





A concise second-order complexity analysis for unconstrained optimization using high-order regularized models

C. Cartis ^a, N. I. M. Gould ^b and Ph. L. Toint^c

^aMathematical Institute, Oxford University, Oxford, UK; ^bComputational Mathematics Group, STFC-Rutherford Appleton Laboratory, Chilton, UK; ^cDepartment of Mathematics, University of Namur, Namur, Belgium

ABSTRACT

An adaptive regularization algorithm is proposed that uses Taylor models of the objective of order $p, p \geq 2$, of the unconstrained objective function, and that is guaranteed to find a first- and second-order critical point in at most

$$O\left(\max\left\{\epsilon_1^{-\frac{p+1}{p}}, \epsilon_2^{-\frac{p+1}{p-1}}\right\}\right)$$

function and derivatives evaluations, where ϵ_1 and ϵ_2 are prescribed first- and second-order optimality tolerances. This is a simple algorithm and associated analysis compared to the much more general approach in Cartis *et al.* [*Sharp worst-case evaluation complexity bounds for arbitrary-order nonconvex optimization with inexpensive constraints*, arXiv:1811.01220, 2018] that addresses the complexity of criticality higher-than-two; here, we use standard optimality conditions and practical subproblem solves to show a same-order sharp complexity bound for second-order criticality. Our approach also extends the method in Birgin *et al.* [*Worst-case evaluation complexity for unconstrained nonlinear optimization using high-order regularized models*, Math. Prog. A 163(1) (2017), pp. 359–368] to finding second-order critical points, under the same problem smoothness assumptions as were needed for first-order complexity.

ARTICLE HISTORY

Received 4 June 2019
Accepted 1 October 2019

KEYWORDS

Nonconvex optimization;
regularization methods;
complexity analysis

1. Introduction

A question of general interest in computational optimization is to know how many evaluations of the functions that define a given problem are needed for an algorithm to find an estimate of a local minimizer. Considerable advances have been made on this topic, both for convex problems [16] and nonconvex ones [7]. Although much of this research has been devoted to the important issue of finding approximate first-order critical points, some authors have addressed the case where higher-order necessary optimality conditions must also be satisfied, and we review relevant literature below. In this paper, we streamline the recent approach and its analysis in [9] that addresses the much more general case of

CONTACT C. Cartis  coralia.cartis@maths.ox.ac.uk  Mathematical Institute, Oxford University, Oxford
OX2 6GG, UK

achieving q th order criticality and where the first improved and sharp complexity bound for second-order criticality was obtained using high-order models. Here, we use standard optimality measures and practical, local subproblem solves, as well as allow different tolerances on first- and second-order optimality criteria, leading to a concise complexity analysis. A more detailed comparison with [9] is given in the penultimate paragraph of this section.

We consider the unconstrained minimization of a twice continuously differentiable objective function $f : \mathbb{R}^n \rightarrow \mathbb{R}$. It is, of course, well known that a finite minimizer x_* of f necessarily satisfies the first- and second-order criticality conditions $\nabla_x f(x_*) = 0$ and $\lambda_{\text{left}}(\nabla_x^2 f(x_*)) \geq 0$, where λ_{left} denotes the leftmost eigenvalue of its symmetric matrix argument. Thus, a reasonable requirement might be to find a point x_k for which

$$\|\nabla_x f(x_k)\| \leq \epsilon_1 \quad \text{and} \quad \lambda_{\text{left}}(\nabla_x^2 f(x_k)) \geq -\epsilon_2 \quad (1)$$

for given, small $\epsilon_1, \epsilon_2 > 0$ and suitable norm $\|\cdot\|$.

The earliest analysis we are aware of that provides both first- and second-order evaluation complexity guarantees considers cubic regularization methods and shows that at most

$$O\left(\max\left\{\epsilon_1^{-3/2}, \epsilon_2^{-3}\right\}\right) \quad (2)$$

evaluations of f are required to satisfy (1) so long as the objective function is bounded from below, and its Hessian is Lipschitz continuous [17]. Cubic regularization variants with adaptive parameter choices and inexact subproblem solves, so-called Adaptive Regularization with Cubics (ARC), have similar complexity guarantees and were proposed in [5,6]. Under similar conditions, many trust-region (TR) algorithms require at most $O(\max\{\epsilon_1^{-2}, \epsilon_2^{-3}\})$ evaluations. Crucially, examples are known for which such order estimates are tight both for trust-region and regularization methods [6]. Of late, more sophisticated trust region methods and quadratic regularization ones have been proposed that echo the order of the ARC estimates [3,11,15]. At the same time, other methods [12,14] have been shown to mirror the TR-like evaluation estimate in a more general or simplified way, respectively.

The fact that the best-known evaluation bound for ARC is essentially tight, suggests that in order to do better, one needs to add further ingredients. A similar picture emerged for evaluation bounds for first-order critical points: improved bounds of order $O\left(\epsilon_1^{-\frac{p+1}{p}}\right)$, $p \geq 2$, were obtained in [2] for p -times continuously differentiable functions using regularization methods that employ higher-order local models. This will be the theme here. In order to improve upon the estimate (2) for second-order criticality, we will use a higher-order model and regularization. The model minimization conditions however, are approximate and local, for both first- and second-order criticality. This is in contrast with [9], where a more general high-order regularization framework is presented that can also achieve the much more challenging requirement of criticality of order higher-than two, using novel optimality conditions that employ the same accuracy requirement on all criticality orders. In particular, if the approach in [9] is applied to finding second-order critical points, the calculation of the criticality measure in [9, (2.6)], and the subproblem solution, require the (exact) global solution of trust-region (quadratic, possibly nonconvex)

subproblems, and it cannot enforce required accuracies on first- and second-order conditions separately. This measure is related to the left-most eigenvalue of the Hessian only at points that are exact first-order critical points (as detailed in [9, (2.7)–(2.9)]). The aim here is to use the criticality measures (1) that are standard and practical, and inexact variants of (1) for the model minimization, in order to present a dedicated high-order regularization algorithm with a simple analysis, that can achieve the sharp and best known evaluation complexity bounds for first- and second-order criticality under minimal requirements on both the objective and the algorithm.

