




SECOND-ORDER OPTIMALITY AND BEYOND:
CHARACTERIZATION AND EVALUATION COMPLEXITY
IN NONCONVEX CONVEXLY-CONSTRAINED OPTIMIZATION

by C. Cartis, N. I. M. Gould and Ph. L. Toint

Report NAXYS-06-2016

15 August 2016



University of Namur, 61, rue de Bruxelles, B5000 Namur (Belgium)

<http://www.unamur.be/sciences/naxys>

Second-order optimality and beyond: characterization and evaluation complexity in convexly-constrained nonlinear optimization

C. Cartis* N. I. M. Gould[†] and Ph. L. Toint[‡]

15 August 2016

Abstract

High-order optimality conditions for convexly-constrained nonlinear optimization problems are analyzed. A corresponding (expensive) measure of criticality for arbitrary order is proposed and extended to define high-order ϵ -approximate critical points. This new measure is then used within a conceptual trust-region algorithm to show that, if derivatives of the objective function up to order $q \geq 1$ can be evaluated and are Lipschitz continuous, then this algorithm applied to the convexly constrained problem needs at most $O(\epsilon^{-(q+1)})$ evaluations of f and its derivatives to compute an ϵ -approximate q -th order critical point. This provides the first evaluation complexity result for critical points of arbitrary order in nonlinear optimization. An example is discussed showing that the obtained evaluation complexity bounds are essentially sharp.

Keywords: nonlinear optimization, high-order optimality conditions, complexity theory, machine learning.

1 Introduction

Recent years have seen a growing interest in the analysis of the worst-case evaluation complexity of nonlinear (possibly nonconvex) smooth optimization (for the nonconvex case only, see [1–7, 10–13, 15–19, 21–30, 34–36, 39, 43–47] among others). In general terms, this analysis aims at giving (sometimes sharp) bounds on the number of evaluations of a minimization problem's functions (objective and constraints, if relevant) and their derivatives that are, in the worst case, necessary for certain algorithms to find an approximate critical point for the unconstrained, convexly-constrained or general nonlinear optimization problem. It is not uncommon that such algorithms may involve possibly extremely costly internal computations, provided the number of calls to the problem functions is kept as low as possible.

*Mathematical Institute, Oxford University, Oxford OX2 6GG, Great Britain. Email: coralia.cartis@maths.ox.ac.uk.

[†]Numerical Analysis Group, Rutherford Appleton Laboratory, Chilton OX11 0QX, Great Britain. Email: 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. Email: philippe.toint@unamur.be.

Most of the research to date focusses on finding first, second and, in one case, third-order critical points. Evaluation complexity for first-order critical point was first investigated, for the unconstrained case, by Nesterov in [38] and for first- and second-order Nesterov and Polyak [39] and by Cartis, Gould and Toint in [12]. Third-order critical points were studied in [1], motivated by highly-nonlinear problems in machine learning. However, the analysis of evaluation complexity for orders higher than three is missing both concepts and results.

The purpose of the present paper is to improve on this situation in two ways. The first is to review optimality conditions of arbitrary orders $q \geq 1$ for convexly-constrained minimization problems, and the second is to describe a theoretical algorithm whose behaviour provides, for this class of problems, the first evaluation complexity bounds for such arbitrary orders of optimality.

The paper is organized as follows. After the present introduction, Section 2 discusses some preliminary results on tensor norms, a generalized Cauchy-Schwarz inequality and high-order error bounds from Taylor series. Section 3 investigates optimality conditions for convexly-constrained optimization, while Section 4 proposes a trust-region based minimization algorithm for solving this class of problems and analyzes its evaluation complexity. An example is introduced in Section 5 to show that the new evaluation complexity bounds are essentially sharp. A final discussion is presented in Section 6.

2 Preliminaries

2.1 Basic notations

In what follows, $y^T x$ denotes the Euclidean inner product of the vectors x and y of \mathbb{R}^n and $\|x\| = (x^T x)^{1/2}$ is the associated Euclidean norm. If T_1 and T_2 are tensors, $T_1 \otimes T_2$ is their tensor product. $\mathcal{B}(x, \Delta)$, the ball of radius Δ centered at x . If \mathcal{X} is a closed set, $\partial\mathcal{X}$ denotes its boundary and \mathcal{X}^0 denotes its interior. The vectors $\{e_i\}_{i=1}^n$ are the coordinate vectors in \mathbb{R}^n . The notation $\lambda_{\min}[M]$ stands for the leftmost eigenvalue of the symmetric matrix M . If $\{a_k\}$ and $\{b_k\}$ are two infinite sequences of non-negative scalars converging to zero, we say that $a_k = o(b_k)$ if and only if $\lim_{k \rightarrow \infty} a_k/b_k = 0$. The normal cone to a general convex set \mathcal{C} at $x \in \mathcal{C}$ is defined by

$$\mathcal{N}_{\mathcal{C}}(x) \stackrel{\text{def}}{=} \{s \in \mathbb{R}^n \mid s^T(z - x) \leq 0 \text{ for all } z \in \mathcal{C}\}$$

and its polar, the tangent cone to \mathcal{F} at x , by

$$\mathcal{T}_{\mathcal{C}}(x) = \mathcal{N}_{\mathcal{C}}^*(x) \stackrel{\text{def}}{=} \{s \in \mathbb{R}^n \mid s^T v \leq 0 \text{ for all } v \in \mathcal{N}_{\mathcal{C}}(x)\}.$$

Note that $\mathcal{C} \subseteq \mathcal{T}_{\mathcal{C}}(x)$ for all $x \in \mathcal{C}$. We also define $P_{\mathcal{C}}[\cdot]$ be the orthogonal projection onto \mathcal{C} and use the Moreau decomposition [37] which states that, for every $x \in \mathcal{C}$ and every $y \in \mathbb{R}^n$

$$y = P_{\mathcal{T}_{\mathcal{C}}(x)}[y] + P_{\mathcal{N}_{\mathcal{C}}(x)}[y] \quad \text{and} \quad (P_{\mathcal{T}_{\mathcal{C}}(x)}[y] - x)^T (P_{\mathcal{N}_{\mathcal{C}}(x)}[y] - x) = 0. \quad (2.1)$$

(See [20, Section 3.5] for a brief introduction of the relevant properties of convex sets and cones, or to [32, Chapter 3] or [42, Part I] for an in-depth treatment.)

2.2 Tensor norms and generalized Cauchy-Schwartz inequality

We will make substantial use of tensors and their norms in what follows, and thus start by establishing some concepts and notation. If the notation $T[v_1, \dots, v_j]$ stands for the tensor of order $q - j$ resulting from the application of the q -th order tensor T to the vectors v_1, \dots, v_j , the (recursively induced⁽¹⁾) Euclidean norm $\|\cdot\|_q$ on the space of q -th order tensors is the given by

$$\|T\|_q \stackrel{\text{def}}{=} \max_{\|v_1\|=\dots=\|v_q\|=1} T[v_1, \dots, v_q]. \quad (2.2)$$

(Observe that this value is always non-negative since we can flip the sign of $T[v_1, \dots, v_q]$ by flipping that of one of the vectors v_i .)

Note that the definition (2.2) implies that

$$\begin{aligned} \|T[v_1, \dots, v_j]\|_{q-j} &= \max_{\|w_1\|=\dots=\|w_{q-j}\|=1} T[v_1, \dots, v_j][w_1, \dots, w_{q-j}] \\ &= \left(\max_{\|w_1\|=\dots=\|w_{q-j}\|=1} T \left[\frac{v_1}{\|v_1\|}, \dots, \frac{v_j}{\|v_j\|}, w_1, \dots, w_{q-j} \right] \right) \left(\prod_{i=1}^j \|v_i\| \right) \\ &\leq \left(\max_{\|w_1\|=\dots=\|w_q\|=1} T[w_1, \dots, w_q] \right) \left(\prod_{i=1}^j \|v_i\| \right) \\ &= \|T\|_q \cdot \prod_{i=1}^j \|v_i\|, \end{aligned} \quad (2.3)$$

a simple generalization of the standard Cauchy-Schwartz inequality for order-1 tensors (vectors) and of $\|Mv\| \leq \|M\| \|v\|$ which is valid for induced norms of matrices (order-2 tensors). Observe also that perturbation theory (see [33, Th. 7]) implies that $\|T\|_q$ is continuous as a function of T .

If T is a symmetric tensor of order q , define the q -kernel of the multilinear q -form

$$T[v]^q \stackrel{\text{def}}{=} T[\underbrace{v, \dots, v}_q]$$

as

$$\ker^q[T] \stackrel{\text{def}}{=} \{v \in \mathbb{R}^n \mid T[v]^q = 0\}$$

(see [8,9]). Note that, in general, $\ker^q[T]$ is a union of cones. Interestingly, the q -kernels are not only unions of cones but also subspaces for $q = 1$. However this is not true for general q -kernels, since both $(0, 1)^T$ and $(1, 0)^T$ belong to the 2-kernel of the symmetric 2-form $x_1 x_2$ on \mathbb{R}^2 , but not their sum.

We also note that, for symmetric tensors of odd order, $T[v]^q = -T[-v]^q$ and thus that

$$- \min_{\|d\| \leq 1} T[d]^q = - \min_{\|d\| \leq 1} (-T[-d]^q) = - \min_{\|d\| \leq 1} (-T[d]^q) = \max_{\|d\| \leq 1} T[d]^q, \quad (2.4)$$

where we used the symmetry of the unit ball with respect to the origin to deduce the second equality.

⁽¹⁾That it is the recursively norm induced by the standard Euclidean norm results from the observation that

$$\max_{\|v_1\|=\dots=\|v_q\|=1} T[v_1, \dots, v_q] = \max_{\|v_q\|=1} \left[\max_{\|v_1\|=\dots=\|v_{q-1}\|=1} T[v_1, \dots, v_{q-1}] \right] [v_q].$$

2.3 High-order error bounds from Taylor series

The tensors considered in what follows are symmetric and arise as high-order derivatives of the objective function f . For the p -th derivative of a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ to be Lipschitz continuous on the set $\mathcal{S} \subseteq \mathbb{R}^n$, we require that there exists a constant $L_{f,p} \geq 0$ such that, for all $x, y \in \mathcal{S}$,

$$\|\nabla_x^p f(x) - \nabla_x^p f(y)\|_p \leq L_{f,p} \|x - y\| \quad (2.5)$$

where $\nabla_x^p h(x)$ is the p -th order symmetric derivative tensor of h at x .

