR OPTIMIZATION USING SECOND-ORDER METHODS*

CORALIA CARTIS†, NICHOLAS I. M. GOULD‡, AND PHILIPPE L. TOINT§

Abstract. When solving the general smooth nonlinear and possibly nonconvex optimization problem involving equality and/or inequality constraints, an approximate first-order critical point of accuracy ϵ can be obtained by a second-order method using cubic regularization in at most $O(\epsilon^{-3/2})$ evaluations of problem functions, the same order bound as in the unconstrained case. This result is obtained by first showing that the same result holds for inequality constrained nonlinear least-squares. As a consequence, the presence of (possibly nonconvex) equality/inequality constraints does not affect the complexity of finding approximate first-order critical points in nonconvex optimization. This result improves on the best known ($O(\epsilon^{-2})$) evaluation-complexity bound for solving general nonconvexly constrained optimization problems.

Key words. evaluation complexity, worst-case analysis, least-squares problems, constrained nonlinear optimization, cubic regularization methods

AMS subject classifications. 68Q25, 90C60, 90C30, 90C26, 65K05

DOI. 10.1137/130915546

1. Introduction.

The past few years have seen several contributions on the worst-case evaluation complexity of solving smooth but possibly nonconvex optimization problems. Here by *evaluation complexity* of an algorithm on a given problem, we mean the number of *evaluations* of problem functions—objective, constraints (when present), and their derivatives (as needed)—required by the algorithm to reach an approximate problem solution. Starting with the contributions of Vavasis (1993), Nesterov (2004), and Gratton, Sartenaer, and Toint (2008) on (essentially) first-order methods for the unconstrained case, a significant step was made with the proposal by Nesterov and Polyak (2006) of a second-order method including cubic regularization terms. The latter paper showed that solving the smooth unconstrained nonconvex optimization problem can be achieved (using this second-order method) in at most $O(\epsilon^{-3/2})$ evaluations if one is happy to terminate the process with an approximate first-order critical point at which the Euclidean norm of the objective function's gradient is at most a user-prescribed threshold $\epsilon \in (0, 1)$. This is in contrast with what can be obtained for first-order methods, which require at most $O(\epsilon^{-2})$ evaluations in a similar context. This remarkable result by Nesterov was subsequently extended by Cartis, Gould, and Toint (2011a, 2011b) to a wider class of algorithms, leading to the conclusion that the class of cubic regularization (ARC) methods and its complexity of order $O(\epsilon^{-3/2})$ are optimal for the smooth unconstrained nonlinear optimization problem in terms of the worst-case number of function evaluations required.

*Received by the editors April 13, 2013; accepted for publication (in revised form) January 20, 2015; published electronically March 26, 2015.
http://www.siam.org/journals/sinum/53-2/91554.html
†The Mathematical Institute, University of Oxford, Woodstock Road, Oxford, OX2 6GG, UK (cartis@maths.ox.ac.uk).
‡Numerical Analysis Group, Rutherford Appleton Laboratory, Chilton, OX11 0QX, UK (nick.gould@stfc.ac.uk).
§Namur Center for Complex Systems (naXys) and Department of Mathematics, University of Namur, 61, rue de Bruxelles, B-5000 Namur, Belgium (philippe.toint@unamur.be).

Following up on these results for the unconstrained problem, the authors of this paper then examined the smooth constrained problem and showed that, somewhat surprisingly, the complexity in $O(\epsilon^{-2})$ evaluations obtained for first-order methods is not affected at all (in order) by the presence of equality and/or inequality constraints (see Cartis, Gould, and Toint (2012a) for the convex inequality case, and Cartis, Gould, and Toint (2014) for the general nonconvex case). Moreover, the first of these papers also showed that a complexity of order $O(\epsilon^{-3/2})$ evaluations can also be achieved under some conditions for the problem involving convex inequalities, while a very similar result was presented for the general nonconvex equality constrained case in Cartis, Gould, and Toint (2013), provided one is ready to solve the primal more accurately than the dual. This left open the central question of whether the general smooth problem involving both nonconvex equality and inequality constraints can be solved in at most $O(\epsilon^{-3/2})$ evaluations using a cubic regularization method under similar conditions.

The purpose of the present paper is to confirm this proposition, thereby providing a complete worst-case analysis for the computation of approximate first-order critical points in smooth optimization. This is achieved by considering, without loss of generality, a general nonlinear optimization problem of the form

$$\min_{x\in\mathbb{R}^n} f(x) \quad \text{subject to} \quad c(x) = 0 \quad \text{and} \quad x \in \mathcal{F}, \tag{1.1}$$

where $f : \mathbb{R}^n \to \mathbb{R}$, $c : \mathbb{R}^n \to \mathbb{R}^m$, and $\mathcal{F} \subseteq \mathbb{R}^n$ is a closed nonempty convex set; also, f and c_i, $i \in \{1, \dots, m\}$, are sufficiently smooth in $\mathcal{F}$ (see AS6–AS8 in section 3). Of particular interest is the case where $\mathcal{F} \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \mid x \geq x_{\text{low}}\}$, where x_{low} is a vector of lower bounds with some components being possible equal to $-\infty$ and the inequality in this definition being understood componentwise. Indeed, it is well known that this more specific formulation of (1.1) covers that involving explicit and possibly nonconvex inequality constraints, and can be obtained from the latter by the incorporation of slack variables. The set $\mathcal{F}$ could be described by finitely many convex inequality constraints or it could be given in a geometric way, as long as we can calculate orthogonal projections onto $\mathcal{F}$.

We present an algorithm for (1.1) for which the $O(\epsilon^{-3/2})$ upper complexity bound holds, and which is inspired by the two-stage methods of Cartis, Gould, and Toint (2014) and Cartis, Gould, and Toint (2013), combined with the projection technique described in Cartis, Gould, and Toint (2012a). Broadly speaking, this algorithm consists of two phases as follows:

- Phase 1: a least-squares formulation of the constraint violation is minimized subject to convex constraints, resulting in either an (approximate) feasible point or a local infeasible minimizer of the constraint violation, an outcome which cannot be excluded from our analysis barring the use of global optimization techniques.
- Phase 2 (entered if an approximately feasible point is obtained in Phase 1): a short-step target-following technique (Cartis, Gould, and Toint (2013, 2014)) is used to reduce the objective function while preserving (approximate) feasibility of the iterates. This involves defining a least-squares merit function involving the objective's deviation from a set target and the constraint violation, and applying (constrained) ARC-like steps to this auxiliary function. Once such successful steps are generated, the target is adjusted to ensure the original problem objective decreases and approximate feasibility is maintained.

Because both phases of our algorithm crucially depend on the solution of a convexly constrained nonlinear least-squares problem, we start by considering this minimization problem and its complexity in section 2. We then turn to the general constrained case, present our two-phase method, and analyze its complexity in section 3. Some conclusions and perspectives are finally discussed in section 4.

