

BLACK-BOX OPTIMIZATION ALGORITHMS FOR REGULARIZED LEAST-SQUARES PROBLEMS*

YANJUN LIU[†], KEVIN H. LAM[‡], AND LINDON ROBERTS[§]

Abstract. We consider the problem of optimizing the sum of a smooth, nonconvex function for which derivatives are unavailable, and a convex, nonsmooth function with easy-to-evaluate proximal operator. Of particular focus is the case where the smooth part has a nonlinear least-squares structure. We adapt two existing approaches for derivative-free optimization of nonsmooth compositions of smooth functions to this setting. Our main contribution is adapting our algorithm to handle inexactly computed stationary measures, where the inexactness is adaptively adjusted as required by the algorithm (where previous approaches assumed access to exact stationary measures, which is not realistic in this setting). Numerically, we provide two extensions of the state-of-the-art DFO-LS solver for nonlinear least-squares problems and demonstrate their strong practical performance.

Key words. derivative-free optimization; nonlinear least-squares; nonsmooth optimization; trust-region methods;

MSC codes. 65K05, 90C30, 90C56

1. Introduction. Derivative-free optimization (DFO), also referred to as black-box optimization, has received growing attention for minimizing a function without accessing its derivative information. The unavailability of derivatives frequently occurs in the field of computer science and engineering with a variety of applications [15]. There are several popular approaches to DFO methods such as direct-search and model-based methods (see [25] for a survey). Here we focus on the derivative-free model-based algorithms based on trust-region methods (see [13, 33]), where at each iteration a trial step is computed (inaccurately) by solving a subproblem of minimizing a model function built through interpolation within a trust region.

In this paper, we consider the minimization of a composite objective function

$$(1.1) \quad \min_{\mathbf{x} \in \mathbb{R}^n} \{\Phi(\mathbf{x}) = f(\mathbf{x}) + h(\mathbf{x})\},$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is smooth and potentially nonconvex and $h : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$ is a convex but possibly nonsmooth regularization term. We assume that the derivatives of f are not accessible. Furthermore, h is Lipschitz continuous in $\text{dom} h := \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}) < \infty\}$ and its proximal operator is assumed to be cheap to evaluate, which might be required for solving the trust-region subproblem. Problems of the form (1.1) are widely studied in data science (e.g. [31]), and in a more traditional DFO context arise for example in image analysis [17].

Motivated by two different algorithmic approaches in [18] and [20], both primarily designed for objectives of the form $f(\mathbf{x}) + h(\mathbf{c}(\mathbf{x}))$ where both f and \mathbf{c} are black boxes,

*Submitted to the editors DATE.

[†]Mathematical Sciences Institute, Building 145, Science Road, The Australian National University, Canberra ACT 2601, Australia. Current address: Department of Operations Research and Financial Engineering, Princeton University, Princeton NJ 08544, USA (yanjun.liu@princeton.edu).

[‡]College of Engineering, Computing and Cybernetics, The Australian National University, 108 North Rd, Acton ACT 2601, Australia (Kevin.Lam@anu.edu.au).

[§]School of Mathematics and Statistics, Carlslaw Building, University of Sydney, Camperdown NSW 2006, Australia (lindon.roberts@sydney.edu.au). This author was supported by the Australian Research Council Discovery Early Career Award DE240100006.

we design two algorithms for solving (1.1) in a derivative-free trust-region framework. Our ultimate aim in this work is to specialize both approaches to (1.1) (i.e. $\mathbf{c}(\mathbf{x}) = \mathbf{x}$) in the specific case of regularized least-squares problems, where $f(\mathbf{x})$ has a nonlinear least-squares structure, $f(\mathbf{x}) = \sum_i r_i(\mathbf{x})^2$ (and derivatives of r_i are not available).

While the approach from [18] transfers readily to this new setting, more work is required for the approach from [20]. Specifically, we must acknowledge that the stationary measure from [20] cannot be computed exactly, despite this quantity being needed in the algorithm. Hence, we need to extend that approach to allow for inexact estimation of the stationary measure. Our new algorithm extending [20] also allows for more standard assumptions on the trust-region subproblem solution and incorporates more sophisticated trust-region management strategies which have been successful in practical DFO codes [30, 6]. We provide a full first-order convergence and worst-case complexity analysis of this, while similar guarantees from [18] follow with limited extra complications. We numerically compare both approaches and provide open-source implementations which extend the smooth least-squares package DFO-LS [6] (see Section 7.1 for details).