In Section 2, we define terminology and propose our new algorithm, while in Section 3, we provide a convergence analysis that indicates an improved complexity bound. We provide further comments and perspectives in Section 4.

2. A regularized p th-order model and algorithm

Let $p \geq 2$. Consider the optimization problem

$$\min_{x \in \mathbb{R}^n} f(x), \quad (3)$$

where we assume that $f \in \mathcal{C}^{p,1}(\mathbb{R}^n)$, namely, that:

- f is p -times continuously differentiable;
- f is bounded below by f_{low}
- the p th derivative of f at x , the p th-order tensor

$$\nabla_x^p f(x) = \left[\frac{\partial^p f(x)}{\partial x_{i_1} \dots \partial x_{i_p}} \right]_{i_j \in \{1, \dots, n\}, j=1, \dots, p}$$

is globally Lipschitz continuous, that is, there exists a constant $L \geq 0$ such that, for all $x, y \in \mathbb{R}^n$,

$$\|\nabla_x^p f(x) - \nabla_x^p f(y)\|_{[p]} \leq (p-1)! L \|x - y\|. \quad (4)$$

In (4), $\|\cdot\|_{[p]}$ is the tensor norm recursively induced by the Euclidean norm $\|\cdot\|$ on the space of p th-order tensors, which is given by

$$\|T\|_{[p]} \stackrel{\text{def}}{=} \max_{\|v_1\|=\dots=\|v_p\|=1} |T[v_1, \dots, v_p]|, \quad (5)$$

where $T[v_1, \dots, v_j]$ stands for the tensor of order $p-j \geq 0$ resulting from the application of the p th-order tensor T to the vectors v_1, \dots, v_j ¹. Let $T_p(x, s)$ be the Taylor series of the function $f(x+s)$ at x truncated at order p

$$T_p(x, s) \stackrel{\text{def}}{=} f(x) + \sum_{j=1}^p \frac{1}{j!} \nabla_x^j f(x) [s]^j, \quad (6)$$

where the notation $T[s]^j$ stands for the tensor T applied j times to the vector s .

We shall use the following crucial bounds.

Lemma 2.1 (See Appendix A.1): Let $f \in C^{p,1}(\mathbb{R}^n)$, and $T_p(x, s)$ be the Taylor approximation of $f(x + s)$ about x . Then for all $x, s \in \mathbb{R}^n$,

$$f(x + s) \leq T_p(x, s) + \frac{L}{p} \|s\|^{p+1}, \quad (7)$$

$$\|\nabla_x^1 f(x + s) - \nabla_s^1 T_p(x, s)\|_{[1]} \leq L \|s\|^p \quad (8)$$

and

$$\|\nabla_x^2 f(x + s) - \nabla_s^2 T_p(x, s)\|_{[2]} \leq (p - 1)L \|s\|^{p-1}. \quad (9)$$

In order to describe our algorithm, we define the regularized Taylor series model

$$m(x, s, \sigma) \stackrel{\text{def}}{=} T_p(x, s) + \frac{\sigma}{p+1} \|s\|^{p+1}, \quad (10)$$

whose gradient and Hessian are

$$\nabla_s^1 m(x, s, \sigma) = \nabla_s^1 T_p(x, s) + \sigma \|s\|^p \frac{s}{\|s\|} \quad (11)$$

and

$$\nabla_s^2 m(x, s, \sigma) = \nabla_s^2 T_p(x, s) + \frac{\sigma}{p+1} \nabla_s^2 (\|s\|^{p+1}), \quad (12)$$

where

$$\nabla_s^2 (\|s\|^{p+1}) = (p+1) \left[(p-1) \|s\|^{p-3} s s^T + \|s\|^{p-1} I \right]. \quad (13)$$

Note that

$$m(x, 0, \sigma) = T_p(x, 0) = f(x). \quad (14)$$

For the objective function f , we define first- and second-order criticality measures as

$$\chi_{f,1}(x) \stackrel{\text{def}}{=} \|\nabla_x^1 f(x)\| \quad (15)$$

and

$$\chi_{f,2}(x) \stackrel{\text{def}}{=} \max \left\{ 0, -\lambda_f(x) \right\} = \max \left\{ 0, -\min_{\|y\|=1} \nabla_x^2 f(x)[y]^2 \right\} \quad (16)$$

where $\lambda_f(x) \stackrel{\text{def}}{=} \lambda_{\text{left}}[\nabla_x^2 f(x)]$. Similarly, for the model (10), we consider the measures

$$\chi_{m,1}(x, s, \sigma) \stackrel{\text{def}}{=} \|\nabla_s^1 m(x, s, \sigma)\| \quad (17)$$

and

$$\chi_{m,2}(x, s, \sigma) \stackrel{\text{def}}{=} \max \left\{ 0, -\lambda_m(x, s, \sigma) \right\} = \max \left\{ 0, -\min_{\|y\|=1} \nabla_s^2 m(x, s, \sigma)[y]^2 \right\} \quad (18)$$

where $\lambda_m(x, s, \sigma) \stackrel{\text{def}}{=} \lambda_{\text{left}}[\nabla_s^2 m(x, s, \sigma)]$.