2. An algorithm for constrained nonlinear least-squares problems.

2.1. A review of the COCARC-S algorithm for minimizing convexly constrained nonconvex objectives. The COCARC-S algorithm (Cartis, Gould, and Toint (2012a)) is designed for minimizing a general nonlinear twice continuously differentiable objective function $f : \mathbb{R}^n \to \mathbb{R}$ within the closed convex set $\mathcal{F}$, that is,

$$\min_{x \in \mathcal{F}} f(x). \tag{2.1}$$

Iteration k of the COCARC-S algorithm proceeds by first checking the approximate first-order criticality of the current iterate. This is achieved by testing if

$$\chi_f(x_k) \le \epsilon, \tag{2.2}$$

where $\epsilon > 0$ is a user-specified accuracy threshold and where

$$\chi_f(x) \stackrel{\text{def}}{=} \left| \min_{x+d \in \mathcal{F}, \|d\| \le 1} \langle \nabla_x f(x), d \rangle \right| \tag{2.3}$$

($\langle \cdot, \cdot \rangle$ denotes the standard Euclidean inner product) is the linearized decrease in f achievable inside a feasible neighborhood of diameter one (Conn et al. (1993); see also Yuan (1985)). If (2.2) fails, a step s_k is computed from the iterate x_k by (approximately) minimizing a cubic model of f of the form

$$m_k(x_k + s) \stackrel{\text{def}}{=} f(x_k) + \langle \nabla_x f(x_k), s \rangle + \tfrac{1}{2} \langle s, B_k s \rangle + \tfrac{1}{3} \sigma_k \|s\|^3 \tag{2.4}$$

subject to $x_k + s \in \mathcal{F}$ for a given "regularization weight" $\sigma_k > 0$, and where B_k is a symmetric approximation of $\nabla_{xx} f(x_k)$. The approximate solution/trial step $x_k^+ = x_k + s_k$ of this constrained model subproblem, namely,

$$\min_{s \in \mathbb{R}^n, x_k + s \in \mathcal{F}} m_k(x_k + s), \tag{2.5}$$

needs to be calculated accurately enough to ensure that

$$\chi_k^{\text{m}}(x_k^+) \le \min(\kappa_{\text{stop}}, \|s_k\|) \, \chi_k, \tag{2.6}$$

where $\kappa_{\text{stop}} \in [0, 1)$ is a constant and where we employ the model-specific first-order criticality measure[1]

$$\chi_k^{\text{m}}(x) \stackrel{\text{def}}{=} \left| \min_{x+d \in \mathcal{F}, \|d\| \le 1} \langle \nabla_x m_k(x), d \rangle \right|. \tag{2.7}$$

Furthermore, it is important (from a complexity point of view and for our results here) to also assume/ensure that this model minimization can be seen as performing

[1] Inequality (2.6) is an adequate stopping condition for the subproblem solution since $\chi_k^{\text{m}}(x_k^*)$ must be identically zero if x_k^* is a local minimizer of (2.5).

a number ℓ_k of successive (possibly incomplete) line minimizations of the model m_k between feasible points (see AS8 in Cartis, Gould, and Toint (2012a)).[2]

Once the trial point $x_k^+ = x_k + s_k$ is computed, the achieved reduction $f(x_k) - f(x_k^+)$ is compared to the predicted one, $f(x_k) - m_k(x_k^+)$. If the ratio of the former to the latter is sufficiently positive, x_k^+ is accepted as the next iterate x_{k+1} and the regularization weight is (possibly) decreased (an iteration where this occurs is called a *successful* iteration). If this ratio is not sufficiently positive, i.e., is below some constant $\eta_1 \in (0,1)$, the trial point is rejected ($x_{k+1} = x_k$) and the regularization weight is increased (by a factor at least $\gamma_1 > 1$).

Algorithm 2.1: COCARC-S Algorithm for (2.1)

A starting point $x_0 \in \mathcal{F}$, a minimum regularization parameter $\sigma_{\min} > 0$, an initial regularization parameter $\sigma_0 \geq \sigma_{\min}$, and algorithmic parameters $\gamma_2 \geq \gamma_1 > 1$ and $1 > \eta_1 > 0$, as well as the tolerance $\epsilon \in (0,1)$ are given.

Step 0: Check for termination. If $\chi_f(x_k) \leq \epsilon$, terminate.

Step 1: Computation of the step. Starting from x_k, approximately minimize $m_k(x_k + s)$ subject to the $x_k + s \in \mathcal{F}$, yielding a trial point $x_k^+ = x_k + s_k$.

Step 2: Acceptance of the trial point. Compute $f(x_k^+)$ and define

$$\rho_k = \frac{f(x_k) - f(x_k^+)}{m_k(x_k) - m_k(x_k^+)}. \tag{2.8}$$

If $\rho_k \geq \eta_1$, then $x_{k+1} = x_k^+$, else set $x_{k+1} = x_k$.

Step 3: Regularization parameter update. If $\rho_k \geq \eta_1$, choose

$$\sigma_{k+1} \in [\sigma_{\min}, \gamma_1 \sigma_k). \tag{2.9}$$

Otherwise, choose $\sigma_{k+1} \in [\gamma_1 \sigma_k, \gamma_2 \sigma_k]$.

Note that a feasible x_0 can be obtained by projection of any user-supplied initial guess onto the convex set $\mathcal{F}$. For future reference, we define $\mathcal{S}$ to be the index set of the successful iterations, that is,

$$\mathcal{S} \stackrel{\text{def}}{=} \{k \geq 0 \mid \rho_k \geq \eta_1\}. \tag{2.10}$$

Denote by $\mathcal{X}$ the closed convex hull of all iterates x_k and trial points x_k^+.

What is the maximum number of COCARC-S iterations that can be necessary before an iterate x_k is found which satisfies (2.2)? In order to answer this question, we now recall the assumptions used to derive that the required complexity results for the COCARC-S algorithm and those of these results that are of interest in our context. More specifically, we assume the following hold.

[2]Note that solving (2.5) to ensure (2.6) and use line minimizations does not require additional objective evaluations, and so it will not worsen the evaluation complexity bound of COCARC-S.

$\boxed{\textbf{AS1}}$ $\mathcal{F}$ is closed, convex, and nonempty.

$\boxed{\textbf{AS2f}}$ f is twice continuously differentiable, its gradient is uniformly Lipschitz continuous on $\mathcal{X}$, and its Hessian is "weakly" uniformly Lipschitz continuous on the segments $[x_k, x_k^+]$, in the sense that there exists a constant $\kappa_{\rm L} \geq 0$ such that, for all k and all $y \in [x_k, x_k^+]$,

$$\|[\nabla_{xx} f(y) - \nabla_{xx} f(x_k)]s_k\| \leq \kappa_{\rm L} \|s_k\|^2.$$

$\boxed{\textbf{AS3f}}$ The Hessian $\nabla_{xx} f(x_k)$ is well approximated by B_k, in the sense that there exists a constant $\kappa_{\rm B} > 0$ such that, for all k,

$$\|[B_k - \nabla_{xx} f(x_k)]s_k\| \leq \kappa_{\rm B} \|s_k\|^2.$$

$\boxed{\textbf{AS4}}$ ℓ_k is bounded above by a constant independent of k and ϵ.

$\boxed{\textbf{AS5}}$ $\mathcal{X}$ is bounded.

Assumptions AS1, AS2f, AS3f, and AS5 are relatively standard, even if AS5 may be judged to be somewhat restrictive. AS4 is really an assumption on the solution of the subproblem of minimizing the model (2.4) subject to convex constraints, which raises the question of a practical algorithm to perform this task. A first-order (steepest-descent-like) variant of the COCARC-S algorithm can be used to minimize the cubic model over the convex set $\mathcal{F}$, with a provable bound on its iteration complexity (see section 4.3 of Cartis, Gould, and Toint (2009)). The computational or iteration complexity of this or other model subproblem minimization procedures is not included in our evaluation complexity bounds (for the second-order COCARC-S algorithm) since it only involves the model and not the objective function, and as such it does not affect the evaluation complexity of the COCARC-S algorithm. The particular way and property of solving the model required in AS4 remains a (reasonable) assumption in our results. The reason AS4 is needed in our complexity results is described in detail in sections 4.2.2–4.2.3 in Cartis, Gould, and Toint (2012a) by means of a simple example (that satisfies AS4).