1.1. Existing Works. Derivative-based algorithms for minimizing a nonsmooth objective function are well-studied. For example, the proximal point method [5] for general nonsmooth optimization, the proximal gradient method [5] and its accelerated variant [2] for minimizing convex but nonsmooth objective function in the form of (1.1). The main ideas behind these proximal algorithms (see [28] for a comprehensive survey) consist of approximating the nonsmooth structure by a smooth function and applying efficient algorithms for smooth optimization. Based on the trust-region framework, the smoothing trust-region method proposed in [10] uses a sequence of parameterized smoothing functions to approximate the original nonsmooth objective function, where the smoothing parameter is updated before applying the trust region method at each iteration to ensure convergence. Another approach in [7] directly applies the trust-region method to solve a class of nonsmooth, nonconvex optimization problems using an appropriate criticality measure. However, the complexity analysis in [7] builds on the exact evaluations of the trust region subproblem and the criticality measure, which are not required in our algorithms. More recently, a derivative-based trust-region method for solving (1.1) with h nonsmooth and possibly nonconvex was introduced [1].

There are also several model-based DFO methods for nonconvex, nonsmooth objectives, for which the survey [4] provides a thorough discussion (and Mesh Adaptive Direct Search [3] is a widely used direct search method for this setting). For problems where the nonsmoothness arises in a structured way (such as (1.1)), the most common model-based DFO setting is composite objectives of the form $h(\mathbf{c}(\mathbf{x}))$ where h is nonsmooth and \mathbf{c} is a black-box function [18, 20, 21, 24].¹ Here, we are particularly interested in the case where f has a nonlinear least-squares structure, which also has a (smooth) composite objective form. The works [34, 9] exploit this composite structure in the least-squares case. Additionally, [12, 22] study model-based DFO methods for optimisation with convex constraints (i.e. where h is the indicator function of a convex set).

Here, our focus is on adapting two existing approaches to the setting (1.1), and we

¹This structure is sometimes called either ‘gray-box’ or ‘glass-box’ optimization, since the structure of h is typically assumed to be known.

describe both these approaches below. Firstly, [20] considers the composite objective

$$(1.2) \quad \Phi(\mathbf{x}) = f(\mathbf{x}) + h(\mathbf{c}(\mathbf{x})),$$

where f and \mathbf{c} are nonconvex, derivative-free functions, and h is convex but potentially nonsmooth. Here, interpolation-based models are formed for both f and \mathbf{c} and used within a trust-region method. Global convergence and a worst-case complexity bound of $\mathcal{O}(|\log \epsilon| \epsilon^{-2})$ objective evaluations to reach ϵ -approximate first-order stationary are then established. However, because of the difficulty of solving the associated subproblems (the trust-region subproblem and calculating approximate stationary estimates), it was only implemented in the case where $f = 0$ and $h(\mathbf{c}) = \max_i c_i$, in which case the subproblems reduce to linear programs.

Secondly, [18], as well as establishing complexity bounds for smooth model-based DFO, considers two nonsmooth problems:

- Where $\Phi(\mathbf{x})$ is nonsmooth and nonconvex with no specific structure; and
- The composite form (1.2).

In the first case, they consider the setting where Φ can be approximated by a sequence of smooth functions (parametrized by some scalar $\mu \rightarrow 0^+$), and introduce a method where a sequence of smooth problems with decreasing values of μ is solved inexactly (using standard smooth model-based DFO methods). They show global convergence and a worst-case complexity bound of $\mathcal{O}(\epsilon^{-3})$ objective evaluations to achieve ϵ -approximate first-order stationary. In the composite case, they propose a very similar method to [20] but with an improved worst-case complexity bound of $\mathcal{O}(\epsilon^{-2})$ objective evaluations.

1.2. Contributions. We introduce two *implementable* methods based on [20] and [18] for solving (1.1), particularly suited for the case where $f(\mathbf{x})$ has a nonlinear least-squares structure.