The minimization algorithm we consider is now described in detail in Algorithm 2.1. Note that if the second-order conditions are removed—namely, the conditions for $i = 2$ in (20) and (22)—then this method reduces to the ARp algorithm in [2]. Furthermore, the high-order regularization algorithm in [9] for criticality of order possibly larger than two uses a combined, different measure in place of (20), and subproblem solves that require global model minimization (albeit over a ‘small’ neighbourhood) instead of the local and inexact common conditions (22).

Algorithm 2.1: ARp

Step 0: Step 0: Initialization. An initial point x_0 and an initial regularization parameter $\sigma_0 > 0$ are given, as well as an accuracy levels ϵ_1 and ϵ_2 . The constants θ , η_1 , η_2 , γ_1 , γ_2 , γ_3 and σ_{\min} are also given and satisfy

$$\theta > 0, \quad \sigma_{\min} \in (0, \sigma_0], \quad 0 < \eta_1 \leq \eta_2 < 1 \quad \text{and} \quad 0 < \gamma_1 < 1 < \gamma_2 < \gamma_3. \quad (19)$$

Compute $f(x_0)$ and set $k = 0$.

Step 1: Test for termination. Evaluate $\{\nabla_x^i f(x_k)\}_{i=1}^2$. If

$$\chi_{f,i}(x_k) \leq \epsilon_i \quad \text{for } i = 1, 2, \quad (20)$$

terminate with the approximate solution $x_\epsilon = x_k$. Otherwise compute derivatives of f from order 3 to p at x_k .

Step 2: Step calculation. Compute the step s_k by approximately minimizing the model $m(x_k, s, \sigma_k)$ with respect to s in the sense that the conditions

$$m(x_k, s_k, \sigma_k) < m(x_k, 0, \sigma_k) \quad (21)$$

and

$$\chi_{m,i}(x_k, s_k, \sigma_k) \leq \theta \|s_k\|^{p+1-i}, \quad (i = 1, 2) \quad (22)$$

hold.

Step 3: Acceptance of the trial point. Compute $f(x_k + s_k)$ and define

$$\rho_k = \frac{f(x_k) - f(x_k + s_k)}{T_p(x_k, 0) - T_p(x_k, s_k)}. \quad (23)$$

If $\rho_k \geq \eta_1$, then define $x_{k+1} = x_k + s_k$; otherwise define $x_{k+1} = x_k$.

Step 4: Regularization parameter update. Set

$$\sigma_{k+1} \in \begin{cases} [\max(\sigma_{\min}, \gamma_1 \sigma_k), \sigma_k] & \text{if } \rho_k \geq \eta_2, \\ [\sigma_k, \gamma_2 \sigma_k] & \text{if } \rho_k \in [\eta_1, \eta_2), \\ [\gamma_2 \sigma_k, \gamma_3 \sigma_k] & \text{if } \rho_k < \eta_1. \end{cases} \quad (24)$$

Increment k by one and go to Step 1 if $\rho_k \geq \eta_1$ or to Step 2 otherwise.

Each iteration of this algorithm requires the approximate minimization of $m(x_k, s, \sigma_k)$, and we note that conditions (21) and (22) are always achievable as they are satisfied at a

second-order critical point of $m(x, s, \sigma)$. Indeed, existing algorithms, such as the standard second-order trust-region method [10, Section 6.6] and ARC [4] will find such a point as the regularized Taylor model is both sufficiently smooth and bounded from below.² Moreover, *this approximate minimization does not involve additional computations of f nor its derivatives at points other than x_k , and therefore the precise method used, and the resulting effort spent, in Step 2 have no impact on the evaluation complexity.*³ Finally note that the second condition in (22) disappears if $\lambda_{\text{left}}(\nabla_x^2 T_p(x, s)) \geq 0$.

Iterations for which $\rho_k \geq \eta_1$ (and hence $x_{k+1} = x_k + s_k$) are called ‘successful’ and we denote by $\mathcal{S}_k \stackrel{\text{def}}{=} \{0 \leq j \leq k \mid \rho_j \geq \eta_1\}$ the index set of all successful iterations between 0 and k . We also denote the complement, \mathcal{U}_k , of \mathcal{S}_k in $\{0, \dots, k\}$, that corresponds to the index set of ‘unsuccessful’ iterations between 0 and k . Note that, before termination, each successful iteration requires the evaluation of f and its first p derivatives, while only the evaluation of f is needed at unsuccessful ones.

3. Complexity analysis

As it is typical for a complexity analysis of (regularization and other) methods, we proceed by showing lower bounds on the Taylor model decrease and on the length of the step at each iteration. The proofs of the next three lemmas is very similar to corresponding results in [2] and hence we defer the proofs to the appendix (but still include them for completeness, as the algorithm has changed).

Lemma 3.1: *The mechanism of Algorithm 2.1 guarantees that, for all $k \geq 0$,*

$$T_p(x_k, 0) - T_p(x_k, s_k) \geq \frac{\sigma_k}{p+1} \|s_k\|^{p+1}, \quad (25)$$

and so (23) is well-defined.

We next deduce a simple upper bound on the regularization parameter σ_k .

Lemma 3.2: *Let $f \in \mathcal{C}^{p,1}(\mathbb{R}^n)$. Then, for all $k \geq 0$,*

$$\sigma_k \leq \sigma_{\max} \stackrel{\text{def}}{=} \max \left\{ \sigma_0, \frac{\gamma_3 L(p+1)}{p(1-\eta_2)} \right\}. \quad (26)$$

Our next move, very much in the line of the theory proposed in [2,5], is to show that the step cannot be arbitrarily small compared with the gradient of the objective function at the trial point $x_k + s_k$.

