

Improved second-order evaluation complexity for unconstrained nonlinear optimization using high-order regularized models

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

July 25th, 2017

Abstract

The unconstrained minimization of a sufficiently smooth objective function $f(x)$ is considered, for which derivatives up to order p , $p \geq 2$, are assumed to be available. An adaptive regularization algorithm is proposed that uses Taylor models of the objective of order p 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 $\epsilon_2 > 0$ are prescribed first- and second-order optimality tolerances. Our approach extends the method in Birgin et al. (2016) to finding second-order critical points, and establishes the novel complexity bound for second-order criticality under identical problem assumptions as for first-order, namely, that the p -th derivative tensor is Lipschitz continuous and that $f(x)$ is bounded from below. The evaluation-complexity bound for second-order criticality improves on all such known existing results.

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 [15] and nonconvex ones [6]. 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.

We consider the unconstrained minimization of a C^2 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

$$\|\nabla_x f(x_k)\| \leq \epsilon_1 \quad \text{and} \quad \lambda_{\text{left}}(\nabla_x^2 f(x_k)) \geq -\epsilon_2 \quad (1.1)$$

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

†Computational Mathematics Group, STFC-Rutherford Appleton Laboratory, Chilton OX11 0QX, England. Email: nick.gould@stfc.ac.uk. The work of this author was supported by EPSRC grant EP/M025179/1

‡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

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

$$O\left(\max\left(\epsilon_1^{-3/2}, \epsilon_2^{-3}\right)\right) \quad (1.2)$$

evaluations of f are required to satisfy (1.1) so long as the objective function is bounded from below, and its Hessian is Lipschitz continuous [16]. Adaptive cubic regularization variants with inexact subproblem solves and similar guarantees were proposed in [4, 5]. Under similar conditions, many trust-region (TR) algorithms require at most $O\left(\max\left(\epsilon_1^{-2}, \epsilon_2^{-3}\right)\right)$ evaluations. Crucially, examples are known for which such order estimates are tight both for trust-region and regularization methods [5]. Of late, more sophisticated trust region methods and quadratic regularization ones have been proposed that echo the order of the ARC estimates [9, 14, 2]. At the same time, other methods [10, 12] 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 $\mathcal{O}\left(\epsilon_1^{-\frac{p+1}{p}}\right)$, $p \geq 2$, were obtained in [1] 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 (1.2) 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.

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

2 A regularized p -th order model and algorithm

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

$$\min_{x \in \mathbb{R}^n} f(x), \quad (2.1)$$

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

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

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

$$\|\nabla_x^p f(x) - \nabla_x^p f(y)\|_{[p]} \leq (p-1)! L \|x - y\|. \quad (2.2)$$

In (2.2), $\|\cdot\|_{[p]}$ is the tensor norm recursively induced by the Euclidean norm $\|\cdot\|$ on the space of p -th order tensors, which is given by

$$\|T\|_{[p]} \stackrel{\text{def}}{=} \max_{\|v_1\|=\dots=\|v_p\|=1} |T[v_1, \dots, v_p]|, \quad (2.3)$$

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

$$T_p(x, s) \stackrel{\text{def}}{=} f(x) + \sum_{j=1}^p \frac{1}{j!} \nabla_x^j f(x) [s]^j, \quad (2.4)$$

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 $x, s \in \mathbb{R}^n$,

$$f(x + s) \leq T_p(x, s) + \frac{L}{p} \|s\|^{p+1}, \quad (2.5)$$

$$\|\nabla_x^1 f(x + s) - \nabla_s^1 T_p(x, s)\|_{[1]} \leq L \|s\|^p \quad (2.6)$$

and

$$\|\nabla_x^2 f(x + s) - \nabla_s^2 T_p(x, s)\|_{[2]} \leq (p - 1)L \|s\|^{p-1}. \quad (2.7)$$

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

$$m(x, s, \sigma) = T_p(x, s) + \frac{\sigma}{p+1} \|s\|^{p+1}, \quad (2.8)$$

whose gradient and Hessian are

$$\nabla_s^1 m(x, s, \sigma) = \nabla_s^1 T_p(x, s) + \sigma \|s\|^p \frac{s}{\|s\|} \quad (2.9)$$

and

$$\nabla_s^2 m(x, s, \sigma) = \nabla_s^2 T_p(x, s) + \frac{\sigma}{p+1} \nabla_s^2 (\|s\|^{p+1}), \quad (2.10)$$

where

$$\nabla_s^2 (\|s\|^{p+1}) = (p+1) [(p-1)\|s\|^{p-3} s s^T + \|s\|^{p-1} I]. \quad (2.11)$$