The above conditions guarantee that the following holds.

LEMMA 2.1. *Suppose that AS1, AS2f, AS3f, AS4, and AS5 hold. Then there exists a positive constant $\kappa_{\rm C} \in (0,1)$ independent of k and ϵ such that*

$$f(x_k) - f(x_{k+1}) \geq \eta_1 \kappa_{\rm C} \chi_f(x_{k+1})^{\frac{3}{2}} \quad \text{for all } k \in \mathcal{S}. \tag{2.11}$$

Moreover, there exists a constant $\sigma_{\rm succ} > 0$ independent of k and ϵ such that, if $\sigma_k \geq \sigma_{\rm succ}$, then iteration k of the COCARC-S algorithm is successful ($k \in \mathcal{S}$). Finally, there exists a constant $\sigma_{\max} \geq \sigma_{\rm succ}$ such that $\sigma_k \leq \sigma_{\max}$ for all $k \geq 0$.

Proof. Relation (2.11) follows from Theorem 4.7 in Cartis, Gould, and Toint (2012a). In order to prove the second statement, we may follow the line of thought of Lemma 5.2 in Cartis, Gould, and Toint (2011a): we note that, using (2.4), AS2f, and AS3f,

$$\begin{aligned} f(x_k^+) - m_k(x_k^+) &\leq \tfrac{1}{2}\|H(\xi_k) - H(x_k)\|\,\|s\|^2 + \tfrac{1}{2}\|(H(x_k) - B_k)s_k\|\,\|s_k\| - \tfrac{1}{3}\sigma_k\|s_k\|^3 \\ &\leq [\tfrac{1}{2}(\kappa_L + \kappa_B) - \tfrac{1}{3}\sigma_k]\,\|s_k\|^3, \end{aligned}$$

where ξ_k belongs to the segment $[x_k, x_k^+]$ and thus $\|\xi_k - x_k\| \leq \|s_k\|$. This relation shows that $f(x_k^+) \leq m_k(x_l^+)$, and hence that $k \in \mathcal{S}$, provided $\sigma_k \geq \kappa_{\rm succ} = \frac{3}{2}(\kappa_L + \kappa_B)$. The final statement of the lemma follows from Lemma 4.3 in Cartis, Gould, and Toint (2012a). $\square$

2.2. Applying COCARC-S to constrained nonlinear least-squares. We consider now the smooth constrained nonlinear least-squares problem given by

$$\min_{x\in\mathcal{F}} \phi(x) \stackrel{\text{def}}{=} \tfrac{1}{2}\|r(x)\|^2, \tag{2.12}$$

where r is a twice-continuously differentiable "residual" function from $\mathbb{R}^n$ into $\mathbb{R}^m$, $\|\cdot\|$ is the standard Euclidean norm, and $\mathcal{F}$ is a nonempty closed convex set. Our objective is to apply the cubically regularized COCARC-S method described in section 2.1 for general nonconvex optimization subject to convex constraints to this problem and to specialize the corresponding complexity results.

In the spirit of Cartis, Gould, and Toint (2013), this involves redefining a suitable termination criterion that exploits the particularity of the least-squares problem. In the latter paper, we have indeed argued that, for the unconstrained case ($\mathcal{F} = \mathbb{R}^n$), it is advisable to replace the standard rule where the algorithm is terminated as soon as an iterate x_k is found such that

$$\|\nabla_x \phi(x_k)\| = \|J(x_k)^T r(x_k)\| \le \epsilon \tag{2.13}$$

(where $J(x)$ denotes the Jacobian of r at x and $\epsilon \in (0,1)$ is a user-defined accuracy threshold) by the rule that iterations are instead terminated as soon as an iterate x_k is found such that, for some user-defined accuracy thresholds $\epsilon_p \in (0,1)$ and $\epsilon_d \in (0,1)$,

$$\|r(x_k)\| \le \epsilon_p \quad \text{or} \quad \|\,\nabla_x \|r(x_k)\|\,\| \le \epsilon_d, \tag{2.14}$$

where

$$\nabla_x \|r(x)\| \stackrel{\text{def}}{=} \begin{cases} \dfrac{J(x)^T r(x)}{\|r(x)\|} & \text{when} \quad r(x) \neq 0, \\ 0 & \text{otherwise.} \end{cases} \tag{2.15}$$

The case corresponding to the first condition in (2.14) is the situation where an ϵ_p-approximate optimal point is found with "zero" residual, while the second corresponds to the case where the residual at the approximate first-order critical point is nonzero.

If we now consider the constrained case ($\mathcal{F} \subset \mathbb{R}^n$), making the gradient of the objective function small, as requested in (2.13), is no longer appropriate, because the solution might lie on the boundary of $\mathcal{F}$. In the spirit of (2.2), an alternative to (2.13) in the constrained case is to stop the COCARC-S algorithm when applied to (2.12) as soon as a point x_k is found such that

$$\chi_\phi(x_k) \le \epsilon, \tag{2.16}$$

where

$$\chi_\phi(x) = \left| \min_{x+d\in\mathcal{F}, \|d\|\le 1} \langle J(x)^T r(x), d\rangle \right| \tag{2.17}$$

(see (2.3)). However, since we are considering a least-squares problem, we may apply the same reasoning as in the unconstrained case, and we therefore suggest terminating the COCARC-S algorithm as soon as an x_k is found in $\mathcal{F}$ such that

$$\|r(x_k)\| \le \epsilon_p \quad \text{or} \quad \chi_{\|r\|}(x_k) = \left| \min_{x_k+d\in\mathcal{F}, \|d\|\le 1} \langle \nabla_x \|r(x_k)\|, d\rangle \right| \le \epsilon_d. \tag{2.18}$$

Note that $\chi_{\|r\|}(x)$ is continuous as a function of x for $r(x) \neq 0$ and is zero if and only if x is a first-order critical point of problem (2.12). Condition (2.18) therefore replaces (2.14) in the constrained case.

In order to adapt this framework of section 2.1 to the constrained nonlinear least-squares case (2.12), we first note that AS1, AS4, and AS5 need not to be modified.[3] AS2f and AS3f must, however, be reformulated in terms of the residual function r.

AS2 Each r_i $(i = 1, \ldots, m)$ is twice continuously differentiable and uniformly Lipschitz continuous on $\mathcal{X}$, and its Hessian is "weakly" uniformly Lipschitz continuous on the segments $[x_k, x_k^+]$ in the sense that there exists a constant $\kappa_{\rm L} \geq 0$ such that, for all k, all $y \in [x_k, x_k^+]$, and all $i = 1, \ldots, m$,

$$\|[\nabla_{xx} r_i(y) - \nabla_{xx} r_i(x_k)] s_k\| \leq \kappa_{\rm L} \|s_k\|^2.$$

Moreover, the Jacobian $J(x)$ is Lipschitz continuous on $\mathcal{X}$ in the sense that there exists a constant $\kappa_{\rm J} \geq 0$ such that, for all $x, y \in \mathcal{X}$,

$$\|J(x) - J(y)\| \leq \kappa_{\rm J} \|x - y\|.$$

AS3 The Hessian $\nabla_{xx}\phi(x_k)$ is well approximated by B_k in the sense that there exists a constant $\kappa_{\rm B} > 0$ such that, for all k,

$$\|[B_k - \nabla_{xx}\phi(x_k)] s_k\| \leq \kappa_{\rm B} \|s_k\|^2.$$

We refer the reader to the technical discussion in Cartis, Gould, and Toint (2013) showing that these assumptions ensure AS2f with ϕ playing the role of f. Observe that, strictly speaking, the fact that the residuals $r_i(x)$ are twice continuously differentiable (AS.2) on a bounded set $\mathcal{X}$ (AS.5) is enough to ensure that r_i and $\nabla_x r_i$ are Lipschitz continuous on $\mathcal{X}$, but we prefer to require these properties explicitly for clarity.