Let $T_{f,p}(x, s)$ denote⁽²⁾ the p -th order Taylor-series approximation to $f(x + s)$ at some $x \in \mathbb{R}^n$ given by

$$T_{f,p}(x, s) \stackrel{\text{def}}{=} f(x) + \sum_{j=1}^p \frac{1}{j!} \nabla_x^j f(x)[s]^j \quad (2.6)$$

and consider the Taylor identity

$$\phi(1) - t_k(1) = \frac{1}{(k-1)!} \int_0^1 (1-\xi)^{k-1} [\phi^{(k)}(\xi) - \phi^{(k)}(0)] d\xi \quad (2.7)$$

involving a given univariate C^k function $\phi(\alpha)$ and its k -th order Taylor approximation $t_k(\alpha) = \sum_{i=0}^k \phi^{(i)}(0) \alpha^i / i!$ expressed in terms of the i th derivatives $\phi^{(i)}$, $i = 1, \dots, k$. Let $x, s \in \mathbb{R}^n$. Then, picking $\phi(\alpha) = f(x + \alpha s)$ and $k = q$, it follows immediately from the fact that $t_p(1) = T_{f,p}(x, s)$, the identity

$$\int_0^1 (1-\xi)^{p-1} d\xi = \frac{1}{p}, \quad (2.8)$$

(2.3), (2.5), (2.6) and (2.7) imply that, for all $x, s \in \mathbb{R}^n$,

$$\begin{aligned} f(x+s) &\leq T_{f,p}(x, s) + \frac{1}{(p-1)!} \int_0^1 (1-\xi)^{p-1} |\nabla_x^p f(x + \xi s)[s]^p - \nabla_x^p f(x)[s]^p| d\xi \\ &\leq T_{f,p}(x, s) + \left[\int_0^1 \frac{(1-\xi)^{p-1}}{(p-1)!} d\xi \right] \max_{\xi \in [0,1]} |\nabla_x^p f(x + \xi s)[s]^p - \nabla_x^p f(x)[s]^p| \\ &\leq T_{f,p}(x, s) + \frac{1}{p!} \|s\|^p \max_{\xi \in [0,1]} \|\nabla_x^p f(x + \xi s) - \nabla_x^p f(x)\|_p \\ &= T_{f,p}(x, s) + \frac{L_{f,p}}{p!} \|s\|^{p+1}. \end{aligned} \quad (2.9)$$

Similarly,

$$f(x+s) \geq T_{f,p}(x, s) - \frac{1}{p!} \|s\|^p \max_{\xi \in [0,1]} \|\nabla_x^p f(x + \xi s) - \nabla_x^p f(x)\|_p \geq T_{f,p}(x, s) - \frac{L_{f,p}}{p!} \|s\|^{p+1}. \quad (2.10)$$

Inequalities (2.9) and (2.10) will be useful in our developments below, but immediately note that they in fact depend only on the weaker requirement that

$$\max_{\xi \in [0,1]} \|\nabla_x^p f(x + \xi s) - \nabla_x^p f(x)\|_p \leq L_{f,p} \|s\|, \quad (2.11)$$

for all x and s of interest, rather than relying on (2.5).

⁽²⁾Unfortunately, double indices are necessary for most of our notation, as we need to distinguish both the function to which the relevant quantity is associated (the first index) and its order (the second index).

3 Unconstrained and convexly constrained problems

The problem we wish to solve is formally described as

$$\min_{x \in \mathcal{F}} f(x) \tag{3.1}$$

where we assume that $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is q -times continuously differentiable and bounded from below, and that f has Lipschitz continuous derivatives of orders 1 to q . We also assume that the feasible set \mathcal{F} is closed, convex and non-empty. Note that this formulation covers unconstrained optimization ($\mathcal{F} = \mathbb{R}^n$), as well as standard inequality (and linear equality) constrained optimization in its different forms: the set \mathcal{F} may be defined by simple bounds, and/or by polyhedral or more general convex constraints. We are tacitly assuming here that the cost of evaluating values and derivatives of the constraint functions possibly involved in the definition of \mathcal{F} is negligible.

3.1 High-order optimality conditions

Given that our ambition is to work with high-order model, it seems natural to aim at finding high-order local minimizers. As is standard, we say that x_* is a local minimizer of f if and only if there exists a (sufficiently small) neighbourhood \mathcal{B}_* of x_* such that

$$f(x) \geq f(x_*) \quad \text{for all } x \in \mathcal{B}_* \cap \mathcal{F}. \tag{3.2}$$

However, we must immediately remember important intrinsic limitations. These are exemplified by the smooth two-dimensional problem

$$\min_{x \in \mathbb{R}^2} f(x) = \begin{cases} x_2 \left(x_2 - e^{-1/x_1^2} \right) & \text{if } x_1 \neq 0, \\ x_2^2 & \text{if } x_1 = 0, \end{cases} \tag{3.3}$$

which is a simplified version of a problem stated by Hancock nearly a century ago [31, p. 36], itself a variation of a famous problem stated even earlier by Peano [41, Nos. 133-136]. The contour lines of its objective function are shown in Figure 3.1 on the following page.

The first conclusion which can be drawn by examining this example is that, in general, assessing that a given point x (the origin in this case) is a local minimizer needs more than verifying that every direction from this point is an ascent direction. Indeed, this latter property holds in the example, but the origin is not a local minimizer (it is a saddle point). This is caused by the fact that objective function decrease may occur along specific curves starting from the point under consideration, and these curves need not be lines (such as $x(\alpha) = 0 + \alpha e_2 + \frac{1}{2} e^{-1/2\alpha^2} e_1$ for $\alpha \geq 0$ in the example).

The second conclusion is that the characterization of a local minimizer cannot always be translated into a set of conditions only involving the Taylor expansion of f at x_* . In our example, the difficulty arises because the coefficients of the Taylor's expansion of e^{-1/x_1^2} at x all vanish as x_1 approaches the origin, and therefore that the (non-)minimizing nature of this point cannot be determined from the values of these coefficients. Thus the gap between necessary and sufficient optimality conditions cannot be closed if one restricts one's attention to using derivatives of the objective function at a putative solution of problem (3.1).




Figure 3.1: The contour lines of the objective function in (3.3)

Note that worse situations may also occur, for instance if we consider the following variation on Hancock simplified example (3.3):

$$\min_{x \in \mathbb{R}^2} f(x) = \begin{cases} x_2 \left(x_2 - \sin(1/x) e^{-1/x_1^2} \right) & \text{if } x_1 \neq 0, \\ x_2^2 & \text{if } x_1 = 0, \end{cases} \quad (3.4)$$

for which no continuous descent path exists in a neighbourhood of the origin despite the origin not being a local minimizer.

3.1.1 Necessary conditions for convexly constrained problems

The above examples show that fully characterizing a local minimizer in terms of general continuous descent paths is in general impossible. However, the fact that no such path exists remains a necessary condition for such points, even if Hancock's example shows that these paths may not be amenable to a characterization using path derivatives. In what follows, we therefore propose derivative-based necessary optimality conditions by focussing on a specific (yet reasonably general) class of descent paths $x(\alpha)$ of the form

$$x(\alpha) = x_* + \sum_{i=1}^q \alpha^i s_i + o(\alpha^q) \quad (3.5)$$

where $\alpha > 0$. Define the q -th order descriptor set of \mathcal{F} at x_* by

$$\mathcal{D}_{\mathcal{F}}^q(x) \stackrel{\text{def}}{=} \bigcup_{\varsigma > 0} \left\{ (s_1, \dots, s_q) \in \mathbb{R}^{n \times q} \mid x + \sum_{i=1}^q \alpha^i s_i \in \mathcal{F} \text{ for all } \alpha \leq \varsigma \right\} \quad (3.6)$$

Note that $\mathcal{D}_{\mathcal{F}}^1(x) = \mathcal{T}_{\mathcal{F}}(x)$, the standard tangent cone to \mathcal{F} at x . We say that a feasible curve⁽³⁾ $x(\alpha)$ is tangent to $\mathcal{D}_{\mathcal{F}}^q(x)$ if (3.5) holds for some $(s_1, \dots, s_q) \in \mathcal{D}_{\mathcal{F}}^q(x)$.

⁽³⁾Or arc, or path.

Note that the definition (3.6) implies that

$$s_i \in \mathcal{T}_{\mathcal{F}}(x_*) \quad \text{for } i \in \{1, \dots, u\} \quad (3.7)$$

where s_u is the first nonzero s_ℓ .

We now consider some conditions that preclude the existence of feasible descent paths of the form (3.5). These conditions involve the index sets $\mathcal{P}(j, k)$ defined, for $k \leq j$, by

$$\mathcal{P}(j, k) \stackrel{\text{def}}{=} \{(\ell_1, \dots, \ell_k) \in \{1, \dots, j\}^k \mid \sum_{i=1}^k \ell_i = j\}. \quad (3.8)$$

For $k \leq j \leq 4$, these are given by Table 3.2.

$j \downarrow$	$k \rightarrow$			
	1	2	3	4
1	{(1)}			
2	{(2)}	{(1,1)}		
3	{(3)}	{(1,2),(2,1)}	{(1,1,1)}	
4	{(4)}	{(1,3),(2,2),(3,1)}	{(1,1,2),(1,2,1),(2,1,1)}	{(1,1,1,1)}

Table 3.2: The sets $\mathcal{P}(j, k)$ for $k \leq j \leq 4$

We now state necessary conditions for x_* to be a local minimizer.

Theorem 3.1 Suppose that f is q times continuously differentiable in an open neighbourhood of x_* and that x_* is a local minimizer for problem (3.1). If $x_* \in \partial\mathcal{F}$, suppose furthermore that a constraint qualification holds in the sense that every feasible arc (3.5) starting from x_* is tangent to $\mathcal{D}_{\mathcal{F}}^q(x_*)$. Then, for $j \in \{1, \dots, q\}$, the inequality

$$\sum_{k=1}^j \frac{1}{k!} \left(\sum_{(\ell_1, \dots, \ell_k) \in \mathcal{P}(j, k)} \nabla_x^k f(x_*)[s_{\ell_1}, \dots, s_{\ell_k}] \right) \geq 0 \quad (3.9)$$

holds for all $(s_1, \dots, s_j) \in \mathcal{D}_{\mathcal{F}}^j(x_*)$ such that, for $i \in \{1, \dots, j-1\}$,

$$\sum_{k=1}^i \frac{1}{k!} \left(\sum_{(\ell_1, \dots, \ell_k) \in \mathcal{P}(i, k)} \nabla_x^k f(x_*)[s_{\ell_1}, \dots, s_{\ell_k}] \right) = 0. \quad (3.10)$$

Proof. Consider feasible paths of the form (3.5). Substituting this relation in the expression $f(x(\alpha)) \geq f(x_*)$ (given by (3.2)) and collecting terms of equal degree in α , we obtain that, for sufficiently small α ,

$$0 \leq f(x(\alpha)) - f(x_*) = \sum_{j=1}^q \alpha^j \sum_{k=1}^j \frac{1}{k!} \left(\sum_{(\ell_1, \dots, \ell_k) \in \mathcal{P}(j, k)} \nabla_x^k f(x_*)[s_{\ell_1}, \dots, s_{\ell_k}] \right) + o(\alpha^q) \quad (3.11)$$

where $\mathcal{P}(i, k)$ is defined in (3.8). For this to be true, we need each coefficient of α^j to be non-negative on the zero set of the coefficients $1, \dots, j-1$, subject to the requirement that the arc (3.5) must be feasible for α sufficiently small. Assume now that $s_1 \in \mathcal{T}_*$ and that (3.10) holds for $i = 1$. This latter condition request s_1 to be in the zero set of the coefficient in α in (3.11), that is

$$s_1 \in \mathcal{T}_* \cap \ker^1[\nabla_x^1 f(x_*)].$$

Then the coefficient of α^2 in (3.11) must be non-negative, which yields, using $\mathcal{P}(2, 1) = \{(2)\}$, $\mathcal{P}(2, 2) = \{(1)\}$ (see Table 3.2), that

$$\nabla_x^1 f(x_*)[s_2] + \frac{1}{2} \nabla_x^2 f(x_*)[s_1]^2 \geq 0. \quad (3.12)$$

which is (3.9) for $q = 2$.

We may then proceed in the same manner for all coefficients up from order 3 to q , each time considering them in the zero set of the previous coefficients (that is (3.10)) and verify that (3.11) directly implies (3.9). \square

Following a long tradition, we say that x_* is a q -th order critical point for problem (3.1) if the conclusions of this theorem hold for $j \in \{1, \dots, q\}$. Of course, a q -th order critical point need not be a local minimizer, but every local minimizer is a q -th order critical point.

