



HAL
open science

Yet another fast variant of Newton's method for nonconvex optimization

Serge Gratton, Sadok Jerad, Philippe L. Toint

► **To cite this version:**

Serge Gratton, Sadok Jerad, Philippe L. Toint. Yet another fast variant of Newton's method for nonconvex optimization. 2023. hal-04269968

HAL Id: hal-04269968

<https://hal.science/hal-04269968>

Preprint submitted on 3 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Yet another fast variant of Newton’s method for nonconvex optimization

S. Gratton*, S. Jerad† and Ph. L. Toint‡

31 VII 2023

Abstract

A class of second-order algorithms is proposed for minimizing smooth nonconvex functions that alternates between regularized Newton and negative curvature steps in an iteration-dependent subspace. In most cases, the Hessian matrix is regularized with the square root of the current gradient and an additional term taking moderate negative curvature into account, a negative curvature step being taken only exceptionally. Practical variants have been detailed where the subspaces are chosen to be the full space, or Krylov subspaces. In the first case, the proposed method only requires the solution of a single linear system at nearly all iterations. We establish that at most $\mathcal{O}(|\log \epsilon| \epsilon^{-3/2})$ evaluations of the problem’s objective function and derivatives are needed for algorithms in the new class to obtain an ϵ -approximate first-order minimizer, and at most $\mathcal{O}(|\log \epsilon| \epsilon^{-3})$ to obtain a second-order one. Encouraging initial numerical experiments with two full-space and two Krylov-subspaces variants are finally presented.

Keywords: Newton’s method, nonconvex optimization, negative curvature, adaptive regularization methods, evaluation complexity.

1 Introduction

It is not an understatement to say that Newton’s method is a central algorithm to solve nonlinear minimization problems, mostly because the method exhibits a quadratic rate of convergence when close to the solution and is affine-invariant. In the worst case, it can however be as slow as a vanilla first-order method [10], [15, Section 3.2] even when globalized with a linesearch [43] or a trust region [18]. This drawback has however been circumvented by the cubic regularization algorithm [44] and its subsequent adaptive variants [11, 12], [15, Section 3.3]. For nonconvex optimization, these latter variants exhibit a worst-case $\mathcal{O}(\epsilon^{-3/2})$ complexity order to find an ϵ -first-order minimizer compared with the $\mathcal{O}(\epsilon^{-2})$ order of second-order trust-region methods [33], [15, Section 3.2]. Adaptive cubic regularization was later extended to handle inexact derivatives [49, 50, 2, 1], probabilistic models [1, 16], and even schemes in which the value of the objective function is never computed [31].

However, as noted in [41], the improvement in complexity has been obtained by trading the simple Newton step requiring only the solution of a single linear system for more complex or slower procedures, such as secular iterations, possibly using Lanczos preprocessing [9, 11] (see also [15, Chapters 8 to 10]) or (conjugate-)gradient descent [36, 6]. In the simpler context of convex

*Université de Toulouse, INP, IRIT, Toulouse, France. Email: serge.gratton@enseeiht.fr. Work partially supported by 3IA Artificial and Natural Intelligence Toulouse Institute (ANITI), French "Investing for the Future - PIA3" program under the Grant agreement ANR-19-PI3A-0004"

†ANITI, Université de Toulouse, INP, IRIT, Toulouse, France. Email: sadok.jerad@toulouse-inp.fr

‡NAXYS, University of Namur, Namur, Belgium. Email: philippe.toint@unamur.be. Partly supported by ANITI.

problems, two recent papers [41, 21] independently proposed another globalization technique. At an iterate x , the step s is computed as

$$s = -(\nabla_x^2 f(x) + \lambda_k I_n)^{-1} \nabla_x^1 f(x) \quad (1.1)$$

where $\lambda_k \sim \sqrt{\|\nabla_x^1 f(x)\|}$. This new approach exhibits the best complexity rate of second-order methods for convex optimization and retains the local superlinear convergence of standard Newton method, while showing remarkable numerical promise [41]. Devising an algorithm for nonconvex functions that can use similar ideas whenever possible appears as a natural extension.

In the nonconvex case, the Hessian may be indefinite and it is well-known that negative curvature can be exploited to ensure progress towards second-order points. Mixing gradient-related (possibly Newton) and negative curvature directions has long been considered and can be traced back to [40], which initiated a line of work using curvilinear search to find a step combining both types of directions. The length of the step is typically tuned using an Armijo-like condition [40, 24, 42]. Improvements were subsequently proposed by incorporating the curvilinear step in a nonmonotone algorithm [23], allowing the resolution of large-scale problems [39] or by choosing between the two steps based on model decrease [26]. Alternatively, negative curvature has also been used to regularize the Hessian matrix, yielding the famous Goldfeld-Quandt-Trotter (GQT) method [25]. Unfortunately, this method also involves more complex computation to find the step and has the same global convergence rate as first-order algorithms [48]. The negative curvature regularization was also the subject of the more recent paper [4], in which various Newton steps are tried at each iteration in order to ensure the optimal ARC [11, 12] global rate of convergence.

One may then wonder if it is possible to devise an adaptive second-order method using a single explicitly regularized Newton step when possible and a negative curvature direction only when necessary, with a near-optimal complexity rate. The objective of this paper is to show that it is indeed possible (and efficient). To this aim, we propose a fast Newton’s method that exploits negative curvature for nonconvex optimization problems and generalizes the method proposed in [41, 21] to the nonconvex case. The new algorithm automatically adjusts the regularization parameter (without knowledge of the Hessian’s Lipschitz constant). The method either uses an appropriately regularized Newton step taking the smallest negative eigenvalue of the Hessian also into account or simply follows the negative curvature otherwise. It first attempts a step along a direction regularized by the square root of the gradient only, as in the convex setting [21, 41]. In that sense, it is inspired by the “convex until proved guilty” strategy advocated by [7]. If this attempt fails, it obtains negative curvature information of the Hessian, which is then used either for regularization or to define a step along a negative-curvature direction. In what follows, all these operations are carried out in a specific, iteration-dependent subspace, whose choice leads to different algorithmic variants. We prove that these methods require at most $\mathcal{O}(|\log \epsilon| \epsilon^{-3/2})$ iterations and evaluations of the problem data to obtain an ϵ -approximate first-order critical point, which is very close to the optimal convergence rate of second-order methods for Lipschitz Hessian functions [13]. We also introduce an further algorithmic variant which is guaranteed to find a second-order critical point in at most $\mathcal{O}(|\log \epsilon| \epsilon^{-3})$ iterations.

The paper is organized as follows. Section 2 describes the general algorithmic framework and compares it with recent work on second-order methods. Section 3 states our assumptions and derives a bound on its worst-case complexity for finding first-order critical points. Section 4 presents the second-order algorithmic variant and states its complexity, the corresponding analysis being detailed in appendix. Section 5 then discusses some choices of the iteration-dependent subspace, including Krylov spaces. Section 6 finally illustrates the numerical behavior of the proposed methods. Conclusions are drawn in Section 7.

Notation The following notations will be used throughout the paper. Let $q \geq 1$. The symbol $\|\cdot\|$ denotes the Euclidean norm for vectors in \mathbb{R}^q and its associated subordinate norm for matrices. $\lambda_{\min}(M)$ and $\lambda_{\max}(M)$ denote the minimum and maximum eigenvalues of a symmetric matrix M , while I_q is the identity matrix in \mathbb{R}^q . For $x \in \mathbb{R}$, we define $[x]_+ = \max(x, 0)$. For two vectors $x, y \in \mathbb{R}^q$, $x^\top y$ denotes their inner product. The i -th column of I_q is denoted by e_i .

2 Adaptive Newton with Negative Curvature

We consider the problem of finding approximate first-order critical points of the smooth unconstrained nonconvex optimization problem

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

and discuss our algorithm called AN2C (for Adaptive Newton with Negative Curvature) on this page. The algorithm, whose purpose is to compute first-order critical points, is presented in the framework of adaptive regularization methods [3, 11] [15, Section 3.3] and proceeds as follows, using two subroutines RegStep and NewtonEigenStep.

Algorithm 2.1: Adaptive Newton with Negative Curvature (AN2C)

Step 0: Initialization An initial point $x_0 \in \mathbb{R}^n$, a regularization parameter $\sigma_0 > 0$ and a gradient accuracy threshold $\epsilon \in (0, 1]$ are given, as well as the parameters

$$\begin{aligned} \sigma_{\min} > 0, \kappa_C, \kappa_\theta > 0, \kappa_a, \kappa_b \geq 1, \varsigma_1 \in (0, 1), \varsigma_2 \in [0, \frac{1}{2}), \varsigma_3 \in [0, 1), \theta \in (0, 1], \\ 0 < \gamma_1 < 1 < \gamma_2 \leq \gamma_3 \text{ and } 0 < \eta_1 \leq \eta_2 < 1. \end{aligned}$$

Set $k = 0$.

Step 1: Check termination Evaluate $g_k \stackrel{\text{def}}{=} \nabla_x^1 f(x_k)$. Terminate if $\|g_k\| \leq \epsilon$.

Step 2: Compute subspace derivatives Choose $p \in \{1, \dots, n\}$ and form $V_p \in \mathbb{R}^{n \times p}$. Compute $\hat{g}_k \stackrel{\text{def}}{=} V_p^\top g_k$ and $\hat{H}_k \stackrel{\text{def}}{=} V_p^\top H_k V_p$ where $H_k \stackrel{\text{def}}{=} \nabla_x^2 f(x_k)$.

Step 3 (Optional): Attempt a regularization step

$$s_k = s_k^{def} = \text{RegStep}(\hat{g}_k, \hat{H}_k, V_p, \|g_k\|, \sigma_k, \kappa_a, \kappa_b, \varsigma_1, \varsigma_2). \quad (2.2)$$

If s_k^{def} has been successfully defined, go to Step 5.

Step 4 : Newton Step Computation

$$s_k = \text{NewtonEigenStep}(\hat{g}_k, \hat{H}_k, V_p, \|g_k\|, \sigma_k, \kappa_C, \kappa_b, \kappa_\theta, \varsigma_3, \theta). \quad (2.3)$$

Step 5: Acceptance ratio computation Evaluate $f(x_k + s_k)$ and compute the acceptance ratio

$$\rho_k = \frac{f(x_k) - f(x_k + s_k)}{-(g_k^\top s_k + \frac{1}{2} s_k^\top H_k s_k)}. \quad (2.4)$$

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

Step 6: Regularization parameter update Set

$$\sigma_{k+1} \in \begin{cases} [\max(\sigma_{\min}, \gamma_1 \sigma_k), \sigma_k] & \text{if } \rho_k \geq \eta_2, \\ [\sigma_k, \gamma_2 \sigma_k] & \text{if } \rho_k \in [\eta_1, \eta_2), \\ [\gamma_2 \sigma_k, \gamma_3 \sigma_k] & \text{if } \rho_k < \eta_1. \end{cases} \quad (2.5)$$

Increment k by one and go to Step 1.

Algorithm 2.2: $\text{RegStep}(\hat{g}_k, \hat{H}_k, V_p, \|g_k\|, \sigma_k, \kappa_a, \kappa_b, \varsigma_1, \varsigma_2)$

Attempt to solve the linear system

$$(\hat{H}_k + \sqrt{\kappa_a \sigma_k \|g_k\|} I_p) y_k^{def} = -\hat{g}_k. \quad (2.6)$$

If a solution y_k^{def} of this system can be obtained such that

$$(y_k^{def})^\top (\hat{H}_k + \sqrt{\kappa_a \sigma_k \|g_k\|} I_p) y_k^{def} > 0, \quad (2.7)$$

$$\|y_k^{def}\| \leq \frac{(1 + \kappa_\theta)}{\varsigma_1} \sqrt{\frac{\|g_k\|}{\kappa_a \sigma_k}}, \quad (2.8)$$

$$\|H_k V_p y_k^{def} + g_k\| \leq \kappa_b \|\hat{H}_k y_k^{def} + \hat{g}_k\|, \quad (2.9)$$

$$\|r_k^{def}\| \leq \min \left(\varsigma_2 \sqrt{\kappa_a \sigma_k \|g_k\|} \|y_k^{def}\|, \kappa_\theta \|\hat{g}_k\| \right) \quad (2.10)$$

where $r_k^{def} = (\hat{H}_k + \sqrt{\kappa_a \sigma_k \|g_k\|} I_p) y_k^{def} + \hat{g}_k$, then return $s_k^{def} \stackrel{\text{def}}{=} V_p y_k^{def}$.

Algorithm 2.3: $\text{NewtonEigenStep}(\hat{g}_k, \hat{H}_k, V_p, \|g_k\|, \sigma_k, \kappa_C, \kappa_b, \kappa_\theta, \varsigma_3, \theta)$

Step 1: Test negative curvature If $\lambda_{\min}(\hat{H}_k) \leq -\kappa_C \sqrt{\sigma_k \|g_k\|}$, go to Step 4.

Step 2: Newton Step Solve

$$\left(\hat{H}_k + (\sqrt{\sigma_k \|g_k\|} + [-\lambda_{\min}(\hat{H}_k)]_+) I_p \right) y_k^{neig} = -\hat{g}_k \quad (2.11)$$

to ensure the residual condition

$$\begin{aligned} \|r_k^{neig}\| &\stackrel{\text{def}}{=} \left\| \left(\hat{H}_k + (\sqrt{\sigma_k \|g_k\|} + [-\lambda_{\min}(\hat{H}_k)]_+) I_p \right) y_k^{neig} + \hat{g}_k \right\| \\ &\leq \min \left(\varsigma_3 \sqrt{\sigma_k \|g_k\|} \|y_k^{neig}\|, \kappa_\theta \|\hat{g}_k\| \right). \end{aligned} \quad (2.12)$$

Step 3: Check global quality of the solution If

$$\|H_k V_p y_k^{neig} + g_k\| \leq \kappa_b \|\hat{H}_k y_k^{neig} + \hat{g}_k\| \quad \text{then set } s_k = s_k^{neig} \stackrel{\text{def}}{=} V_p y_k^{neig}. \quad (2.13)$$

Else, go back AN2C[Step 2].

Step 4: Eigenvector direction Compute u_k such that

$$\hat{g}_k^\top u_k \leq 0, \quad \|u_k\| = 1 \quad \text{and} \quad u_k^\top \hat{H}_k u_k \leq \theta \lambda_{\min}(\hat{H}_k) \quad (2.14)$$

and set

$$s_k = s_k^{curv} \stackrel{\text{def}}{=} \frac{\theta \kappa_C \sqrt{\sigma_k \|g_k\|}}{\sigma_k} V_p u_k. \quad (2.15)$$

The selection of the iteration-dependent subspace defined as the range of V_p in Step 2 is of course crucial for the algorithm. At this stage of the algorithm description, cycling may possibly occur between Step 2 and (2.13) in Step 4, should the choice of the subspace be consistently inadequate. We will however discuss some practical choices in Section 5, for which this situation cannot happen. For our subsequent analysis, we therefore assume the following.

AS.0 For each iteration k , condition (2.13) is satisfied after finitely many choices of V_p . Moreover, there exists a constant $V_{\max} \geq 1$ such that

$$\|V_p\| \leq V_{\max} \text{ for all } p \in \{1, \dots, n\}. \quad (2.16)$$

After selecting the subspace⁽¹⁾ and projecting the current gradient and Hessian, we first attempt a step that avoids computing negative curvature information. Indeed, the s_k^{def} notation, where *def* stands for “definite”, in (2.6) makes the connection with the two conditions (2.6) and (2.7). The condition (2.7) is significantly less restrictive than checking the positive-definiteness of the regularized matrix in (2.6). This is at variance with the work of [5] where a factorization is required at each step, and coherent with the ‘capped-CG’ subroutine proposed at [51]. Should the problem be (locally) convex, (2.7) would automatically hold (see [41, 21]). The test (2.8) is required as to avoid steps whose magnitude is too large compared to the gradient (the motivation for its particular form of the test will become clear in Section 3).