Having reformulated our assumptions, we are now entitled to deduce that (2.11) (with $f = \phi$) holds as AS1–AS5 are satisfied. The next step is then to modify this lower bound in the spirit of Lemma 3.1 in Cartis, Gould, and Toint (2013). Assume first that, for a given $\beta \in (0, 1)$, $\|r(x_{k+1})\| \leq \beta \|r(x_k)\|$. Then $k \in \mathcal{S}$ and

$$\|r(x_k)\| - \|r(x_{k+1})\| \geq (1 - \beta)\|r(x_k)\| \tag{2.19}$$

and

$$\|r(x_k)\|^{\frac{1}{2}} - \|r(x_{k+1})\|^{\frac{1}{2}} \geq (1 - \sqrt{\beta})\|r(x_k)\|^{\frac{1}{2}} \geq \frac{(1 - \sqrt{\beta})}{\sqrt{\beta}} \|r(x_{k+1})\|^{\frac{1}{2}}. \tag{2.20}$$

If, on the other hand, $\|r(x_{k+1})\| > \beta \|r(x_k)\|$, we nevertheless know that, for $k \in \mathcal{S}$, $\|r(x_{k+1})\| = \|r(x_k^+)\| < \|r(x_k)\|$, and thus that

$$\begin{aligned} \|r(x_k)\| - \|r(x_{k+1})\| &= \frac{\|r(x_k)\|^2 - \|r(x_{k+1})\|^2}{\|r(x_k)\| + \|r(x_{k+1})\|} \\ &\geq \frac{\phi(x_k) - \phi(x_{k+1})}{\|r(x_k)\|} \\ &\geq \eta_1 \kappa_{\rm C} \left(\frac{\|r(x_{k+1})\|}{\|r(x_k)\|} \right)^{\frac{3}{2}} \|r(x_k)\|^{\frac{1}{2}} \left(\frac{\chi_\phi(x_{k+1})}{\|r(x_{k+1})\|} \right)^{\frac{3}{2}} \\ &> \eta_1 \kappa_{\rm C} \beta^{\frac{3}{2}} \|r(x_k)\|^{\frac{1}{2}} \left(\chi_{\|r\|}(x_{k+1}) \right)^{\frac{3}{2}}, \end{aligned} \tag{2.21}$$

[3]Except for the obvious change from ϵ to ϵ_p and ϵ_d in AS4.

where we have used (2.11) to obtain the penultimate inequality, and (2.15), (2.3), and the inequality $\|r(x_{k+1})\| > \beta\|r(x_k)\|$ to obtain the last. Hence, using this latter inequality again, we have that

$$\sqrt{\|r(x_k)\|} - \sqrt{\|r(x_{k+1})\|} = \frac{\|r(x_k)\| - \|r(x_{k+1})\|}{\sqrt{\|r(x_k)\|} + \sqrt{\|r(x_{k+1})\|}} \geq \tfrac{1}{2}\eta_1\kappa_{\rm C}\beta^{\frac{3}{2}} \left(\chi_{\|r\|}(x_{k+1})\right)^{\frac{3}{2}}. \tag{2.22}$$

As a consequence, we conclude from (2.19) and (2.21) that

$$\|r(x_k)\|-\|r(x_{k+1})\| \geq \min[\eta_1\kappa_{\rm C}\beta^{\frac{3}{2}}, (1-\beta)]\cdot\min[\|r(x_k)\|^{\frac{1}{2}} \left(\chi_{\|r\|}(x_{k+1})\right)^{\frac{3}{2}}, \|r(x_k)\|]. \tag{2.23}$$

Similarly, we deduce from (2.20) and (2.22) that

$$\|r(x_k)\|^{\frac{1}{2}} - \|r(x_{k+1})\|^{\frac{1}{2}} \geq \kappa_\phi \min\left[\,\left(\chi_{\|r\|}(x_{k+1})\right)^{\frac{3}{2}},\ \|r(x_{k+1})\|^{\frac{1}{2}}\,\right],$$

where $\kappa_\phi \stackrel{\rm def}{=} \min[\frac{1}{2}\eta_1\kappa_{\rm C}\beta^{\frac{3}{2}}, \beta^{-\frac{1}{2}} - 1]$. Thus, as long as the COCARC-S algorithm applied to problem (2.12) is not terminated, i.e., as long as (2.18) is violated, we have that, for $k \in \mathcal{S}$,

$$\|r(x_k)\|^{\frac{1}{2}} - \|r(x_{k+1})\|^{\frac{1}{2}} \geq \kappa_\phi \min[\,\epsilon_d^{\frac{3}{2}},\, \epsilon_p^{\frac{1}{2}}\,]. \tag{2.24}$$

Because, obviously, $0 \leq \|r(x_k)\|^{\frac{1}{2}} \leq \|r(x_0)\|^{\frac{1}{2}}$ for all k, we deduce that, provided AS1–AS5 hold, there are at most

$$\left\lfloor \frac{\|r(x_0)\|^{\frac{1}{2}}}{\kappa_\phi \min[\,\epsilon_d^{\frac{3}{2}},\, \epsilon_p^{\frac{1}{2}}\,]} \right\rfloor \stackrel{\rm def}{=} N_{\mathcal{S}}(\epsilon_p, \epsilon_d)$$

successful iterations until the COCARC-S algorithm applied to the constrained nonlinear least-squares problem (2.12) finds an iterate x_k such that (2.18) holds at x_{k+1}. Theorem 2.1 (equation (2.14)) in Cartis, Gould, and Toint (2011a) gives us that the number of unsuccessful iterations that occur up to an(y) iteration k is bounded above by

$$\left\lfloor [1 + N_{\mathcal{S}}(\epsilon_p, \epsilon_d)]\, \frac{1}{\log\gamma_1} \log\frac{\sigma_{\max}}{\sigma_{\min}} \right\rfloor,$$

where γ_1 is the algorithm parameter defined in (2.9), and $\sigma_{\max}$ is the uniform upper bound on the regularization weight σ_k derived in Lemma 2.1. Thus we conclude that, for $\epsilon_p, \epsilon_d \in (0,1)$, the total number of (successful and unsuccessful) iterations required by Algorithm COCARC-S to find x_k is bounded above by

$$\left\lfloor \frac{\kappa_{\rm CNLS}}{\min[\,\epsilon_d^{\frac{3}{2}},\, \epsilon_p^{\frac{1}{2}}\,]} \right\rfloor \tag{2.25}$$

with

$$\kappa_{\rm CNLS} \stackrel{\rm def}{=} \frac{\|r(x_0)\|^{\frac{1}{2}}}{\kappa_\phi} + \left(1 + \frac{\|r(x_0)\|^{\frac{1}{2}}}{\kappa_\phi}\right) \frac{\log(\sigma_{\max}/\sigma_{\min})}{\log\gamma_1}. \tag{2.26}$$

We summarize our findings in the form of the following theorem.