Our adaptation of [20] (as specialized to the case $\mathbf{c}(\mathbf{x}) = \mathbf{x}$ in (1.2)) is different in several notable respects:

- It allows inexact calculation of stationary measures throughout the algorithm, where the level of allowable inexactness is adapted to the algorithm's progress and requirements. When a regularization term is present, computing the stationarity measure typically requires solving a convex optimization problem, and so exact evaluations (as required by [20]) are not available, but arbitrarily accurate approximations are available through suitable iterative methods.
- The sufficient decrease condition for the trust-region subproblem in [20] is proportional to the global optimality gap. Here, we propose a new, simpler sufficient decrease condition to compute the trial step, analogous to the standard Cauchy decrease condition for the smooth case (but again relying on inexactly computed stationary estimates).
- The overall algorithmic framework is based on that of DFO-GN [9], which has a more sophisticated trust-region mechanism which has been demonstrated to work well in practice.

We show global convergence and a worst-case complexity analysis of this approach, matching the improved $\mathcal{O}(\epsilon^{-2})$ complexity bound from [18].

Adapting the approach [18] to our setting requires less effort. We demonstrate how to derive suitable smooth approximations to (1.1) based on the Moreau envelope for h , and show that the convergence and complexity results from [18] apply to this setting. We extend the global convergence theory from [18] to additionally show Clarke stationary of all accumulation points of the algorithm.

Both algorithms require the solution of specific subproblems in each iteration. For all subproblems encountered by both methods, we demonstrate how a smoothed variant of FISTA [5, Chapter 10.8.4] can be applied to solve these in practice.

Finally, we implement both techniques for the regularized nonlinear least-squares setting by extending the state-of-the-art solver DFO-LS [6], which only handles the cases where $h = 0$ or h is the indicator function of a convex set [22]. Both methods outperform NOMAD [26] (which implements the mesh adaptive direct search method) on this problem class.

Code Availability. Both adaptations of DFO-LS are available on Github.²

Structure of paper. We first introduce our variant of [20] in Section 2. Our new convergence and complexity analysis of this approach is given in Section 3, and we then explicitly adapt this method to the nonlinear least-squares case in Section 4. We then describe our smoothing variant of [18] in Section 5. Our approach for calculating approximate subproblem solutions is given in Section 6, and numerical results are provided in Section 7.

Notation. We use $\|\cdot\|$ to be the Euclidean norm of vectors and the operator 2-norm of matrices, and $B(\mathbf{x}, \Delta) := \{\mathbf{y} \in \mathbb{R}^n : \|\mathbf{y} - \mathbf{x}\| \leq \Delta\}$ to be the closed ball centered at $\mathbf{x} \in \mathbb{R}^n$ with radius $\Delta > 0$.

2. Algorithmic Framework. In this section, we outline the general model-based DFO algorithmic framework for (1.1). Our framework follows the structure given in [9], except the modifications for including a regularization term based on [20]. We first introduce the first-order criticality measure for (1.1), and then present our algorithm.

2.1. Criticality Measure. Similar to [7], we linearize $f(\mathbf{x})$ and the argument of h around any \mathbf{x} to give an approximation of Φ

$$(2.1) \quad l(\mathbf{x}, \mathbf{s}) := f(\mathbf{x}) + \nabla f(\mathbf{x})^T \mathbf{s} + h(\mathbf{x} + \mathbf{s}), \mathbf{s} \in \mathbb{R}^n.$$

Then we define the quantity

$$(2.2) \quad \Psi_r(\mathbf{x}) := l(\mathbf{x}, \mathbf{0}) - \min_{\|\mathbf{s}\| \leq r} l(\mathbf{x}, \mathbf{s}).$$

Following [32], $\Psi_r(\mathbf{x})$ is non-negative and continuous for all $\mathbf{x} \in \mathbb{R}^n$, and we say that \mathbf{x}^* is a critical point of Φ if

$$(2.3) \quad \Psi_1(\mathbf{x}^*) = 0.$$

The condition (2.3) is equivalent to other first-order optimal necessary conditions (see [32, Lemma 2.1] or the discussion in [7, Section 2], for example).

However, in the DFO setting, we cannot calculate Ψ_1 because ∇f is not accessible. Instead, in our algorithm we follow [14] and approximate f with a model constructed by interpolation to points near the current iterate. Let \mathbf{x}_0 denote the initial iterate and suppose that at k -th iteration, we form the model through sampling within a closed ball $B(\mathbf{x}_k, \Delta_k)$, where Δ_k is bounded above by Δ_{\max} . Our notion of model accuracy is the following [14]:

²From <https://github.com/yanjunliu-regina/dfols> and <https://github.com/khflam/dfols/>. See Section 7.1 for details of both implementations.