When computing a vector s_k^{def} satisfying (2.7) to (2.10) is not possible, we (approximately) solve a linear system in \mathbb{R}^p (2.11) whose definition involves $[-\lambda_{\min}(\widehat{H}_k)]_+$. Even if an exact solution can be obtained at a marginal cost for small p , we still allow an approximate solution satisfying (2.12). We note that $[-\lambda_{\min}(\widehat{H}_k)]_+$ could have been replaced in (2.11) by $\kappa_C \sqrt{\sigma_k \|g_k\|}$ and the remainder of the complexity analysis would remain valid. An interesting connection can also be established between the regularization in (2.11) for $\widehat{H}_k = H_k$ and the GQT method [25], as the regularization parameter $(\sqrt{\sigma_k \|g_k\|} + [-\lambda_{\min}(H_k)]_+)$ is very similar in spirit to that used in this method. In the closely related algorithm of [4], a term μ is added to $[-\lambda_{\min}(H_k)]_+$ and multiple μ ’s are tested as to ensure ‘cubic’ descent. In our case, $\sqrt{\sigma_k \|g_k\|}$ directly yields a regularization of the desired order. Also observe that, in most cases, the “approximate minimum curvature direction” u_k is already available when computing $\lambda_{\min}(\widehat{H}_k)$. It can be also retrieved via a Lanczos procedure as proposed in [47, Lemma 9].

We now provide some comments that apply to the definition of both s_k^{def} and s_k^{neig} . Specifically, focusing on the latter, condition (2.13) serves to ensure the appropriateness of the subspace spanned by V_p . This condition guarantees that the projected residual (2.12) is sufficiently small compared to both the projected and unprojected gradients. In a more standard setting, where $V_p = I_n$, this condition simplifies to

$$\|(H_k + (\sqrt{\sigma_k \|g_k\|} + [-\lambda_{\min}(H_k)]_+)I_n)s_k + g_k\| \leq \min(\varsigma_3 \sqrt{\sigma_k \|g_k\|} \|s_k\|, \kappa_\theta \|g_k\|), \quad (2.17)$$

where the $\kappa_\theta \|g_k\|$ term is standard when devising truncated CG algorithms. The other term ensures the typical condition required for the approximate minimization of the cubic model m_k , namely that

$$\|\nabla_s^1 m_k(s_k)\| \leq \mathcal{O}(\|s_k\|^2). \quad (2.18)$$

This condition is typically used to derive the optimal complexity rate $\mathcal{O}(\epsilon^{-3/2})$, see [11, 3] and the references therein.

Finally we note that the condition of Step 1 in the `NewtonEigenStep` algorithm, which forces the negative curvature step (2.15), can be interpreted as the comparison of the minimal curvature of the quadratic $(\lambda_{\min}(V_p^\top H_k V_p))$ with the quantity $\sigma_k \sqrt{\|g_k\|/\sigma_k}$, which itself can be viewed as the curvature of the regularization term $\frac{1}{6}\sigma_k \|s\|^3$ for some s whose length $\sqrt{\|g_k\|/\sigma_k}$ is of the order of a standard regularized step (see [15, Lemma 3.3], for instance). The test thus ensures a “regularization-like” step when the quadratic’s negative curvature is strong enough to dominate that of the regularization for too small steps (see (2.29) below).

⁽¹⁾Since we do not specify at this point how to make this selection, AN2C may be viewed as a *class of algorithms* depending on the choice of V_p .

Once the step has been computed, the mechanisms for accepting/rejecting the new iterate (Step 5) and updating the regularization parameter (Step 6) are typical of adaptive regularization algorithms (see [3, 11] or [15, Section 3.3.1], for instance).

Before delving into the complexity analysis of AN2C, we further explore its fundamental properties and discuss its relationships with closely related nonconvex optimization algorithms. The method presented in [19] differs from AN2C in that it employs a gradient step followed by a negative curvature step. On the other hand, [38] adopts a condition-based approach to choose between gradient descent and negative curvature directions, relying on known smoothness parameters, while our methods remain fully adaptive. Another related approach is presented in [47], which, unlike AN2C, examines various conditions to select a specific direction (gradient, Newton, negative curvature) and performs a linesearch. Furthermore, [20] proposes a trust-region algorithm (in contrast to adaptive regularization) that tackles the trust-region subproblem using a combination of conjugate gradients and negative curvature. Notably, their condition on the residuals of this subproblem [20, Inequality (3.2)] can be related to (2.17).

Following well-established practice, we now define

$$\mathcal{S} \stackrel{\text{def}}{=} \{k \geq 0 \mid x_{k+1} = x_k + s_k\} = \{k \geq 0 \mid \rho_k \geq \eta_1\},$$

the set of indexes of “successful iterations”, and

$$\mathcal{S}_k \stackrel{\text{def}}{=} \mathcal{S} \cap \{0, \dots, k\},$$

the set of indexes of successful iterations up to iteration k . We further partition \mathcal{S}_k in three subsets depending on the nature of the step taken, so that

$$\mathcal{S}_k^{\text{neig}} \stackrel{\text{def}}{=} \mathcal{S}_k \cap \{s_k = s_k^{\text{neig}}\}, \quad \mathcal{S}_k^{\text{curv}} \stackrel{\text{def}}{=} \mathcal{S}_k \cap \{s_k = s_k^{\text{curv}}\}, \quad \mathcal{S}_k^{\text{def}} \stackrel{\text{def}}{=} \mathcal{S}_k \cap \{s_k = s_k^{\text{def}}\}.$$

We also recall a well-known result bounding the total number of iterations of adaptive regularization methods in terms of the number of successful ones.

Lemma 2.1 [3, Lemma 2.4], [15, Lemma 2.4.1] Suppose that the AN2C algorithm is used and that $\sigma_k \leq \sigma_{\max}$ for some $\sigma_{\max} > 0$. Then

$$k \leq |\mathcal{S}_k| \left(1 + \frac{|\log \gamma_1|}{\log \gamma_2} \right) + \frac{1}{\log \gamma_2} \log \left(\frac{\sigma_{\max}}{\sigma_0} \right). \quad (2.19)$$

Proof. Since Steps 5 and 6 of the AN2C algorithm are identical to Steps 3 and 4 of the algorithm presented on page 43 of [15] and since Lemma 2.4.1 in this reference only depends on these steps, the conclusion of the lemma remain valid for AN2C algorithm. \square

This result implies that the overall complexity of the algorithm can be estimated once bounds on σ_k and $|\mathcal{S}_k|$ are known, as we will show in the next section.

We now state a simple relation between $\|s_k^{\text{neig}}\|$, σ_k and $\|g_k\|$ inspired by [41].

Lemma 2.2 For all iterations k where s_k^{neig} is computed, we have that

$$\|s_k^{neig}\| \leq V_{\max} \|y_k^{neig}\| \quad (2.20)$$

$$\hat{g}_k = - \left(\hat{H}_k + (\sqrt{\sigma_k \|g_k\|} + [-\lambda_{\min}(\hat{H}_k)]_+) I_p \right) y_k^{neig} + r_k^{neig} \quad (2.21)$$

and

$$\|s_k^{neig}\| \leq (1 + \kappa_\theta) V_{\max}^{\frac{3}{2}} \sqrt{\frac{\|g_k\|}{\sigma_k}}. \quad (2.22)$$

Similarly, when s_k^{def} is computed,

$$\hat{g}_k = - \left(\hat{H}_k + \sqrt{\kappa_a \sigma_k \|g_k\|} I_p \right) y_k^{def} + r_k^{def}, \quad (2.23)$$

and

$$\|s_k^{def}\| \leq V_{\max} \|y_k^{def}\|. \quad (2.24)$$

At last, when s_k^{curv} is computed,

$$\|s_k^{curv}\| \leq V_{\max} \frac{\theta \kappa_C \sqrt{\sigma_k \|g_k\|}}{\sigma_k}. \quad (2.25)$$

Proof. First note from the second part of (2.13) and (2.16),

$$\|s_k^{neig}\| = \|V_p y_k^{neig}\| \leq V_{\max} \|y_k^{neig}\|$$

yielding (2.20). A similar proof can be followed to derive (2.24).

Equation (2.21) results from (2.11), (2.13) and the definition of the residual (2.12). Let us rewrite now (2.21) in function of \hat{g}_k and \hat{H}_k ,

$$\hat{g}_k = - \left(\hat{H}_k + (\sqrt{\sigma_k \|g_k\|} + [-\lambda_{\min}(\hat{H}_k)]_+) I_p \right) y_k^{neig} + r_k^{neig}.$$

From (2.21), the facts that $\hat{H}_k + (\sqrt{\sigma_k \|g_k\|} + [-\lambda_{\min}(\hat{H}_k)]_+) I_p$ is a positive definite matrix with

$$\lambda_{\min}(\hat{H}_k + (\sqrt{\sigma_k \|g_k\|} + [-\lambda_{\min}(\hat{H}_k)]_+) I_p) \geq \sqrt{\sigma_k \|g_k\|}$$

and that $\|r_k^{neig}\| \leq \kappa_\theta \|\hat{g}_k\|$ because of (2.12). We thus obtain that

$$\|y_k^{neig}\| \leq (1 + \kappa_\theta) \sqrt{\frac{\|\hat{g}_k\|}{\sigma_k}} \leq (1 + \kappa_\theta) \sqrt{\frac{V_{\max} \|g_k\|}{\sigma_k}}, \quad (2.26)$$

where the last inequality follows from (2.16). This last inequality and (2.20) give (2.22).

If $k \in \mathcal{S}_k^{def}$, (2.23) is obtained from (2.6) and the definition of the residual r_k^{def} .

Else if $k \in \mathcal{S}_k^{curv}$, (2.16), the fact that $\|u_k\| = 1$ and (2.15) give (2.25). \square

The next lemma gives a lower bound on the decrease of the local quadratic approximation. In standard adaptive regularization algorithms, this decrease automatically results from the minimization of the model (See [3] for instance). In our case, we need to use the properties of s_k^{def} , s_k^{curv} and s_k^{neig} to obtain the desired result.

Lemma 2.3 Let k be a successful an iteration of AN2C. If $k \in \mathcal{S}_k^{def}$, we have that

$$-\left(g_k^\top s_k + \frac{1}{2}s_k^\top H_k s_k\right) \geq \frac{1-2\varsigma_2}{2}\sqrt{\kappa_a\sigma_k\|g_k\|}\|y_k^{def}\|^2 \geq \frac{1-2\varsigma_2}{2}\sqrt{\kappa_a\sigma_k\|g_k\|}\frac{\|s_k\|^2}{V_{\max}^2}. \quad (2.27)$$

If $k \in \mathcal{S}_k^{neig}$, then

$$-\left(g_k^\top s_k + \frac{1}{2}s_k^\top H_k s_k\right) \geq (1-\varsigma_3)\sqrt{\sigma_k\|g_k\|}\|y_k^{neig}\|^2 \geq (1-\varsigma_3)\sqrt{\sigma_k\|g_k\|}\frac{\|s_k\|^2}{V_{\max}^2}. \quad (2.28)$$

Else, if $k \in \mathcal{S}_k^{curv}$,

$$-\left(g_k^\top s_k + \frac{1}{2}s_k^\top H_k s_k\right) \geq \frac{1}{2}\theta^3\kappa_C^3\frac{\|g_k\|^{\frac{3}{2}}}{\sqrt{\sigma_k}} \geq \frac{1}{2}\sigma_k\frac{\|s_k\|^3}{V_{\max}^3}. \quad (2.29)$$

Proof. Suppose first that $k \in \mathcal{S}_k^{def}$. We then obtain from (2.2), (2.10) and (2.8) that

$$\begin{aligned} g_k^\top s_k^{def} + \frac{1}{2}(s_k^{def})^\top H_k s_k^{def} &= (V_p^\top g_k)^\top y_k^{def} + \frac{1}{2}(y_k^{def})^\top V_p^\top H_k V_p y_k^{def} \\ &= (r_k^{def})^\top y_k^{def} - (y_k^{def})^\top (\widehat{H}_k + \sqrt{\kappa_a\sigma_k\|g_k\|}I_p) y_k^{def} \\ &\quad + \frac{1}{2}(y_k^{def})^\top \widehat{H}_k y_k^{def} \\ &= -\sqrt{\kappa_a\sigma_k\|g_k\|}\|y_k^{def}\|^2 + (r_k^{def})^\top y_k^{def} - \frac{1}{2}(y_k^{def})^\top \widehat{H}_k y_k^{def} \\ &\leq -\sqrt{\kappa_a\sigma_k\|g_k\|}\|y_k^{def}\|^2 + \varsigma_2\sqrt{\kappa_a\sigma_k\|g_k\|}\|y_k^{def}\|^2 \\ &\quad + \frac{1}{2}\sqrt{\kappa_a\sigma_k\|g_k\|}\|y_k^{def}\|^2. \end{aligned}$$

Hence (2.27) follows from (2.24).

Suppose now that $k \in \mathcal{S}_k^{neig}$. By using (2.21) and the fact that $\widehat{H}_k + [-\lambda_{\min}(\widehat{H}_k)]_+ I_p \succeq 0$,

$$\begin{aligned} g_k^\top s_k^{neig} + \frac{1}{2}(s_k^{neig})^\top H_k s_k^{neig} &= (V_p^\top g_k)^\top y_k^{neig} + \frac{1}{2}(y_k^{neig})^\top V_p^\top H_k V_p y_k^{neig} \\ &= (r_k^{neig})^\top y_k^{neig} - (y_k^{neig})^\top (\widehat{H}_k + [-\lambda_{\min}(\widehat{H}_k)]_+ I_p) y_k^{neig} \\ &\quad + \frac{1}{2}(y_k^{neig})^\top (\widehat{H}_k + [-\lambda_{\min}(\widehat{H}_k)]_+ I_p) y_k^{neig} \\ &\quad - \frac{1}{2}[-\lambda_{\min}(\widehat{H}_k)]_+ \|y_k^{neig}\|^2 - \sqrt{\sigma_k\|g_k\|}\|y_k^{neig}\|^2 \\ &= (r_k^{neig})^\top y_k^{neig} - \frac{1}{2}(y_k^{neig})^\top (\widehat{H}_k + [-\lambda_{\min}(\widehat{H}_k)]_+ I_p) y_k^{neig} \\ &\quad - \frac{1}{2}[-\lambda_{\min}(\widehat{H}_k)]_+ \|y_k^{neig}\|^2 - \sqrt{\sigma_k\|g_k\|}\|y_k^{neig}\|^2 \\ &\leq \varsigma_3\sqrt{\sigma_k\|g_k\|}\|y_k^{neig}\|^2 - \frac{1}{2}[-\lambda_{\min}(\widehat{H}_k)]_+ \|y_k^{neig}\|^2 - \sqrt{\sigma_k\|g_k\|}\|y_k^{neig}\|^2, \end{aligned}$$

where we have used (2.12) to obtain the last inequality. Rearranging, ignoring the $\frac{1}{2}[-\lambda_{\min}(\widehat{H}_k)]_+ \|y_k^{neig}\|^2$ term and using (2.20) yield (2.28).

Suppose finally that $k \in \mathcal{S}_k^{curv}$. As (2.14) and (2.15) hold and that Step 4 of `NewtonEigenStep` is

taken when $\lambda_{\min}(\widehat{H}_k) \leq -\kappa_C \sqrt{\sigma_k \|g_k\|}$, we deduce that

$$\begin{aligned} g_k^\top s_k^{curv} + \frac{1}{2} (s_k^{curv})^\top H_k s_k^{curv} &= g_k^\top V_p u_k + \frac{1}{2} (s_k^{curv})^\top H_k s_k^{curv} \\ &\leq \frac{1}{2} \frac{\theta^2 \kappa_C^2 \|g_k\|}{\sigma_k} u_k^\top \widehat{H}_k u_k \leq \frac{1}{2} \frac{\theta^3 \kappa_C^2 \|g_k\|}{\sigma_k} \lambda_{\min}(\widehat{H}_k) \\ &\leq -\frac{1}{2} \theta^3 \kappa_C^3 \frac{\|g_k\|^{\frac{3}{2}}}{\sqrt{\sigma_k}}, \end{aligned} \quad (2.30)$$

yielding the first inequality in (2.29). For the second inequality, remark that from (2.25), we derive that

$$\theta^3 \kappa_C^3 \frac{\|g_k\|^{\frac{3}{2}}}{\sqrt{\sigma_k}} = \sigma_k \theta^3 \kappa_C^3 \frac{\|g_k\|^{\frac{3}{2}}}{\sigma_k^{\frac{3}{2}}} \geq \sigma_k \frac{\|s_k\|^3}{V_{\max}^3},$$

injecting the last bound in (2.30) gives the second inequality in (2.29). \square

3 Complexity analysis for the AN2C algorithm

We now turn to analyzing the worst-case complexity of the AN2C algorithm. Our analysis is conducted under AS.0 and the following assumptions.

AS.1 The function f is two times continuously differentiable in \mathbb{R}^n .

AS.2 There exists a constant f_{low} such that $f(x) \geq f_{\text{low}}$ for all $x \in \mathbb{R}^n$.