THEOREM 2.2. *Assume that AS1–AS5 hold. Consider $\epsilon_p, \epsilon_d \in (0,1)$. Then there is a constant $\kappa_{\text{CNLS}} > 0$ whose expression is given in* (2.26) *such that the COCARC-S algorithm applied to problem* (2.12) *requires at most*

$$\left\lfloor \kappa_{\text{CNLS}} \max[\,\epsilon_d^{-\frac{3}{2}},\, \epsilon_p^{-\frac{1}{2}}\,] \right\rfloor \tag{2.27}$$

iterations (and evaluations of r and possibly its derivatives) to find an iterate x_k such that

$$\|r(x_k)\| \le \epsilon_p \quad \text{or} \quad \chi_{\|r\|}(x_k) \le \epsilon_d. \tag{2.28}$$

Note that this bound is identical (in order) to that obtained by Cartis, Gould, and Toint (2013, Theorem 3.2) for the *unconstrained* nonlinear least-squares problem ($\mathcal{F} = \mathbb{R}^n$). Moreover, provided $\epsilon_p \ge \epsilon_d^3$, then the number of COCARC-S iterations is bounded above by $O(\epsilon_d^{-3/2})$, which is the same complexity as that of solving the general *unconstrained* nonlinear optimization problem with the ARC cubic regularization algorithm (see Nesterov and Polyak (2006) and Cartis, Gould, and Toint (2012b)).

3. The general nonlinear optimization problem. Having considered the constrained nonlinear least-squares case, we now turn to the general nonlinear optimization problem (1.1).

Both phases of our proposed analysis critically depend on applying the COCARC-S algorithm, first to the squared norm of the constraint violation

$$\theta(x) \stackrel{\text{def}}{=} \tfrac{1}{2}\|c(x)\|^2 \quad \text{for} \quad x \in \mathcal{F}, \tag{3.1}$$

terminating the computation as soon as a point $x_1 \in \mathcal{F}$ is found such that, for some user-defined accuracy thresholds $\epsilon_p \in (0,1)$ and $\epsilon_d \in (0,1)$,

$$\|c(x_1)\| \le \delta\epsilon_p \quad \text{or} \quad \chi_{\|c\|}(x_1) \le \epsilon_d \tag{3.2}$$

(for some $\delta \in (0,1)$), and subsequently to a sequence of suitably defined least-squares problems whose objective function is denoted by

$$\mu(x, t_k) \stackrel{\text{def}}{=} \tfrac{1}{2}\|r(x,t_k)\|^2 \stackrel{\text{def}}{=} \tfrac{1}{2}\left\| \begin{pmatrix} c(x) \\ f(x) - t_k \end{pmatrix} \right\|^2 \quad \text{for} \quad x \in \mathcal{F} \tag{3.3}$$

for some monotonically decreasing sequence of "targets" t_k ($k = 1, \ldots$).

We now describe our two-phase algorithm as Algorithm 3.1, where we use the symbol $\mathcal{P}$ to denote the projection onto the set $\mathcal{F}$. Note that the iterations in step 2(a) of Phase 2 correspond to applying the COCARC-S algorithm (ignoring the termination test in step 0) for each new value of the target t_k until the first successful iteration occurs. Also observe that $\chi_{\|r(\cdot,t)\|}(x_{k+1}, t_k)$ (in step 2(b)) is well-defined when $\|r(x_{k+1}, t_k)\| > \delta\epsilon_p > 0$.

Analyzing the worst-case behavior of this algorithm once more requires specifying the necessary assumptions. As above, we denote by $\mathcal{X} \subseteq \mathcal{F}$ the closed convex hull of all Phase 1 and Phase 2 iterates and trials points, and by $\mathcal{X}_2 \subseteq \mathcal{X}$ that of all Phase 2 iterates and trials points.

AS6 The function c is twice continuously differentiable on an open neighborhood of $\mathcal{X}$, and f is twice continuously differentiable in an open neighborhood of $\mathcal{X}_2$.

Algorithm 3.1: Short-Step ARC Algorithm for (1.1)

A starting point x_0, a minimum regularization parameter $\sigma_{\min} > 0$, an initial regularization parameter $\sigma_0 \geq \sigma_{\min}$, a parameter $\delta \in (0, 1)$, as well as the tolerances $\epsilon_p \in (0, 1)$ and $\epsilon_d \in (0, 1)$, are given.

Phase 1:

Starting from $\mathcal{P}(x_0)$, apply the COCARC-S algorithm to minimize $\theta(x)$ subject to $x \in \mathcal{F}$ until a point $x_1 \in \mathcal{F}$ is found at which (3.2) holds. If $\|c(x_1)\| > \delta\epsilon_p$, terminate.

Phase 2:

1. Set $t_1 = f(x_1) - \sqrt{\epsilon_p^2 - \|c(x_1)\|^2}$ and $k = 1$.
2. For $k = 1, 2, \ldots$, do:
   (a) Loop on steps 1 to 3 of the COCARC-S algorithm to minimize $\mu(x, t_k)$ as a function of $x \in \mathcal{F}$ until a successful iteration is obtained, yielding a new iterate $x_{k+1} \in \mathcal{F}$ and a new value of the regularization parameter $\sigma_{k+1} \geq \sigma_{\min}$.
   (b) If $\|r(x_{k+1}, t_k)\| > \delta\epsilon_p$, then terminate if $\chi_{\|r(\cdot,t)\|}(x_{k+1}, t_k) \leq \epsilon_d$. Otherwise (that is, if either $\|r(x_{k+1}, t_k)\| \leq \delta\epsilon_p$ or $\chi_{\|r(\cdot,t)\|}(x_{k+1}, t_k) > \epsilon_d$ and $\|r(x_{k+1}, t_k)\| > \delta\epsilon_p$), set

$$t_{k+1} = f(x_{k+1}) - \sqrt{\|r(x_k, t_k)\|^2 - \|r(x_{k+1}, t_k)\|^2 + (f(x_{k+1}) - t_k)^2}. \tag{3.4}$$

AS7 The components c_i $(i = 1, \ldots, m)$ and the Jacobian $J(x)$ are globally Lipschitz continuous in $\mathcal{X}$ with Lipschitz constants $L_{c_i} > 0$ and $L_J > 0$, respectively. The components $\nabla^2 c_i(x)$ are weakly Lipschitz continuous on the segments $[x_k, x_k^+]$ (for both phases) with Lipschitz constant L_{H,c_i} for $i \in \{1, \ldots, m\}$.

AS8 $f(x)$ and $g(x)$ are Lipschitz continuous in $\mathcal{X}_2$ with Lipschitz constants L_f and $L_{g,f} > 0$, respectively. Moreover, $\nabla^2 f(x)$ is weakly Lipschitz continuous on all Phase 2 segments $[x_k, x_k^+]$ with Lipschitz constant $L_{H,f}$.

AS.9 The objective $f(x)$ is bounded above and below in a neighborhood of the feasible set; that is, there exist constants $\alpha > 0$, f_{low}, and $f_{\text{up}} \geq f_{\text{low}} + 1$ such that

$$f_{\text{low}} \leq f(x) \leq f_{\text{up}} \qquad \text{for all } x \in \mathcal{F} \cap \mathcal{C}_\alpha,$$

where

$$\mathcal{C}_\alpha = \{x \in \mathbb{R}^n \mid \|c(x)\| \leq \alpha\}. \tag{3.5}$$

AS10 ℓ_k is bounded above by a constant independent of k and ϵ in all constrained cubic model minimizations (in both phases of the algorithm).

AS11 $\mathcal{X}$ is bounded.