Note that the constraint qualification assumption automatically holds if \mathcal{F} is defined by a set of explicit polynomial inequalities and/or linear equations. Also note that, as the order j grows, (3.9) may be interpreted as imposing a condition on s_j (via $\nabla_x^1 f(x_*)[s_j]$), given the directions $\{s_i\}_{i=1}^{j-1}$ satisfying (3.10).

In more general situations, the fact that conditions (3.9) and (3.10) not only depends on the behaviour of the objective function in some well-chosen subspace, but involves the geometry of the all possible feasible arcs makes the second-order condition (3.12) difficult to use, particular in the case where $\mathcal{F} \subset \mathbb{R}^n$. In what follows we discuss, as far as we currently can, two resulting questions of interest.

1. Are there cases where these conditions reduce to checking homogeneous polynomials involving the objective function's derivatives on a subspace?
2. If that is not the case, are there circumstances in which not only the complete left-hand side of (3.10) vanishes, but also each term of this left-hand side?

We start by deriving useful consequences of Theorem 3.1.

Corollary 3.2 Suppose that the assumptions of Theorem 3.1 hold and let \mathcal{N}_* be the normal cone to \mathcal{F} at x_* and \mathcal{T}_* the corresponding tangent cone. Then

$$-\nabla_x^1 f(x_*) \in \mathcal{N}_* \quad (3.13)$$

and

$$s_1 \in \mathcal{T}_* \cap \text{span} \{ \nabla_x^1 f(x_*) \}^\perp \subseteq \partial \mathcal{T}_* \quad \text{and} \quad s_2 \in \mathcal{T}_*.$$

Moreover

$$\nabla_x^1 f(x_*)[s_i] \geq 0 \quad (i = 1, 2). \quad (3.14)$$

Proof. First observe that (3.9) for $j = 1$ reduces to $\nabla_x^1 f(x_*)[s_1] = 0$ for all $s_1 \in \mathcal{T}_*$. Thus (3.13) holds. Also note that (3.9)-(3.10) impose that

$$s_1 \in \mathcal{T}_* \cap \ker^1[\nabla_x^1 f(x_*)] = \mathcal{T}_* \cap \text{span} \{ \nabla_x^1 f(x_*) \}^\perp, \quad (3.15)$$

which, because of (3.13) and the polarity of \mathcal{N}_* and \mathcal{T}_* , yields that s_1 belongs to $\partial\mathcal{T}_*$. Assume now that $s_2 \notin \mathcal{T}_*$. Then, for all α sufficiently small, $\alpha s_1 + \alpha^2 s_2$ does not belong to \mathcal{T}_* and thus $x(\alpha) = x_* + \alpha s_1 + \alpha^2 s_2 + o(\alpha^2)$ cannot belong to \mathcal{F} , which is a contradiction. Hence $s_2 \in \mathcal{T}_*$ and (3.14) follows for $i = 2$, while it follows from $s_1 \in \mathcal{T}_*$ and (3.13) for $i = 1$. \square

The first-order necessary condition (3.13) is well-known for general first-order minimizers (see [40, Th. 12.9, p. 353] for instance).

Consider now the second-order conditions (3.12). If $\mathcal{F} = \mathbb{R}^n$ (or if the convex constraints are inactive at x_*), then $\nabla_x^1 f(x_*) = 0$ because of (3.13) and (3.12) is nothing but the familiar condition that the Hessian of the objective function must be positive semi-definite. If x_* happens to lie on the boundary of \mathcal{F} and $\nabla_x^1 f(x_*) \neq 0$, (3.12) indicates that the effect of the curvature of the boundary of \mathcal{F} may be represented by the term $\nabla_x^1 f(x_*)[s_2]$, which is non-negative because of (3.14). Consider, for example, the problem

$$\min_{x \in \mathcal{F} \subset \mathbb{R}^2} x_1 \quad \text{where} \quad \mathcal{F} = \{x \in \mathbb{R}^n \mid x_1 \geq \frac{1}{2}x_2^2\},$$

whose global solution is at the origin. In this case it is easy to check that $-\nabla_x^1 f(0) = -e_1 \in \mathcal{N}_* = \text{span} \{-e_1\}$, that $\nabla_x^2 f(0) = 0$, and that second-order feasible arcs of the form (3.5) with $x(0) = 0$ may be chosen with $s_1 = \pm e_2$ and $s_2 = \beta e_1$ where $\beta \geq \frac{1}{2}$. This imposes $\nabla_x^2 f(0)[s_1]^2 \geq -1$, which (unsurprisingly) holds.

Interestingly, there are cases where the geometry of the set of locally feasible arcs is simple and manageable. In particular, suppose that the boundary of \mathcal{F} is locally polyhedral. Then, given $\nabla_x^1 f(x_*)$, either $\mathcal{T}_* \cap \text{span} \{ \nabla_x^1 f(x_*) \}^\perp = \emptyset$, in which case conditions (3.9) and (3.10) are void, or there exists $d \neq 0$ in that subspace. It then possible to define a locally feasible arc with $s_1 = d$ and $s_2 = \dots = s_q = 0$. As a consequence, the smallest possible value of $\nabla_x^1 f(x_*)[s_2]$ for feasible arcs starting from x_* is identically zero and this term therefore vanishes from (3.9)-(3.10). Moreover, because of the definition of $\mathcal{P}(k, j)$ (see Table 3.2), all terms but that in $\nabla_x^j f(x_*)[s_1]^j$ also vanish from these conditions, which then simplify to

$$\nabla_x^j f(x_*)[s_1]^j \geq 0 \quad \text{for all} \quad s_1 \in \mathcal{T}_* \cap \left(\bigcap_{i=1}^{j-1} \ker^i[\nabla_x^i f(x_*)] \right) \quad (3.16)$$

for $j = 2, \dots, q$, which is a condition only involving subspaces and (for $i \geq 2$) cones (for $i \geq 3$).

3.1.2 Necessary conditions for unconstrained problems

Consider now the case where x_* belongs to \mathcal{F}^0 , which is obviously the case if the problem is unconstrained. Then we have that the path (3.5) is unrestricted, $\mathcal{D}_{\mathcal{F}}^q(x_*) = \mathbb{R}^n$, and one is then free to choose the vectors $\{s_i\}_{i=1}^q$ (and their sign) arbitrarily. Note first that, since $\mathcal{N}_* = \{0\}$, (3.13) implies that, unsurprisingly,

$$\nabla_x^1 f(x_*) = 0.$$

For the second-order condition, we obtain from (3.9), again unsurprisingly, that, because $\ker^1[\nabla_x^1 f(x_*)] = \mathbb{R}^n$,

$$\nabla_x^2 f(x_*) \text{ is positive semi-definite on } \mathbb{R}^n.$$

Hence, if there exists a vector $s_1 \in \ker^2[\nabla_x^2 f(x_*)]$, we have that $\|[\nabla_x^2 f(x_*)]^{\frac{1}{2}} s_1\| = 0$ and therefore that $\nabla_x^2 f(x_*)[s_1, s_2] = 0$ for all $s_2 \in \mathbb{R}^n$. Thus the term for $k = 1$ vanishes from (3.9), as well as all terms involving $\nabla_x^2 f(x_*)$ applied to a vector $s_1 \in \ker^2[\nabla_x^2 f(x_*)]$. This implies in particular that the third-order condition may now be written as

$$\nabla_x^3 f(x_*)[s_1]^3 = 0 \quad \text{for all } s_1 \in \ker^2[\nabla_x^2 f(x_*)], \quad (3.17)$$

where the equality is obtained by considering both s_1 and $-s_1$.

Unfortunately, complications arise with fourth-order conditions, even when the objective function is a polynomial. Consider the following variant of Peano's [41] problem:

$$\min_{x \in \mathbb{R}^2} f(x) = x_2^2 - \kappa_1 x_1^2 x_2 + \kappa_2 x_1^4, \quad (3.18)$$

where κ_1 and κ_2 are parameters. Then one can verify that

$$\nabla_x^1 f(0) = 0, \quad \nabla_x^2 f(0) = \begin{pmatrix} 0 & 0 \\ 0 & 2 \end{pmatrix}$$

$$[\nabla_x^3 f(0)]_{ijk} = \begin{cases} -2\kappa_1 & \text{for } (i, j, k) \in \{(1, 2, 1), (1, 1, 2), (2, 1, 1)\} \\ 0 & \text{otherwise,} \end{cases}$$

and

$$[\nabla_x^4 f(0)]_{ijkl} = \begin{cases} 24\kappa_2 & \text{for } (i, j, k, \ell) = (1, 1, 1, 1) \\ 0 & \text{otherwise.} \end{cases}$$

Hence

$$\ker^1[\nabla_x^1 f(0)] = \mathbb{R}^2, \quad \ker^2[\nabla_x^2 f(0)] = \text{span}\{e_1\}, \quad \text{and} \quad \ker^3[\nabla_x^3 f(0)] = \text{span}\{e_1\} \cup \text{span}\{e_2\}.$$

The necessary condition (3.9) then state that, if the origin is a minimizer, then, using the arc defined by $s_1 = e_1$ and $s_2 = \frac{1}{2}\kappa_1 e_2$,

$$\frac{1}{2}\nabla_x^2 f(0)[s_2]^2 + \frac{1}{2}\nabla_x^3 f(0)[s_1, s_1, s_2] + \frac{1}{24}\nabla_x^4 f(0)[s_1]^4 = \frac{1}{4}\kappa_1^2 - \frac{1}{2}\kappa_1^2 + \kappa_2 = -\frac{1}{4}\kappa_1^2 + \kappa_2 \geq 0.$$

This shows that the condition $\nabla_x^4 f(x_*)[s_1]^4 \geq 0$ on $\cap_{i=1}^3 \ker^i[\nabla_x^i f(x_*)]$, although necessary, is arbitrarily far away from the weaker necessary condition

$$\frac{1}{2}\nabla_x^2 f(0)[s_2]^2 + \frac{1}{2}\nabla_x^3 f(0)[s_1, s_1, s_2] + \frac{1}{24}\nabla_x^4 f(0)[s_1]^4 \geq 0 \quad (3.19)$$

when κ_1 grows. As was already the case for problem (3.3), the example for $\kappa_1 = 1$ and $\kappa_2 = 2$, say, shows that a function may admit a saddle point ($x_* = 0$) which is a maximum ($x_* = 0$) along a curve ($x_2 = \frac{3}{2}x_1^2$ in this case) while at the same time be minimal along every line passing through x_* . Figure 3.2 shows the contour lines of the objective function of (3.18) for increasing values of κ_2 , keeping $\kappa_1 = 3$.

One may attribute the problem that not every term in (3.9) vanishes to the fact that switching signs of s_1 or s_2 does imply that any of the terms in (3.19) is zero (as we have verified) because of the terms $\nabla_x^2 f(0)[s_2]^2$ and $\nabla_x^4 f(0)[s_1]^4$. Is this a feature of even orders




Figure 3.2: The contour lines of (3.18) for $\kappa_1 = 3$ and $\kappa_2 = 2$ (left), $\kappa_2 = 2.25$ (center) and $\kappa_2 = 2.5$ (right)

only? Unfortunately, this is not the case for $q = 7$. Indeed it is not difficult to verify that the terms whose multi-index (ℓ_1, \dots, ℓ_k) is a permutation of $(1, 2, 2, 2)$ belong to $\mathcal{P}(7, 4)$ and those whose multi-index is a permutation of $(1, 1, 1, 1, 1, 2)$ belong to $\mathcal{P}(7, 6)$. Moreover, the contribution of these terms to the sum (3.9) cannot be distinguished by varying s_1 or s_2 , for instance by switching their signs as this technique yields only one equality in two unknowns. In general, we may therefore conclude that (3.9) must involve a mixture of terms with derivative tensors of various degrees.