AS.3 The Hessian of f is globally Lipschitz continuous, that is, there exists a non-negative constant L_H such that

$$\|\nabla_x^2 f(x) - \nabla_x^2 f(y)\| \leq L_H \|x - y\| \quad \text{for all } x, y \in \mathbb{R}^n. \quad (3.1)$$

AS.4 There exists a constant $\kappa_B > 0$ such that

$$\max(0, -\lambda_{\min}(\nabla_x^2 f(x))) \leq \kappa_B \quad \text{for all } x \in \{y \in \mathbb{R}^n \mid f(y) \leq f(x_0)\}.$$

AS.1-AS.3 are standard assumptions when analyzing algorithms that utilize second-order information [12, 3]. AS.4 is weaker than assuming bounded Hessians, a condition often used when theoretically analyzing second-order methods that combines negative curvature and gradient based directions [47, 19, 38]. The left-hand side of the inequality is sometimes called the ‘‘convex deviation’’ or ‘‘modulus of nonconvexity’’ [37]. As it turns out, AS.4 is only needed for x being any iterate x_k produced by the algorithm and these iterates all belong to the level associated with the starting point x_0 because the acceptance condition in Step 5 ensures that the sequence $\{f(x_k)\}$ is non-increasing. If this level set is bounded or if the sequence $\{x_k\}$ remains bounded for any other reason, we immediately obtain that

$$\max(0, -\lambda_{\min}(H_k)) \leq \kappa_B \quad \text{for all } k \geq 0 \quad (3.2)$$

for some $\kappa_B \geq 0$, and both AS.3 and AS.4 automatically hold.

Having established a lower bound on the decrease ratio in Lemma 2.3, we next proceed to derive an upper bound on the regularization parameter. This is a crucial step when analyzing adaptive regularization methods.

Lemma 3.1 Suppose that AS.1 and AS.3 hold. Then, for all $k \geq 0$,

$$\sigma_k \leq \sigma_{\max} \stackrel{\text{def}}{=} \gamma_3 \max \left(\sigma_0, \varsigma_{\max} \frac{L_H}{6(1-\eta_2)} \right), \quad (3.3)$$

where

$$\varsigma_{\max} \stackrel{\text{def}}{=} \max \left(\frac{(1+\kappa_\theta)V_{\max}^{\frac{7}{2}}}{(1-\varsigma_3)}, \frac{2(1+\kappa_\theta)V_{\max}^3}{\kappa_a \varsigma_1 (1-2\varsigma_2)}, 2V_{\max}^3 \right). \quad (3.4)$$

Proof. Let us compute the ratio ρ_k for $k \in \mathcal{S}_k^{\text{neig}}$. By using AS.3 and the standard error bound for Lipschitz approximation of the function (see [14, Lemma 2.1]), that $\varsigma_3 < 1$, (2.28) and (2.22), we obtain that

$$\begin{aligned} 1 - \rho_k &= \frac{f(x_k + s_k) - f(x_k) - g_k^\top s_k - \frac{1}{2} s_k^\top H_k s_k}{-(g_k^\top s_k + \frac{1}{2} s_k^\top H_k s_k)} \\ &\leq \frac{L_H V_{\max}^2 \|s_k^{\text{neig}}\|^3}{6(1-\varsigma_3) \sqrt{\sigma_k} \|g_k\| \|s_k^{\text{neig}}\|^2} \\ &\leq \frac{L_H V_{\max}^2 \|s_k^{\text{neig}}\|}{6(1-\varsigma_3) \sqrt{\sigma_k} \|g_k\|} \\ &\leq \frac{L_H (1+\kappa_\theta) V_{\max}^{\frac{7}{2}}}{6(1-\varsigma_3) \sigma_k}. \end{aligned} \quad (3.5)$$

Hence, if $\sigma_k \geq \frac{L_H (1+\kappa_\theta) V_{\max}^{\frac{7}{2}}}{6(1-\varsigma_3)(1-\eta_2)}$, then $\rho_k \geq \eta_2$, which implies that iteration k is successful and $\sigma_{k+1} \leq \sigma_k$ because of (2.5). The mechanism of (2.5) in the algorithm then ensures that

$$\sigma_k \leq \gamma_3 \max \left(\sigma_0, \frac{L_H (1+\kappa_\theta) V_{\max}^{\frac{7}{2}}}{6(1-\varsigma_3)(1-\eta_2)} \right). \quad (3.6)$$

Similarly, if $k \in \mathcal{S}_k^{\text{def}}$, we use AS.3, the Lipschitz approximation error bound, the fact that $\varsigma_2 < \frac{1}{2}$, (2.27), (2.8) and (2.24) to deduce that

$$1 - \rho_k \leq \frac{L_H \|s_k^{\text{def}}\| V_{\max}^2}{3(1-2\varsigma_2) \sqrt{\kappa_a \sigma_k} \|g_k\|} \leq \frac{L_H (1+\kappa_\theta) V_{\max}^3}{3\kappa_a \varsigma_1 (1-2\varsigma_2) \sigma_k}.$$

Using the same argument as above, we now obtain that

$$\sigma_k \leq \gamma_3 \max \left(\sigma_0, \frac{L_H (1+\kappa_\theta) V_{\max}^3}{3\kappa_a \varsigma_1 (1-2\varsigma_2)(1-\eta_2)} \right). \quad (3.7)$$

Consider finally the case where $k \in \mathcal{S}_k^{\text{curv}}$. Again using AS.3, the Lipschitz approximation error bound and (2.29) lower-bound, we derive that

$$1 - \rho_k = \frac{f(x_k + s_k) - f(x_k) - g_k^\top s_k - \frac{1}{2} s_k^\top H_k s_k}{-g_k^\top s_k - \frac{1}{2} s_k^\top H_k s_k} \leq \frac{L_H \|s_k^{\text{curv}}\|^3 V_{\max}^3}{6\frac{1}{2} \sigma_k \|s_k^{\text{curv}}\|^3} = \frac{L_H V_{\max}^3}{3\sigma_k},$$

so that

$$\sigma_k \leq \gamma_3 \max \left(\sigma_0, \frac{L_H V_{\max}^3}{3(1-\eta_2)} \right). \quad (3.8)$$

Combining (3.6), (3.7) and (3.8) gives (3.3) with ς_{\max} defined by (3.4). \square

We now prove a lower bound on the decrease at a successful iteration k using negative curvature. We will also bound the change in the norm $\|g_{k+1}\|$ in term of $\|g_k\|$, which will be useful later to bound the cardinal of a subset of $\mathcal{S}_k^{neg} \cup \mathcal{S}_k^{curv}$.

Lemma 3.2 Suppose that AS.1, AS.3 and AS.4 hold and that $k \in \mathcal{S}_k^{curv}$ before termination. Then

$$f(x_k) - f(x_{k+1}) \geq \frac{\eta_1 \theta^3 \kappa_C^3}{2\sqrt{\sigma_{\max}}} \epsilon^{\frac{3}{2}}, \quad (3.9)$$

and

$$\|g_{k+1}\| \leq \left(\frac{L_H V_{\max}^2}{2\sigma_k} \kappa_C^2 \theta^2 + \frac{\theta^2 \kappa_B \kappa_C}{\sqrt{\epsilon \sigma_k}} + 1 \right) \|g_k\|. \quad (3.10)$$

Proof. Let $k \in \mathcal{S}_k^{curv}$. From (2.4) and (2.29), we obtain that

$$f(x_k) - f(x_{k+1}) \geq \eta_1 \left(-g_k^\top s_k - \frac{1}{2} s_k^\top H_k s_k \right) \geq \frac{\eta_1 \theta^3 \kappa_C^3}{2\sqrt{\sigma_k}} \|g_k\|^{\frac{3}{2}}.$$

Since $\|g_k\| \geq \epsilon$ before termination and that $\sigma_k \leq \sigma_{\max}$ by Lemma 3.1, we obtain (3.9).

Let us now prove (3.10). By using the Lipschitz error bound for the gradient ([14, Lemma 2.1]), the triangular inequality, the fact that $k \in \mathcal{S}_k^{curv}$, (2.14), (2.15), and (2.25), we obtain that

$$\begin{aligned} \|g_{k+1}\| &\leq \|g_{k+1} - g_k - H_k s_k\| + \|H_k s_k + g_k\| \\ &\leq \frac{L_H}{2} \|s_k\|^2 + \|g_k\| + \|H_k s_k\| \\ &= \frac{L_H}{2} \|s_k^{curv}\|^2 + \|g_k\| + \|H_k s_k^{curv}\| \\ &\leq \frac{L_H V_{\max}^2}{2\sigma_k} \kappa_C^2 \theta^2 \|g_k\| + \|g_k\| + \|H_k s_k^{curv}\|. \end{aligned} \quad (3.11)$$

Now, using (2.14), (2.15) again,

$$\begin{aligned} \|H_k s_k^{curv}\| &= \theta \kappa_C \sqrt{\frac{\|g_k\|}{\sigma_k}} \|H_k V_D u_k\| = \theta \kappa_C \sqrt{\frac{\|g_k\|}{\sigma_k}} \sqrt{u_k^\top \hat{H}_k^2 u_k} \\ &\leq \theta^2 \kappa_C \sqrt{\frac{\|g_k\|}{\sigma_k}} |\lambda_{\min}(\hat{H}_k)| \leq \theta^2 \kappa_C \sqrt{\frac{\|g_k\|}{\sigma_k}} |\lambda_{\min}(H_k)|. \end{aligned}$$

Hence (3.11) together with AS.4 and the fact $\|g_k\| \geq \epsilon$ before termination, give that

$$\begin{aligned} \|g_{k+1}\| &\leq \frac{L_H V_{\max}^2}{2\sigma_k} \kappa_C^2 \theta^2 \|g_k\| + \|g_k\| + \theta^2 \kappa_B \kappa_C \sqrt{\frac{\|g_k\|}{\sigma_k}} \\ &= \left(\frac{L_H V_{\max}^2}{2\sigma_k} \kappa_C^2 \theta^2 + \frac{\theta^2 \kappa_B \kappa_C}{\sqrt{\sigma_k \|g_k\|}} + 1 \right) \|g_k\| \\ &\leq \left(\frac{L_H V_{\max}^2}{2\sigma_k} \kappa_C^2 \theta^2 + \frac{\theta^2 \kappa_B \kappa_C}{\sqrt{\sigma_k \epsilon}} + 1 \right) \|g_k\|, \end{aligned}$$

yielding (3.10). \square

This lemma is the only result requiring AS.4 or its weaker formulation (3.2). Note that this assumption is only required along directions of negative curvature, which we expect to occur rarely in practice for suitably large choices of κ_C .

After proving a lower bound on the quadratic's decrease when $k \in \mathcal{S}_k^{def}$, we now exhibit a relationship between the decrease on the objective function decrease and gradient both at iteration k and $k + 1$ for $k \in \mathcal{S}_k^{neig} \cup \mathcal{S}_k^{def}$. This is also where the two global conditions (2.13) and (2.9) on the subspace V_p will be useful. Moreover, we also prove an inequality between the norms of the gradient at two successive iterations, similar to (3.10).

Lemma 3.3 Suppose that AS.1 and AS.3 hold and that $k \in \mathcal{S}_k^{neig} \cup \mathcal{S}_k^{def}$ before termination. Then

$$\|g_{k+1}\| \leq \left(\frac{L_H V_{\max}^3 (1 + \kappa_\theta)}{2\varsigma_1^2 \sigma_k} + \frac{2\kappa_b \sqrt{V_{\max}}}{\varsigma_1} + \kappa_b \kappa_C \sqrt{V_{\max}} \right) (1 + \kappa_\theta) \|g_k\| \quad (3.12)$$

and

$$f(x_k) - f(x_{k+1}) \geq \eta_1 \varsigma_{\min} \sqrt{\sigma_k \|g_k\|} \left(\frac{-(2 + \kappa_C) \kappa_b \sqrt{\kappa_a \sigma_k \|g_k\|} + \sqrt{(\kappa_b (2 + \kappa_C))^2 \kappa_a \sigma_k \|g_k\|} + 2 V_{\max}^2 L_H \|g_{k+1}\|}{L_H V_{\max}^2} \right)^2 \quad (3.13)$$

where

$$\varsigma_{\min} \stackrel{\text{def}}{=} \min \left(\frac{1 - 2\varsigma_2}{2}, 1 - \varsigma_3 \right). \quad (3.14)$$

Proof. Consider first the case where $k \in \mathcal{S}_k^{neig}$. By using the Lipschitz error bound for the gradient ([14, Lemma 2.1]), that (2.13) holds, r_k^{neig} expression (2.21), the condition on $\|r_k^{neig}\|$ (2.12) and the fact that $[-\lambda_{\min}(\widehat{H}_k)]_+ \leq \kappa_C \sqrt{\sigma_k \|g_k\|}$ for $k \in \mathcal{S}_k^{neig}$, we deduce that

$$\begin{aligned} \|g_{k+1}\| &\leq \|g_{k+1} - H_k s_k^{neig} - g_k\| + \|H_k s_k^{neig} + g_k\| \\ &\leq \frac{L_H}{2} \|s_k^{neig}\|^2 + \kappa_b \|\widehat{H}_k y_k^{neig} + \widehat{g}_k\| \\ &\leq \frac{L_H}{2} \|s_k^{neig}\|^2 + \kappa_b (\sqrt{\sigma_k \|g_k\|} + [-\lambda_{\min}(\widehat{H}_k)]_+) \|y_k^{neig}\| + \kappa_b \|r_k^{neig}\| \\ &\leq \frac{L_H}{2} \|s_k^{neig}\|^2 + \kappa_b (1 + \kappa_C) \sqrt{\sigma_k \|g_k\|} \|y_k^{neig}\| + \kappa_b \varsigma_3 \sqrt{\sigma_k \|g_k\|} \|y_k^{neig}\|. \end{aligned} \quad (3.15)$$

Using now (2.26) and (2.22) in the last inequality

$$\|g_{k+1}\| \leq \left(\frac{L_H V_{\max}^3}{2\sigma_k} (1 + \kappa_\theta) + \kappa_b (1 + \kappa_C) \sqrt{V_{\max}} + \varsigma_3 \kappa_b \sqrt{V_{\max}} \right) (1 + \kappa_\theta) \|g_k\|. \quad (3.16)$$

Consider now $k \in \mathcal{S}_k^{def}$. By arguments similar to those used for (3.15), this time with (2.23), (2.9) and (2.10), we obtain that

$$\begin{aligned} \|g_{k+1}\| &\leq \|g_{k+1} - H_k s_k^{def} - g_k\| + \|H_k s_k^{def} + g_k\| \\ &\leq \|g_{k+1} - H_k s_k^{def} - g_k\| + \kappa_b \|\widehat{H}_k y_k^{def} + \widehat{g}_k\| \\ &\leq \frac{L_H}{2} \|s_k^{def}\|^2 + \kappa_b \sqrt{\kappa_a \sigma_k \|g_k\|} \|y_k^{def}\| + \kappa_b \|r_k^{def}\| \\ &\leq \frac{L_H}{2} \|s_k^{def}\|^2 + \kappa_b \sqrt{\kappa_a \sigma_k \|g_k\|} \|y_k^{def}\| + \kappa_b \varsigma_2 \sqrt{\kappa_a \sigma_k \|g_k\|} \|y_k^{def}\|. \end{aligned} \quad (3.17)$$

Bounding $\|s_k^{def}\|$ with (2.24) and utilizing (2.8) yields that

$$\|g_{k+1}\| \leq \left(\frac{L_H(1 + \kappa_\theta)V_{\max}^2}{2\varsigma_1^2\kappa_a\sigma_k} + \frac{\kappa_b(1 + \varsigma_2)}{\varsigma_1} \right) (1 + \kappa_\theta) \|g_k\|, \quad (3.18)$$

so that taking the larger bound for both (3.16) and (3.18) and using the bounds $\varsigma_1 < 1$, $\varsigma_2 < \frac{1}{2}$, $\varsigma_3 < 1$, $V_{\max} \geq 1$ and $\kappa_b \geq 1$ gives (3.12).