DEFINITION 2.1. A model $p_k \in C^1$ for $f \in C^1$ is fully linear in $B(\mathbf{x}_k, \Delta_k)$ if for any $\mathbf{y} \in B(\mathbf{x}_k, \Delta_k)$,

$$(2.4) \quad |f(\mathbf{y}) - p_k(\mathbf{y})| \leq \kappa_{\text{ef}} \Delta_k^2,$$

$$(2.5) \quad \|\nabla f(\mathbf{y}) - \nabla p_k(\mathbf{y})\| \leq \kappa_{\text{eg}} \Delta_k,$$

where κ_{ef} and κ_{eg} are independent of \mathbf{y} , \mathbf{x}_k and Δ_k .

The local fully linear model p_k (for f) induces an approximate criticality measure below, where in (2.1) we replace ∇f with ∇p_k , as used in [20].

DEFINITION 2.2. Given a continuously differentiable model $p_k : \mathbb{R}^n \rightarrow \mathbb{R}$ for f , for each $\mathbf{x} \in \mathbb{R}^n$ we define

$$(2.6) \quad \tilde{l}(\mathbf{x}, \mathbf{s}) := f(\mathbf{x}) + \nabla p_k(\mathbf{x})^T \mathbf{s} + h(\mathbf{x} + \mathbf{s}), \mathbf{s} \in \mathbb{R}^n,$$

and for all $r > 0$, we define

$$(2.7) \quad \eta_r(\mathbf{x}) := \tilde{l}(\mathbf{x}, \mathbf{0}) - \min_{\|\mathbf{s}\| \leq r} \tilde{l}(\mathbf{x}, \mathbf{s}).$$

We choose η_1 as the *approximate criticality measure* for our algorithm. When $h \equiv 0$, this η_1 reduces to $\|\mathbf{g}_k\|$ as expected, and when $h \equiv I_C$ is an indicator function for a convex set C , η_1 reduces to the approximate criticality measure chosen in [22].

Fortunately, if p_k is a fully linear model for f , the error between the true criticality measure Ψ_1 and our approximation η_1 is controlled.

LEMMA 2.3. Suppose that $f \in C^1$ and h is continuous. Assume that $p_k : \mathbb{R}^n \rightarrow \mathbb{R}$ is a fully linear model of f with respect to constant κ_{ef} and κ_{eg} on the ball $B(\mathbf{x}_k, \Delta_k)$. Then

$$(2.8) \quad |\Psi_1(\mathbf{y}) - \eta_1(\mathbf{y})| \leq \kappa_{\text{eg}} \Delta_k,$$

for any $\mathbf{y} \in B(\mathbf{x}_k, \Delta_k)$.

Proof. This proof is based on [20, Theorem 1]. Take any $\mathbf{y} \in B(\mathbf{x}_k, \Delta_k)$. Since p_k is a fully linear model of f on $B(\mathbf{x}_k, \Delta_k)$, it follows from (2.5) that

$$(2.9) \quad \|\nabla f(\mathbf{y}) - \nabla p_k(\mathbf{y})\| \leq \kappa_{\text{eg}} \Delta_k.$$

Since $l(\mathbf{y}, \mathbf{s})$ is continuous with respect to \mathbf{s} on $B(\mathbf{0}, 1)$, consider $\tilde{\mathbf{s}} \in B(\mathbf{0}, 1)$ such that

$$(2.10) \quad \min_{\|\mathbf{s}\| \leq 1} l(\mathbf{y}, \mathbf{s}) = l(\mathbf{y}, \tilde{\mathbf{s}}).$$

Then from (2.2), (2.7) and (2.9), it follows that

$$\begin{aligned} \Psi_1(\mathbf{y}) - \eta_1(\mathbf{y}) &= \left(l(\mathbf{y}, \mathbf{0}) - \min_{\|\mathbf{s}\| \leq 1} l(\mathbf{y}, \mathbf{s}) \right) - \left(\tilde{l}(\mathbf{y}, \mathbf{0}) - \min_{\|\mathbf{s}\| \leq 1} \tilde{l}(\mathbf{y}, \mathbf{s}) \right), \\ &= \min_{\|\mathbf{s}\| \leq 1} \tilde{l}(\mathbf{y}, \mathbf{s}) - \min_{\|\mathbf{s}\| \leq 1} l(\mathbf{y}, \mathbf{s}), \\ &= \min_{\|\mathbf{s}\| \leq 1} \tilde{l}(\mathbf{y}, \mathbf{s}) - l(\mathbf{y}, \tilde{\mathbf{s}}), \\ &\leq \tilde{l}(\mathbf{y}, \tilde{\mathbf{s}}) - l(\mathbf{y}, \tilde{\mathbf{s}}), \\ &= (\nabla p(\mathbf{y}) - \nabla f(\mathbf{y}))^T \tilde{\mathbf{s}}, \\ &\leq \|\nabla p(\mathbf{y}) - \nabla f(\mathbf{y})\|, \\ (2.11) \quad &\leq \kappa_{\text{eg}} \Delta_k. \end{aligned}$$