3.1.3 A sufficient condition

Despite the limitations we have seen when considering the simplified Hancock example, we may still derive a sufficient condition for x_* to be an isolated minimizer, which is inspired by the standard second-order case (see Theorem 2.4 in Nocedal and Wright [40] for instance).

Theorem 3.3 Suppose that f is q times continuously differentiable in an open neighbourhood of $x_* \in \mathcal{F}$. If $x_* \in \partial\mathcal{F}$, suppose also that a constraint qualification holds in the sense that every feasible arc starting from x_* is tangent to $\mathcal{D}_{\mathcal{F}}^p(x_*)$. Let \mathcal{T}_* be the tangent cone to \mathcal{F} at x_* . If there exists an $q \in [1, p - 1]$ such that, for all $s \in \mathcal{T}_*$,

$$\nabla_x^i f(x_*)[s]^i = 0 \quad (i = 1, \dots, q) \quad \text{and} \quad \nabla_x^{q+1} f(x_*)[s]^{q+1} > 0, \quad (3.20)$$

then x_* is an isolated minimizer for problem (3.1).

Proof. The second part of condition (3.20) and the continuity of the $(q+1)$ -th derivative imply that

$$\nabla_x^{q+1} f(z)[s]^{q+1} > 0 \quad (3.21)$$

for all $s \in \mathcal{T}_*$ and all z is a sufficiently small feasible neighbourhood of x_* . Now, using Taylor's expansion, we obtain that, for all $s \in \mathcal{T}_*$ and all $\tau \in (0, 1)$,

$$f(x_* + \tau s) - f(x_*) = \sum_{i=1}^q \frac{\tau^i}{i!} \nabla_x^i f(x_*)[s]^i + \frac{\tau^{q+1}}{(q+1)!} \nabla_x^{q+1} f(z)[s]^{q+1}$$

for some $z \in [x_*, x_* + \tau s]$. If τ is sufficiently small, then this equality, the first part (3.20) and (3.21) ensure that $f(x_* + \tau s) > f(x_*)$. Since this strict inequality holds for all $s \in \mathcal{T}_*$ and all τ sufficiently small, x_* must be a feasible isolated minimizer. \square

Observe that, in Peano's example (see (3.18) with $\kappa_1 = 3$ and $\kappa_2 = 2$), we have that the curvature of the objective function is positive along every line passing through the origin, but that the order of the curvature varies with s (second order along $s = e_2$ and fourth order along $s = e_1$), which precludes applying Theorem 3.3.

3.1.4 An approach using Taylor models

As already noted, the conditions expressed in Theorem 3.1 may, in general, be very complicated to verify in an algorithm, due to their dependence on the geometry of the set of feasible paths. To avoid this difficulty, we now explore a different approach. Let us now define, for some $\Delta \in (0, 1]$ and some $j \in \{1, \dots, p\}$,

$$\phi_{f,j}^\Delta(x) \stackrel{\text{def}}{=} f(x) - \underset{\substack{x+d \in \mathcal{F} \\ \|d\| \leq \Delta}}{\text{globmin}} T_{f,j}(x, d), \quad (3.22)$$

the smallest value of the j -th order Taylor model $T_{f,j}(x, s)$ achievable by a feasible point at distance at most Δ from x . Note that $\phi_{f,j}^\Delta(x)$ is a continuous function of x and Δ for given \mathcal{F} and f (see [33, Th. 7]). The introduction of this quantity is in part motivated by the following theorem.

Theorem 3.4 Suppose that f is q times continuously differentiable in an open neighbourhood of x . If $x \in \partial\mathcal{F}$, suppose furthermore that a constraint qualification holds in the sense that every feasible arc starting from x is tangent to $\mathcal{D}_{\mathcal{F}}^q(x)$. Define

$$\mathcal{Z}_{\mathcal{F}}^{f,j}(x) \stackrel{\text{def}}{=} \{(s_1, \dots, s_j) \in \mathcal{D}_{\mathcal{F}}^j(x) \mid (s_1, \dots, s_i) \text{ satisfy (3.10) (at } x) \text{ for } i \in \{1, \dots, j-1\}\}.$$

Then

$$\lim_{\Delta \rightarrow 0} \frac{\phi_{f,j}^\Delta(x)}{\Delta^j} = 0 \quad \text{implies that (3.9) holds (at } x) \text{ for all } (s_1, \dots, s_j) \in \mathcal{Z}_{\mathcal{F}}^{f,j}(x).$$

Proof. We start by rewriting the power series (3.11) for degree j , for any given arc $x(\alpha)$ tangent to $\mathcal{D}_{\mathcal{F}}^j(x)$ in the form

$$f(x(\alpha)) - f(x) = \sum_{i=1}^j c_i \alpha^i + o(\alpha^j) = T_{f,j}(x, s(\alpha)) - f(x) \quad (3.23)$$

where $s(\alpha) \stackrel{\text{def}}{=} x(\alpha) - x$ and

$$c_i \stackrel{\text{def}}{=} \sum_{k=1}^i \frac{1}{k!} \sum_{(\ell_1, \dots, \ell_k) \in \mathcal{P}(i,k)} \nabla_x^k f(x)[s_{\ell_1}, \dots, s_{\ell_k}],$$

and where the last equality in (3.23) holds because f and $T_{f,j}$ share the first j derivatives at x . This reformulation allows us to write that, for $i \in \{1, \dots, j\}$,

$$c_i = \frac{1}{i!} \left. \frac{d^i}{d\alpha^i} [T_{f,j}(x, s(\alpha)) - f(x)] \right|_{\alpha=0}. \quad (3.24)$$

Assume now there exists an $(s_1, \dots, s_j) \in \mathcal{Z}_{\mathcal{F}}^{f,j}(x)$ such that (3.9) does not hold. In the notation just introduced, this means that, for this particular (s_1, \dots, s_j) ,

$$c_i = 0 \quad \text{for } i \in \{1, \dots, j-1\} \quad \text{and} \quad c_j < 0.$$

Then, from (3.24),

$$\left. \frac{d^i}{d\alpha^i} [T_{f,j}(x, s(\alpha)) - f(x)] \right|_{\alpha=0} = 0 \quad \text{for } i \in \{1, \dots, j-1\}, \quad (3.25)$$

and thus the first $(j-1)$ coefficients of the polynomial $T_{f,j}(x, s(\alpha)) - f(x)$ vanish. Thus, using (3.23),

$$\left. \frac{d^j}{d\alpha^j} [T_{f,j}(x, s(\alpha)) - f(x)] \right|_{\alpha=0} = j! \lim_{\alpha \rightarrow 0} \frac{T_{f,j}(x, s(\alpha)) - f(x)}{\alpha^j}. \quad (3.26)$$

Now let i_0 be the index of the first nonzero s_i . Note that $i_0 \in \{1, \dots, j\}$ since otherwise the structure of the sets $\mathcal{P}(i, k)$ implies that $c_j = 0$. Observe also that we may redefine the α parameter to $\alpha \|s_{i_0}\|^{1/i_0}$ so that we may assume, without loss of generality that $\|s_{i_0}\| = 1$. As a consequence, we obtain that, for sufficiently small α ,

$$\|s(\alpha)\| \leq \frac{3}{2} \alpha^{i_0} \leq \frac{3}{2} \alpha. \quad (3.27)$$

Hence, successively using the facts that $c_j < 0$, that (3.24) and (3.26) hold for all arcs $x(\alpha)$ tangent to $\mathcal{D}_{\mathcal{F}}^q(x)$, and that (3.27) and (3.22) hold, we may deduce that

$$\begin{aligned} 0 < |c_j| &\leq \frac{j!}{j!} \lim_{\alpha \rightarrow 0} \frac{f(x) - T_{f,j}(x, s(\alpha))}{\alpha^j} \\ &\leq \left(\frac{3}{2}\right)^j \lim_{\alpha \rightarrow 0} \frac{f(x) - T_{f,j}(x, s(\alpha))}{\|s(\alpha)\|^j} \\ &= \left(\frac{3}{2}\right)^j \lim_{\|s(\alpha)\| \rightarrow 0} \frac{f(x) - T_{f,j}(x, s(\alpha))}{\|s(\alpha)\|^j} \\ &\leq \left(\frac{3}{2}\right)^j \lim_{\Delta \rightarrow 0} \frac{\phi_{f,j}^{\Delta}(x)}{\Delta^j}. \end{aligned}$$

The conclusion of the theorem immediately follows since $\lim_{\Delta \rightarrow \infty} \frac{\phi_{f,j}^{\Delta}(x)}{\Delta^j} = 0$. \square

This theorem has a useful consequence.

Corollary 3.5 Suppose that f is q times continuously differentiable in an open neighbourhood of x . If $x \in \partial\mathcal{F}$, suppose furthermore that a constraint qualification holds in the sense that every feasible arc starting from x is tangent to $\mathcal{D}_{\mathcal{F}}^q(x)$. Then x is a q -th order critical point for problem (3.1) if

$$\lim_{\Delta \rightarrow 0} \frac{\phi_{f,j}^{\Delta}(x)}{\Delta^j} = 0 \quad \text{for } j \in \{1, \dots, q\}. \quad (3.28)$$

Proof. We successively apply Theorem 3.4 q times and deduce that x is a j -th order critical point for $j = 1, \dots, q$. \square

This last result says that we may avoid the difficulty of dealing the the possibly complicated geometry of $\mathcal{D}_{\mathcal{F}}^q(x)$ if we are ready to perform the global optimization occurring in (3.22) exactly and find a way to compute or overestimate the limit in (3.28). Although this is a positive conclusion, these two remaining challenges remain daunting. However, it worthwhile noting that the standard approach to computing first- second- and third-order criticality measures for unconstrained problems follows the exact same approach. In the first-order case, it is easy to verify that

$$\|\nabla_x^1 f(x)\| = \frac{1}{\Delta} \left[- \min_{\|d\| \leq \Delta} \nabla_x^1 f(x)[d] \right] = \frac{1}{\Delta} \left[f(x) - \operatorname{globmin}_{\|d\| \leq \Delta} \left(f(x) + \nabla_x^1 f(x)[d] \right) \right]$$

where the first equality is justified by the convexity of $\nabla_x^1 f(x)[d]$ as a function of d . Because the left-hand side of the above relation is independent of Δ , the computation of the limit (3.28) for Δ tending to zero is trivial when $j = 1$ and the limiting value is $\|\nabla_x f(x)\|$. For the second-order case, assuming $\|\nabla_x^1 f(x)\| = 0$,

$$\begin{aligned} \left| \min \left[0, \lambda_{\min}[\nabla_x^2 f(x)] \right] \right| &= \frac{1}{2\Delta^2} \left[- \operatorname{globmin}_{\|d\| \leq \Delta} \nabla_x^2 f(x)[d]^2 \right] \\ &= \frac{1}{\Delta^2} \left[f(x) - \operatorname{globmin}_{\|d\| \leq \Delta} \left(f(x) + \nabla_x^1 f(x)[d] + \frac{1}{2} \nabla_x^2 f(x)[d]^2 \right) \right], \end{aligned} \quad (3.29)$$

the first global optimization problem being easily solvable by a trust-region-type calculation [20, Section 7.3] or directly by an equivalent eigenvalue analysis. As for the first-order case, the left-hand side of the equation is independent of Δ and obtaining the limit for Δ tending to zero is trivial.