Finally, from (3.17), (3.15), (2.20), (2.24), the bounds $\max(\varsigma_3, \varsigma_2) < 1$ and $\kappa_a \geq 1$, we obtain that, for $k \in \mathcal{S}_k^{def} \cup \mathcal{S}_k^n$,

$$\frac{L_H V_{\max}^2}{2} \|y_k\|^2 + \kappa_b(2 + \kappa_C) \sqrt{\kappa_a \sigma_k} \|g_k\| \|y_k\| - \|g_{k+1}\| \geq 0.$$

Hence $\|y_k\|$ is larger than the positive root of this quadratic and therefore

$$\|y_k\| \geq \frac{-\kappa_b(2 + \kappa_C) \sqrt{\kappa_a \sigma_k} \|g_k\| + \sqrt{\kappa_b^2(2 + \kappa_C)^2 \kappa_a \sigma_k \|g_k\| + 2L_H V_{\max}^2 \|g_{k+1}\|}}{L_H V_{\max}^2} > 0.$$

We then deduce (3.13) from this inequality, (2.4), the lower bounds on the quadratic decrease for $k \in \mathcal{S}_k^{neig}$ or $k \in \mathcal{S}_k^{def}$ ((2.28) and (2.27) respectively) and the definition of ς_{\min} in (3.14). \square

The bound (3.13) is not sufficient for deriving the required $\mathcal{O}(\epsilon^{-3/2})$ optimal complexity rate because the decrease depends on both $\|g_{k+1}\|$ and $\|g_k\|$. Indeed, when $\|g_{k+1}\| \ll \|g_k\|$, the right-hand side of (3.13) tends to zero. To circumvent this difficulty, the next lemma borrows some elements of [41, Theorem 1] and partitions $\mathcal{S}_k^{neig} \cup \mathcal{S}_k^{def}$ in two further subsets. The minimum decrease on the objective function is of the required magnitude in the first one while no meaningful information can be derived on the decrease on the function value in the second, albeit the magnitude of the gradient at the next iteration is halved. The bounds (3.12) and (3.10) are then used to bound the cardinal of the latter set.

Lemma 3.4 Suppose that AS.1, AS.3 and AS.4 hold and that $\mathcal{S}_k^{neig} \cup \mathcal{S}_k^{def}$ is partitioned as

$$\mathcal{S}_k^{neig} \cup \mathcal{S}_k^{def} = \mathcal{S}_k^{decr} \cup \mathcal{S}_k^{divgrad} \quad (3.19)$$

where

$$\mathcal{S}_k^{decr} \stackrel{\text{def}}{=} \{k \in \mathcal{S}_k^{neig} \cup \mathcal{S}_k^{def}, \sigma_k \|g_k\| \leq \kappa_m 2L_H \|g_{k+1}\|\}, \quad (3.20)$$

$$\mathcal{S}_k^{divgrad} \stackrel{\text{def}}{=} \{k \in \mathcal{S}_k^{neig} \cup \mathcal{S}_k^{def}, \sigma_k \|g_k\| > \kappa_m 2L_H \|g_{k+1}\|\} \quad (3.21)$$

with

$$\kappa_m \stackrel{\text{def}}{=} \gamma_3 \max \left(\frac{\sigma_0}{L_H}, \frac{\varsigma_{\max}}{6(1-\eta_2)} \right). \quad (3.22)$$

Then, for all $k \in \mathcal{S}_k^{decr}$,

$$f(x_k) - f(x_{k+1}) \geq \frac{\eta_1 \varsigma_{\min}(\sigma_k \|g_k\|)^{\frac{3}{2}}}{\left(\kappa_m L_H \left(\kappa_b(2 + \kappa_C) \sqrt{\kappa_a} + \sqrt{(\kappa_b(2 + \kappa_C))^2 \kappa_a + \frac{V_{\max}^2}{\kappa_m}} \right) \right)^2}. \quad (3.23)$$

Moreover,

$$|\mathcal{S}_k^{divgrad}| \leq \kappa_n |\mathcal{S}_k^{decr}| + \left(\frac{1}{2 \log(2)} |\log(\epsilon)| + \kappa_{curv} \right) |\mathcal{S}_k^{curv}| + \frac{|\log(\epsilon)| + \log(\|g_0\|)}{\log(2)} + 1, \quad (3.24)$$

where

$$\kappa_n \stackrel{\text{def}}{=} \frac{1}{\log(2)} \log \left(\frac{L_H(1 + \kappa_\theta) V_{\max}^3}{2\varsigma_1^2 \sigma_{\min}} + \frac{2\sqrt{V_{\max} \kappa_b}}{\varsigma_1} + \sqrt{V_{\max} \kappa_C \kappa_b} \right) + \frac{\log(1 + \kappa_\theta)}{\log(2)}, \quad (3.25)$$

$$\kappa_{curv} \stackrel{\text{def}}{=} \frac{1}{\log(2)} \log \left(\frac{L_H V_{\max}^2}{2\sigma_{\min}} \kappa_C^2 \theta^2 + \frac{\theta^2 \kappa_B \kappa_C}{\sqrt{\sigma_{\min}}} + 1 \right). \quad (3.26)$$

Proof. Let $k \in \mathcal{S}_k^{decr}$. Injecting the definition of \mathcal{S}_k^{decr} (3.20) in (3.13), we obtain that

$$f(x_k) - f(x_{k+1}) \geq \eta_1 \varsigma_{\min}(\sigma_k \|g_k\|)^{\frac{3}{2}} \left(\frac{-(2 + \kappa_C) \kappa_b \sqrt{\kappa_a} + \sqrt{(2 + \kappa_C)^2 \kappa_b^2 \kappa_a + \frac{V_{\max}^2}{\kappa_m}}}{L_H V_{\max}^2} \right)^2.$$

Taking the conjugate both at the denominator and numerator yields (3.23).

Let $k \in \mathcal{S}_k^{divgrad}$. Using the definition of κ_m in (3.22) and that of $\mathcal{S}_k^{divgrad}$ in (3.21) gives that

$$\|g_{k+1}\| < \frac{\sigma_k}{\kappa_m L_H} \frac{\|g_k\|}{2} \leq \frac{\sigma_k}{\gamma_3 \max \left(\frac{\sigma_0}{L_H}, \frac{\varsigma_{\max}}{6(1-\eta_2)} \right) L_H} \frac{\|g_k\|}{2} \leq \frac{\|g_k\|}{2}, \quad (3.27)$$

where the last inequality results from the upper bound on σ_k in (3.3).

Successively using the fact that $\mathcal{S}_k = \mathcal{S}_k^{decr} \cup \mathcal{S}_k^{divgrad} \cup \mathcal{S}_k^{curv}$, the relationship between $\|g_{k+1}\|$ and $\|g_k\|$ in the three cases ((3.27), (3.12) and (3.10)), the fact that $\sigma_k \geq \sigma_{\min}$ in (3.12) and

(3.10), we then deduce that

$$\begin{aligned}
 \frac{\epsilon}{\|g_0\|} &\leq \frac{\|g_k\|}{\|g_0\|} = \prod_{i \in \mathcal{S}_k \setminus \{k\}} \frac{\|g_{i+1}\|}{\|g_i\|} \\
 &= \prod_{i \in \mathcal{S}_k^{decr} \setminus \{k\}} \frac{\|g_{i+1}\|}{\|g_i\|} \prod_{i \in \mathcal{S}_k^{divgrad} \setminus \{k\}} \frac{\|g_{i+1}\|}{\|g_i\|} \prod_{i \in \mathcal{S}_k^{curv} \setminus \{k\}} \frac{\|g_{i+1}\|}{\|g_i\|} \\
 &\leq \left[\left(\frac{L_H(1 + \kappa_\theta)V_{\max}^3}{2\varsigma_1^2\sigma_{\min}} + \frac{2\kappa_b\sqrt{V_{\max}}}{\varsigma_1} + \kappa_C\kappa_b\sqrt{V_{\max}} \right) (1 + \kappa_\theta) \right]^{|\mathcal{S}_k^{decr} \setminus \{k\}|} \times \\
 &\quad \frac{1}{2^{|\mathcal{S}_k^{divgrad} \setminus \{k\}|}} \times \left[\frac{L_H V_{\max}^2}{2\sigma_{\min}} \kappa_C^2 \theta^2 + \frac{\theta^2 \kappa_B \kappa_C}{\sqrt{\epsilon \sigma_{\min}}} + 1 \right]^{|\mathcal{S}_k^{curv} \setminus \{k\}|}.
 \end{aligned}$$

Now $\varsigma_1 \leq 1$ and thus both terms in brackets are larger than one. Moreover, obviously, $|\mathcal{S}_k^{decr} \setminus \{k\}| \leq |\mathcal{S}_k^{decr}|$ and $|\mathcal{S}_k^{curv} \setminus \{k\}| \leq |\mathcal{S}_k^{curv}|$, so that

$$\begin{aligned}
 \frac{2^{|\mathcal{S}_k^{divgrad} \setminus \{k\}|} \epsilon}{\|g_0\|} &\leq \left[\left(\frac{L_H(1 + \kappa_\theta)V_{\max}^3}{2\varsigma_1^2\sigma_{\min}} + \frac{2\kappa_b\sqrt{V_{\max}}}{\varsigma_1} + \sqrt{V_{\max}}\kappa_C\kappa_b \right) (1 + \kappa_\theta) \right]^{|\mathcal{S}_k^{decr}|} \\
 &\quad \times \left[\frac{L_H V_{\max}^2}{2\sigma_{\min}} \kappa_C^2 \theta^2 + \frac{\theta^2 \kappa_B \kappa_C}{\sqrt{\epsilon \sigma_{\min}}} + 1 \right]^{|\mathcal{S}_k^{curv}|}.
 \end{aligned}$$

Taking logarithms gives that

$$\begin{aligned}
 |\mathcal{S}_k^{divgrad} \setminus \{k\}| \log(2) &\leq \log \left[\left(\frac{L_H(1 + \kappa_\theta)V_{\max}^3}{2\varsigma_1^2\sigma_{\min}} + \frac{\kappa_b\sqrt{V_{\max}}}{\varsigma_1} + \kappa_C\kappa_b\sqrt{V_{\max}} \right) (1 + \kappa_\theta) \right]^{|\mathcal{S}_k^{decr}|} \\
 &\quad + \log(\|g_0\|) + |\log(\epsilon)| + \log \left[\frac{L_H V_{\max}^2}{2\sigma_{\min}} \kappa_C^2 \theta^2 + \frac{\theta^2 \kappa_B \kappa_C}{\sqrt{\epsilon \sigma_{\min}}} + 1 \right]^{|\mathcal{S}_k^{curv}|}.
 \end{aligned}$$

We then obtain (3.24) with the values of κ_n and κ_{curv} stated in (3.25) and (3.26) by dividing this last inequality by $\log(2)$ and using the facts that $|\mathcal{S}_k^{divgrad} \setminus \{k\}| \geq |\mathcal{S}_k^{divgrad}| - 1$ and $\frac{1}{\sqrt{\epsilon}} \geq 1$. \square

Combining the previous lemmas, we are now able to state the complexity of the AN2C algorithm. Our theorem statement relies on the observation that the objective function is evaluated once per iteration, and its derivatives once per successful iteration.

Theorem 3.5 Suppose that AS.1- AS.4 hold. Then the AN2C algorithm requires at most

$$|\mathcal{S}_k| \leq \left(\kappa_\star + \frac{\kappa_{\text{negdecr}}}{2 \log(2)} |\log(\epsilon)| \right) \epsilon^{-\frac{3}{2}} + \frac{|\log(\epsilon)| + \log(\|g_0\|)}{\log(2)} + 1$$

successful iterations and evaluations of the gradient and the Hessian and at most

$$\left(1 + \frac{|\log \gamma_1|}{\log \gamma_2} \right) \left[\left(\kappa_\star + \frac{\kappa_{\text{negdecr}}}{2 \log(2)} |\log(\epsilon)| \right) \epsilon^{-\frac{3}{2}} + \frac{|\log(\epsilon)| + \log(\|g_0\|)}{\log(2)} + 1 \right] + \frac{1}{\log \gamma_3} \log \left(\frac{\sigma_{\max}}{\sigma_0} \right)$$

evaluations of f to produce a vector x_ϵ such that $\|g(x_\epsilon)\| \leq \epsilon$, where κ_\star is defined by

$$\kappa_\star \stackrel{\text{def}}{=} \kappa_{\text{decr}} (1 + \kappa_n) + \kappa_{\text{negdecr}} (1 + \kappa_{\text{curv}}), \quad (3.28)$$

with

$$\kappa_{\text{decr}} \stackrel{\text{def}}{=} \frac{\left(L_H \kappa_m (\sqrt{\kappa_a} \kappa_b (2 + \kappa_C) + \sqrt{\kappa_a (\kappa_b (2 + \kappa_C))^2 + \frac{V_{\max}^2}{\kappa_m}}) \right)^2}{\eta_1 \zeta_{\min} \sigma_{\min}^{\frac{3}{2}}} \quad (3.29)$$

and

$$\kappa_{\text{negdecr}} \stackrel{\text{def}}{=} \frac{2(f(x_0) - f_{\text{low}}) \sqrt{\sigma_{\max}}}{\eta_1 \kappa_C^3 \theta^3}, \quad (3.30)$$

and where κ_n and κ_{curv} are given by (3.25) and (3.26).

Proof. First note that we only need to prove an upper bound on $|\mathcal{S}_k^{\text{decr}}|$ and $|\mathcal{S}_k^{\text{curv}}|$ to derive a bound on $|\mathcal{S}_k|$ since

$$|\mathcal{S}_k| = |\mathcal{S}_k^{\text{decr}}| + |\mathcal{S}_k^{\text{curv}}| + |\mathcal{S}_k^{\text{divgrad}}| \quad (3.31)$$

and a bound on $|\mathcal{S}_k^{\text{divgrad}}|$ is given by (3.24). We start by proving an upper bound on $|\mathcal{S}_k^{\text{curv}}|$. Using AS.2, the lower bound on the decrease of the function values (3.9) and that $\sigma_k \leq \sigma_{\max}$ as stated in Lemma 3.1, we derive that, for $k \in \mathcal{S}_k^{\text{curv}}$,

$$f(x_0) - f_{\text{low}} \geq \sum_{i \in \mathcal{S}_k} f(x_i) - f(x_{i+1}) \geq \sum_{i \in \mathcal{S}_k^{\text{curv}}} f(x_i) - f(x_{i+1}) \geq |\mathcal{S}_k^{\text{curv}}| \frac{\eta_1 \kappa_C^3 \theta^3}{2\sqrt{\sigma_{\max}}} \epsilon^{\frac{3}{2}}$$

and hence that

$$|\mathcal{S}_k^{\text{curv}}| \leq \frac{2(f(x_0) - f_{\text{low}}) \sqrt{\sigma_{\max}}}{\eta_1 \kappa_C^3 \theta^3} \epsilon^{-\frac{3}{2}} = \kappa_{\text{negdecr}} \epsilon^{-\frac{3}{2}}. \quad (3.32)$$

Similarly for $k \in \mathcal{S}_k^{\text{decr}}$, using AS.2, (3.23), the fact that $\sigma_k \geq \sigma_{\min}$ and $\|g_k\| \geq \epsilon$ before termination yields that

$$f(x_0) - f_{\text{low}} \geq \sum_{i \in \mathcal{S}_k^{\text{decr}}} f(x_i) - f(x_{i+1}) \geq \frac{|\mathcal{S}_k^{\text{decr}}| \eta_1 \zeta_{\min} (\sigma_{\min} \epsilon)^{\frac{3}{2}}}{\left(L_H \kappa_m (\sqrt{\kappa_a} (2 + \kappa_C) \kappa_b + \sqrt{\kappa_a (\kappa_b (2 + \kappa_C))^2 + \frac{V_{\max}^2}{\kappa_m}}) \right)^2}$$

where κ_m is defined in (3.22). Rearranging the last inequality yields that

$$|\mathcal{S}_k^{\text{decr}}| \leq \frac{\left(L_H \kappa_m (\sqrt{\kappa_a} \kappa_b (2 + \kappa_C) + \sqrt{\kappa_a (\kappa_b (2 + \kappa_C))^2 + \frac{V_{\max}^2}{\kappa_m}}) \right)^2}{\eta_1 \zeta_{\min} \sigma_{\min}^{\frac{3}{2}}} \epsilon^{-\frac{3}{2}} = \kappa_{\text{decr}} \epsilon^{-\frac{3}{2}}. \quad (3.33)$$

Combining now (3.32) and (3.33) with the upper-bound (3.24) on $|\mathcal{S}_k^{divgrad}|$, we deduce that

$$|\mathcal{S}_k^{divgrad}| \leq \kappa_n \kappa_{decr} \epsilon^{-\frac{3}{2}} + \left(\frac{|\log(\epsilon)|}{2 \log(2)} + \kappa_{curv} \right) \kappa_{negdecr} \epsilon^{-\frac{3}{2}} + \frac{|\log(\epsilon)| + \log(\|g_0\|)}{\log(2)} + 1. \quad (3.34)$$

By summing equations (3.32), (3.33), and (3.34) to bound $|\mathcal{S}_k|$ in (3.32), while also isolating the terms based on their different orders with respect to ϵ , we obtain that

$$|\mathcal{S}_k| \leq \left(\kappa_\star + \frac{\kappa_{negdecr}}{2 \log(2)} |\log(\epsilon)| \right) \epsilon^{-\frac{3}{2}} + \frac{|\log(\epsilon)| + \log(\|g_0\|)}{\log(2)} + 1, \quad (3.35)$$

where κ_\star is defined in (3.28), thus proving the first part of the theorem. The second part is then deduced from (3.35) combined with Lemma 2.1. \square

Regrouping all the problem's dependent constant of the last theorem and keeping the worst dependency w.r.t ϵ , we derive a $\mathcal{O}(|\log(\epsilon)|\epsilon^{-3/2})$ complexity order in ϵ that only differs by the factor $|\log(\epsilon)|$ from the optimal order for nonconvex second-order methods [13], a factor which is typically small for practical values of ϵ . The AN2C algorithm thus enjoys a better complexity order than that of past hybrid algorithms [19, 38, 25] for which the order is $\mathcal{O}(\epsilon^{-2})$. However, it is marginally worse than that of the more complex second-order linesearch of [47] which attains the optimal order. Moreover, we see in the proof of Theorem 3.5 that the $|\log \epsilon|$ term appears because of (3.24) and (3.32) and we may hope that the number of s_k^{curv} iterations is typically much less than its worst-case $\mathcal{O}(\epsilon^{-3/2})$ in practice. The trust-region algorithm of [20] has the same total complexity as AN2C although their method requires only $\mathcal{O}(\epsilon^{-3/2})$ gradient and Hessian calls whereas our algorithm suffers from an additional $|\log(\epsilon)|$ term.