From here on, our analysis is nearly identical to that presented in Cartis, Gould, and Toint (2013). We know from the previous section that Phase 1 of Algorithm 3.1 will terminate in a number of iterations (and function evaluations) as given by (2.27).

Let us now consider Phase 2, and exploit again the least-squares structure of the minimizations carried on in step 2. We obtain the following properties.

LEMMA 3.1. *Suppose that $\epsilon_p \le \alpha$. Then in every Phase 2 iteration $k \ge 1$ of Algorithm 3.1 we have that*

$$t_k \ge t_{k+1}, \tag{3.6}$$

$$f(x_k) - t_k \ge 0, \tag{3.7}$$

$$\|r(x_k, t_k)\| = \epsilon_p, \tag{3.8}$$

$$\|c(x_k)\| \le \epsilon_p \;\; and \;\; |f(x_k) - t_k| \le \epsilon_p, \tag{3.9}$$

and so $x_k \in \mathcal{F} \cap \mathcal{C}_\alpha$. In addition, if AS6–AS8 hold, each Phase 2 iteration requires at most $\kappa_{\mathcal{S}}$ evaluations of problem functions, where $\kappa_{\mathcal{S}} \ge 1$ is a constant independent of k.

Proof. Note that the monotonicity property of the COCARC-S iterates in Phase 2 of Algorithm 3.1 provides

$$\|r(x_k, t_k)\| \ge \|r(x_{k+1}, t_k)\| \text{ for all } k \ge 1, \tag{3.10}$$

and so the updating procedure for t_k in (3.4) is well-defined and gives

$$t_k - t_{k+1} = -(f(x_{k+1}) - t_k) + \sqrt{\|r(x_k, t_k)\|^2 - \|r(x_{k+1}, t_k)\|^2 + (f(x_{k+1}) - t_k)^2} \tag{3.11}$$

for any successful $k \ge 1$. Due to (3.11), (3.6) follows immediately in the case when $f(x_{k+1}) \le t_k$. Otherwise, when $f(x_{k+1}) > t_k$, conjugacy properties and (3.11) give

$$t_k - t_{k+1} = \frac{\|r(x_k, t_k)\|^2 - \|r(x_{k+1}, t_k)\|^2}{f(x_{k+1}) - t_k + \sqrt{\|r(x_k, t_k)\|^2 - \|r(x_{k+1}, t_k)\|^2 + (f(x_{k+1}) - t_k)^2}} \ge 0,$$

where in the last inequality we also used (3.10).

Note that (3.7) holds at $k = 1$ due to the particular choice of t_1 and at $k > 1$ due to (3.4) and (3.10). Also, (3.9) follows straightforwardly from (3.8), which also provides that $x_k \in \mathcal{C}_1$ due to (3.5). It remains to prove (3.8), by induction on k. Again, the particular choice of t_1 gives (3.8) at $k = 1$. Assume now that (3.8) holds at $k > 1$, namely,

$$\|r(x_k, t_k)\| = \epsilon_p. \tag{3.12}$$

If k is an unsuccessful iteration, then $x_{k+1} = x_k$ and $t_{k+1} = t_k$, and so (3.8) is satisfied at $k + 1$. Otherwise, we have

$$\begin{aligned}(f(x_{k+1}) - t_{k+1})^2 &= \|r(x_k, t_k)\|^2 - \|r(x_{k+1}, t_k)\|^2 + (f(x_{k+1}) - t_k)^2 \\ &= \|r(x_k, t_k)\|^2 - \|c(x_{k+1})\|^2,\end{aligned}$$

where (3.7) and (3.4) give the first identity, while the second equality follows from (3.3). Thus we deduce, also using (3.3), that

$$\|r(x_{k+1}, t_{k+1})\|^2 = \|r(x_k, t_k)\|^2,$$

which concludes our induction step due to (3.12).

In each Phase 2 iteration, steps 1 to 3 of the COCARC-S Algorithm 2.1 are iterated upon until an (inner) successful iteration is obtained in step 2. Because the regularization parameter is increased at least by a factor $\gamma_1 > 1$ at unsuccessful (inner) iterations, because $\sigma_k \geq \sigma_{\min}$, and because of Lemma 2.1, we know that at most

$$\kappa_{\mathcal{S}} \stackrel{\text{def}}{=} \max\left[\frac{\log(\sigma_{\text{succ}}/\sigma_{\min})}{\log \gamma_1}, 1\right]$$

inner iterations of this algorithm are necessary to compute x_{k+1}. Since each of these iterations involves a single problem function evaluation (in step 2), we obtain the desired result. $\square$

We may then pursue our analysis exactly as in Cartis, Gould, and Toint (2013) and deduce the following important result on the decrease in t_k, directly inspired by Lemma 5.3 in that reference.

LEMMA 3.2. *Suppose that AS*3 *(with $\phi = r(\cdot, t_k)$) and AS*6*–AS*11 *hold and that*

$$\epsilon_p \leq \alpha \quad \textit{and} \quad \epsilon_d \leq \epsilon_p^{1/3}. \tag{3.13}$$

Then, for every Phase 2 *iteration k of Algorithm* 3.1*, we have that*

$$t_k - t_{k+1} \geq \kappa_t \epsilon_d^{3/2} \epsilon_p^{1/2} \tag{3.14}$$

for some constant $\kappa_t \in (0, 1)$ independent of k, ϵ_d, and ϵ_p.

Proof. We consider two cases. The first case is when $\|r(x_{k+1}, t_k)\| \leq \delta\epsilon_p$. Then we have, from (3.4), (3.3), and (3.8), that

$$\begin{aligned} t_k - t_{k+1} &= -(f(x_{k+1}) - t_k) + \sqrt{\|r(x_k, t_k)\|^2 - \|c(x_{k+1})\|^2} \\ &= -(f(x_{k+1}) - t_k) + \sqrt{\epsilon_p^2 - \|c(x_{k+1})\|^2}. \end{aligned} \tag{3.15}$$

It also follows from (3.3) that

$$(f(x_{k+1}) - t_k)^2 + \|c(x_{k+1})\|^2 = \|r(x_{k+1}, t_k)\|^2 \leq \delta^2\epsilon_p^2. \tag{3.16}$$

We can regard the right-hand side of the second equality in (3.15) as well as (3.16) as functions of unknowns $f \stackrel{\text{def}}{=} f(x_{k+1}) - t_k$ and $c \stackrel{\text{def}}{=} c(x_{k+1})$ and look for the solution of the following optimization problem in (f, c):

$$\min_{(f,c)\in\mathbb{R}^2} F(f, c) \stackrel{\text{def}}{=} -f + \sqrt{\epsilon_p^2 - c^2} \text{ subject to } f^2 + c^2 \leq \delta^2\epsilon_p^2. \tag{3.17}$$

It is easy to show that the global minimum of (3.17) is attained at $(f_*, c_*) = (\delta\epsilon_p, 0)$, and it is given by $F(f_*, c_*) = -\delta\epsilon_p + \epsilon_p$ (see Lemma 5.2 and its proof in Cartis, Gould, and Toint (2013)). This implies, due to (3.15), that

$$t_k - t_{k+1} \geq (1 - \delta)\epsilon_p \geq (1 - \delta)\epsilon_d^{\frac{3}{2}} \epsilon_p^{\frac{1}{2}}, \tag{3.18}$$

where we have used the second part of (3.13) to deduce the last inequality.