Finally, if $\mathcal{M}(x) \stackrel{\text{def}}{=} \ker[\nabla_x^1 f(x)] \cap \ker[\nabla_x^2 f(x)]$ and $P_{\mathcal{M}(x)}$ is the orthogonal projection onto that subspace,

$$\begin{aligned} \|P_{\mathcal{M}(x)}(\nabla_x^3 f(x))\| &= \frac{1}{6\Delta^3} \left[- \min_{\|d\| \leq \Delta} \nabla_x^1 f(x)[d] \right] \\ &= \frac{1}{\Delta^3} \left[f(x) - \operatorname{globmin}_{\|d\| \leq \Delta} \left(f(x) + \nabla_x^1 f(x)[d] + \frac{1}{2} \nabla_x^2 f(x)[d]^2 + \frac{1}{6} \nabla_x^3 f(x)[d]^3 \right) \right] \end{aligned} \quad (3.30)$$

where the first equality results from (2.2). In this case, the global optimization in the subspace $\mathcal{M}(x)$ is potentially harder to solve exactly (a randomization argument is used in [1] to derive

a upper bound on its value), although it still involves a subspace⁽⁴⁾.

While we are unaware of a technique for making the global minimization in (3.22) easy in the even more complicated general case, we may think of approximating the limit in (3.28) by choosing a (user-supplied) value of $\Delta > 0$ small enough⁽⁵⁾ and consider the size of the quantity

$$\frac{1}{\Delta^j} \phi_{f,j}^\Delta(x). \quad (3.31)$$

Unfortunately, it is easy to see that, if Δ is fixed at some positive value, a zero value of $\phi_{f,j}^\Delta(x)$ alone is not a necessary condition for x being a local minimizer. Indeed consider the univariate problem of minimizing $f(x) = x^2(1 - \alpha x)$ for $\alpha > 0$. One verifies that, for any $\Delta > 0$, the choice $\alpha = 2/\Delta$ yields that

$$\phi_{f,1}^\Delta(0) = 0, \quad \phi_{f,2}^\Delta(0) = 0 \quad \text{but} \quad \phi_{f,3}^\Delta(0) = \frac{4}{\alpha^2} > 0, \quad (3.32)$$

despite 0 being a local (but not global) minimizer. As a matter of fact, $\phi_{f,j}^\Delta(x)$ gives more information than the mere potential proximity of a j -th order critical point: it is able to see beyond an infinitesimal neighbourhood of x and provides information on possible further descent beyond such a neighbourhood. Rather than a true criticality measure, it can be considered, for fixed Δ , as an indicator of further progress, but its use for terminating at a *local* minimizer is clearly imperfect.

Despite this drawback, the above arguments would suggest that it is reasonable to consider a (conceptual) minimization algorithm whose objective is to find a point x_ϵ such that

$$\phi_{f,j}^\Delta(x_\epsilon) \leq \epsilon \Delta^j \quad \text{for} \quad j = 1, \dots, q \quad (3.33)$$

for some $\Delta \in (0, 1]$ sufficiently small and some $q \in \{1, \dots, p\}$. This condition implies an approximate minimizing property which we make more precise by the following result.

Theorem 3.6 Suppose that f is q times continuously differentiable and that $\nabla_x^q f$ is Lipschitz continuous with constant $L_{f,q}$ (in the sense of (2.5)) in an open neighbourhood of x_ϵ of radius larger than Δ . Suppose also (3.33) holds for $j = q$. Then

$$f(x_\epsilon + d) \geq f(x_\epsilon) - 2\epsilon \Delta^q \quad \text{for all } x_\epsilon + d \in \mathcal{F} \text{ such that } \|d\| \leq \min \left(\frac{p! \epsilon \Delta^q}{L_{f,p}} \right)^{\frac{1}{q+1}}. \quad (3.34)$$

Proof. Consider $x + d \in \mathcal{F}$. Using the triangle inequality, we have that

$$\begin{aligned} f(x_\epsilon + d) &= f(x_\epsilon + d) - T_{f,q}(x_\epsilon, d) + T_{f,q}(x_\epsilon, d) \\ &\geq -|f(x_\epsilon + d) - T_{f,q}(x_\epsilon, d)| + T_{f,q}(x_\epsilon, d). \end{aligned} \quad (3.35)$$

Now, condition (3.33) for $j = q$, implies that, if $\|d\| \leq \Delta$,

$$T_{f,q}(x_\epsilon, d) \geq T_{f,q}(x_\epsilon, 0) - \epsilon \Delta^q = f(x_\epsilon) - \epsilon \Delta^q. \quad (3.36)$$

⁽⁴⁾We saw in Section 3.1.2 that $q = 3$ is the highest order for which this is possible.

⁽⁵⁾Note that a small Δ has the advantage of limiting the global optimization effort.

Hence, substituting (2.10) and (3.36) into (3.35) and using again that $\|d\| \leq \Delta < 1$, we deduce that

$$f(x_\epsilon + d) \geq f(x_\epsilon) - \frac{L_{f,p}}{p!} \|d\|^{q+1} - \epsilon \Delta^q,$$

and the desired result follows. \square

The size of the neighbourhood of x_ϵ where f is “locally smallest” — in that the first part (3.34) holds — therefore increases with the criticality order q , a feature potentially useful in various contexts such as global optimization.

Before turning to more algorithmic aspects, we briefly compare the results of Theorem 3.6 which what can be deduced on the local behaviour of the Taylor series $T_{f,q}(x_*, s)$ if, instead of requiring the exact necessary condition (3.9) to hold exactly, this condition is relaxed to

$$\sum_{k=1}^j \frac{1}{k!} \left(\sum_{(\ell_1, \dots, \ell_k) \in \mathcal{P}(j,k)} \nabla_x^k f(x_*)[s_{\ell_1}, \dots, s_{\ell_k}] \right) \geq -\epsilon \quad (3.37)$$

while insisting that (3.10) should hold exactly. If $j = q = 1$, it is easy to verify that (3.37) for $s_1 \in \mathcal{T}_*$ is equivalent to the condition that

$$\|P_{\mathcal{T}_*}[\nabla_x^1 f(x_*)]\| \leq \epsilon, \quad (3.38)$$

from which we deduce, using the Cauchy-Schwarz inequality, that

$$T_{f,1}(x_*, s) \geq T_{f,1}(x_*, 0) - \epsilon \Delta \quad (3.39)$$

for all $s \in \mathcal{T}_*$ with $\|d\| \leq \Delta$, that is (3.33) for $j = 1$. Thus, by Theorem 3.6, we obtain that (3.34) holds for $j = 1$.

4 A trust-region minimization algorithm

Aware of the optimality conditions and their limitations, we may now consider an algorithm to achieve (3.33). This objective naturally suggests a trust-region⁽⁶⁾ formulation with adaptive model degree, in which the user specifies a desired criticality order q , assuming that derivatives of order $1, \dots, q$ are available when needed. We made this idea explicit in Algorithm 4.1 on the next page.

We first state a useful property of Algorithm 4.1, which ensures that a fixed fraction of the iterations $1, 2, \dots, k$ must be either successful or very successful. Indeed, if we define

$$\mathcal{S}_k \stackrel{\text{def}}{=} \{\ell \in \{1, \dots, k\} \mid \rho_\ell \geq \eta_1\},$$

the following bound holds.

⁽⁶⁾A detailed account and a comprehensive bibliography on trust-region methods can be found in [20].

Algorithm 4.1: Trust-region algorithm using adaptive order models for convexly-constrained problems (TR q)

Step 0: Initialization. A criticality order q , an accuracy threshold $\epsilon \in (0, 1]$, a starting point x_0 and an initial trust-region radius $\Delta_1 \in [\epsilon, 1]$ are given, as well as algorithmic parameters $\Delta_{\max} \in [\Delta_1, 1]$, $\gamma_1 \leq \gamma_2 < 1 \leq \gamma_3$ and $0 < \eta_1 \leq \eta_2 < 1$. Compute $x_1 = P_{\mathcal{F}}[x_0]$, evaluate $f(x_1)$ and set $k = 1$.

Step 1: Step computation. For $j = 1, \dots, q$,

1. Evaluate $\nabla^j f(x_k)$ and compute $\phi_{f,j}^{\Delta_k}(x_k)$ from (3.22).
2. If $\phi_{f,j}^{\Delta_k}(x_k) > \epsilon \Delta_k^j$, go to Step 3 with $s_k = d$, where d is the argument of the global minimum in the computation of $\phi_{f,j}^{\Delta_k}(x_k)$.

Step 2: Termination. Terminate with $x_\epsilon = x_k$ and $\Delta_\epsilon = \Delta_k$.

Step 3: Accept the new iterate. Compute $f(x_k + s_k)$ and

$$\rho_k = \frac{f(x_k) - f(x_k + s_k)}{T_{f,j}(x_k, 0) - T_{f,j}(x_k, s_k)}. \quad (4.1)$$

If $\rho_k \geq \eta_1$, set $x_{k+1} = x_k + s_k$. Otherwise set $x_{k+1} = x_k$.

Step 4: Update the trust-region radius. Set

$$\Delta_{k+1} \in \begin{cases} [\gamma_1 \Delta_k, \gamma_2 \Delta_k] & \text{if } \rho_k < \eta_1, \\ [\gamma_2 \Delta_k, \Delta_k] & \text{if } \rho_k \in [\eta_1, \eta_2), \\ [\Delta_k, \min(\Delta_{\max}, \gamma_3 \Delta_k)] & \text{if } \rho_k \geq \eta_2, \end{cases} \quad (4.2)$$

increment k by one and go to Step 1.

Lemma 4.1 Assume that $\Delta_k \geq \Delta_{\min}$ for some $\Delta_{\min} > 0$ independent of k . Then Algorithm 4.1 ensures that, whenever $\mathcal{S}_k \neq \emptyset$,

$$k \leq \kappa_u |\mathcal{S}_k|, \quad \text{where } \kappa_u \stackrel{\text{def}}{=} \left(1 + \frac{\log \gamma_3}{|\log \gamma_2|}\right) + \frac{1}{|\log \gamma_2|} \log \left(\frac{\Delta_1}{\Delta_{\min}}\right). \quad (4.3)$$

Proof. The trust-region update (4.2) ensures that

$$\Delta_k \leq \Delta_1 \gamma_2^{|\mathcal{U}_k|} \gamma_3^{|\mathcal{S}_k|},$$

where $\mathcal{U}_k = \{1, \dots, k\} \setminus \mathcal{S}_k$. This inequality then yields (4.3) by taking logarithms and using that $|\mathcal{S}_k| \geq 1$ and $k = |\mathcal{S}_k| + |\mathcal{U}_k|$. \square

4.1 Evaluation complexity for Algorithm 4.1

We start our worst-case analysis by formalizing our assumptions. Let

$$\mathcal{L}_f \stackrel{\text{def}}{=} \{x + z \in \mathbb{R}^n \mid x \in \mathcal{F}, f(x) \leq f(x_1) \text{ and } \|z\| \leq \Delta_{\max}\}.$$

AS.1 The feasible set \mathcal{F} is closed, convex and non-empty.

AS.2 The objective function f is q times continuously differentiable on an open set containing \mathcal{L}_f .

AS.3 For $j \in \{1, \dots, q\}$, the j -th derivative of f is Lipschitz continuous on \mathcal{L}_f (in the sense of (2.5)) with Lipschitz constant $L_{f,j} \geq 1$.