Similarly, considering $\bar{\mathbf{s}} \in B(\mathbf{0}, 1)$ such that

$$(2.12) \quad \min_{\|\mathbf{s}\| \leq 1} \tilde{l}(\mathbf{y}, \mathbf{s}) = l(\mathbf{y}, \bar{\mathbf{s}}),$$

We obtain the inequality

$$(2.13) \quad \eta_1(\mathbf{y}) - \Psi_1(\mathbf{y}) \leq \kappa_{\text{eg}} \Delta_k.$$

Hence, combining (2.11) and (2.13), we conclude that (2.8) holds. \square

2.2. Main Algorithm. Our algorithm for solving (1.1) is based on a trust-region framework [13]. At each iteration, we construct a model m_k to approximate our objective function Φ around the current iterate \mathbf{x}_k , and in particular within a trust region $B(\mathbf{x}_k, \Delta_k)$. Then we find a tentative new point $\mathbf{x}_k + \mathbf{s}_k$ by approximately solving the trust-region subproblem

$$(2.14) \quad \mathbf{s}_k \approx \arg \min_{\|\mathbf{s}\| \leq \Delta_k} m_k(\mathbf{x}_k + \mathbf{s}).$$

We measure the sufficient decrease in the objective function using the ratio

$$(2.15) \quad R_k := \frac{\Phi(\mathbf{x}_k) - \Phi(\mathbf{x}_k + \mathbf{s}_k)}{m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k)},$$

which is used to determine the next iterate \mathbf{x}_{k+1} and update the trust region radius $\Delta_k > 0$. If R_k is sufficiently large, we accept this step (i.e. $\mathbf{x}_{k+1} = \mathbf{x}_k + \mathbf{s}_k$) and increase Δ_k , otherwise we reject this step (i.e. $\mathbf{x}_{k+1} = \mathbf{x}_k$) and decrease Δ_k if necessary.

To construct our model, we first take a quadratic approximation p_k for f

$$(2.16) \quad f(\mathbf{x}_k + \mathbf{s}) \approx p_k(\mathbf{x}_k + \mathbf{s}) = f(\mathbf{x}_k) + \mathbf{g}_k^T \mathbf{s} + \frac{1}{2} \mathbf{s}^T H_k \mathbf{s},$$

for some choices $\mathbf{g}_k \in \mathbb{R}^n$ and $H_k \in \mathbb{R}^{n \times n}$. We can then construct our model m_k for Φ as

$$(2.17) \quad \Phi(\mathbf{x}_k + \mathbf{s}) \approx m_k(\mathbf{x}_k + \mathbf{s}) := p_k(\mathbf{x}_k + \mathbf{s}) + h(\mathbf{x}_k + \mathbf{s}).$$

This construction relies on our assumption that the regularizer h is known and easy to evaluate. We note the similarity between (2.17) and the more general model used for composite nonsmooth optimization in [20].

In practice, the model p_k is constructed using the techniques from [15]: we maintain a set of interpolation points Y_k containing \mathbf{x}_k , and approximate \mathbf{g}_k and H_k by enforcing the interpolation conditions

$$(2.18) \quad f(\mathbf{y}) = p_k(\mathbf{y}), \quad \forall \mathbf{y} \in Y_k.$$

This model construction allows p_k to be fully linear provided Y_k satisfies certain geometric conditions:

DEFINITION 2.4. Let $\{\Lambda_1(\mathbf{x}), \dots, \Lambda_{|Y_k|}(\mathbf{x})\}$ be the set of Lagrange polynomials for Y_k . For some $\Lambda \geq 1$, Y_k is Λ -poised in some ball $B(\mathbf{x}, \Delta)$ if $Y_k \subset B(\mathbf{x}, \Delta)$ and

$$(2.19) \quad \max_{t=1, \dots, |Y_k|} \max_{\mathbf{y} \in B(\mathbf{x}, \Delta)} |\Lambda_t(\mathbf{y})| \leq \Lambda.$$

Algorithm 2.1 DFO-LSR: a model-based DFO method for (1.1).