4 Finding second-order critical points

Can the AN2C algorithm be strengthened to ensure it will compute second-order critical points? We show in this section under the same assumptions as that used for its first-order analysis that approximate second order points can be reached.

The resulting modified algorithm, which we call SOAN2C (for Second-Order AN2C) makes extensive use of AN2C, and is detailed on the next page.

Prior to reaching an approximate first-order point, we utilize only the `RegStep` and `NewtonEigenStep` subroutines to generate tentative steps, hence the 'fo' (first-order) superscripts in (4.2) and (4.3). Similar to Section 2, AS.0 is necessary to obtain a valid step when `NewtonEigenStep` is invoked. Once an approximate first-order point is reached, further progress towards second-order stationarity is obtained by exploiting the negative-curvature direction (4.4)-(4.5), thereby justifying the 'so' (second-order) superscript.

An upper bound on the evaluation complexity of the SOAN2C algorithm is given by the following theorem.

Algorithm 4.1: Second-Order Adaptive Newton with Negative Curvature (SOAN2C)

Step 0: Initialization Identical to AN2C[Step 0] with $\epsilon \in (0, 1]$ now replaced by $\epsilon_1 \in (0, 1]$ and $\epsilon_2 \in (0, 1]$.

Step 1: Compute current derivatives Evaluate g_k and H_k . Terminate if

$$\|g_k\| \leq \epsilon_1 \text{ and } \lambda_{\min}(H_k) \geq -\epsilon_2. \quad (4.1)$$

Step 2: Compute subspace derivatives Form \hat{g}_k and \hat{H}_k as in AN2C[Step 2].

Step 3: Step calculation If $\|g_k\| > \epsilon_1$,

$$s_k = s_k^{fo} = \text{RegStep}(\hat{g}_k, \hat{H}_k, V_p, \|g_k\|, \sigma_k, \kappa_a, \kappa_b, \varsigma_1, \varsigma_2), \quad (\text{Optional}). \quad (4.2)$$

If s_k^{fo} has been successfully defined, go to Step 4. Else, compute

$$s_k = s_k^{fo} = \text{NewtonEigenStep}(\hat{g}_k, \hat{H}_k, V_p, \|g_k\|, \sigma_k, \kappa_C, \kappa_b, \kappa_\theta, \varsigma_3, \theta). \quad (4.3)$$

Else ($\|g_k\| \leq \epsilon_1$), compute u_k such that

$$g_k^\top u_k \leq 0, \quad \|u_k\| = 1 \text{ and } H_k u_k = \lambda_{\min}(H_k) u_k, \quad (4.4)$$

and set

$$s_k = s_k^{so} \stackrel{\text{def}}{=} \frac{-\lambda_{\min}(H_k)}{\sigma_k} u_k. \quad (4.5)$$

Step 4: Acceptance ratio computation Identical to AN2C[Step 5].

Step 5: Regularization parameter update Identical to AN2C[Step 6].

Theorem 4.1 Suppose that AS.1–AS.4 hold. Then the SOAN2C algorithm requires at most

$$|\mathcal{S}_k| \leq \kappa_\star \epsilon_1^{-\frac{3}{2}} + \kappa_{\text{so}} \epsilon_2^{-3} + \frac{|\log(\epsilon_1)|}{2 \log(2)} \kappa_{\text{negdecr}} \epsilon_1^{-\frac{3}{2}} + \left(\frac{|\log(\epsilon_1)| + \log(\kappa_{\text{gpi}})}{\log(2)} + 1 \right) (\kappa_{\text{so}} \epsilon_2^{-3} + 1)$$

successful iterations and evaluations of the gradient and the Hessian and at most

$$\begin{aligned} & \left(1 + \frac{|\log \gamma_1|}{\log \gamma_2} \right) \left[\kappa_\star \epsilon_1^{-\frac{3}{2}} + \kappa_{\text{so}} \epsilon_2^{-3} + \frac{|\log(\epsilon_1)|}{2 \log(2)} \kappa_{\text{negdecr}} \epsilon_1^{-\frac{3}{2}} \right. \\ & \quad \left. + \left(\frac{|\log(\epsilon_1)| + \log(\kappa_{\text{gpi}})}{\log(2)} + 1 \right) (\kappa_{\text{so}} \epsilon_2^{-3} + 1) \right] + \frac{1}{\log \gamma_3} \log \left(\frac{\sigma_{\text{max}}}{\sigma_0} \right) \end{aligned}$$

evaluations of f to produce a vector x_ϵ such that $\|g(x_\epsilon)\| \leq \epsilon_1$ and $\lambda_{\min}(H_{x_\epsilon}) \geq -\epsilon_2$, where

$$\kappa_{\text{so}} \stackrel{\text{def}}{=} \frac{2\sigma_{\text{max}}^2(f(x_0) - f_{\text{low}})}{\eta_1} \quad (4.6)$$

κ_{gpi} is defined in (A.5) and κ_\star , κ_{negdecr} and σ_{max} (defined by (3.28), (3.30) and (3.3), respectively) depend solely on the problem .

As for Theorem 3.5, the bound, in which the ϵ_2^{-3} term is likely to dominate, differs from standard one for second-order algorithms seeking second-order points (in $\mathcal{O}(\max(\epsilon_1^{-3/2}, \epsilon_2^{-3}))$) [15, Theorems 3.3.9 and 3.4.6] by a (modest) factor $|\log(\epsilon_1)|$.

To prove Theorem 4.1, we need to take two main issues into account. The first is that, because the step may be computed using (4.2), (4.3) but also (4.5), we need to complete the partition of $|\mathcal{S}_k|$ by introducing subsets relevant to this new type of steps. The second is clearly that negative curvature information must be exploited in order to guarantee a sufficient decrease of the objective function when it is discovered close to a first-order critical point. This leads to a development which broadly follows the lines of Section 3, extending the proofs when necessary to handle the more complicated situation. The details of this development are given in appendix.

5 Choosing the subspace

In practice, the algorithm crucially depends on how one chooses the matrix V_p spanning the iteration-dependent subspace, and we discuss two options. Each of the choices presented below can be included in both AN2C and SOAN2C, defined in Section 2 and Section 4, respectively. For conciseness, we only consider AN2C.

5.1 A full-space variant

A simple choice of V_p is to consider $V_p \stackrel{\text{def}}{=} I_n$, that is the subspace is in fact the whole space. We note that, in this case, conditions (2.9) or (2.13) automatically hold.

We define two variants in this context. The first is called AN2CER (for AN2C Exact using RegStep) exploits the RegStep algorithm in order to limit the need of possibly costly second-order information. The second, potentially more costly, is called AN2CE and does not use the optional RegStep algorithm, therefore making no attempt to avoid eigenvalue computations.

These variants may be useful for problems in which systems (2.6) and (2.11) may effectively be solved (for instance using Cholesky factorizations). As we will see below, they require on average a single such solution/factorization per iteration. AN2CER and AN2CE may thus be attractive in the large class of applications for which off-the-shelf linear solvers are available. The computation

of $\lambda_{\min}(H_k)$ also needs to be feasible but, due to Algorithm `RegStep`, this occurs only rarely in `AN2CER`.

5.2 A Krylov variant

When the dimension of the problem grows and factorizations become impractical, one can turn to exploiting Krylov subspaces, as we now show. The resulting algorithmic variant will be called `AN2CK`, where `K` stands here for Krylov, and is obtained by replacing Steps 3 and 4 of the `AN2C` algorithm by Algorithm `AN2CKStep` on the following page. In this variant, the subspace generation and step computation are combined in order to best exploit the structure of the resulting subproblem. As is common in Krylov-based methods, we assume the availability of a 'preconditioner', that is a positive-definite matrix M_k approximating the Hessian H_k in the sense that $M_k^{-1}H_k$ is close to the identity. For clarity, we ignore the iteration subscript k in what follows.

Each iteration of the `AN2CKStep` algorithm has a moderate cost (a few vector assignments, one matrix-vector product and –possibly– the computation of the smallest eigenvalue of a tridiagonal matrix, see [17] and the references therein for details). We observe that (5.1)-(5.2) amounts to using the standard preconditioned Lanczos process for building an orthonormal (in the $\langle \cdot, M \cdot \rangle$ inner product) basis V_p of successive Krylov subspaces generate by the preconditioned gradient and Hessian. We therefore build on existing theory for this process (see [18, Section 5.2], for instance). We note that the use of the full Lanczos basis V_p is only requested at the end of the process (in (5.7) and (5.9)). As a consequence two options are available for its detailed implementation: one can store the Lanczos basis vectors as the iterations proceed and use them at the end of the step computation, or one can forget them but re-run the necessary Lanczos process to re-generate them (as has been done in the `GALAHAD` library [28] for the `GLTR` and `GLRT` algorithms for trust-region and regularization subproblems, respectively). Obviously, V_p and T_p may be updated incrementally in (5.3) and (5.4). When updating T_p , it is also easy to check if it remains positive definite by recurring the pivots of its Cholesky factorization, which are given by

$$\pi_1 = \delta_1 \quad \text{and} \quad \pi_p = \delta_p - \alpha_p^2 / \pi_{p-1} \quad (p > 1).$$

As long as π_p stays positive, it is thus unnecessary to compute $\lambda_{\min}(T_p)$ since $[-\lambda_{\min}(T_p)]_+$ is then identically zero in (5.5). Finally, should a preconditioner M be unavailable, setting $M = I_n$ is possible, in which case w_p and z_p can be dispensed of because they are identical to r_p and v_p , respectively.

We now verify that, as stated, Algorithm `AN2CK` is a correct instantiation of Algorithm `AN2C` (without the optional Step 3).

Theorem 5.1 Suppose that

$$\mu_1 \leq \lambda_{\min}(M) \quad \text{and} \quad \lambda_{\max}(M) \leq \mu_2 \tag{5.10}$$

for some $\mu_2 \geq \mu_1 > 0$. Then the definitions and conditions (5.8), (5.6) and (5.5) of Algorithm `AN2CKStep` are equivalent to (2.14), (2.13) (with κ_b redefined as $\max(1, \kappa_b \sqrt{\mu_2})$) and (2.11) of Algorithm 2.3, respectively. Moreover, `AS.0` holds and (5.4) is valid.

Proof. If Z_p is the matrix whose columns are z_1, \dots, z_p , we deduce from (5.1) and (5.2) that

$$HV_p = Z_p T_p + \alpha_{p+1} z_{p+1} e_p^T = M V_p T_p + \alpha_{p+1} M v_{p+1} e_p^T. \tag{5.11}$$

Using that $V_p^T M v_{p+1} = 0$ yields (5.4). Note also that as $v_1 = \frac{w_1}{\alpha_1} = M^{-1} z_1$ from (5.1) and $V_p^T M V_p = I_p$,

$$V_p^T g = \alpha_1 V_p^T z_1 = \alpha_1 V_p^T M v_1 = \alpha_1 e_1. \tag{5.12}$$

Algorithm 5.1: AN2CKStep($g, H, \sigma, M, \kappa_C, \kappa_b, \theta$)

Step 0: Initialization Set $p = 1$, $r_1 = g$, $w_1 = M^{-1}r_1$, $\alpha_1 = \sqrt{w_1^\top r_1}$ and $z_0 = 0$.

Step 1: Form the orthonormal basis Compute

$$z_p = \frac{r_p}{\alpha_p}, \quad v_p = \frac{w_p}{\alpha_p}, \quad \delta_p = v_p^\top H v_p, \quad (5.1)$$

$$r_{p+1} = H v_p - \delta_p z_p - \alpha_p z_{p-1}, \quad w_{p+1} = M^{-1}r_{p+1}, \quad \alpha_{p+1} = \sqrt{w_{p+1}^\top r_{p+1}}, \quad (5.2)$$

and define

$$V_p = (v_1, v_2, \dots, v_p) \in \mathbb{R}^{n \times p}. \quad (5.3)$$

Step 2: Newton step computation Form the subspace Hessian

$$T_p \stackrel{\text{def}}{=} V_p^\top H V_p = \begin{pmatrix} \delta_1 & \alpha_2 & & & & \\ \alpha_2 & \delta_2 & \alpha_3 & & & \\ & \ddots & \ddots & \ddots & & \\ & & & \delta_{p-1} & \alpha_p & \\ & & & \alpha_p & \delta_p & \end{pmatrix} \quad (5.4)$$

and compute its minimum eigenvalue.

If $\lambda_{\min}(T_p) \leq -\kappa_C \sqrt{\sigma \|g\|}$, go to Step 4.

Otherwise, solve

$$\left(T_p + (\sqrt{\sigma \|g\|} + [-\lambda_{\min}(T_p)]_+) I_p \right) y_p = -\alpha_1 e_1. \quad (5.5)$$

Step 3: Check global quality of the solution If

$$\sqrt{\alpha_{p+1}^2 (e_p^\top y_p)^2 + \|T_p y_p + \alpha_1 e_1\|^2} \leq \kappa_b \|T_p y_p + \alpha_1 e_1\|, \quad (5.6)$$

then return

$$s = s^{neig} = V_p y_p. \quad (5.7)$$

Else increment p by one and go back to Step 1.

Step 4: Eigenvector direction Compute u such that

$$e_1^\top u \leq 0, \quad \|u\| = 1 \quad \text{and} \quad u^\top T_p u \leq \theta \lambda_{\min}(T_p). \quad (5.8)$$

Return

$$s = s^{curv} = \theta \kappa_C \sqrt{\frac{\|g\|}{\sigma}} V_p u. \quad (5.9)$$

The last identity with the fact that $T_p = V_p^\top H V_p$ ensures that (5.8) and (5.5) are reformulations of (2.14) and (2.11). We now prove that (5.6) implies (2.13). Using (5.5), (5.12), (5.11), we obtain that

$$\begin{aligned} Hs + g &= H V_p y_p + \alpha_1 M v_1 = H V_p y_p + \alpha_1 M V_p e_1 \\ &= H V_p y_p - M V_p T_p y_p - (\sqrt{\sigma \|g\|} + [-\lambda_{\min}(T_p)]_+) M V_p y_p \\ &= \alpha_{p+1} (e_p^\top y_p) M v_{p+1} - (\sqrt{\sigma \|g\|} + [-\lambda_{\min}(T_p)]_+) M V_p y_p. \end{aligned}$$

Since $V_p^\top M V_p = I_p$ and $V_p^\top M v_{p+1} = 0$, we deduce, using (5.5) and (5.6), that

$$\begin{aligned} \|Hs + g\|^2 &\leq \lambda_{\max}(M) (Hs + g)^\top M^{-1} (Hs + g) \\ &= \lambda_{\max}(M) [\alpha_{p+1}^2 (e_p^\top y_p)^2 + (\sqrt{\sigma \|g\|} + [-\lambda_{\min}(T_p)]_+)^2 \|y_p\|^2] \\ &= \lambda_{\max}(M) [\alpha_{p+1}^2 (e_p^\top y_p)^2 + \|T_p y_p + \alpha_1 e_1\|^2] \\ &\leq \kappa_b^2 \lambda_{\max}(M) \|T_p y_p + \alpha_1 e_1\|^2, \end{aligned}$$

and (2.13) follows with the redefined κ_b . We finally verify that AS.0 holds. Because

$$1 = \|M^{\frac{1}{2}} V_p\| \geq \lambda_{\min}(M^{\frac{1}{2}}) \|V_p\| = \sqrt{\lambda_{\min}(M)} \|V_p\|$$

(2.16) holds with $V_{\max} = 1/\sqrt{\lambda_{\min}(M)} \leq \mu_1^{-1/2}$, where we again used (5.10) to derive the last inequality. Moreover, given that $\kappa_b \geq 1$, termination necessarily occurs when $p = n$, $V_n^\top M V_n = I_n$, V_n spans the whole space and $\alpha_{p+1} = 0$ in (5.6). \square

The optional Step 3 of Algorithm 2.1 is in fact implicitly contained in Algorithm 5.1 since convexity along the current step (condition (2.7)) is verified at each step of the Lanczos process by checking the positive-definiteness of T_p .