For simplicity of notation, define $L_f \stackrel{\text{def}}{=} \max_{j \in \{1, \dots, q\}} L_{f,j}$.

Algorithm 4.1 is required to start from a feasible $x_1 \in \mathcal{F}$, which, together with the fact that the subproblem solution in Step 2 involves minimization over \mathcal{F} , leads to AS.1. Note that AS.3 automatically holds if f is $q+1$ times continuously differentiable and \mathcal{F} is bounded. It is also important to note that we could replace AS.3 by the condition that (2.11) holds on the path of iterates $\cup_{k \geq 1} [x_k, x_{k+1}]$ without altering any of the proofs below. While this weaker formulation may be useful, we prefer to use AS.3 in the sequel of this paper because it is independent of the sequence of iterates and may be easier to verify a priori, given the problem (3.1).

Lemma 4.2 Suppose that AS.2 and AS.3 hold. Then, for all $\ell \in \{1, \dots, k\}$,

$$\Delta_\ell \geq \kappa_\Delta \epsilon \quad (4.4)$$

where

$$\kappa_\Delta \stackrel{\text{def}}{=} \min \left[1, \frac{\gamma_1(1 - \eta_2)}{L_f} \right]. \quad (4.5)$$

Proof. Assume that, for some $\ell \in \{1, \dots, k\}$

$$\Delta_\ell \leq \frac{1 - \eta_2}{L_f} \epsilon. \quad (4.6)$$

From (4.1), we obtain that, for some $j \in \{1, \dots, q\}$,

$$|1 - \rho_\ell| = \frac{f(x_\ell + s_\ell) - T_{f,j}(x_\ell, s_\ell)}{T_{f,j}(x_\ell, 0) - T_{f,j}(x_\ell, s_\ell)} < \frac{L_f \|s_\ell\|^{j+1}}{j! \epsilon \Delta_\ell^j} \leq \frac{L_f \Delta_\ell}{j! \epsilon} \leq (1 - \eta_2),$$

where we used (2.9) and the fact that $\phi_{f,j}^{\Delta_\ell}(x_\ell) > \epsilon \Delta_\ell^j$ to deduce the first inequality, the bound $\|s_\ell\| \leq \Delta_\ell$ to deduce the second, and (4.6) with $j \geq 1$ to deduce the third. Thus $\rho_\ell \geq \eta_2$ and $\Delta_{\ell+1} \geq \Delta_\ell$. The mechanism of the algorithm and the inequality $\Delta_1 \geq \epsilon$ then ensures that, for all $\ell \in k$,

$$\Delta_\ell \geq \min \left[\Delta_1, \frac{\gamma_1(1 - \eta_2)\epsilon}{L_f} \right] \geq \kappa_\Delta \epsilon.$$

□

We now derive a simple lower bound on the objective function decrease at successful iterations.

Lemma 4.3 Suppose that AS.1–AS.3 hold. Then, if k is the index of a successful iteration before termination,

$$f(x_k) - f(x_{k+1}) \geq \eta_1 \kappa_\Delta \epsilon^{q+1}. \quad (4.7)$$

Proof. We have, using (4.1), the fact that $\phi_{f,j}^{\Delta_k}(x_k) > \epsilon \Delta_k^j$ for some $j \in \{1, \dots, q\}$ and (4.4) successively, that

$$f(x_k) - f(x_{k+1}) \geq \eta_1 [T_{f,j}(x_k, 0) - T_{f,j}(x_k, s_k)] = \eta_1 \phi_{f,j}^{\Delta_k}(x_k) > \eta_1 \kappa_\Delta \epsilon^{j+1} \geq \eta_1 \kappa_\Delta \epsilon^{q+1}.$$

□

Our worst-case evaluation complexity results can now be proved by summing the decreases guaranteed by this last lemma.

Theorem 4.4 Suppose that AS.1–AS.3 hold. Let f_{low} be a lower bound on f within \mathcal{F} . Then, given $\epsilon \in (0, 1]$, Algorithm 4.1 applied to problem (3.1) needs at most

$$\left\lceil \kappa_{\mathcal{S}}^f \frac{f(x_0) - f_{\text{low}}}{\epsilon^{q+1}} \right\rceil \quad (4.8)$$

successful iterations (each possibly involving one evaluation of f and its q first derivatives) and at most

$$\left\lceil \kappa_u \kappa_{\mathcal{S}}^f \frac{f(x_0) - f_{\text{low}}}{\epsilon^{q+1}} \right\rceil + 1 \quad (4.9)$$

iterations in total to terminate with an iterate x_ϵ such that (3.33) holds, where

$$\kappa_{\mathcal{S}}^f = \frac{1}{\eta_1} \max \left[1, \frac{L_f}{\gamma_1(1 - \eta_2)} \right], \quad (4.10)$$

and κ_u is given by (4.3). Moreover, if Δ_ϵ is the value of Δ_k at termination,

$$f(x_\epsilon + d) \geq f(x_\epsilon) - 2\epsilon \Delta_\epsilon^q \quad (4.11)$$

for all d such that

$$x_\epsilon + d \in \mathcal{F} \quad \text{and} \quad \|d\| \leq (\epsilon \Delta_\epsilon^q)^{\frac{1}{q+1}} \left(\frac{L_f}{q!} \right)^{-\frac{1}{q+1}}. \quad (4.12)$$

Observe that, because of (4.2) and (4.4), $\Delta_\epsilon \in [\kappa_\delta \epsilon, \Delta_{\text{max}}]$.

Proof. Let k be the index of an arbitrary iteration before termination. Using the definition of f_{low} , the nature of successful iterations, (4.10) and Lemma 4.3, we deduce that

$$f(x_0) - f_{\text{low}} \geq f(x_0) - f(x_{k+1}) = \sum_{i \in \mathcal{S}_k} [f(x_i) - f(x_{i+1})] \geq |\mathcal{S}_k| [\kappa_{\mathcal{S}}^f]^{-1} \epsilon^{q+1} \quad (4.13)$$

which proves (4.8).

We next call upon Lemma 4.1 to compute the upper bound on the total number of iterations before termination (obviously, there must be a least one successful iteration unless termination occurs for $k = 1$) and add one for the evaluation at termination. Finally, (4.11)-(4.12) result from Theorem 3.6, AS.3 and the fact that $\phi_{f,q}^{\Delta_k}(x_\epsilon) \leq \epsilon \Delta_{k_\epsilon}^q$ at termination. \square

Theorem 4.4 generalizes the known bounds for the cases where $\mathcal{F} = \mathbb{R}$ and $q = 1$ [38], $q = 2$ [12, 39] and $q = 3$ [1]. The results for $q = 2$ with $\mathcal{F} \subset \mathbb{R}^n$ and for $q > 3$ appear to be new. The latter provide the first evaluation complexity bounds for general criticality order q . Note that, if $q = 1$, bounds of the type $O(\epsilon^{(p+1)/p})$ exist if one is ready to minimize models of degree $p > q$ (see [6]). Whether similar improvements can be obtained for $q > 1$ remains an open question at this stage.

We also observe that the above theory remains valid if the termination rule

$$\phi_{k,j}^{\Delta_k}(x_k) \leq \epsilon \Delta_k^j \quad \text{for } j \in \{1, \dots, q\} \quad (4.14)$$

used in Step 1 is replaced by a more flexible one, involving other acceptable termination circumstances, such as if (4.14) hold *or* some other condition holds. We conclude this section by noting that the global optimization effort involved in the computation of $\phi_{j,j}^{\Delta_k}(x_k)$ ($j \in \{1, \dots, q\}$) in Algorithm 4.1 might be limited by choosing Δ_{\max} small enough.

5 Sharpness

It is interesting that an example was presented in [14] showing that the bound in $O(\epsilon^{-3})$ evaluations for $q = 2$ is essentially sharp for both the trust-region and a regularization algorithm. This is significant, because requiring $\phi_{f,2}^{\Delta}(x) \leq \epsilon \Delta^2$ is slightly stronger, for small Δ , than the standard condition

$$\|\nabla_x^1 f(x)\| \leq \epsilon \quad \text{and} \quad \min \left[0, \lambda_{\min}[\nabla_x^2 f(x)] \right] \geq -\epsilon \quad (5.1)$$

(used in [39] and [12] for instance). Indeed, for one-dimensional problems and assuming $\nabla_x^2 f(x) \leq 0$, the former condition amounts to requiring that

$$\frac{1}{2} \left(-\nabla_x^2 f(x) + 2 \frac{|\nabla_x^1 f(x)|}{\Delta} \right) \leq \epsilon \quad (5.2)$$

where the absolute value reflects the fact that $s = \pm \Delta$ depending on the sign of g . In the remainder of this section, we show that the example proposed in [14] can be extended to arbitrary order q , and thus that the complexity bounds (4.8)-(4.9) are essentially sharp for our trust-region algorithm.

The idea of our generalized example is to apply Algorithm 4.1 to a unidimensional objective function f for some fixed $q \geq 1$ and $\mathcal{F} = \mathbb{R}_+$ (hence guaranteeing AS.1), generating a sequence of iterates $\{x_k\}_{k \geq 0}$ starting from the origin, i.e. $x_0 = x_1 = 0$. We first choose the sequences of derivatives values up to order q to be, for all $k \geq 1$,

$$\nabla_x^j f(x_k) = 0 \quad \text{for } j \in \{1, \dots, q-1\} \quad \text{and} \quad \nabla_x^q f(x_k) = -q! \left(\frac{1}{k+1} \right)^{\frac{1}{q+1} + \delta} \quad (5.3)$$

where $\delta \in (0, 1)$ is a (small) positive constant. This means that, at iterate x_k , the q -th order Taylor model is given by

$$T_{f,q}(x_k, s) = f(x_k) - \left(\frac{1}{k+1} \right)^{\frac{1}{q+1} + \delta} s^q$$

where the value of $f(x_k)$ remains unspecified for now. The step is then obtained by minimizing this model in a trust-region of radius

$$\Delta_k = \left(\frac{1}{k+1} \right)^{\frac{1}{q+1} + \delta},$$

yielding that

$$s_k = \Delta_k = \left(\frac{1}{k+1} \right)^{\frac{1}{q+1} + \delta} \in (0, 1). \quad (5.4)$$

As a consequence, the model decrease is given by

$$T_{f,q}(x_k, 0) - T_{f,q}(x_k, s_k) = -\frac{1}{q!} \nabla_x^q f(x_k) s_k^q = \left(\frac{1}{k+1} \right)^{1+(q+1)\delta}. \quad (5.5)$$

For our example, we define the objective function decrease at iteration k to be

$$\Delta f_k \stackrel{\text{def}}{=} f(x_k) - f(x_k + s_k) = \frac{1}{2}(\eta_1 + \eta_2)[T_{f,q}(x_k, 0) - T_{f,q}(x_k, s_k)], \quad (5.6)$$

thereby ensuring that $\rho_k \in [\eta_1, \eta_2]$ and $x_{k+1} = x_k + s_k$ for each k . Summing up function decreases, we may then specify the objective function's values at the iterates by

$$f(x_0) = \frac{\eta_1 + \eta_2}{2} \zeta(1 + (q+1)\delta) \quad \text{and} \quad f(x_{k+1}) = f(x_k) - \frac{\eta_1 + \eta_2}{2} \left(\frac{1}{k+1} \right)^{1+(q+1)\delta}. \quad (5.7)$$

where $\zeta(t) \stackrel{\text{def}}{=} \sum_{k=1}^{\infty} k^{-t}$ is the Riemann zeta function. This function is finite for all $t > 1$ (and thus also for $t = 1 + (q+1)\delta$), thereby ensuring that $f(x_k) \geq 0$ for all $k \geq 0$. We also verify that

$$\frac{\Delta_{k+1}}{\Delta_k} = \left(\frac{k+1}{k+2} \right)^{\frac{1}{q+1} + \delta} \in [\gamma_2, 1]$$

in accordance with (4.2), provided $\gamma_2 \leq \left(\frac{2}{3}\right)^{\frac{1}{q+1} + \delta}$. Observe also that (5.3) and (5.5) ensure that, for each $k \geq 1$,

$$\phi_{f,j}^{\Delta_k}(x_k) = 0 \quad \text{for } j \in \{1, \dots, q-1\} \quad (5.8)$$

and

$$\phi_{f,q}^{\Delta_k}(x_k) = \left(\frac{1}{k+1} \right)^{1+(q+1)\delta} = \left(\frac{1}{k+1} \right)^{\frac{1}{q+1} + \delta} \Delta_k^q. \quad (5.9)$$