Lemma 3.3: *Let $f \in \mathcal{C}^{p,1}(\mathbb{R}^n)$. Then, for all $k \geq 0$,*

$$\|s_k\| \geq \left(\frac{\chi_{f,1}(x_k + s_k)}{L + \theta + \sigma_k} \right)^{\frac{1}{p}}. \quad (27)$$

Next we show that the step cannot also be arbitrarily small compared to the second-order criticality measure (16) at the trial point $x_k + s_k$. This is the crucial novel ingredient of the paper, that is essential to the improved second-order complexity results.

Lemma 3.4: Let $f \in \mathcal{C}^{p,1}(\mathbb{R}^n)$. Then, for all $k \geq 0$,

$$\|s_k\| \geq \left(\frac{\chi_{f,2}(x_k + s_k)}{(p-1)L + \theta + p\sigma_k} \right)^{\frac{1}{p-1}}. \quad (28)$$

Proof: Using (10) and the fact that $\min_z[a(z) + b(z)] \geq \min_z[a(z)] + \min_z[b(z)]$, we find that

$$\begin{aligned} & \lambda_f(x_k + s_k) \\ &= \min_{\|y\|=1} \nabla_x^2 f(x_k + s_k)[y]^2 \\ &= \min_{\|y\|=1} \left(\nabla_x^2 f(x_k + s_k) - \nabla_s^2 T_p(x_k, s_k) - \frac{\sigma_k}{p+1} \nabla_s^2 \|s_k\|^{p+1} + \nabla_s^2 m(x_k, s_k, \sigma_k) \right) [y]^2 \\ &\geq \min_{\|y\|=1} \left(\nabla_x^2 f(x_k + s_k) - \nabla_s^2 T_p(x_k, s_k) \right) [y]^2 + \frac{\sigma_k}{p+1} \min_{\|y\|=1} \left(-\nabla_s^2 \|s_k\|^{p+1} \right) [y]^2 \\ &\quad + \min_{\|y\|=1} \nabla_s^2 m(x_k, s_k, \sigma_k) [y]^2. \end{aligned}$$

Considering each term in turn, and using (5) and (9), we see that

$$\begin{aligned} & \min_{\|y\|=1} \left(\nabla_x^2 f(x_k + s_k) - \nabla_s^2 T_p(x_k, s_k) \right) [y]^2 \\ &\geq \min_{\|y_1\|=\|y_2\|=1} \left(\nabla_x^2 f(x_k + s_k) - \nabla_s^2 T_p(x_k, s_k) \right) [y_1, y_2] \\ &\geq - \max_{\|y_1\|=\|y_2\|=1} \left| \left(\nabla_x^2 f(x_k + s_k) - \nabla_s^2 T_p(x_k, s_k) \right) [y_1, y_2] \right| \\ &= -\| \nabla_x^2 f(x_k + s_k) - \nabla_s^2 T_p(x_k, s_k) \|_{[2]} \\ &\geq -(p-1)L\|s_k\|^{p-1}, \end{aligned}$$

and using (13), we find that $\nabla_s^2(\|s_k\|^{p+1})[y]^2 = (p+1)[(p-1)\|s_k\|^{p-3}(s_k^T y)^2 + \|s_k\|^{p-1}\|y\|^2]$, and so

$$\min_{\|y\|=1} \left(-\nabla_s^2(\|s_k\|^{p+1}) \right) [y]^2 = - \max_{\|y\|=1} \nabla_s^2(\|s_k\|^{p+1}) [y]^2 = -p(p+1)\|s_k\|^{p-1}.$$

Recalling (18), we have $\min_{\|y\|=1} \nabla_s^2 m(x_k, s_k, \sigma_k) [y]^2 = \lambda_m(x_k, s_k, \sigma_k)$. This, and the last two displayed equations imply that

$$-\lambda_f(x_k + s_k) \leq (p-1)L\|s_k\|^{p-1} + p\sigma_k\|s_k\|^{p-1} - \min\{0, \lambda_m(x_k, s_k, \sigma_k)\}. \quad (29)$$

As the right hand side of (29) is nonnegative, the bound (29) can be re-written as

$$\max\{0, -\lambda_f(x_k + s_k)\} \leq [(p-1)L + p\sigma_k] \|s_k\|^{p-1} + \max\{0, -\lambda_m(x_k, s_k, \sigma_k)\}.$$

Combining the above with (16) and (18), and with (22) with $i = 2$, we conclude

$$\begin{aligned}\chi_{f,2}(x_k + s_k) &\leq ((p-1)L + p\sigma_k)\|s_k\|^{p-1} + \chi_{m,2}(x_k, s_k, \sigma_k) \\ &\leq ((p-1)L + \theta + p\sigma_k)\|s_k\|^{p-1},\end{aligned}$$

and so (28) follows. ■

We now bound the number of unsuccessful iterations as a function of the number of successful ones and include a proof in the Appendix.

Lemma 3.5 ([5, Theorem 2.1]): *The mechanism of Algorithm 2.1 guarantees that, if*

$$\sigma_k \leq \sigma_{\max}, \quad (30)$$

for some $\sigma_{\max} > 0$, then

$$k+1 \leq |\mathcal{S}_k| \left(1 + \frac{|\log \gamma_1|}{\log \gamma_2}\right) + \frac{1}{\log \gamma_2} \log \left(\frac{\sigma_{\max}}{\sigma_0}\right). \quad (31)$$

Using all the above results, we are now in position to state our main evaluation complexity result.