Note that

$$m(x, 0, \sigma) = T_p(x, 0) = f(x). \quad (2.12)$$

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

$$\chi_{f,1}(x) \stackrel{\text{def}}{=} \|\nabla_x^1 f(x)\| \quad (2.13)$$

⁽¹⁾Note that $\|\cdot\|_{[1]} = \|\cdot\|$, the usual Euclidean vector norm.

and

$$\chi_{f,2}(x) \stackrel{\text{def}}{=} \max \left[0, -\lambda_f(x) \right] = \max \left[0, -\min_{\|y\|=1} \nabla_x^2 f(x)[y]^2 \right] \quad (2.14)$$

where $\lambda_f(x) \stackrel{\text{def}}{=} \lambda_{\text{left}}[\nabla_x^2 f(x)]$. Similarly, for the model (2.8), we consider the measures

$$\chi_{m,1}(x, s, \sigma) \stackrel{\text{def}}{=} \|\nabla_s^1 m(x, s, \sigma)\| \quad (2.15)$$

and

$$\chi_{m,2}(x, s, \sigma) \stackrel{\text{def}}{=} \max \left[0, -\lambda_m(x, s, \sigma) \right] = \max \left[0, -\min_{\|y\|=1} \nabla_s^2 m(x, s, \sigma)[y]^2 \right] \quad (2.16)$$

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 (2.18) and (2.20) – then this method reduces to the ARp algorithm in [1].

Algorithm 2.1: ARp

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, ϵ_2 and ϵ_3 . The constants $\theta, \eta_1, \eta_2, \gamma_1, \gamma_2, \gamma_3$ and σ_{\min} are also given and satisfy

$$\theta > 0, \quad \sigma_{\min} \in (0, \sigma_0], \quad 0 < \eta_1 \leq \eta_2 < 1 \quad \text{and} \quad 0 < \gamma_1 < 1 < \gamma_2 < \gamma_3. \quad (2.17)$$

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

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

$$\chi_{f,i}(x_k) \leq \epsilon_i \quad \text{for } i = 1, 2, \quad (2.18)$$

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

$$m(x_k, s_k, \sigma_k) < m(x_k, 0, \sigma_k) \quad (2.19)$$

and

$$\chi_{m,i}(x_k, s_k, \sigma_k) \leq \theta \|s_k\|^{p+1-i}, \quad (i = 1, 2) \quad (2.20)$$

hold.

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

$$\rho_k = \frac{f(x_k) - f(x_k + s_k)}{T_p(x_k, 0) - T_p(x_k, s_k)}. \quad (2.21)$$

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

$$\sigma_{k+1} \in \begin{cases} [\max(\sigma_{\min}, \gamma_1 \sigma_k), \sigma_k] & \text{if } \rho_k \geq \eta_2, \\ [\sigma_k, \gamma_2 \sigma_k] & \text{if } \rho_k \in [\eta_1, \eta_2), \\ [\gamma_2 \sigma_k, \gamma_3 \sigma_k] & \text{if } \rho_k < \eta_1. \end{cases} \quad (2.22)$$

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 (2.19) and (2.20) 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 [8, §6.6] and ARC [3] 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 (2.20) 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 [1] and hence we defer the proofs to the appendix (but still include them for completeness, as the algorithm has changed).

⁽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.

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

$$T_p(x_k, 0) - T_p(x_k, s_k) \geq \frac{\sigma_k}{p+1} \|s_k\|^{p+1}, \quad (3.1)$$

and so (2.21) 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 $k \geq 0$,

$$\sigma_k \leq \sigma_{\max} \stackrel{\text{def}}{=} \max \left[\sigma_0, \frac{\gamma_3 L(p+1)}{p(1-\eta_2)} \right]. \quad (3.2)$$

Our next move, very much in the line of the theory proposed in [4, 1], 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 $k \geq 0$,

$$\|s_k\| \geq \left(\frac{\chi_{f,1}(x_k + s_k)}{L + \theta + \sigma_k} \right)^{\frac{1}{p}}. \quad (3.3)$$

Next we show that the step cannot also be arbitrarily small compared to the second order criticality measure (2.14) 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 $k \geq 0$,

$$\|s_k\| \geq \left(\frac{\chi_{f,2}(x_k + s_k)}{(p-1)L + \theta + p\sigma_k} \right)^{\frac{1}{p-1}}. \quad (3.4)$$