Input: Starting point $\mathbf{x}_0 \in \text{dom } h$ and the initial trust-region radius $\Delta_0^{\text{init}} > 0$.

Parameters: maximum trust-region radius $\Delta_{\max} \geq \Delta_0^{\text{init}}$ with $\Delta_{\max} > 1$, trust region radius scaling factors $0 < \gamma_{\text{dec}} < 1 < \gamma_{\text{inc}} \leq \bar{\gamma}_{\text{inc}}$ and $0 < \alpha_1 < \alpha_2 < 1$, criticality constants $\epsilon_C, \mu > 0$, acceptance thresholds $0 < \beta_1 \leq \beta_2 < 1$, trust region subproblem constant $c_1 := \min\{1, \Delta_{\max}^{-2}\}/2$, safety constants $0 < \omega_S < 1$ and $0 < \gamma_S < 2e_3c_1/(1 + \sqrt{1 + 2e_3c_1})$, poisedness constant $\Lambda \geq 1$, accuracy level $0 < e_1 < 1$, $e_2 > 0$ and $0 < e_3 < 1$.

- 1: Build an initial interpolation set Y_0 of size $n + 1$, with $\mathbf{x}_0 \in Y_0$. Set $\rho_0^{\text{init}} = \Delta_0^{\text{init}}$.
- 2: **for** $k = 0, 1, 2, \dots$ **do**
- 3: Given \mathbf{x}_k and Y_k , solve the interpolation problem in (2.18) to get J_k^{init} and build the model m_k^{init} in (2.17).
- 4: Compute $\eta_1^{\text{init}}(\mathbf{x}_k) = \bar{l}^{\text{init}}(\mathbf{x}_k, \mathbf{0}) - \bar{l}^{\text{init}}(\mathbf{x}_k, \mathbf{d}_k)$, where

$$(2.20) \quad \mathbf{d}_k \approx \arg \min_{\|\mathbf{d}\| \leq 1} \bar{l}^{\text{init}}(\mathbf{x}_k, \mathbf{d}),$$

and

$$(2.21) \quad \bar{l}^{\text{init}}(\mathbf{x}_k, \mathbf{d}) = f(\mathbf{x}_k) + (\mathbf{g}_k^{\text{init}})^T \mathbf{d} + h(\mathbf{x}_k + \mathbf{d}).$$

Here \mathbf{d}_k is calculated approximately so that the approximate estimate $\bar{\eta}_1^{\text{init}}(\mathbf{x}_k)$ satisfies

$$(2.22) \quad \eta_1^{\text{init}}(\mathbf{x}_k) - \bar{\eta}_1^{\text{init}}(\mathbf{x}_k) < \min\{(1 - e_1)\epsilon_C, e_2\Delta_k^{\text{init}}\}$$

- 5: **if** $\bar{\eta}_1^{\text{init}}(\mathbf{x}_k) \leq e_1\epsilon_C$ **then**
- 6: **Criticality Phase:** using Algorithm 2.2, modify Y_k and find $\Delta_k \leq \Delta_k^{\text{init}}$ and $\bar{\eta}_1(\mathbf{x}_k)$ such that Y_k is Λ -poised in $B(\mathbf{x}_k, \Delta_k)$, $\Delta_k \leq \mu\bar{\eta}_1(\mathbf{x}_k)$ and $\eta_1(\mathbf{x}_k) - \bar{\eta}_1(\mathbf{x}_k) < e_2\Delta_k$. Set $\rho_k = \min(\rho_k^{\text{init}}, \Delta_k)$.
- 7: **else**
- 8: Set $m_k = m_k^{\text{init}}$, $\Delta_k = \Delta_k^{\text{init}}$, $\rho_k = \rho_k^{\text{init}}$ and $\bar{\eta}_1(\mathbf{x}_k) = \bar{\eta}_1^{\text{init}}(\mathbf{x}_k)$
- 9: **end if**
- 10: Approximately solve (2.14) to get a step \mathbf{s}_k satisfying $\|\mathbf{s}_k\| \leq \Delta_k$, $m_k(\mathbf{x}_k + \mathbf{s}_k) \leq m_k(\mathbf{x}_k)$ and

$$(2.23) \quad m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k) \geq e_3c_1\bar{\eta}_1(\mathbf{x}_k) \min \left\{ \Delta_k, \frac{\bar{\eta}_1(\mathbf{x}_k)}{\max\{1, \|H_k\|\}} \right\}.$$