Theorem 3.6: *Let $f \in \mathcal{C}^{p,1}(\mathbb{R}^n)$. Then, given $\epsilon_1 > 0$ and $\epsilon_2 > 0$, Algorithm 2.1 needs at most*

$$\left\lceil \kappa_s(f(x_0) - f_{\text{low}}) \max \left\{ \epsilon_1^{-\frac{p+1}{p}}, \epsilon_2^{-\frac{p+1}{p-1}} \right\} \right\rceil + 1$$

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

$$\left\lceil \kappa_s(f(x_0) - f_{\text{low}}) \max \left\{ \epsilon_1^{-\frac{p+1}{p}}, \epsilon_2^{-\frac{p+1}{p-1}} \right\} \right\rceil \left(1 + \frac{|\log \gamma_1|}{\log \gamma_2}\right) + \frac{1}{\log \gamma_2} \log \left(\frac{\sigma_{\max}}{\sigma_0}\right) + 1$$

iterations in total to produce an iterate x_ϵ such that $\|\nabla_x^1 f(x_\epsilon)\| \leq \epsilon_1$ and $\lambda_{\text{left}}(\nabla_x^2 f(x_\epsilon)) \geq -\epsilon_2$, where σ_{\max} is given by (26) and where

$$\kappa_s \stackrel{\text{def}}{=} \frac{p+1}{\eta_1 \sigma_{\min}} \max \left\{ (L + \theta + \sigma_{\max})^{\frac{p+1}{p}}, ((p-1)L + \theta + p\sigma_{\max})^{\frac{p+1}{p-1}} \right\}.$$

Proof: At each successful iteration k before termination, either the first-order or the second-order approximate optimality condition must fail (at the next iteration), namely,

$$\chi_{f,1}(x_{k+1}) > \epsilon_1 \text{ or } \chi_{f,2}(x_{k+1}) > \epsilon_2, \quad (32)$$

and we also have the guaranteed decrease

$$f(x_k) - f(x_{k+1}) \geq \eta_1 (T_p(x_k, 0) - T_p(x_k, s_k)) \geq \frac{\eta_1 \sigma_{\min}}{p+1} \|s_k\|^{p+1} \quad (33)$$

where we used (23), (25) and (24). For any successful iteration for which the first condition in (32) holds, we deduce from (33), (27) and (26) that

$$f(x_k) - f(x_{k+1}) \geq \kappa_1 \epsilon_1^{\frac{p+1}{p}} \text{ where } \kappa_1 \stackrel{\text{def}}{=} \frac{\eta_1 \sigma_{\min}}{p+1} \left(\frac{1}{L + \theta + \sigma_{\max}} \right)^{\frac{p+1}{p}}. \quad (34)$$

Similarly, for any successful iteration for which the second condition in (32) holds, we deduce from (33), (28) and (26) that

$$f(x_k) - f(x_{k+1}) \geq \kappa_2 \epsilon_2^{\frac{p+1}{p-1}} \text{ where } \kappa_2 \stackrel{\text{def}}{=} \frac{\eta_1 \sigma_{\min}}{p+1} \left(\frac{1}{(p-1)L + \theta + p\sigma_{\max}} \right)^{\frac{p+1}{p-1}}. \quad (35)$$

Thus on any successful iteration until termination, we can guarantee the minimal of the two decreases in (34) and (35), and hence, since $\{f(x_k)\}$ decreases monotonically,

$$f(x_0) - f(x_{k+1}) \geq \min\{\kappa_1, \kappa_2\} \min \left\{ \epsilon_1^{\frac{p+1}{p}}, \epsilon_2^{\frac{p+1}{p-1}} \right\} \cdot |\mathcal{S}_k|.$$

Using that f is bounded below by f_{low} , we conclude

$$|\mathcal{S}_k| \leq \frac{f(x_0) - f_{\text{low}}}{\min\{\kappa_1, \kappa_2\}} \max \left\{ \epsilon_1^{-\frac{p+1}{p}}, \epsilon_2^{-\frac{p+1}{p-1}} \right\}$$

until termination, from which the desired bound on the number of successful iterations follows. Lemma 3.5 is then invoked to compute the upper bound on the total number of iterations. ■

Observe that we may modify the algorithm to seek only first-order points by restricting (22) to $i = 1$. The corresponding complexity is then

$$O \left(\epsilon_1^{-\frac{p+1}{p}} \right),$$

which coincides with the bound in [2]. Moreover, the same complexity result holds if, by chance, $\lambda_{\text{left}}(\nabla_x^2 f(x_k)) \geq -\epsilon_2$ for all iterations. By contrast, if ϵ_1 is so large that $\|\nabla_x^1 f(x_k)\| \leq \epsilon_1$ at every iteration, the complexity is

$$O \left(\epsilon_2^{-\frac{p+1}{p-1}} \right)$$

to find a point with a sufficiently large leftmost eigenvalue.

4. Final comments

Our goal has been to devise a simple algorithm that can be guaranteed to find an approximate first- and second-order critical point in fewer evaluations than the best-known current champions and with the simplest requirements on the problem and the algorithm. The new algorithm we have designed finds such a point in at most

$$O\left(\max\left\{\epsilon_1^{-\frac{p+1}{p}}, \epsilon_2^{-\frac{p+1}{p-1}}\right\}\right)$$

function and derivative evaluations under suitable differentiability and Lipschitz continuity conditions. When $p = 2$, we recover the standard best bound (2), while for $p = 3$, this improves to $O(\max\{\epsilon_1^{-4/3}, \epsilon_2^{-2}\})$ function and derivative evaluations, and approaches $O(\max\{\epsilon_1^{-1}, \epsilon_2^{-1}\})$ evaluations as p increases to infinity. Of course, this comes at an increased cost of requiring derivatives of order up to p , and of needing to approximately solve a potentially harder step subproblem. Note though, that the conditions (21) and (22) for model minimization are only local ones, and that the improved second-order approximate criticality result is achieved under the same problem assumptions as the first-order one (in [2] and here). Furthermore, the approach here offers a simple alternative to the more generally applicable approach in [9].