Returning now to the complete sequence of minimization iterates, we see that, whenever the AN2CK algorithm is used with iteration-dependent preconditioners $M_k \neq I_n$, Theorems 3.5 and 4.1 remain valid provided (5.10) holds uniformly for all iterations.

6 Numerical illustration

We now illustrate the behavior of our proposed algorithms on three sets of test problems from the freely available OPM collection⁽²⁾ [34]. The first set contains 119 small-dimensional problems, the second contains 74 medium-size ones, while the third contains 59 ‘‘largish’’ ones. The list of problems and their dimensions are listed in Tables A.3, A.4 and A.5 in appendix.

6.1 Using the full-space variants

We use Matlab implementations of AN2CE and AN2CER where the involved linear systems are solved by using the Matlab sparse Cholesky factorization, and where we have set

$$\kappa_C = 10^3, \kappa_a = 50 \text{ (AN2CE) or } 100 \text{ (AN2CER)}, \kappa_\theta = 1, \varsigma_1 = \frac{1}{2}, \varsigma_2 = \varsigma_3 = 10^{-10},$$

$$\sigma_0 = 1, \sigma_{\min} = 10^{-8}, \gamma_1 = \frac{1}{2}, \gamma_2 = \gamma_3 = 10, \eta_1 = 10^{-4} \text{ and } \eta_2 = 0.95.$$

The values of κ_C and κ_a were obtained from a hyper-parameter search⁽³⁾ on the set of small problems. The values of ς_2 and ς_3 are given here for consistency, but are irrelevant since factorizations are used to solve the linear systems. Other parameters values are typical of regularization algorithms.

⁽²⁾This collection is a subset of the CUTEest [29] collection where the test problems are described in Matlab.

⁽³⁾Covering the choice $\{10^{30}, 10^8, 10^5, 10^3, 10^2, 10\}$ for κ_C and $\{100, 50, 10\}$ for κ_a .

We compare AN2CE and AN2CER with implementations of the standard adaptive regularization AR2 and trust-region TR2M, two well-regarded methods. All these algorithms use quadratic approximations of the objective function (i.e. gradients and Hessians). The first three also use the same acceptance thresholds η_1 and η_2 and values of γ_1 , γ_2 and γ_3 . The TR2M methods shrinks the trust-region radius by a factor $\sqrt{10}$ and expands it by a factor 2 (see [15, Section 11.2] for a discussion of the coherence of these factors between trust-region and adaptive regularization methods). The authors are aware that further method-dependent tuning would possibly result in improved performance, but the values chosen here appear to work reasonably well for each method. The step computation is performed in AR2 following [15, page 67] or [3] using an (unpreconditioned) Lanczos approach while a standard Moré-Sorensen method⁽⁴⁾ is used in TR2M (see [15, Chapter 9] for details). For AR2, the step computation is terminated as soon as

$$\|g_k + H_k s_k\| \leq \frac{1}{2} \theta_{sub} \sigma_k \|s_k\|^2 \quad (6.1)$$

which slightly differs from the test $\|\nabla_s^1 m_k(s_k)\| \leq \frac{1}{2} \theta_{sub} \sigma_k \|s_k\|^2$ used in [15, page 65] and [3] while maintaining the desired $\mathcal{O}(\epsilon^{-3/2})$ evaluation complexity bound (see [35] for a justification of (6.1) –including the fact that it more often allows the pure Newton step to be accepted– or [15, page 67]). The Moré-Sorensen iterations in TR2M are terminated as soon as $\|s_k\| \in [(1-\theta_{sub})\Delta_k, (1+\theta_{sub})\Delta_k]$, where, in both cases, $\theta_{sub} = 10^{-3}$ for $n \leq 100$ and 10^{-2} for $n > 100$. All experiments were run on a Dell Precision computer with Matlab 2022b.

We discuss our experiments from the efficiency and reliability points of view. Efficiency is measured, in accordance with the complexity theory, in number of iterations (or, equivalently, function and possibly derivatives' evaluations): the fewer the more efficient the algorithm. In addition to presenting the now standard performance profiles [22] for our four algorithms in Figure 1, we follow [45, 32] and consider the derived “global” measure π_{algo} to be $\frac{1}{10}$ of the area below the curve corresponding to `algo` in the performance profile, for abscissas in the interval $[1, 10]$. The larger this area and the closer π_{algo} to one, the closer the curve to the left and top borders of the plot and the better the global performance.

When reporting reliability, we say that the run of an algorithmic variant on a specific test problem is successful if the gradient norm tolerance $\epsilon = 10^{-6}$ has been achieved in the allotted cpu-time (1h) and before the maximum number of iterations (5000) is reached. The ρ_{algo} statistic denotes the percentage of successful runs taken on all problems in each of the three classes.




Figure 1: Full-space variants: iteration performance profiles for OPM problems (left: small, center: medium, right: largish). We report on the vertical axis the proportion of problems for which the number of iterations of each algorithm is at most a fraction (given by the horizontal axis) of the smallest across all algorithms (see [22]).

Figure 1 and Table 1 suggest that the reliability of AN2CE and AN2CER is comparable to that of AR2 and TR2M for all problem sizes. They also indicate that AN2CER is somewhat slower iteration-wise than AR2 and TR2M, but AN2CE is very comparable. The fact that the computationally

⁽⁴⁾Given that our version of AN2C uses matrix factorizations, it seems more natural to compare it with a Moré-Sorensen-based trust-region than to one using truncated conjugate gradients.

	small pbs.		medium pbs.		largish pbs.	
algo	π_{algo}	ρ_{algo}	π_{algo}	ρ_{algo}	π_{algo}	ρ_{algo}
AN2CER	0.88	96.64	0.85	93.24	0.85	94.92
AN2CE	0.91	96.64	0.91	95.95	0.81	86.44
AR2	0.92	97.48	0.85	93.24	0.84	93.22
TR2M	0.91	94.96	0.86	93.24	0.83	91.53

Table 1: Efficiency and reliability statistics for the OPM problems (full-space variants)

more expensive AN2CE is often faster than AN2CER in terms of iteration numbers is not surprising. Indeed, the regularization term in (2.11) becomes $\sqrt{\sigma_k} \|g_k\|$ in convex regions, recovering the analysis of [41, 21], whereas AN2CE regularizes the problem more strongly in (2.6) (by a factor 10 in our numerical settings) and therefore may further restricts the steplength. AN2CE may however be computationally more intensive⁽⁵⁾ than AN2CER. Which of the two algorithms is preferable in practice is likely to depend on the CPU cost of calculating the Hessian’s smallest eigenvalue.

As expected, the call to `NewtonEigenStep` in AN2CER is typically performed on very few iterations (for less 6.4% of them for the small-problems testset) and, when used, results in a negative-curvature step (2.15) even more exceptionally (less than 1%). This means in particular that a single linear-system solve was necessary for approximately 93% of all iterations. The AN2CE variant of course called `NewtonEigenStep` at every iteration, but (2.15) was never actually used.

We also ran the SOAN2CE and SOAN2CER variants with $\epsilon_1 = 10^{-6}$ and $\epsilon_2 = 10^{-4}$, but their results are undistinguishable (for our test sets) from those obtained with AN2CE and AN2CER, except for a final eigenvalue analysis at the found approximate first-order point, which confirmed in all cases that the second-order condition (4.1) did also hold at this point. No step of the form (4.5) was ever taken in our runs, despite the fact that such steps are necessary in theory (think of starting the minimization at a first-order saddle point).

6.2 Using the Krylov-based variants

We ran two variants of the AN2CK algorithm on our three problem sets, which differ in how the vector u is chosen in (5.8). In the first, called AN2CKU, u is chosen as the eigenvector associated with the eigenvalue $\lambda_{\min}(T_p)$. In the second, called AN2CKYU, u is chosen as the sum of the current vector y_p plus a multiple of the eigenvector associated with $\lambda_{\min}(T_p)$ chosen to ensure that the last inequality in (5.8) holds as an equality. An hyper-parameter search on a subset of the medium-sized test set yielded the values

$$\kappa_C = 3, \quad \kappa_b = 50 \text{ and } \theta = \frac{1}{2}.$$

None of the tested methods used preconditioning (that is the choice $M = I_n$ was made throughout). The matrices V_p were stored explicitly.

We again compared these two variants with AR2 and with TR2K, an implementation of the trust-region close to TR2M, but in which the step is computed by minimizing the quadratic model in the intersection of the trust-region and the successive Krylov spaces until

$$\|g_k + H_k s_k\| \leq \frac{1}{10} \|g_k\|. \tag{6.2}$$

The results of our comparison (using the same metrics as in the previous subsection) are given in Figure 2 and Table 2.

We observe that AN2CKU significantly trails the other variants and is in particular both less efficient and less reliable than AN2CKYU, which we explain by the fact that, should a negative curvature step occur, the former strategy does not exploit the decrease of the quadratic model

⁽⁵⁾Most failures of this algorithm on large problems occurred because the time limit was reached.




Figure 2: Krylov-space variants: iteration performance profiles for OPM problems (left: small, center: medium, right: largish). We report on the vertical axis the proportion of problems for which the number of iterations of each algorithm is at most a fraction (given by the horizontal axis) of the smallest across all algorithms (see [22]).

algo	small pbs.		medium pbs.		largish pbs.	
	π_{algo}	ρ_{algo}	π_{algo}	ρ_{algo}	π_{algo}	ρ_{algo}
AN2CKU	0.86	96.64	0.81	93.24	0.77	86.44
AN2CKYU	0.91	96.64	0.90	95.95	0.85	91.53
AR2	0.92	97.48	0.87	93.24	0.89	93.22
TR2K	0.94	96.64	0.85	87.84	0.77	84.75

Table 2: Efficiency and reliability statistics for the OPM problems (Krylov-space variants)

already obtained by the “convex step” y_p . By contrast, AN2CKYU appears to be competitive with both AR2 and TR2K, irrespective of problem size.

For the AN2CKYU variant, the average ratio of the number of matrix-vector products divided by the product of the number of iterations and the problem size (a ratio which is one if every Lanczos process takes n iterations) is below 0.5 for small problems, below 0.15 for medium ones and below 0.03 for large ones. Negative curvature directions (5.9) are also used, for this variant, by 0.25% of the iterations for small problems, 0.23% of iterations for medium ones and never for large ones.

Finally, we also tested SOAN2CKU and SOAN2CKYU, the versions of AN2CKU and AN2CKYU which enforce second-order optimality. As for full-space methods, the results obtained are undistinguishable from those for AN2CKU and AN2CKYU, except for a final eigenvalue analysis confirming the approximate second-order optimality of the computed solution.

These early results are encouraging but the authors are aware that only further experiments will allow a proper assessment of the method’s true potential, both from the number of function/derivatives evaluations and CPU-usage points of view. Several further algorithmic developments within the new algorithms are also of interest, including a possibly better balance between NewtonEigenStep and RegStep in the full-space version, as well as refinements of the regularization parameter update (2.5), possibly in the spirit of [30].

7 Conclusions and Perspectives

We have proposed AN2C and AN2CK, two second-order minimization methods for nonconvex problems that alternate, in an iteration dependent subspace, between Newton and negative-curvature directions. These methods differ from the more standard trust-region and adaptive-regularization techniques in that the involved step computation is free of further inner iterative processes and only requires the approximate solution of at most two (but typically one) linear systems per iter-

ation. We have also proved that these algorithms require at most $\mathcal{O}(|\log(\epsilon)|\epsilon^{-3/2})$ iterations to obtain an ϵ -approximate first-order critical point. Our proof builds on some elements of [41, 21] for the convex case and arguments for adaptive regularization [3] and other nonconvex optimization methods [19, 47]. At each iteration, the algorithms either take an explicit Newton step or negative curvature when it is sufficiently large compared to the square root of the gradient. The norm of the residuals of the Newton step are adjusted dynamically and different types of solvers can be used to solve the linear systems, depending on how subspaces are chosen.

An extension of the algorithmic framework ensuring approximate second-order optimality has also been introduced, and we have proved that the resulting methods require at most $\mathcal{O}(|\log(\epsilon)|\epsilon^{-3})$ iterations to achieve its objective.

A first set of numerical experiments with full-space variants AN2CE and AN2CER as well as Krylov-subspaces iterative ones AN2CKU and AN2CKYU indicates that they are very reliable and competitive with standard techniques in terms of number of iterations.

The reader may wonder why we haven't considered selecting iteration-dependent low-dimensional random subspaces, as has been advocated in [8, 46] for instance. The main reason is that using the Johnson-Lindenstrauss lemma (the basic tool in such an approach) is possible for defining a probabilistically accurate approximate gradient in the subspace. However, as far as we know, using this idea is problematic for the full Hessian matrix unless it is assumed to be of low rank. We could therefore attempt to follow the Cauchy-point-based analysis of [8, 46], and hopefully obtain a probabilistic complexity bound in $\mathcal{O}(\epsilon^{-2})$. However, we do not see at this point how to design a low-dimensional random-subspace algorithm with an $\mathcal{O}(|\log(\epsilon)|\epsilon^{-3/2})$ probabilistic complexity bound for minimizing functions with general (possibly full-rank) Hessians.

Promising lines for future work include inexact derivatives as in [50, 51], estimating the regularization parameter without evaluating the objective function (as in [31]), stochastic variants and the handling of simple constraints such as bounds on the variables in the spirit of [15, Section 14.2].

Acknowledgment

The authors are indebted to two excellent referees for their thoughtful comments and suggestions.

References

- [1] S. Bellavia, G. Gurioli, and B. Morini. (2021) 'Adaptive cubic regularization methods with dynamic inexact Hessian information and applications to finite-sum minimization', *IMA Journal of Numerical Analysis*, 41(1), 764–799.
- [2] S. Bellavia, G. Gurioli, B. Morini, and Ph. L. Toint. (2019) 'Adaptive regularization algorithms with inexact evaluations for nonconvex optimization', *SIAM Journal on Optimization*, 29(4), 2881–2915.
- [3] E. G. Birgin, J. L. Gardenghi, J. M. Martínez, S. A. Santos, and Ph. L. Toint. (2017) 'Worst-case evaluation complexity for unconstrained nonlinear optimization using high-order regularized models', *Mathematical Programming, Series A*, 163(1):359–368.
- [4] E. G. Birgin, and J. M. Martínez. (2017) 'The Use of Quadratic Regularization with a Cubic Descent Condition for Unconstrained Optimization', *SIAM Journal on Optimization*, 27(2): 1049–1074.
- [5] E. G. Birgin and J. M. Martínez. (2019) 'A Newton-like method with mixed factorizations and cubic regularization for unconstrained minimization', *Computational Optimization and Applications*, 73(3): 707–753.
- [6] Y. Carmon and J. C. Duchi. (2021) 'Gradient descent efficiently finds the cubic-regularized non-convex Newton step', *SIAM Journal on Optimization*, 29(3):2146–2178.
- [7] Y. Carmon, J. C. Duchi, O. Hinder, and A. Sidford. (2017) ' "Convex until proven guilty": Dimension-free acceleration of gradient descent on non-convex functions', *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 654–663.
- [8] C. Cartis, J. Fowkes and Z. Shao. (2022) 'Randomised subspace methods for non-convex optimization, with applications to nonlinear least-squares', *arXiv*, 2211.09873.
- [9] C. Cartis, N. I. M. Gould, and Ph. L. Toint. (2009), 'Trust-region and other regularization of linear least-squares problems', *BIT*, 49(1):21–53.
- [10] C. Cartis, N. I. M. Gould, and Ph. L. Toint. (2010), ' On the complexity of steepest descent, Newton's and regularized Newton's methods for nonconvex unconstrained optimization', *SIAM Journal on Optimization*, 20(6):2833–2852.