We now use Hermite interpolation to construct the objective function f on the successive intervals $[x_k, x_{k+1}]$, and define

$$f(x) = p_k(x - x_k) + f(x_k) \quad \text{for } x \in [x_k, x_{k+1}] \text{ and } k \geq 1, \quad (5.10)$$

where p_k is the polynomial

$$p_k(s) = \sum_{i=0}^{2q+1} c_{i,k} s^i, \quad (5.11)$$

with coefficients defined by the interpolation conditions

$$\begin{aligned} p_k(0) &= f(x_k) - f(x_{k+1}), & p_k(s_k) &= 0; \\ \nabla_s^j p_k(0) &= 0 = \nabla_s^j p_k(s_k) & \text{for } j &\in \{1, \dots, q-1\}, \\ \nabla_s^q p_k(0) &= \nabla_x^q f(x_k), & \nabla_s^q p_k(s_k) &= \nabla_x^q f(x_{k+1}). \end{aligned} \quad (5.12)$$

These conditions ensure that $f(x)$ is q times continuously differentiable on \mathbb{R}_+ and thus that AS.2 holds. They also impose the following values for the first $q + 1$ coefficients

$$c_{0,k} = f(x_k) - f(x_{k+1}), \quad c_{j,k} = 0 \quad (j \in \{1, \dots, q-1\}), \quad c_{q,k} = -\nabla_x^q f(x_k); \quad (5.13)$$

and the remaining $q + 1$ coefficients are solutions of the linear system

$$\begin{pmatrix} s_k^{q+1} & s_k^{q+2} & \cdots & s_k^{2q+1} \\ (q+1)s_k^q & (q+2)s_k^{q+1} & \cdots & (2q+1)s_k^{2q} \\ \vdots & \vdots & & \vdots \\ \frac{(q+1)!}{1!} s_k & \frac{(q+2)!}{2!} s_k^2 & \cdots & \frac{(2q+1)!}{(q+1)!} s_k^{q+1} \end{pmatrix} \begin{pmatrix} c_{q+1,k} \\ c_{q+2,k} \\ \vdots \\ c_{2q+1,k} \end{pmatrix} = r_k, \quad (5.14)$$

where the right-hand side is given by

$$r_k = \begin{pmatrix} -\Delta f_k - \frac{1}{q!} \nabla_x^q f(x_k) s_k^q \\ -\frac{1}{(q-1)!} \nabla_x^q f(x_k) s_k^{q-1} \\ \vdots \\ -\nabla_x^p f(x_k) s_k \\ \nabla_x^q f(x_{k+1}) - \nabla_x^q f(x_k) \end{pmatrix}. \quad (5.15)$$

Observe now that the coefficient matrix of this linear system may be written as

$$\begin{pmatrix} s_k^{q+1} & & & \\ & s_k^q & & \\ & & \ddots & \\ & & & s_k \end{pmatrix} M_q \begin{pmatrix} 1 & & & \\ & s_k & & \\ & & \ddots & \\ & & & s_k^q \end{pmatrix}$$

where

$$M_q \stackrel{\text{def}}{=} \begin{pmatrix} 1 & 1 & \cdots & 1 \\ q+1 & q+2 & \cdots & 2q+1 \\ \vdots & \vdots & & \vdots \\ \frac{(q+1)!}{1!} & \frac{(q+2)!}{2!} & \cdots & \frac{(2q+1)!}{(q+1)!} \end{pmatrix} \quad (5.16)$$

is an invertible matrix independent of k (see Appendix). Hence

$$\begin{pmatrix} c_{q+1,k} \\ c_{q+2,k} \\ \vdots \\ c_{2q+1,k} \end{pmatrix} = \begin{pmatrix} 1 & & & \\ & s_k^{-1} & & \\ & & \ddots & \\ & & & s_k^{-q} \end{pmatrix} M_q^{-1} \begin{pmatrix} s_k^{-(q+1)} & & & \\ & s_k^{-q} & & \\ & & \ddots & \\ & & & s_k^{-1} \end{pmatrix} r_k \quad (5.17)$$

Observe now that, because of (5.4), (5.6), (5.5) and (5.3),

$$|\Delta f_k| = O(s_k^{q+1}), \quad |\nabla_x^q f(x_k) s_k^{q-j}| = O(s_k^{q+1-j}) \quad (j \in \{0, \dots, q-1\})$$

and, since $\nabla_x^q f(x_k) < \nabla_x^q f(x_{k+1}) < 0$,

$$|\nabla_x^q f(x_{k+1}) - \nabla_x^q f(x_k)| \leq |\nabla_x^q f(x_k)| \leq q! s_k.$$

These bounds and (5.15) imply that $[r_k]_i$, the i -th component of r_k , satisfies

$$|[r_k]_i| = O(s_k^{q+2-i}) \quad \text{for } i \in \{1, \dots, q+1\}.$$

Hence, using (5.17) and the non-singularity of M_q , we obtain that there exists a constant $\kappa_q \geq 1$ independent of k such that

$$|c_{i,k}| s_k^{i-q-1} \leq \kappa_q \quad \text{for } i \in \{q+1, \dots, 2q+1\}, \quad (5.18)$$

and thus that

$$|\nabla_s^{q+1} p_k(s)| \leq \sum_{i=q+1}^{2q+1} i! |c_{i,k}| s_k^{i-q-1} \leq \left(\sum_{i=q+1}^{2q+1} i! \right) \kappa_q$$

Moreover, using successively (5.11), the triangle inequality, (5.13), (5.3), (5.4), (5.18) and $\kappa_q \geq 1$, we obtain that, for $j \in \{1, \dots, q\}$,

$$\begin{aligned} |\nabla_s^j p_k(s)| &\leq \sum_{i=j}^{2q+1} \frac{i!}{(i-j)!} |c_{i,k}| s^{i-j} \\ &= \frac{q!}{(q-j)!} |c_{q,k}| s^{q-j} + \sum_{i=q+1}^{2q+1} \frac{i!}{(i-j)!} |c_{i,k}| s^{i-q-1} s^{q+1-j} \\ &\leq \frac{q!}{(q-j)!} + \sum_{i=q+1}^{2q+1} \frac{i!}{(i-j)!} |c_{i,k}| s^{i-q-1} \\ &\leq \left(\sum_{i=q}^{2q+1} \frac{i!}{(i-j)!} \right) \kappa_q \end{aligned}$$

and thus all derivatives of order one up to q remain bounded on $[0, s_k]$. Because of (5.10), we therefore obtain that AS.3 holds. Moreover (5.13), (5.18), the inequalities $|\nabla_x^q f(x_k)| \leq q!$ and $f(x_k) \geq 0$, (5.10) and (5.4) also ensure that $f(x)$ is bounded below.

We have therefore shown that the bounds of Theorem 4.4 are essentially sharp, in that, for every $\delta > 0$, Algorithm 4.1 applied to the problem of minimizing the lower-bounded objective function f just constructed and satisfying AS.1-AS.3 will take, because of (5.8) and (5.9),

$$\left[\frac{1}{e^{\frac{q+1}{1+(q+1)\delta}}} \right]$$

iterations and evaluation of f and its q first derivatives to find an iterate x_k such that condition (4.14) holds. Moreover, it is clear that, in the example presented, the global rate of convergence is driven by the term of degree q in the Taylor series.

6 Discussion

We have analyzed the optimality conditions of order 2 and above, and proposed a measure of criticality for arbitrary order for convexly constrained nonlinear optimization problems. As this measure can be extended to define ϵ -approximate critical points of high-order, we have then used it in a conceptual trust-region algorithm to show that, if derivatives of the objective

function up to order $q \geq 1$ can be evaluated and are Lipschitz continuous, then this algorithm applied to the convexly constrained problem (3.1) needs at most $O(\epsilon^{-(q+1)})$ evaluations of f and its derivatives to compute an ϵ -approximate q -order critical point. Moreover, we have shown by an example that this bound is essentially sharp.

In the purely unconstrained case, this result recovers known results for $q = 1$ (first-order criticality for Lipschitz gradients) [38], $q = 2$ (second-order criticality⁽⁷⁾ with Lipschitz Hessians) [14, 39] and $q = 3$ (third-order criticality⁽⁸⁾ with Lipschitz continuous third derivative) [1], but extends them to arbitrary order. The results for the convexly constrained case appear to be new and provide in particular the first complexity bound for second- and third-order criticality for such inequality constrained problems.

Because the condition (4.14) measure different orders of criticality, we could choose to use a different ϵ for every order (as in [14]), complicating the expression of the bound accordingly. However, as shown by our example, the worst-case behaviour of Algorithm 4.1 is dominated by that of $\nabla_x^q f$, which makes the distinction of the various ϵ -s less crucial.

Because of the global optimization occurring in the definition of the criticality measure $\phi_{f,j}^\Delta(x)$, the algorithm discussed in the present paper remains, in general, of a theoretical nature. However there may be cases where this computation is tractable for small enough Δ , for instance if the derivative tensors of the objective function are strongly structured. Such approaches may hopefully be of use for small dimensional or structured highly nonlinear problems, such as those occurring in machine learning using deep learning techniques (see [1]).

The present framework for handling convex constraints is not free of limitations, resulting from our choice to transfer difficulties associated with the original problem to the subproblem solution, thereby sparing precious evaluations of f and its derivatives. In particular, the cost of evaluating any constraint function/derivative possibly defining the convex feasible set \mathcal{F} is neglected by the present approach, which must therefore be seen as a suitable framework to handle “cheap inequality constraint” such as simple bounds.

It is known that the complexity of obtaining ϵ -approximate first-order criticality for unconstrained and convexly-constrained problem can be reduced to $O(\epsilon^{-(p+1)/p})$ if one is ready to define the step by using a regularization model of order $p \geq 1$. In the unconstrained case, this was shown for $p = 2$ in [12, 39] and for general $p \geq 1$ in [6], while the convexly constrained case was analyzed (for $p = 2$) in [13]. Whether this methodology and the associated improvements in evaluation complexity bounds can be extended to order above one is the subject of ongoing research.

Acknowledgements

The authors would like to thank Oliver Stein for suggesting reference [33], as well as Jim Burke and Adrian Lewis for interesting discussions. The third author would also like to acknowledge the support provided by the Belgian Fund for Scientific Research (FNRS), the Leverhulme Trust (UK), Balliol College (Oxford), the Department of Applied Mathematics of the Hong Kong Polytechnic University, ENSEEIHT (Toulouse, France) and INDAM (Florence, Italy).

References

- [1] A. Anandkumar and R. Ge. Efficient approaches for escaping high-order saddle points in nonconvex optimization. [arXiv.1602.05908](https://arxiv.org/abs/1602.05908), 2016.