In practice, the test (20) for termination in Step 1 of Algorithm 2.1 would be arranged to check one of the pair of required inequalities, and only to check the other if the first holds (the order is immaterial). One could imagine a variant of the algorithm in which failure of one (but not both) of (20) might influence the requirement for the next step calculation/model minimization. Specifically, if $\chi_{f,1}(x_k) > \epsilon_1$, one might simply require that $\chi_{m,1}(x_k, s_k, \sigma_k) \leq \theta \|s_k\|^p$ rather than (22) as this alone would aim to improve first-order criticality. However, though this decoupling is possible both in practice and in the analysis, it is not as straightforward as in the case of say, trust-region methods [14], as the lower bounds on the step in (27) and (28) depend on the objective's gradient and Hessian value at the next trial point/iterate, not the current x_k . Also, one might modify the ARp algorithm to check the optimality measures (20) at every trial point, not just successful ones. This may allow earlier termination but possibly at an unsuccessful step and at increased first- and second-derivatives evaluation cost.

Extending the approach here to the constrained case, even convex constraints, seems challenging as the connection between model eigenvalues and function eigenvalues in a set is no longer straightforward. Another aspect for future work is quantifying the cost of the subproblem solution in a similar vein to recent works [1,13], where there is particular interest due to large scale applications, in quantifying the number of derivative actions required per iteration as derivatives cannot be stored/called explicitly. More generally, finding efficient ways to solve higher order polynomial models would bring ARp methods closer to practical use.

Notes

1. Note that $\|\cdot\|_{[1]} = \|\cdot\|$, the usual Euclidean vector norm.
2. When p is even, $m(x, s, \sigma)$ is smooth everywhere but at the origin, but a step from $s = 0$ in the steepest-descent/eigen direction will move to a region for which the model is always smooth.
3. We implicitly assume here that derivatives at x_k can be stored explicitly.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

The work of the first author was supported by the Alan Turing Institute, UK. The work of the second author was supported by EPSRC [grant number EP/M025179/1].

Notes on contributors

C. Cartis holds a PhD in mathematics from the University of Cambridge, and is currently Associate Professor in numerical optimization at the University of Oxford. She was awarded a Leslie Fox Prize in numerical analysis (second place). Her research interests cover numerical optimization with emphasis on methods and analysis for nonconvex problems, and applications in compressed sensing and climate modelling.

N. I. M. Gould, an STFC Senior Fellow and Professor of Numerical Optimization, has interests in theoretical and practical optimization, and links to linear algebra. He is a SIAM Fellow, a winner of the Leslie Fox and Beale-Orchard Hays prizes, and his most recent book is *Trust-region Methods* (SIAM, 2000), written jointly with Andy Conn and Philippe Toint.

Ph. L. Toint received his PhD from mathematics from the University of Namur (Belgium) and Cambridge (UK), after which he was appointed as assistant professor and then full professor in mathematics at the University of Namur (Belgium), where he is now emeritus. He was awarded the Beale-Orchard-Hays prize and the Lagrange prize for continuous optimization. He is an inaugural SIAM Fellow, and former chair of the Mathematical Optimisation Society (2010–2013). His research interests are smooth nonlinear optimization, with an emphasis on the algorithmic viewpoint, ranging from convergence theory to numerical considerations and software development, as well as practical and multidisciplinary applications of optimization techniques such as in transportation systems and planning.

ORCID

C. Cartis  <http://orcid.org/0000-0002-0963-5550>

N. I. M. Gould  <http://orcid.org/0000-0002-1031-1588>

References

- [1] N. Agarwal, Z. Allen-Zhu, B. Bullins, E. Hazan and T. Ma, *Finding Approximate Local Minima Faster Than Gradient Descent*. Princeton University, Technical Report, ArXiv: 1611.01146. 2016.
- [2] 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*, Math. Program. Ser. A 163(1) (2017), pp. 359–368.
- [3] E.G. Birgin and J.M. Martínez, *The use of quadratic regularization with a cubic descent condition for unconstrained optimization*, SIAM J. Optim. 27(2) (2017), pp. 1049–1074.
- [4] C. Cartis, N.I.M. Gould and Ph.L. Toint, *Adaptive cubic regularisation methods for unconstrained optimization. Part I: motivation, convergence and numerical results*, Math. Prog. Ser. A 127(2) (2011), pp. 245–295.
- [5] C. Cartis, N.I.M. Gould and Ph.L. Toint, *Adaptive cubic regularisation methods for unconstrained optimization. Part II: worst-case function and derivative-evaluation complexity*, Math. Prog. Ser. A 130(2) (2011), pp. 295–319.