The second case is when $\|r(x_{k+1}, t_k)\| > \delta\epsilon_d$ and $\chi_{\|r(\cdot,t)\|}(x_{k+1}, t_k) > \epsilon_d$. Then, from (2.23) and the second part of (3.13), we have that

$$t_k - t_{k+1} \geq \kappa_1 \epsilon_d^{3/2} \epsilon_p^{1/2}$$

for some $\kappa_1 \in (0, 1)$. Combining this last bound with (3.18) then gives (3.14) with $\kappa_t = \min[\kappa_1, (1 - \delta)]$. $\square$

It is then easy to combine the complexity analysis we already mentioned for Phase 1 of Algorithm 3.1 with the second part of (3.9) and AS9.

LEMMA 3.3. *Suppose that AS6–AS11 and (3.13) hold. Then Algorithm 3.1 generates an iterate $x_j \in \mathcal{F}$ such that either*

$$\chi_{\|c\|}(x_j) \leq \epsilon_d \ \ \textit{and} \ \ \|c(x_j)\| > \delta\epsilon_p \tag{3.19}$$

or

$$\|r(x_j, t_{j-1})\| > \delta\epsilon_p, \quad \chi_{\|r(\cdot,t)\|}(x_j, t_{j-1}) \leq \epsilon_d, \ \ \textit{and} \ \ \|c(x_j)\| \leq \epsilon_p \tag{3.20}$$

in at most

$$\left\lfloor \kappa_{\text{NLO}}\, \epsilon_p^{-\frac{1}{2}} \epsilon_d^{-\frac{3}{2}} \right\rfloor \tag{3.21}$$

evaluations of f and c (and their derivatives), where $\kappa_{\text{NLO}} > 0$ is a problem-dependent constant independent of ϵ_p, ϵ_d, and x_0.

Proof. We have already discussed above (see Theorem 2.2) the fact that Phase 1 of Algorithm 3.1 will terminate in at most

$$\left\lfloor \kappa_{\text{CNLS}} \max[(\delta\epsilon_p)^{-\frac{1}{2}} \epsilon_d^{-\frac{3}{2}}, (\delta\epsilon_p)^{-1}] \right\rfloor \leq \left\lfloor \kappa_{\text{CNLS}} \delta^{-1} \max[\epsilon_p^{-\frac{1}{2}} \epsilon_d^{-\frac{3}{2}}, \epsilon_p^{-1}] \right\rfloor \leq \left\lfloor \kappa_{\text{CNLS}} \delta^{-1}\, \epsilon_p^{-\frac{1}{2}} \epsilon_d^{-\frac{3}{2}} \right\rfloor \tag{3.22}$$

iterations (and problem evaluations), where we have taken into account our change of primal accuracy from ϵ_p to $\delta\epsilon_p$ specified in (3.2) and (3.13). If the algorithm terminates at this stage, then (3.19) must hold, as required. Assume now that Phase 2 of Algorithm 3.1 is entered. We start by observing that AS9 implies that, for every k,

$$f_{\text{low}} \leq f(x_k) \leq t_k + \epsilon_p \leq t_1 - k\, \kappa_t \epsilon_d^{\frac{3}{2}} \epsilon_p^{\frac{1}{2}} + \epsilon_p \leq f(x_1) - k\, \kappa_t \epsilon_d^{\frac{3}{2}} \epsilon_p^{\frac{1}{2}} + \epsilon_p,$$

where we have also used (3.14) and the definition of t_1 in Algorithm 3.1. Hence, we obtain from the inequality $f(x_1) \leq f_{\text{up}}$ (itself implied by AS9 again), the second part of (3.9), and $\epsilon_p \in (0, 1)$ that at most

$$\frac{f_{\text{up}} - f_{\text{low}} + 1}{\kappa_t \epsilon_d^{\frac{3}{2}} \epsilon_p^{\frac{1}{2}}}$$

Phase 2 iterations may occur before $\chi_{\|r(\cdot,t)\|}(x_{k+1}, t_k) \leq \epsilon_d$ with $\|r(x_{k+1}, t_k)\| > \delta\epsilon_p$. Since the last part of Lemma 3.1 states that at most $\kappa_{\mathcal{S}}$ evaluations of f and c (and their derivatives) occur for each such iteration, we therefore deduce that at most

$$\kappa_{\text{phase2}} \stackrel{\text{def}}{=} \left\lfloor \frac{\kappa_{\mathcal{S}}(f_{\text{up}} - f_{\text{low}} + 1)}{\kappa_t \epsilon_d^{\frac{3}{2}} \epsilon_p^{\frac{1}{2}}} \right\rfloor$$

Phase 2 iterations are needed to satisfy the termination test in step 2(b) of Algorithm 3.1. Combining this result with (3.22) and the first part of (3.9), we obtain the desired conclusion with $\kappa_{\text{NLO}} = \max[\kappa_{\text{CNLS}}\delta^{-1}, \kappa_{\text{phase2}}]$. $\square$

To complete our analysis, we now comment on the meaning of the termination test (3.20); we already discussed (3.19) in section 2.2. This meaning is best expressed by using

$$\ell(x,y) \stackrel{\text{def}}{=} f(x) + \langle y, c(x)\rangle, \tag{3.23}$$

the Lagrangian of the original problem where only equality constraints are kept, that is,

$$\min_{x\in\mathbb{R}^n} f(x) \quad \text{such that} \quad c(x) = 0. \tag{3.24}$$

Lemma 3.4. *Assume that, at iteration k in Phase 2 of Algorithm 3.1,*

$$\|r(x_{k+1},t_k)\| > \delta\epsilon_p \quad \textit{and} \quad \chi_{\|r(\cdot,t)\|}(x_{k+1},t_k) \le \epsilon_d. \tag{3.25}$$

Then either, for some vector $y_{k+1} \in \mathbb{R}^m$,

$$\chi_\ell(x_{k+1},y_{k+1}) \le \epsilon_d\|(y_{k+1},1)\| \quad \textit{and} \quad \|c(x_{k+1})\| \le \epsilon_p, \tag{3.26}$$

where $\ell(x,y)$ is the Lagrangian of the equality constrained problem given by (3.23)*, or*

$$\chi_{\|c\|}(x_{k+1}) \le \epsilon_d \quad \textit{and} \quad \|c(x_{k+1})\| \in (\delta\epsilon_p, \epsilon_p). \tag{3.27}$$

Proof. Assume that $f(x_{k+1}) \ne t_k$. Then the second part of (3.25) can be rewritten as

$$\begin{aligned}\epsilon_d &\ge \left|\min_{x+d\in\mathcal{F},\|d\|\le 1}\left\langle \frac{J(x_{k+1})^T c(x_{k+1}) + (f(x_{k+1})-t_k)g_{k+1}}{\|r(x_{k+1},t_k)\|}, d\right\rangle\right| \\ &= \left|\min_{x+d\in\mathcal{F},\|d\|\le 1}\left\langle \frac{J(x_{k+1})^T y_{k+1} + g_{k+1}}{\|(y_{k+1},1)\|}, d\right\rangle\right|,\end{aligned}$$

where $y_{k+1} = c(x_{k+1})/(f(x_{k+1}) - t_k)$. Thus, given the definition of $\ell(x,y)$ in (3.23) and that of χ in (2.3), we obtain that the first part of (3.26) holds. The second part of this statement results from the first inequality in (3.9). Suppose now that $f(x_{k+1}) = t_k$. The second part of (3.27) is easily deduced from the observation that, in this case,

$$\|c(x_{k+1})\| = \|r(x_{k+1},t_k)\| \in (\delta\epsilon_p, \|r(x_k,t_k)\|) = (\delta\epsilon_p, \epsilon_p), \tag{3.28}$$

where we successively used (3.3), the first part of (3.25), the monotonic nature of the COCARC-S algorithm, and (3.8). The first part of (3.27) then follows directly from (3.25) and the relation

$$\chi_{\|r(\cdot,t)\|}(x_{k+1},t_k) = \left|\min_{x+d\in\mathcal{F},\|d\|\le 1}\left\langle \frac{J(x_{k+1})^T c(x_{k+1})}{\|r(x_{k+1},t_k)\|}, d\right\rangle\right| = \chi_{\|c\|}(x_{k+1}),$$

where we used (3.28) to deduce the second equality. $\square$

Condition (3.26) expresses the approximate first-order criticality of (x_{k+1}, y_{k+1}) by assessing that the maximum feasible linearized decrease in the Lagrangian corresponding to the problem (3.24) only involving equalities is small compared to the size of the multiplier. The use of a scaled measure of criticality of this type was already argued in Cartis, Gould, and Toint (2013).