- 11: Calculate $\tau_k := \min\{\bar{\eta}_1(\mathbf{x}_k)/(\|\mathbf{g}_k\| + L_h), 1\}$.
- 12: **if** $\|\mathbf{s}_k\| < \tau_k\gamma_S\rho_k$ **then**
- 13: **Safety Phase:** Set $\mathbf{x}_{k+1} = \mathbf{x}_k$ and $\Delta_{k+1}^{\text{init}} = \max\{\rho_k, \omega_S\Delta_k\}$, and form Y_{k+1} by making Y_k Λ -poised in $B(\mathbf{x}_{k+1}, \Delta_{k+1}^{\text{init}})$.
- 14: **If** $\Delta_{k+1}^{\text{init}} = \rho_k$, set $(\rho_{k+1}^{\text{init}}, \Delta_{k+1}^{\text{init}}) = (\alpha_1\rho_k, \alpha_2\rho_k)$, otherwise set $\rho_{k+1}^{\text{init}} = \rho_k$.
- 15: **goto** Line 2.
- 16: **end if**
- 17: Evaluate $f(\mathbf{x}_k + \mathbf{s}_k)$ and calculate ratio R_k in (2.15)
- 18: Accept/reject step and update trust region radius: set

$$\mathbf{x}_{k+1} = \begin{cases} \mathbf{x}_k + \mathbf{s}_k & R_k \geq \beta_1, \\ \mathbf{x}_k & R_k < \beta_1, \end{cases} \quad \text{and } \Delta_{k+1}^{\text{init}} = \begin{cases} \min(\max(\gamma_{\text{inc}}\Delta_k, \bar{\gamma}_{\text{inc}}\|\mathbf{s}_k\|), \Delta_{\max}), & R_k \geq \beta_2, \\ \max(\gamma_{\text{dec}}\Delta_k, \|\mathbf{s}_k\|, \rho_k), & \beta_1 \leq R_k < \beta_2, \\ \max(\min(\gamma_{\text{dec}}\Delta_k, \|\mathbf{s}_k\|)/\tau_k, \rho_k), & R_k < \beta_1. \end{cases}$$

- 19: **if** $R_k \geq \beta_1$ **then**
 - 20: Form $Y_{k+1} = Y_k \cup \{\mathbf{x}_{k+1}\} \setminus \{\mathbf{y}_t\}$ for some $\mathbf{y}_t \in Y_k$ and set $\rho_{k+1}^{\text{init}} = \rho_k$.
 - 21: **else if** Y_k is not Λ -poised in $B(\mathbf{x}_k, \Delta_k)$ **then**
 - 22: **Model Improvement Phase:** Form Y_{k+1} by making Y_k Λ -poised in $B(\mathbf{x}_{k+1}, \Delta_{k+1}^{\text{init}})$ and set $\rho_{k+1}^{\text{init}} = \rho_k$.
 - 23: **else** [$R_k < \beta_1$ and Y_k is Λ -poised in $B(\mathbf{x}_k, \Delta_k)$]
 - 24: **Unsuccessful Phase:** Set $Y_{k+1} = Y_k$, and if $\Delta_{k+1}^{\text{init}} = \rho_k$, set $(\rho_{k+1}^{\text{init}}, \Delta_{k+1}^{\text{init}}) = (\alpha_1\rho_k, \alpha_2\rho_k)$, otherwise set $\rho_{k+1}^{\text{init}} = \rho_k$.
 - 25: **end if**
 - 26: **end for**
-

Intuitively, the set Y_k has ‘good geometry’ if Y_k is Λ -poised with a small Λ . The exact form of interpolation used to construct p_k is left open for now, and will be revisited in the context of regularized nonlinear least-squares problems in Section 4.

For now, in light of the techniques given in [15, Chapter 6] we simply assume the existence of procedures which can verify whether or not a set Y_k is Λ -poised, and if it is not, add and remove points from Y_k until it becomes Λ -poised.

Remark 2.5. If p_k is a linear model, the techniques from [22, Section 4] allow Definition 2.4 to be weakened to only maximizing over $\mathbf{y} \in B(\mathbf{x}, \Delta) \cap C$ in (2.19) for any closed convex set C with nonempty interior. In our case, we might take $C = \text{dom } h$, for example.