- [11] C. Cartis, N. I. M. Gould, and Ph. L. Toint. (2011), 'Adaptive cubic overestimation methods for unconstrained optimization. Part I: motivation, convergence and numerical results', *Mathematical Programming, Series A*, 127(2):245–295.
- [12] C. Cartis, N. I. M. Gould, and Ph. L. Toint. (2011), 'Adaptive cubic overestimation methods for unconstrained optimization. Part II: worst-case function-evaluation complexity', *Mathematical Programming, Series A*, 130(2):295–319.
- [13] C. Cartis, N. I. M. Gould, and Ph. L. Toint. (2018), 'Worst-case evaluation complexity and optimality of second-order methods for nonconvex smooth optimization', *Invited Lectures, Proceedings of the 2018 International Conference of Mathematicians (ICM 2018), Rio de Janeiro*, (4):3729–3768.
- [14] C. Cartis, N. I. M. Gould, and Ph. L. Toint. (2020), 'Sharp worst-case evaluation complexity bounds for arbitrary-order nonconvex optimization with inexpensive constraints', *SIAM Journal on Optimization*, 30(1):513–541.
- [15] C. Cartis, N. I. M. Gould, and Ph. L. Toint. (2022), 'Evaluation complexity of algorithms for nonconvex optimization', *Number 30 in MOS-SIAM Series on Optimization. SIAM*.
- [16] C. Cartis and K. Scheinberg. (2018), 'Global convergence rate analysis of unconstrained optimization methods based on probabilistic models', *Mathematical Programming, Series A*, 159(2):337–375.
- [17] Ed. S. Coakley and V. Rokhlin. (2013), 'A fast divide-and-conquer algorithm for computing the spectra of real symmetric tridiagonal matrices', *Applied and Computational Harmonic Analysis*, 34(3): 379–414.
- [18] A. R. Conn, N. I. M. Gould, and Ph. L. Toint. (2000), 'Trust-Region Methods', *Number 1 in MOS-SIAM Optimization Series. SIAM*.
- [19] F. E. Curtis and D. P. Robinson. (2018), 'Exploiting negative curvature in deterministic and stochastic optimization', *Mathematical Programming, Series A*, 176(1-2):69–94.
- [20] F. E. Curtis, D. P. Robinson, C. W. Royer, and S. J. Wright. (2021), 'Trust-region Newton-CG with strong second-order complexity guarantees for nonconvex optimization', *SIAM Journal on Optimization*, 31(1):518–544.
- [21] N. Doikov and Yu. Nesterov. (2023), 'Gradient regularization of Newton method with Bregman distances', *Mathematical Programming*
- [22] E. D. Dolan, J. J. Moré, and T. S. Munson. (2006), 'Optimality measures for performance profiles', *SIAM Journal on Optimization*, 16(3):891–909.
- [23] M. Ferris, S. Lucidi, and M. Roma. (1996), 'Nonmonotone curvilinear line search methods for unconstrained optimization', *Computational Optimization and Applications*, 6(2):117–136.
- [24] D. Goldfarb. The use of negative curvature in minimization algorithms. (1980), *Technical Report TR80-412, Department of Computer Science, Cornell University*.
- [25] S. M. Goldfeldt, R. E. Quandt, and H. F. Trotter. (1966), 'Maximization by quadratic hill-climbing', *Econometrica*, 34:541–551.
- [26] N. I. M. Gould, S. Lucidi, M. Roma, and Ph. L. Toint. (2000), 'Exploiting negative curvature in a linesearch algorithm for unconstrained optimization', *Optimization Methods and Software*, 14(1-2):75–98.
- [27] N. I. M. Gould, S. Lucidi, M. Roma, and Ph. L. Toint. (1999), 'Solving the trust-region subproblem using the Lanczos method', *SIAM Journal on Optimization*, 9(2):504–525.
- [28] N. I. M. Gould, D. Orban and Ph. L. Toint. (2003), 'GALAHAD—a library of thread-safe Fortran 90 packages for large-scale nonlinear optimization', *ACM Transactions on Mathematical Software*, 29(4):353–372.
- [29] N. I. M. Gould, D. Orban and Ph. L. Toint. (2015), 'CUTEst: a Constrained and Unconstrained Testing Environment with safe threads for mathematical optimization', *Computational Optimization and Applications*, 60(3):545–557.
- [30] N. I. M. Gould, M. Porcelli, and Ph. L. Toint. (2012), 'Updating the regularization parameter in the adaptive cubic regularization algorithm', *Computational Optimization and Applications*, 53(1):1–22.
- [31] S. Gratton, S. Jerad, and Ph. L. Toint. (2023), 'Convergence properties of an objective-function-free optimization regularization algorithm, including an $\mathcal{O}(\epsilon^{-3/2})$ complexity bound', *SIAM Journal on Optimization*, 33(3):1621–1646.
- [32] S. Gratton, S. Jerad, and Ph. L. Toint. (2022), 'First-order objective-function-free optimization algorithms and their complexity', *arXiv:2203.01757, 2022*.
- [33] S. Gratton, A. Sartenaer, and Ph. L. Toint. (2008), 'Recursive trust-region methods for multiscale nonlinear optimization', *SIAM Journal on Optimization*, 19(1):414–444.
- [34] S. Gratton and Ph. L. Toint. (2021), 'OPM, a collection of optimization problems in Matlab', *arXiv:2112.05636*.
- [35] S. Gratton and Ph. L. Toint. (2022), 'Adaptive regularization minimization algorithms with non-smooth norms', *IMA Journal of Numerical Analysis*, 43(2):920–949.
- [36] A. Griewank. (1981), 'The modification of Newton's method for unconstrained optimization by bounding cubic terms', *Technical Report NA/12, Department of Applied Mathematics and Theoretical Physics, University of Cambridge*.

- [37] S. Kong and A. S. Lewis. (2022), 'The cost of nonconvexity in deterministic nonsmooth optimization', *arXiv:2210.00652*.
- [38] M. Liu, Z. Li, X. Wang, J. Yi, and T. Yang. (2018), 'Adaptive negative curvature descent with applications in non-convex optimization', In *Advances in Neural Information Processing Systems, vol. 31*.
- [39] S. Lucidi, F. Rochetich, and M. Roma.(1988), 'Curvilinear stabilization techniques for truncated Newton methods in large scale unconstrained optimization', *SIAM Journal on Optimization*, 8(4):916–939.
- [40] G. P. McCormick. (1977), 'A modification of Armijo's step-size rule for negative curvature', *Mathematical Programming*, 13(1):111–115.
- [41] K. Mischenko. (2023) 'Regularized Newton method with global $\mathcal{O}(1/k^2)$ convergence', *SIAM Journal on Optimization.*, 33(3):1440–1462.
- [42] J. J. Moré and D. C. Sorensen. (1979), 'On the use of directions of negative curvature in a modified Newton method', *Mathematical Programming*, 16(1):1–20.
- [43] Yu. Nesterov. (2018), 'Lectures on Convex Optimization', *Springer Verlag, Heidelberg, Berlin, New York*.
- [44] Yu. Nesterov and B. T. Polyak. (2007), 'Cubic regularization of Newton method and its global performance', *Mathematical Programming, Series A*, 108(1):177–205.
- [45] M. Porcelli and Ph. L. Toint. (2019), 'A note on using performance and data profiles for training algorithms', *ACM Transactions on Mathematical Software*, 45(2):1–25.
- [46] Z. Shao. (2021), 'On Random Embeddings and Their Application to Optimization', *University of Oxford*
- [47] C. W. Royer and S. J. Wright. (2018), 'Complexity analysis of second-order line-search algorithms for smooth nonconvex optimization', *SIAM Journal on Optimization*, 28(2):1448–1477.
- [48] K. Ueda and N. Yamashita. (2014), 'A regularized Newton method without line search for unconstrained optimization', *Computational Optimization and Applications*, 59(1-2):321–351.
- [49] P. Xu, F. Roosta-Khorasani, and M. W. Mahoney. (2020), 'Newton-type methods for non-convex optimization under inexact Hessian information', *Mathematical Programming, Series A*, 184((1-2)):35–70.
- [50] Z. Yao, P. Xu, F. Roosta-Khorasani, and M. W. Mahoney. (2021), 'Inexact non-convex Newton-type methods', *INFORMS Journal on Optimization*, 3(2):154–182.
- [51] Z. Yao, P. Xu, F. Roosta-Khorasani, M. W. Mahoney, and S. J. Wright (2021), 'Inexact Newton-CG algorithms with complexity guarantees', *IMA Journal of Numerical Analysis*, 43(3):1855–1897.

Proof of Theorem 4.1

As we noted in Section 4, the step in the SOAN2C algorithm may be computed using (4.2), (4.3) or (4.5). The notations defining the partition of $|\mathcal{S}_k|$ remain relevant, but we complete them by introducing

$$\mathcal{S}^{so} \stackrel{\text{def}}{=} \mathcal{S} \cap \{s_k = s_k^{so}\}, \quad \mathcal{S}_k^{so} \stackrel{\text{def}}{=} \mathcal{S}_k \cap \{s_k = s_k^{so}\}, \quad \mathcal{S}^{fo} \stackrel{\text{def}}{=} \mathcal{S} \setminus \mathcal{S}^{so} \quad \text{and} \quad \mathcal{S}_k^{fo} \stackrel{\text{def}}{=} \mathcal{S}_k \setminus \mathcal{S}_k^{so}.$$

In addition, for $m \geq \ell \geq 0$, we define

$$\mathcal{S}_{\ell,m} \stackrel{\text{def}}{=} \mathcal{S} \cap \{\ell, \dots, m\}$$

and we naturally extend this notation using superscripts identifying the subsets of $\mathcal{S}_{\ell,m}$ corresponding to the different iteration types identified above. We also introduce two index sequences whose purpose is to keep track of when $s_k = s_k^{fo}$ (4.2)-(4.3) or $s_k = s_k^{so}$ (4.5) are used, in the sense that

$$s_k = s_k^{fo} \text{ for } k \in \bigcup_{i \geq 0, p_i \geq 0} \{p_i, \dots, q_i - 1\} \text{ and } s_k = s_k^{so} \text{ for } k \in \bigcup_{i \geq 0} \{q_i, \dots, p_{i+1} - 1\}.$$

Formally,

$$p_0 = \begin{cases} 0 & \text{if } \|g_0\| > \epsilon_1 \\ -1 & \text{if } \|g_0\| \leq \epsilon_1, \end{cases} \quad \text{and} \quad q_0 = \begin{cases} \inf\{k > 0 \mid \|g_k\| \leq \epsilon_1\} & \text{if } \|g_0\| > \epsilon_1 \\ 0 & \text{if } \|g_0\| \leq \epsilon_1. \end{cases} \quad (\text{A.1})$$

Then

$$p_i \stackrel{\text{def}}{=} \inf\{k > q_{i-1} \mid \|g_k\| > \epsilon_1\} \quad \text{and} \quad q_i \stackrel{\text{def}}{=} \inf\{k > p_i \mid \|g_k\| \leq \epsilon_1\} \quad \text{for } i \geq 1. \quad (\text{A.2})$$

The following lemma states an important decrease property holding when (4.5) is used. We also verify that the bound on the regularization parameter derived in Section 3 still applies.

Lemma A.1 Suppose that AS.1 and AS.3 hold. Let $k \in \mathcal{S}^{so}$. Then

$$-g_k^\top s_k - \frac{1}{2} s_k^\top H_k s_k \geq \frac{1}{2} \sigma_k \|s_k\|^3. \quad (\text{A.3})$$

Moreover, the upper bound (3.3) still holds for all $k \geq 0$.

Proof. We obtain from (4.4) and (4.5) that

$$g_k^\top s_k^{so} + \frac{1}{2} (s_k^{so})^\top H_k s_k^{so} \leq \frac{1}{2} \|s_k^{so}\|^2 u_k^\top H_k u_k = \frac{1}{2} \|s_k^{so}\|^2 \lambda_{\min}(H_k) \leq -\frac{1}{2} \sigma_k \|s_k^{so}\|^3,$$

which gives (A.3). As in Lemma 3.1, we now use AS.3, the standard Lipschitz error bound for the function (see [14, Lemma 2.1]) and (A.3) to deduce that

$$1 - \rho_k = \frac{f(x_k + s_k) - f(x_k) - g_k^\top s_k - \frac{1}{2} s_k^\top H_k s_k}{-g_k^\top s_k - \frac{1}{2} s_k^\top H_k s_k} \leq \frac{L_H \|s_k^{so}\|^3}{6(\frac{1}{2} \sigma_k \|s_k^{so}\|^3)} = \frac{L_H}{3\sigma_k}.$$

Thus, if $\sigma_k \geq \frac{L_H}{3(1-\eta_2)}$, we have that $\rho_k \geq \eta_2$ and k is a successful iteration. We may then use the argument of Lemma 3.1 and the fact that ζ_{\max} introduced in (3.4) is larger than two as $V_{\max} \geq 1$. Therefore, we deduce that (3.3) also holds for the SOAN2C algorithm. \square

We now prove an analogue of Lemma 3.1, now using the negative-curvature step as described in (4.4)-(4.5). We also bound the sequence of $\|g_{p_i}\|$.

Lemma A.2 Suppose that AS.1, AS.3 and AS.4 hold. Then, for $k \in \mathcal{S}^{so}$,

$$f(x_k) - f(x_{k+1}) \geq \frac{\eta_1}{2\sigma_{\max}^2} \epsilon_2^3. \quad (\text{A.4})$$

We also have that

$$\|g_{p_i}\| \leq \kappa_{g_{p_i}} \stackrel{\text{def}}{=} \max \left[\|g_0\|, \left(\frac{L_H \kappa_B^2}{2\sigma_{\min}^2} + \frac{\kappa_B^2}{\sigma_{\min}} + 1 \right) \right], \quad (\text{A.5})$$

for all $p_i \geq 0$ as defined in (A.1)-(A.2).

Proof. Let $k \in \mathcal{S}^{so}$. From (2.4) and (A.3), we obtain that

$$f(x_k) - f(x_{k+1}) \geq \eta_1 \left(-g_k^\top s_k - \frac{1}{2} s_k^\top H_k s_k \right) \geq \frac{\eta_1}{2} \sigma_k \|s_k^{so}\|^3.$$

Using now that $\|s_k^{so}\|^3 = \frac{|\lambda_{\min}(H_k)|^3}{\sigma_k^3}$ (see (4.5)) in the previous inequality gives that

$$f(x_k) - f(x_{k+1}) \geq \frac{\eta_1}{2\sigma_k^2} |\lambda_{\min}(H_k)|^3.$$

Now $|\lambda_{\min}(H_k)| \geq \epsilon_2$ when s_k^{so} is computed and $\sigma_k \leq \sigma_{\max}$ by Lemma A.1, from which (A.4) follows. Observe now that (A.5) trivially holds if $p_i = p_0 = 0$. Consider now $p_i > 0$. From the definition of p_i and q_i in (A.2), we see that $p_i - 1 \in \mathcal{S}^{so}$. Using the Lipschitz error bound

for the gradient ([14, Lemma 2.1]), the triangular inequality, (4.4), (4.5), (3.2) (resulting from AS.4), we obtain that

$$\begin{aligned}
 \|g_{p_i}\| &\leq \|g_{p_i} - g_{p_i-1} - H_{p_i-1}s_{p_i-1}^{so}\| + \|H_{p_i-1}s_{p_i-1}^{so} + g_{p_i-1}\| \\
 &\leq \frac{L_H}{2}\|s_{p_i-1}^{so}\|^2 + \|g_{p_i-1}\| + \|H_{p_i-1}s_{p_i-1}^{so}\| \\
 &\leq \frac{L_H|\lambda_{\min}(H_{p_i-1})|^2}{2\sigma_{p_i-1}^2} + \|g_{p_i-1}\| + \frac{|\lambda_{\min}(H_{p_i-1})|\|H_{p_i-1}u_{p_i-1}\|}{\sigma_{p_i-1}} \\
 &\leq \frac{L_H|\lambda_{\min}(H_{p_i-1})|^2}{2\sigma_{p_i-1}^2} + \|g_{p_i-1}\| + \frac{|\lambda_{\min}(H_{p_i-1})|^2}{\sigma_{p_i-1}} \\
 &= \frac{L_H(-\lambda_{\min}(H_{p_i-1}))^2}{2\sigma_{p_i-1}^2} + \|g_{p_i-1}\| + \frac{(-\lambda_{\min}(H_{p_i-1}))^2}{\sigma_{p_i-1}} \\
 &\leq \frac{L_H\kappa_B^2}{2\sigma_{p_i-1}^2} + \|g_{p_i-1}\| + \frac{\kappa_B^2}{\sigma_{p_i-1}}.
 \end{aligned}$$

But $\|g_{p_i-1}\| \leq \epsilon_1 \leq 1$ since $p_i - 1 \in \mathcal{S}^{so}$ and $\sigma_k \geq \sigma_{\min}$ for all $k \geq 0$, which then implies (A.5). \square

In addition to this lemma, all properties of the different steps derived in Section 3 remain valid because these steps are only computed for $\|g_k\| > \epsilon_1$. In particular, (3.10) still applies with $\epsilon = \epsilon_1$. However, (3.24) in Lemma 3.4 may no longer hold because its proof relies on the fact that $\|g_k\| \geq \epsilon_1$, which is no longer true. The purpose of the next lemma is to provide an analogue of (3.24) for the case where SOAN2C is used.