- [6] C. Cartis, N.I.M. Gould and Ph.L. Toint, *Complexity bounds for second-order optimality in unconstrained optimization*, J. Complex. 28 (2012), pp. 93–108.
- [7] C. Cartis, N.I.M. Gould and Ph.L. Toint, *How much patience do you have? a worst-case perspective on smooth nonconvex optimization*, Optima 88 (2012), pp. 1–10.
- [8] C. Cartis, N.I.M. Gould and Ph.L. Toint, *Second-order optimality and beyond: characterization and evaluation complexity in nonconvex convexly-constrained optimization*, Found. Comput. Math 18(5) (2018), pp. 1073–1107.
- [9] C. Cartis, N.I.M. Gould and Ph.L. Toint, *Sharp worst-case evaluation complexity bounds for arbitrary-order nonconvex optimization with inexpensive constraints*. arXiv:1811.01220, 2018.
- [10] A.R. Conn, N.I.M. Gould and Ph.L. Toint, *Trust-Region Methods*, SIAM, Philadelphia, PA, 2000.
- [11] 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*, Math. Prog. Ser. A 162(1) (2017), pp. 1–32.
- [12] G.N. Grapiglia, J. Yuan and Y. Yuan, *Nonlinear stepsize control algorithms: Complexity bounds for first and second-order optimality*, J. Optim. Theory. Appl. 17(3) (2016), pp. 980–997.
- [13] Y. Carmon and J.C. Duchi, *Gradient Descent Efficiently Finds the Cubic-Regularized Nonconvex Newton Step*. Stanford University, Technical Report. ArXiv:1612.00547. 2016.
- [14] S. Gratton, C.W. Royer and L.N. Vicente, *A Decoupled First/Second-Order Steps Technique for Nonconvex Nonlinear Unconstrained Optimization With Improved Complexity Bounds*. Technical Report 17-21, Department of Mathematics, University of Coimbra. 2017.
- [15] J.M. Martínez and M. Raydan, *Cubic-regularization counterpart of a variable-norm trust-region method for unconstrained minimization*. J. Global Opt. 2016. DOI: 10.1007/s10898-016-0475-8.
- [16] Y. Nesterov, *Introductory Lectures on Convex Optimization*, Kluwer Academic Publishers, Dordrecht, 2004.
- [17] Yu. Nesterov and B.T. Polyak, *Cubic regularization of Newton method and its global performance*, Math. Prog. Ser. A 108(1) (2006), pp. 177–205.

Appendix

A.1 Proof of Lemma 2.1

As in [8], consider the Taylor identity

$$\phi(1) - \tau_k(1) = \frac{1}{(k-1)!} \int_0^1 (1-\xi)^{k-1} [\phi^{(k)}(\xi) - \phi^{(k)}(0)] d\xi \quad (\text{A1})$$

involving a given univariate C^k function $\phi(\alpha)$ and its k th-order Taylor approximation

$$\tau_k(\alpha) = \sum_{i=0}^k \phi^{(i)}(0) \frac{\alpha^i}{i!}$$

expressed in terms of the value $\phi^{(0)} = \phi$ and i th derivatives $\phi^{(i)}$, $i = 1, \dots, k$. Then, picking $\phi(\alpha) = f(x + \alpha s)$ and $k = p$, the identity

$$\int_0^1 (1-\xi)^{k-1} d\xi = \frac{1}{k}, \quad (\text{A2})$$

(4), (5) and (A1) imply that, for all $x, s \in \mathbb{R}^n$,

$$f(x+s) \leq T_p(x, s) + \frac{L}{p} \|s\|^{p+1}$$

[8, (2.8) with $L_{f,p} = (p-1)!L$] since $\tau_p(1) = T_p(x, s)$, which is the required (7).

Likewise, for an arbitrary unit vector v , selecting instead $\phi(\alpha) = \nabla_x^1 f(x + \alpha s)[v]$ and $k = p-1$, it follows from (A1) that

$$\begin{aligned} & (\nabla_x^1 f(x + s) - \nabla_s^1 T_p(x, s))[v] \\ &= \frac{1}{(p-2)!} \int_0^1 (1-\xi)^{p-2} (\nabla_x^p f(x + \xi s) - \nabla_x^p f(x))[s]^{p-1}[v] \, d\xi \end{aligned} \quad (\text{A3})$$

since $\tau_{p-1}(1) = \nabla_s^1 T_p(x, s)[v]$. Thus, using the symmetry of the derivative tensors, picking v to maximize the absolute value of the left-hand side of (A3) and using (A2), (5) and (4) successively, we obtain that

$$\begin{aligned} & \|\nabla_x^1 f(x + s) - \nabla_s^1 T_p(x, s)\|_{[1]} \\ &= \frac{1}{(p-2)!} \left| \int_0^1 (1-\xi)^{p-2} (\nabla_x^p f(x + \xi s) - \nabla_x^p f(x))[v] \left[\frac{s}{\|s\|} \right]^{p-1} \|s\|^{p-1} \, d\xi \right| \\ &\leq \frac{1}{(p-2)!} \left[\int_0^1 (1-\xi)^{p-2} \, d\xi \right] \max_{\xi \in [0,1]} \left| (\nabla_x^p f(x + \xi s) - \nabla_x^p f(x))[v] \left[\frac{s}{\|s\|} \right]^{p-1} \right| \|s\|^{p-1} \\ &\leq \frac{1}{(p-1)!} \max_{\xi \in [0,1]} \max_{\|w_1\|=\dots=\|w_p\|=1} \left| (\nabla_x^p f(x + \xi s) - \nabla_x^p f(x))[w_1, \dots, w_p] \right| \|s\|^{p-1} \\ &= \frac{1}{(p-1)!} \max_{\xi \in [0,1]} \|\nabla_x^p f(x + \xi s) - \nabla_x^p f(x)\|_{[p]} \|s\|^{p-1} \\ &\leq L \|s\|^p \end{aligned}$$

which gives (8).