Combining Lemmas 3.3 and 3.4 then gives our final complexity result.

THEOREM 3.5. *Suppose that AS*3 *(with* $\phi = r(\cdot, t_k)$*), AS*6*–AS*11*, and* (3.13) *hold. Then Algorithm* 3.1 *generates an iterate* $x_k \in \mathcal{F}$ *such that either*

$$\chi_\ell(x_k, y_k) \le \epsilon_d \|(y_k, 1)\| \quad \textit{and} \quad \|c(x_k)\| \le \delta\epsilon_p, \tag{3.29}$$

where $\ell(x, y)$ *is the Lagrangian of the equality constrained problem* (3.23) *and* $y_k \in \mathbb{R}^m$ *is an approximate Lagrange multiplier for problem* (3.24)*, or*

$$\chi_{\|c\|}(x_k) \le \epsilon_p \quad \textit{and} \quad \|c(x_k)\| > \delta\epsilon_p \tag{3.30}$$

in at most

$$\left\lfloor \kappa_{\rm NLO}\, \epsilon_p^{-\frac{1}{2}}\, \epsilon_d^{-\frac{3}{2}} \right\rfloor \tag{3.31}$$

evaluations of f *and* c *(and their derivatives), where* $\kappa_{\rm NLO} > 0$ *is a problem-dependent constant independent of* ϵ_p*,* ϵ_d*, and* x_0*.*

Again, we note that if $\epsilon_d = \epsilon_p^{2/3}$, Lemma 3.3 implies an overall complexity bound of $O(\epsilon_p^{-3/2})$ iterations and problem evaluations for applying Algorithm 3.1 to the general nonlinear optimization problem (1.1).

4. Conclusions and perspectives. We have examined the worst-case complexity of finding approximate first-order critical points for the nonlinear least-squares problem with convex inequality constraints and the general nonlinear optimization problem (involving both nonconvex equality and inequality constraints). We have shown that, under acceptable assumptions, both of these problems can be approximately solved using a second-order method of cubic regularization type in a number of problem (objective, constraints, and derivatives) evaluations that is at most $O(\epsilon_d^{-3/2})$ (for the constrained nonlinear least-squares) and $O(\epsilon_d^{-3/2}\epsilon_p^{-1/2})$ (for the general problem), where ϵ_d and ϵ_p are the dual and primal accuracy thresholds, respectively. The latter bound reduces to $O(\epsilon_p^{-3/2})$ problem evaluations if the dual threshold is chosen such that $\epsilon_d = \epsilon_p^{2/3}$. It is also known that this last bound is sharp (and hence can be attained on a somewhat contrived example) for methods using cubic regularization (see Cartis, Gould, and Toint (2010)) and optimal in a large class of second-order methods (see Cartis, Gould, and Toint (2011b)).

This result remains surprising because it shows that the inclusion of nonlinear equality and inequality constraints in the problem does not affect its worst-case evaluation analysis. Indeed, the worst-case complexity of the general nonlinear optimization problem is identical (in order) to that of the unconstrained case.

We end with a disclaimer that the algorithms proposed in this paper have been developed entirely for theoretical purposes, and any attempt to use them in practice would probably require substantial modifications. In particular, the short-step strategy of staying close to the constraints while attempting to reduce the objective by a small amount each time is unduly conservative and would be inefficient in practice. Being careful in the short-step strategy does help us avoid worst-case pitfalls, hence allowing us to obtain the improved/optimal evaluation complexity bound.

REFERENCES

C. Cartis, N. I. M. Gould, and Ph. L. Toint (2009), *An Adaptive Cubic Regularization Algorithm for Nonconvex Optimization with Convex Constraints and Its Function-Evaluation Complexity*, Technical Report, Optimization Online; http://www.optimization-online.org/DB_FILE/2009/02/2236.pdf.

C. Cartis, N. I. M. Gould, and Ph. L. Toint (2010), *On the complexity of steepest descent, Newton's and regularized Newton's methods for nonconvex unconstrained optimization*, SIAM J. Optim., 20, pp. 2833–2852.

C. Cartis, N. I. M. Gould, and Ph. L. Toint (2011a), *Adaptive cubic overestimation methods for unconstrained optimization. Part* II: *Worst-case function-evaluation complexity*, Math. Program. Ser. A, 130, pp. 295–319.

C. Cartis, N. I. M. Gould, and Ph. L. Toint (2011b), *Optimal Newton-type Methods for Nonconvex Optimization*, Technical Report, Optimization Online; http://www.optimization-online.org/DB_FILE/2011/06/3070.pdf.

C. Cartis, N. I. M. Gould, and Ph. L. Toint (2012a), *An adaptive cubic regularization algorithm for nonconvex optimization with convex constraints and its function-evaluation complexity*, IMA J. Numer. Anal., 32, pp. 1662–1645.

C. Cartis, N. I. M. Gould, and Ph. L. Toint (2012b), *Complexity bounds for second-order optimality in unconstrained optimization*, J. Complexity, 28, pp. 93–108.

C. Cartis, N. I. M. Gould, and Ph. L. Toint (2013), *On the evaluation complexity of cubic regularization methods for potentially rank-deficient nonlinear least-squares problems and its relevance to constrained nonlinear optimization*, SIAM J. Optim., 23, pp. 1553–1574.

C. Cartis, N. I. M. Gould, and Ph. L. Toint (2014), *On the complexity of finding first-order critical points in constrained nonlinear optimization*, Math. Program. Ser. A, 144, pp. 93–106.

A. R. Conn, N. I. M. Gould, A. Sartenaer, and Ph. L. Toint (1993), *Global convergence of a class of trust region algorithms for optimization using inexact projections on convex constraints*, SIAM J. Optim., 3, pp. 164–221.

S. Gratton, A. Sartenaer, and Ph. L. Toint (2008), *Recursive trust-region methods for multiscale nonlinear optimization*, SIAM J. Optim., 19, pp. 414–444.

Yu. Nesterov (2004), *Introductory Lectures on Convex Optimization*, Appl. Optim., Kluwer Academic Publishers, Dordrecht, The Netherlands.

Yu. Nesterov and B. T. Polyak (2006), *Cubic regularization of Newton method and its global performance*, Math. Program. Ser. A, 108, pp. 177–205.

S. A. Vavasis (1993), *Black-box complexity of local minimization*, SIAM J. Optim., 3, pp. 60–80.

Y. Yuan (1985), *Conditions for convergence of trust region algorithms for nonsmooth optimization*, Math. Program., 31, pp. 220–228.