Lemma A.3 Suppose that AS.1, AS.3 and AS.4 hold and the SOAN2C algorithm is used. Consider the partition of $\mathcal{S}_k^{neig} \cup \mathcal{S}_k^{def}$ into $\mathcal{S}_k^{decr} \cup \mathcal{S}_k^{divgrad}$ defined in Lemma 3.4 with the same κ_m (defined in (3.22)). Then (3.23) holds for all $k \in \mathcal{S}_k^{decr}$. Moreover,

$$\begin{aligned}
 |\mathcal{S}_k^{divgrad}| &\leq \kappa_n |\mathcal{S}_k^{decr}| + \left(\frac{1}{2\log(2)} |\log(\epsilon_1)| + \kappa_{curv} \right) |\mathcal{S}_k^{curv}| \\
 &\quad + \left(\frac{|\log(\epsilon_1)| + \log(\kappa_{gpi})}{\log(2)} + 1 \right) (|\mathcal{S}_k^{so}| + 1)
 \end{aligned} \tag{A.6}$$

where κ_n and κ_{curv} are defined in (3.25) and (3.26) and κ_{gpi} is given by (A.5).

Proof. The proof of (3.23) is identical to that used in Lemma 3.4. Moreover, we still obtain (3.27) for $k \in \mathcal{S}_k^{divgrad}$, because the definition of κ_m in (3.22) is unchanged and Lemma A.1 ensures that (3.3) continues to hold for the SOAN2C algorithm.

We now prove (A.6). If \mathcal{S}_k^{fo} is empty, then so is its subset $\mathcal{S}_k^{divgrad}$ and (A.6) trivially holds. If \mathcal{S}_k^{fo} is not empty, we see from the definitions (A.1)-(A.2) that, for some $m \geq 0$ depending on k ,

$$\mathcal{S}_k^{fo} = \{0, \dots, k\} \cap \{\|g_k\| > \epsilon_1\} = \left(\bigcup_{i=0, p_i \geq 0}^{m-1} \{p_i, \dots, q_i - 1\} \right) \cup \{p_m, \dots, k\}. \tag{A.7}$$

Note that the last set in this union is empty unless $k \in \mathcal{S}^{fo}$, in which case $p_m \geq 0$. Suppose first that the set of indices corresponding to the union in brackets is non-empty and let i be an index in this set. Moreover, suppose also that $p_i < q_i - 1$. Using (A.5) and the

facts that $\|g_{q_i-1}\| > \epsilon_1$, that the gradient only changes at successful iterations and that $\mathcal{S}_{p_i, q_i-2} = \mathcal{S}_{p_i, q_i-2}^{curv} \cup \mathcal{S}_{p_i, q_i-2}^{divgrad} \cup \mathcal{S}_{p_i, q_i-2}^{decr}$, we now derive that

$$\begin{aligned} \frac{\epsilon_1}{\kappa_{gpi}} &\leq \frac{\|g_{q_i-1}\|}{\|g_{p_i}\|} = \prod_{j=p_i}^{q_i-2} \frac{\|g_{j+1}\|}{\|g_j\|} = \prod_{j \in \mathcal{S}_{p_i, q_i-2}} \frac{\|g_{j+1}\|}{\|g_j\|} \\ &= \prod_{j \in \mathcal{S}_{p_i, q_i-2}^{decr}} \frac{\|g_{j+1}\|}{\|g_j\|} \prod_{j \in \mathcal{S}_{p_i, q_i-2}^{curv}} \frac{\|g_{j+1}\|}{\|g_j\|} \prod_{j \in \mathcal{S}_{p_i, q_i-2}^{divgrad}} \frac{\|g_{j+1}\|}{\|g_j\|} \\ &\leq \left(\left(\frac{L_H(1 + \kappa_\theta)V_{\max}^3}{2\varsigma_1^2\sigma_{\min}} + \frac{2\kappa_b\sqrt{V_{\max}}}{\varsigma_1} + \kappa_C\kappa_b\sqrt{V_{\max}} \right) (1 + \kappa_\theta) \right)^{|\mathcal{S}_{p_i, q_i-2}^{decr}|} \times \\ &\quad \frac{1}{2^{|\mathcal{S}_{p_i, q_i-2}^{divgrad}|}} \times \left(\frac{L_H V_{\max}^2}{2\sigma_{\min}} \kappa_C^2 \theta^2 x + \frac{\theta^2 \kappa_B \kappa_C}{\sqrt{\epsilon_1 \sigma_{\min}}} + 1 \right)^{|\mathcal{S}_{p_i, q_i-2}^{curv}|} \end{aligned}$$

where we used (3.12), (3.10) and (3.27) to derive the last inequality. Rearranging terms, taking the log, using the inequality $|\mathcal{S}_{p_i, q_i-2}^{divgrad}| \geq |\mathcal{S}_{p_i, q_i-1}^{divgrad}| - 1$ and dividing by $\log(2)$ then gives that

$$\left(|\mathcal{S}_{p_i, q_i-1}^{divgrad}| - 1 \right) + \frac{\log(\epsilon_1) - \log(\kappa_{gpi})}{\log(2)} \leq \kappa_n |\mathcal{S}_{p_i, q_i-2}^{decr}| + \left(\frac{|\log(\epsilon_1)|}{2\log(2)} + \kappa_{curv} \right) |\mathcal{S}_{p_i, q_i-2}^{curv}|$$

with κ_n and κ_{curv} given by (3.25) and (3.26). Further rearranging this inequality and using the fact that $|\mathcal{S}_{p_i, q_i-2}| \leq |\mathcal{S}_{p_i, q_i-1}|$ for the different types of step, we obtain that

$$|\mathcal{S}_{p_i, q_i-1}^{divgrad}| \leq \kappa_n |\mathcal{S}_{p_i, q_i-1}^{decr}| + \left(\frac{|\log(\epsilon_1)|}{2\log(2)} + \kappa_{curv} \right) |\mathcal{S}_{p_i, q_i-1}^{curv}| + \frac{|\log(\epsilon_1)| + \log(\kappa_{gpi})}{\log(2)} + 1. \quad (\text{A.8})$$

If now $p_i = q_i - 1$, then clearly $|\mathcal{S}_{p_i, q_i-1}^{divgrad}| \leq 1$ and (A.8) also holds. Using the same reasoning when $\{p_m, \dots, k\}$ is non-empty, we derive that,

$$|\mathcal{S}_{p_m, k}^{divgrad}| \leq \kappa_n |\mathcal{S}_{p_m, k}^{decr}| + \left(\frac{|\log(\epsilon_1)|}{2\log(2)} + \kappa_{curv} \right) |\mathcal{S}_{p_m, k}^{curv}| + \frac{|\log(\epsilon_1)| + \log(\kappa_{gpi})}{\log(2)} + 1, \quad (\text{A.9})$$

and this inequality also holds if $\{p_m, \dots, k\} = \emptyset$ since $\mathcal{S}_{p_m, k}^{divgrad} \subseteq \{p_m, \dots, k\}$. Adding now (A.8) for $i \in \{0, \dots, m\}$ and (A.9) to take (A.7) into account gives that

$$|\mathcal{S}_k^{divgrad}| \leq \kappa_n |\mathcal{S}_k^{decr}| + \left(\frac{|\log(\epsilon_1)|}{2\log(2)} + \kappa_{curv} \right) |\mathcal{S}_k^{curv}| + \left(\frac{|\log(\epsilon_1)| + \log(\kappa_{gpi})}{\log(2)} + 1 \right) (m + 1).$$

As (A.7) divides \mathcal{S}_k^{fo} into $m + 1$ consecutive sequences, these sequences are then separated by at least a second-order step, so that $m \leq \mathcal{S}_k^{so}$ and (A.6) follows. \square

Equipped with this last lemma and the results of Sections 2 and 3, we may finally establish the worst-case iteration/evaluation complexity of the SOAN2C algorithm and prove Theorem 4.1 itself.

Proof. Note that the bounds (3.32) and (3.33) derived in the proof of Theorem 3.5 are still valid because they only cover steps computed using AN2C, so that we now need to focus on bounding \mathcal{S}_k^{so} . Using AS.2 and the lower bound on the decrease of the function values (A.4), we derive that, for $k \in \mathcal{S}^{so}$,

$$f(x_0) - f_{\text{low}} \geq \sum_{i \in \mathcal{S}_k} f(x_i) - f(x_{i+1}) \geq \sum_{i \in \mathcal{S}_k^{so}} f(x_i) - f(x_{i+1}) \geq |\mathcal{S}_k^{so}| \frac{\eta_1}{2\sigma_{\max}^2} \epsilon_2^3,$$

and therefore that

$$|\mathcal{S}_k^{so}| \leq \frac{2\sigma_{\max}^2(f(x_0) - f_{\text{low}})}{\eta_1} \epsilon_2^{-3} = \kappa_{so} \epsilon_2^{-3}. \quad (\text{A.10})$$

Injecting now (A.10), (3.33) and (3.32) in the bound (A.6) on $\mathcal{S}_k^{divgrad}$ yields that

$$\begin{aligned} |\mathcal{S}_k^{divgrad}| &\leq \kappa_n \kappa_{decr} \epsilon_1^{-\frac{3}{2}} + \left(\frac{|\log(\epsilon_1)|}{2 \log(2)} + \kappa_{curv} \right) \kappa_{negdecr} \epsilon_1^{-\frac{3}{2}} \\ &\quad + \left(\frac{|\log(\epsilon_1)| + \log(\kappa_{gpi})}{\log(2)} + 1 \right) (\kappa_{so} \epsilon_2^{-3} + 1). \end{aligned}$$

Combining the last inequality with (A.10), (3.33) and (3.32) in $|\mathcal{S}_k| = |\mathcal{S}_k^{divgrad}| + |\mathcal{S}_k^{curv}| + |\mathcal{S}_k^{so}| + |\mathcal{S}_k^{decr}|$ and the definition of (3.28) gives that

$$|\mathcal{S}_k| \leq \kappa_* \epsilon_1^{-\frac{3}{2}} + \kappa_{so} \epsilon_2^{-3} + \frac{|\log(\epsilon_1)|}{2 \log(2)} \kappa_{negdecr} \epsilon_1^{-\frac{3}{2}} + \left(\frac{|\log(\epsilon_1)| + \log(\kappa_{gpi})}{\log(2)} + 1 \right) (\kappa_{so} \epsilon_2^{-3} + 1).$$

This proves the first part of the theorem. The second part follows from the last inequality and Lemma 2.1. \square

The factor $|\log(\epsilon_1)|$ by which the bound of Theorem 4.1 differs from $\mathcal{O}(\max(\epsilon_1^{-3/2}, \epsilon_2^{-3}))$ occurs as a consequence of (A.6), (A.10) and (3.32) and one expects that, in practice, (A.10) is smaller than $\mathcal{O}(\epsilon_2^{-3})$ so that Newton steps are taken most often.

The test problems and their dimensions

Problem	n	Problem	n	Problem	n	Problem	n	Problem	n	Problem	n
argauss	3	chebyqad	10	dixmaanl	12	heart8ls	8	msqrtals	16	scurlly10	10
arglina	10	cliff	2	dixon	10	helix	3	msqrtbls	16	scosine	10
arglinb	10	clplatea	16	dqartic	10	hilbert	10	morebv	12	sisser	2
arglinc	10	clplateb	16	edensch	10	himln3	2	nlminsurf	16	spmsqrt	10
argtrig	10	clustr	2	eg2	10	himm25	2	nondquar	10	tcontact	49
arwhead	10	cosine	10	eg2s	10	himm27	2	nzfl	13	tquartic	10
bard	3	crglyv	4	eigenals	12	himm28	2	osbornea	5	trigger	7
bdarwhd	10	cube	2	eigenbls	12	himm29	2	osborneb	11	tridia	10
beale	2	curly10	10	eigencls	12	himm30	3	penalty1	10	tlminsurfx	16
biggs5	5	dixmaana	12	engvall	10	himm32	4	penalty2	10	tlminsurfx	16
biggs6	6	dixmaanb	12	engval2	3	himm33	2	penalty3	10	vardim	10
brownnden	4	dixmaanc	12	expfit	2	hypcir	2	powellbs	2	vibrbeam	8
booth	2	dixmaand	12	extrosnb	10	indef	10	powellsg	12	watson	12
box3	3	dixmaane	12	fminsurf	16	integreq	10	powellsq	2	wmsqrtals	16
brkmcc	2	dixmaanf	12	freuroth	4	jensmp	2	powr	10	wmsqrtbls	16
brownal	10	dixmaang	12	genhumps	5	kowosb	4	recipe	2	woods	12
browns	2	dixmaanh	12	gottfr	2	lminsurf	16	rosenbr	10	yfitu	3
broyden3d	10	dixmaani	12	gulf	4	mancino	10	s308	2	zangwill2	2
broydenbd	10	dixmaanj	12	hairy	2	mexhat	2	sensors	10	zangwill3	3
chandheu	10	dixmaank	12	heart6ls	6	meyer3	3	schmvett	3		

Table A.3: The OPM small test problems and their dimension

Problem	n	Problem	n	Problem	n	Problem	n	Problem	n	Problem	n
arglina	400	crglyv	400	dixmaanl	600	fminsurf	400	ncb20c	500	tcontact	400
arglinb	50	cube	500	dixmaank	600	freuroth	500	nlminsurf	400	tquartic	500
arglinc	50	curly10	500	dixmaanl	600	helix	500	nondquar	500	tridia	500
argtrig	50	deconvu	51	dixon	500	hilbert	500	nzfl	520	tlminsurfx	400
arwhead	500	dixmaana	600	dqartic	500	hydc20ls	99	penalty1	500	tlminsurfx	400
bdarwhd	500	dixmaanb	600	edensch	500	indef	500	penalty2	100	vardim	500
brownal	500	dixmaanc	600	eg2	400	integreq	500	penalty3	500	wmsqrtals	400
broyden3d	500	dixmaand	600	eg2s	400	lminsurf	400	powellsg	500	wmsqrtbls	400
broydenbd	500	dixmaane	600	eigenals	110	msqrtals	400	powr	500	woods	500
chandheu	500	dixmaanf	600	eigenbls	110	msqrtbls	400	rosenbr	100		
chebyqad	150	dixmaang	600	eigencls	110	morebv	500	sensors	100		
clplatea	400	dixmaanh	600	engvall	500	ncb20	500	scosine	500		
clplateb	400	dixmaani	600	extrosnb	500	ncb20b	500	spmsqrt	997		

Table A.4: The OPM medium-size test problems and their dimension

Problem	n	Problem	n	Problem	n	Problem	n	Problem	n
arwhead	2000	dixmaand	2400	eg2	1600	integreq	2000	powellsg	2000
bdarwhd	2000	dixmaane	2400	eg2s	1600	lminsurf	4900	powr	2000
broyden3d	2000	dixmaanf	2400	eg2s	2550	msqrtals	1600	rosenbr	2000
broydenbd	2000	dixmaang	2400	eigenals	2550	msqrtbls	1600	spmsqrt	1498
clplatea	4900	dixmaanh	2400	eigenbls	2550	morebv	5000	tcontact	4900
clplateb	4800	dixmaani	2400	eigencls	2550	ncb20b	2000	tquartic	2000
crglyv	4000	dixmaanj	2400	engvall	2000	ncb20c	2000	tridia	2000
cube	2000	dixmaank	2400	fminsurf	4900	nlminsurf	4900	tlminsurfx	4900
curly10	1000	dixmaanl	2400	freuroth	2000	nondquar	2000	tlminsurfx	4900
dixmaana	2400	dixon	2000	helix	2000	nzfl	2600	vardim	2000
dixmaanb	2400	dqartic	2000	hilbert	2000	penalty1	2000	woods	2000
dixmaanc	2400	edensch	2000	indef	2000	penalty3	2000		

Table A.5: The OPM largish test problems and their dimension