Finally, for arbitrary unit vectors v_1 and v_2 , choosing $\phi(\alpha) = \nabla_x^2 f(x + \alpha s)[v_1, v_2]$ and $k = p-2$, the identity $\tau_{p-2}(1) = \nabla_s^2 T_p(x, s)[v_1, v_2]$ and (A1) together show that

$$\begin{aligned} & (\nabla_x^2 f(x + s) - \nabla_s^2 T_p(x, s))[v_1, v_2] \\ &= \frac{1}{(p-3)!} \int_0^1 (1-\xi)^{p-3} (\nabla_x^p f(x + \xi s) - \nabla_x^p f(x))[v_1, v_2][s]^{p-2} \, d\xi. \end{aligned} \quad (\text{A4})$$

As before, picking v_1 and v_2 to maximize the absolute value of the left-hand side of (A4),

$$\begin{aligned} & \|\nabla_x^2 f(x + s) - \nabla_s^2 T_p(x, s)\|_{[2]} \\ &= \frac{1}{(p-3)!} \left| \int_0^1 (1-\xi)^{p-3} (\nabla_x^p f(x + \xi s) - \nabla_x^p f(x))[v_1, v_2] \left[\frac{s}{\|s\|} \right]^{p-2} \|s\|^{p-2} \, d\xi \right| \\ &\leq \frac{1}{(p-3)!} \left[\int_0^1 (1-\xi)^{p-3} \, d\xi \right] \max_{\xi \in [0,1]} \left| (\nabla_x^p f(x + \xi s) - \nabla_x^p f(x))[v_1, v_2] \left[\frac{s}{\|s\|} \right]^{p-2} \right| \|s\|^{p-2} \\ &\leq \frac{1}{(p-2)!} \max_{\xi \in [0,1]} \max_{\|w_1\|=\dots=\|w_p\|=1} \left| (\nabla_x^p f(x + \xi s) - \nabla_x^p f(x))[w_1, \dots, w_p] \right| \|s\|^{p-2} \\ &= \frac{1}{(p-2)!} \max_{\xi \in [0,1]} \|\nabla_x^p f(x + \xi s) - \nabla_x^p f(x)\|_{[p]} \|s\|^{p-2} \\ &\leq (p-1)L \|s\|^{p-1} \end{aligned}$$

again using (4), (5) and (A2), which provides (9).

A.2 Proof of Lemmas in Section 3

Proof of Lemma 3.1: (See [2, Lemma 2.1]) Observe that, because of (21) and (10),

$$0 < m(x_k, 0, \sigma_k) - m(x_k, s_k, \sigma_k) = T_p(x_k, 0) - T_p(x_k, s_k) - \frac{\sigma_k}{p+1} \|s_k\|^{p+1}$$

which implies the desired bound. Note that $s_k \neq 0$ as long as we can satisfy condition (21), and so (25) implies (23) is well defined. ■

Proof of Lemma 3.2: (See [2, Lemma 2.2]) Assume that

$$\sigma_k \geq \frac{L(p+1)}{p(1-\eta_2)}. \quad (\text{A5})$$

Using (7) and (25), we may then deduce that

$$|\rho_k - 1| \leq \frac{|f(x_k + s_k) - T_p(x_k, s_k)|}{|T_p(x_k, 0) - T_p(x_k, s_k)|} \leq \frac{L(p+1)}{p\sigma_k} \leq 1 - \eta_2$$

and thus that $\rho_k \geq \eta_2$. Then iteration k is very successful in that $\rho_k \geq \eta_2$ and $\sigma_{k+1} \leq \sigma_k$. As a consequence, the mechanism of the algorithm ensures that (26) holds. ■

Proof of Lemma 3.3: (See [2, Lemma 2.3]) Using the triangle inequality, (8), (11) and (22) for $i = 1$, we obtain that

$$\begin{aligned} \chi_{f,1}(x_k + s_k) &\leq \|\nabla_x^1 f(x_k + s_k) - \nabla_s^1 T_p(x_k, s_k)\| + \left\| \nabla_s^1 T_p(x_k, s_k) + \sigma_k \|s_k\|^p \frac{s_k}{\|s_k\|} \right\| \\ &\quad + \sigma_k \|s_k\|^p \\ &= \|\nabla_x^1 f(x_k + s_k) - \nabla_s^1 T_p(x_k, s_k)\|_{[1]} + \chi_{m,1}(x_k, s_k, \sigma_k) + \sigma_k \|s_k\|^p \\ &\leq L \|s_k\|^p + \chi_{m,1}(x_k, s_k, \sigma_k) + \sigma_k \|s_k\|^p \\ &\leq [L + \theta + \sigma_k] \|s_k\|^p \end{aligned}$$

and (27) follows. ■

Proof of Lemma 3.5: The regularization parameter update (24) gives that, for each k ,

$$\gamma_1 \sigma_j \leq \max[\gamma_1 \sigma_j, \sigma_{\min}] \leq \sigma_{j+1}, \quad j \in \mathcal{S}_k, \quad \text{and} \quad \gamma_2 \sigma_j \leq \sigma_{j+1}, \quad j \in \mathcal{U}_k.$$

Thus we deduce inductively that

$$\sigma_0 \gamma_1^{|\mathcal{S}_k|} \gamma_2^{|\mathcal{U}_k|} \leq \sigma_k.$$

We therefore obtain, using (30), that

$$|\mathcal{S}_k| \log \gamma_1 + |\mathcal{U}_k| \log \gamma_2 \leq \log \left(\frac{\sigma_{\max}}{\sigma_0} \right),$$

which then implies that

$$|\mathcal{U}_k| \leq -|\mathcal{S}_k| \frac{\log \gamma_1}{\log \gamma_2} + \frac{1}{\log \gamma_2} \log \left(\frac{\sigma_{\max}}{\sigma_0} \right),$$

since $\gamma_2 > 1$. The desired result (31) then follows from the equality $k+1 = |\mathcal{S}_k| + |\mathcal{U}_k|$ and the inequality $\gamma_1 < 1$ given by (19). ■