Proof. Using (2.8) 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 (2.3) and (2.7), we see that

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

and using (2.11), 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 (2.16), 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

$$-\lambda_f(x_k + s_k) \leq (p-1)L \|s_k\|^{p-1} + p\sigma_k \|s_k\|^{p-1} - \min[0, \lambda_m(x_k, s_k, \sigma_k)]. \quad (3.5)$$

As the right hand side of (3.5) is nonnegative, the bound (3.5) 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 (2.14) and (2.16), and with (2.20) 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 (3.4) follows. \square

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 [4, Theorem 2.1] The mechanism of Algorithm 2.1 guarantees that, if

$$\sigma_k \leq \sigma_{\max}, \quad (3.6)$$

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

$$k + 1 \leq |\mathcal{S}_k| \left(1 + \frac{|\log \gamma_1|}{\log \gamma_2} \right) + \frac{1}{\log \gamma_2} \log \left(\frac{\sigma_{\max}}{\sigma_0} \right). \quad (3.7)$$

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 (3.2) 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,

$$\chi_{f,1}(x_{k+1}) > \epsilon_1 \quad \text{or} \quad \chi_{f,2}(x_{k+1}) > \epsilon_2, \quad (3.8)$$

and we also have the guaranteed decrease

$$f(x_k) - f(x_{k+1}) \geq \eta_1 (T_p(x_k, 0) - T_p(x_k, s_k)) \geq \frac{\eta_1 \sigma_{\min}}{p+1} \|s_k\|^{p+1} \quad (3.9)$$

where we used (2.21), (3.1) and (2.22). For any successful iteration for which the first condition in (3.8) holds, we deduce from (3.9), (3.3) and (3.2) that

$$f(x_k) - f(x_{k+1}) \geq \kappa_1 \epsilon_1^{\frac{p+1}{p}} \quad \text{where} \quad \kappa_1 \stackrel{\text{def}}{=} \frac{\eta_1 \sigma_{\min}}{p+1} \left(\frac{1}{L + \theta + \sigma_{\max}} \right)^{\frac{p+1}{p}}. \quad (3.10)$$

Similarly, for any successful iteration for which the second condition in (3.8) holds, we deduce from (3.9), (3.4) and (3.2) that

$$f(x_k) - f(x_{k+1}) \geq \kappa_2 \epsilon_2^{\frac{p+1}{p-1}} \quad \text{where} \quad \kappa_2 \stackrel{\text{def}}{=} \frac{\eta_1 \sigma_{\min}}{p+1} \left(\frac{1}{(p-1)L + \theta + p\sigma_{\max}} \right)^{\frac{p+1}{p-1}}. \quad (3.11)$$

Thus on any successful iteration until termination we can guarantee the minimal of the two decreases in (3.10) and (3.11), 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. \square

Observe that we may modify the algorithm to seek only first-order points by restricting (2.20) to $i = 1$. The corresponding complexity is then

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

which coincides with the bound in [1]. 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 an algorithm that can guaranteed to find an approximate first- and second-order critical point in fewer evaluations than the best known current champions. 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 (1.2), while for $p = 3$, this improves to $O\left(\max\left(\epsilon_1^{-4/3}, \epsilon_2^{-2}\right)\right)$ function and derivative evaluations, and approaches $O\left(\max\left(\epsilon_1^{-1}, \epsilon_2^{-1}\right)\right)$ 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 (2.19) and (2.20) 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 [1] and here).

In practice, the test (2.18) 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 (2.18) 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 (2.20) 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 [12], as the lower bounds on the step in (3.3) and (3.4) 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 (2.18) 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.

While one might be tempted to try to provide bounds for an algorithm that guarantees approximate third- and higher-order necessary optimality conditions, we have not, as yet, been able to do so. The main sticking point has been that third-order necessary conditions involve the behaviour of the third-order term of the Taylor series in the nullspace of the Hessian (if it exists) and that this (typically proper) subspace of \mathbb{R}^n is highly sensitive. Its use or the use of an approximating set is therefore open to miss-diagnosis. Higher-order criticality becomes successively trickier; the critical spaces are then no longer subspaces but cones [7].

Extending the approach here to the constrained case, even convex constraints, also 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 [13, 11], 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.