⁽⁷⁾Using (3.29).

⁽⁸⁾Using (3.30).

- [2] W. Bian and X. Chen. Worst-case complexity of smoothing quadratic regularization methods for non-Lipschitzian optimization. *SIAM Journal on Optimization*, 23(3):1718–1741, 2013.
- [3] W. Bian and X. Chen. Linearly constrained non-Lipschitzian optimization for image restoration. Technical report, Department of Applied Mathematics, Polytechnic University of Hong Kong, Hong Kong, 2015.
- [4] W. Bian, X. Chen, and Y. Ye. Complexity analysis of interior point algorithms for non-Lipschitz and nonconvex minimization. *Mathematical Programming, Series A*, 149:301–327, 2015.
- [5] E. G. Birgin, J. L. Gardenghi, J. M. Martínez, S. A. Santos, and Ph. L. Toint. Evaluation complexity for nonlinear constrained optimization using unscaled KKT conditions and high-order models. *SIAM Journal on Optimization*, 26(2):951–967, 2016.
- [6] E. G. Birgin, J. L. Gardenghi, J. M. Martínez, S. A. Santos, and Ph. L. Toint. Worst-case evaluation complexity for unconstrained nonlinear optimization using high-order regularized models. *Mathematical Programming, Series A*, (to appear), 2016.
- [7] N. Boumal, P.-A. Absil, and C. Cartis. Global rates of convergence for nonconvex optimization on manifolds. Technical Report Technical Report arXiv:1605.08101, Oxford University, Oxford, UK, 2016.
- [8] O. A. Brezhneva and A. Tret'yakov. Optimality conditions for degenerate extremum problems with equality constraints. *SIAM Journal on Control and Optimization*, 42(2):729–743, 2003.
- [9] O. A. Brezhneva and A. Tret'yakov. The p th-order optimality conditions for inequality constrained optimization problems. *Nonlinear Analysis*, 63:1357–1366, 2005.
- [10] C. Cartis, N. I. M. Gould, and Ph. L. Toint. On the complexity of steepest descent, Newton's and regularized Newton's methods for nonconvex unconstrained optimization. *SIAM Journal on Optimization*, 20(6):2833–2852, 2010.
- [11] C. Cartis, N. I. M. Gould, and Ph. L. Toint. Adaptive cubic overestimation methods for unconstrained optimization. Part I: motivation, convergence and numerical results. *Mathematical Programming, Series A*, 127(2):245–295, 2011.
- [12] C. Cartis, N. I. M. Gould, and Ph. L. Toint. Adaptive cubic overestimation methods for unconstrained optimization. Part II: worst-case function-evaluation complexity. *Mathematical Programming, Series A*, 130(2):295–319, 2011.
- [13] C. Cartis, N. I. M. Gould, and Ph. L. Toint. An adaptive cubic regularization algorithm for nonconvex optimization with convex constraints and its function-evaluation complexity. *IMA Journal of Numerical Analysis*, 32(4):1662–1695, 2012.
- [14] C. Cartis, N. I. M. Gould, and Ph. L. Toint. Complexity bounds for second-order optimality in unconstrained optimization. *Journal of Complexity*, 28:93–108, 2012.
- [15] C. Cartis, N. I. M. Gould, and Ph. L. Toint. On the complexity of finding first-order critical points in constrained nonlinear optimization. *Mathematical Programming, Series A*, 144(1):93–106, 2013.
- [16] C. Cartis, N. I. M. Gould, and Ph. L. Toint. On the evaluation complexity of cubic regularization methods for potentially rank-deficient nonlinear least-squares problems and its relevance to constrained nonlinear optimization. *SIAM Journal on Optimization*, 23(3):1553–1574, 2013.
- [17] C. Cartis, N. I. M. Gould, and Ph. L. Toint. On the evaluation complexity of constrained nonlinear least-squares and general constrained nonlinear optimization using second-order methods. *SIAM Journal on Numerical Analysis*, 53(2):836–851, 2015.
- [18] C. Cartis, Ph. R. Sampaio, and Ph. L. Toint. Worst-case complexity of first-order non-monotone gradient-related algorithms for unconstrained optimization. *Optimization*, 64(5):1349–1361, 2015.
- [19] C. Cartis and K. Scheinberg. Global convergence rate analysis of unconstrained optimization methods based on probabilistic models. Technical Report TR 1887 (Maths E-Print Archive), Mathematical Institute, University of Oxford, Oxford, UK, 2015.
- [20] A. R. Conn, N. I. M. Gould, and Ph. L. Toint. *Trust-Region Methods*. MPS-SIAM Series on Optimization. SIAM, Philadelphia, USA, 2000.
- [21] F. E. Curtis, D. P. Robinson, and M. Samadi. A trust-region algorithm with a worst-case iteration complexity of $O(\epsilon^{-3/2})$ for nonconvex optimization. Technical Report COR@L 14T-009, Lehigh University, Bethlehem, USA, 2014.

- [22] M. Dodangeh, L. N. Vicente, and Z. Zhang. On the optimal order of worst case complexity of direct search. *Optimization Letters*, pages 1–10, June 2015.
- [23] J. P. Dussault. Simple unified convergence proofs for the trust-region and a new ARC variant. Technical report, University of Sherbrooke, Sherbrooke, Canada, 2015.
- [24] R. Garmanjani, D. Júdice, and L. N. Vicente. Trust-region methods without using derivatives: Worst case complexity and the non-smooth case. Technical Report 15-03, Department of mathematics, University of Coimbra, Coimbra, Portugal, 2015.
- [25] D. Ge, X. Jiang, and Y. Ye. A note on the complexity of L_p minimization. *Mathematical Programming, Series A*, 21:1721–1739, 2011.
- [26] S. Ghadimi and G. Lan. Accelerated gradient methods for nonconvex nonlinear and stochastic programming. *Mathematical Programming, Series A*, 156(1-2):59–100, 2016.
- [27] G. N. Grapiglia, J. Yuan, and Y. Yuan. Nonlinear stepsize control algorithms: Complexity bounds for first and second-order optimality. Technical report, Departamento de mathematica, Universidade Federal do Paraná, Centro Politécnico, Curitiba, Paraná, Brazil, 2015.
- [28] G. N. Grapiglia, J. Yuan, and Y. Yuan. On the convergence and worst-case complexity of trust-region and regularization methods for unconstrained optimization. *Mathematical Programming, Series A*, 152:491–520, 2015.
- [29] S. Gratton, C. W. Royer, L. N. Vicente, and Z. Zhang. Direct search based on probabilistic descent. *SIAM Journal on Optimization*, 25(3):1515–1541, 2015.
- [30] S. Gratton, A. Sartenaer, and Ph. L. Toint. Recursive trust-region methods for multiscale nonlinear optimization. *SIAM Journal on Optimization*, 19(1):414–444, 2008.
- [31] H. Hancock. *The Theory of Maxima and Minima*. The Athenaeum Press, Ginn & Co, New York, USA, 1917. Available on line at <https://archive.org/details/theoryofmaximami00hancuoft>.
- [32] J.-B. Hiriart-Urruty and C. Lemaréchal. *Convex Analysis and Minimization Algorithms. Part 1: Fundamentals*. Springer Verlag, Heidelberg, Berlin, New York, 1993.
- [33] W. Hogan. Point-to-set maps in mathematical programming. *SIAM Review*, 15(3):591–603, 1973.
- [34] F. Jarre. On Nesterov’s smooth Chebyshev-Rosenbrock function. *Optimization Methods and Software*, 28(3):478–500, 2013.
- [35] S. Lu, Z. Wei, and L. Li. A trust-region algorithm with adaptive cubic regularization methods for nonsmooth convex minimization. *Computational Optimization and Applications*, 51:551–573, 2012.
- [36] J. M. Martínez and M. Raydan. Cubic-regularization counterpart of a variable-norm trust-region method for unconstrained minimization. Technical report, Department of Mathematics, IMECC-UNICAMP, University of Campinas, Campinas, Brazil, November 2015.
- [37] J. J. Moreau. Décomposition orthogonale d’un espace hilbertien selon deux cônes mutuellement polaires. *Comptes-Rendus de l’Académie des Sciences (Paris)*, 255:238–240, 1962.
- [38] Yu. Nesterov. *Introductory Lectures on Convex Optimization*. Applied Optimization. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2004.
- [39] Yu. Nesterov and B. T. Polyak. Cubic regularization of Newton method and its global performance. *Mathematical Programming, Series A*, 108(1):177–205, 2006.
- [40] J. Nocedal and S. J. Wright. *Numerical Optimization*. Series in Operations Research. Springer Verlag, Heidelberg, Berlin, New York, 1999.
- [41] G. Peano. *Calcolo differenziale e principii di calcolo integrale*. Fratelli Bocca, Roma, Italy, 1884.
- [42] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, USA, 1970.
- [43] K. Scheinberg and X. Tang. Complexity in inexact proximal Newton methods. Technical report, Lehigh University, Bethlehem, USA, 2014.
- [44] K. Ueda and N. Yamashita. Convergence properties of the regularized Newton method for the unconstrained nonconvex optimization. *Applied Mathematics & Optimization*, 62(1):27–46, 2009.
- [45] K. Ueda and N. Yamashita. Convergence properties of the regularized Newton method for the unconstrained nonconvex optimization. *Applied Mathematics and Optimization*, 62(1):27–46, 2010.
- [46] K. Ueda and N. Yamashita. On a global complexity bound of the Levenberg-Marquardt method. *Journal of Optimization Theory and Applications*, 147:443–453, 2010.
- [47] L. N. Vicente. Worst case complexity of direct search. *EURO Journal on Computational Optimization*, 1:143–153, 2013.

Appendix: Non-singularity of M_q

We prove the non-singularity of the matrix M_q introduced in (5.16). Assume, for the purpose of a contradiction that there exists a nonzero vector $v = (c_{q+1,k}, \dots, c_{2q+1,k})^T \in \mathbb{R}^{q+1}$ such that $M_q v = 0$. From the argument of Section 5, this amounts to saying that there exists a polynomial of the form (5.11) with one of the coefficients $c_{q+1,k}, \dots, c_{2q+1,k}$ being nonzero and which satisfies the interpolation conditions (5.12) (i.e. (5.13) and (5.14)) with the restriction that r_k given by (5.15) is identically zero. Since $s_k > 0$, the fact that components 2 to q of r_k are zero implies that $\nabla_x^q f(x_k) = q! c_{q,k} = 0$, and hence (from the first component) that $\Delta f_k = 0$. The interpolation conditions thus specify that

$$\begin{aligned} p_k(0) &= \Delta f_k = 0, & p_k(s_k) &= 0; \\ \nabla_s^j p_k(0) &= 0 = \nabla_s^j p_k(s_k) & \text{for } j &\in \{1, \dots, q-1\}, \\ \nabla_s^q p_k(0) &= c_{q,k} = 0, & \nabla_s^q p_k(s_k) &= 0, \end{aligned}$$

where the last equality results from the fact that the last component of r_k is zero. Because $p_k(s)$ is nonzero, this implies that $p_k(s)$ must be of the form $A s^{q+1} (s - s_k)^{q+1} p_1(s)$ where A is a constant and $p_1(s)$ is a polynomial in s . But, since $p_k(s)$ is of degree $(2q+1)$ and $s^{q+1}(s - s_k)^{q+1}$ of degree $2q+2$, one must have that $p_1(s) = 0 = p_k(s)$, which is impossible. Hence M_q is non-singular.