References

- [1] 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*, 163(1):359–368, 2017.
- [2] E.G. Birgin and J.M. Martínez. The use of quadratic regularization with a cubic descent condition for unconstrained optimization. *SIAM Journal on Optimization*, 27(2):1049–1074, 2017.
- [3] C. Cartis, N. I. M. Gould, and Ph. L. Toint. Adaptive cubic regularisation methods for unconstrained optimization. Part I: motivation, convergence and numerical results. *Mathematical Programming, Series A*, 127(2):245–295, 2011.
- [4] 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. *Mathematical Programming, Series A*, 130(2):295–319, 2011.
- [5] 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.
- [6] 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*, 88:1–10, 2012.
- [7] C. Cartis, N. I. M. Gould, and Ph. L. Toint. Second-order optimality and beyond: characterization and evaluation complexity in nonconvex convexly-constrained optimization. Preprint RAL-P-2016-008, Rutherford Appleton Laboratory, Chilton, Oxfordshire, England, 2016.
- [8] A. R. Conn, N. I. M. Gould, and Ph. L. Toint. *Trust-Region Methods*. SIAM, Philadelphia, 2000.
- [9] 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. *Mathematical Programming, Series A*, 162(1):1–32, 2017.
- [10] G.N. Grapiglia, J. Yuan, and Y. Yuan. Nonlinear stepsize control algorithms: Complexity bounds for first and second-order optimality. *Journal of Optimization Theory and Applications*, 17(3):980–997, 2016.
- [11] Y. Carmon and J. C. Duchi. Gradient descent efficiently finds the cubic-regularized nonconvex Newton step. Stanford University, Technical Report. ArXiv:1612.00547, 2016.
- [12] S. Gratton, C. W. Royer, and L. N. Vicente. A decoupled first/second-order steps technique for non-convex nonlinear unconstrained optimization with improved complexity bounds. Technical Report 17-21, Department of Mathematics, University of Coimbra, 2017.
- [13] 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.

- [14] J. M. Martínez and M. Raydan. Cubic-regularization counterpart of a variable-norm trust-region method for unconstrained minimization. *Journal of Global Optimization*, 2016. DOI: 10.1007/s10898-016-0475-8.
- [15] Y. Nesterov. *Introductory lectures on convex optimization*. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2004.
- [16] Yu. Nesterov and B. T. Polyak. Cubic regularization of Newton method and its global performance. *Mathematical Programming, Series A*, 108(1):177–205, 2006.

Appendix A

A.1 Proof of Lemma 2.1

As in [7], consider the Taylor identity

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

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

$$\int_0^1 (1-\xi)^{k-1} d\xi = \frac{1}{k}, \quad (\text{A.2})$$

(2.2), (2.3) and (A.1) imply that, for all $x, s \in \mathbb{R}^n$,

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

[7, (2.8) with $L_{f,p} = (p-1)!L$] since $\tau_p(1) = T_p(x, s)$, which is the required (2.5).

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 (A.1) that

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

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 (A.3) and using (A.2), (2.3) and (2.2)

successively, we obtain that

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

which gives (2.6).

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 (A.1) together show that

$$\begin{aligned}
 & (\nabla_x^2 f(x+s) - \nabla_s^2 T_p(x, s))[v_1, v_2] \\
 &= \frac{1}{(p-3)!} \int_0^1 (1-\xi)^{p-3} (\nabla_x^p f(x+\xi s) - \nabla_x^p f(x))[v_1, v_2] [s]^{p-2} d\xi. \tag{A.4}
 \end{aligned}$$

As before, picking v_1 and v_2 to maximize the absolute value of the left-hand side of (A.4),

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

again using (2.2), (2.3) and (A.2). which provides (2.7).

A.2 Proof of Lemmas in Section 3

Proof of Lemma 3.1 (See [1, Lemma 2.1]) Observe that, because of (2.19) and (2.8),

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

which implies the desired bound. Note that $s_k \neq 0$ as long as we can satisfy condition (2.19), and so (3.1) implies (2.21) is well defined. \square

Proof of Lemma 3.2 (See [1, Lemma 2.2]) Assume that

$$\sigma_k \geq \frac{L(p+1)}{p(1-\eta_2)}. \tag{A.5}$$

Using (2.5) and (3.1), 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 (3.2) holds. \square

Proof of Lemma 3.3 (See [1, Lemma 2.3]) Using the triangle inequality, (2.6), (2.9) and (2.20) 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 (3.3) follows. \square

Proof of Lemma 3.5. The regularization parameter update (2.22) 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 (3.6), 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 (3.7) then follows from the equality $k + 1 = |\mathcal{S}_k| + |\mathcal{U}_k|$ and the inequality $\gamma_1 < 1$ given by (2.17). \square