Our main algorithm for solving (1.1) is presented in Algorithm 2.1. The overall structure is based on DFO-GN [9, Algorithm 1], which is designed for nonlinear least-squares minimization (c.f. Section 4). Motivated by the practical performance of DFO-GN, we keep important algorithmic features not present in other similar methods (e.g. [15, Chapter 10]) such as the safety phase and the maintenance of a lower bound ρ_k on the trust region radius Δ_k .

In particular, we extend the safety phase from DFO-GN (which originates in [29]), which provides a way to detect insufficient decrease generated by the step size $\|\mathbf{s}_k\|$ before evaluating $f(\mathbf{x}_k)$ and R_k (2.15). Specifically, to accommodate the proof of Lemma 3.8, we introduce a new variable τ_k in the generalized safety phase. This τ_k is calculated in line 11 and then used in the entry condition (line 12) and the update rule for Δ_k (lines 18 and 24). If $h \equiv 0$, then $\tau_k = 1$ and the generalised safety phase is equivalent to the original safety phase from DFO-GN.

Of course, we also use our new criticality measure $\eta_1(\mathbf{x}_k)$ instead of $\|\mathbf{g}_k\|$ in DFO-GN algorithm, just as in [20]. Unfortunately, to calculate the value of $\eta_1(\mathbf{x}_k)$, we need to solve a minimization subproblem as defined in (2.7). So to make our framework practical, we therefore extend [20] to allow an approximate estimate of $\eta_1(\mathbf{x}_k)$ up to a predetermined accuracy level (line 4). Note that if we compute \mathbf{d}_k in (2.20) inexactly, then we automatically have $\bar{\eta}_1(\mathbf{x}_k) \leq \eta_1(\mathbf{x}_k)$.

With the nonsmooth term in our model m_k (2.17), our trust-region subproblem (2.14) and criticality estimation subproblem (2.20) are not straightforward to solve to the required accuracy. We discuss this issue further in Sections 4 and 6.

We describe the geometry-improvement step used in the criticality phase below, which is an adaptation of [9, Algorithm 2]. Compared to Algorithm 2.1, the criticality measure is evaluated approximately up to a different accuracy in line 5.

3. Convergence and Complexity Analysis. Throughout this section, we consider the following standard assumptions:

ASSUMPTION 3.1. *The function f (1.1) is bounded below by f_{low} and continuously differentiable in the convex hull of $\mathcal{B} := \cup_k B(\mathbf{x}_k, \Delta_{\text{max}})$, where ∇f is $L_{\nabla f}$ -Lipschitz continuous in \mathcal{B} .*

ASSUMPTION 3.2. *$h : \mathbb{R}^n \rightarrow \mathbb{R}$ is convex and possibly nonsmooth. We also assume that $h(\mathbf{x})$ is bounded below by h_{low} and Lipschitz continuous with Lipschitz constant L_h in $\text{dom } h$.*

We also need the assumption that the model Hessians are uniformly bounded above for the trust-region subproblem (2.14).

ASSUMPTION 3.3. *There exists $\kappa_H \geq 1$ such that $\|H_k\| \leq \kappa_H - 1$ for all k .*

Lastly, we use the below result to link the Λ -poisedness of Y_k with the accuracy of the model p_k .

LEMMA 3.4. *If Assumption 3.1 holds and Y_k is Λ -poised in $B(\mathbf{x}_k, \Delta_k)$, then p_k is a fully linear model for f .*

Algorithm 2.2 Geometry-Improvement for Criticality Phase

Input: Iterate \mathbf{x}_k , initial set Y_k and trust region radius Δ_k^{init} .

Parameters: $\mu > 0$, $\omega_C \in (0, 1)$ and poisedness constant $\Lambda > 0$.

- 1: Set $Y_k^{(0)} = Y_k$.
- 2: **for** $i=1, 2, \dots$ **do**
- 3: Form $Y_k^{(i)}$ by modifying $Y_k^{(i-1)}$ until it is Λ -poised in $B(\mathbf{x}_k, \omega_C^{i-1} \Delta_k^{\text{init}})$.
- 4: Solve the interpolation system for $Y_k^{(i)}$ to get $J_k^{(i)}$, and form $m_k^{(i)}$ in (2.17).
- 5: Compute $\eta_1^{(i)}(\mathbf{x}_k) = \bar{l}^{(i)}(\mathbf{x}_k, \mathbf{0}) - \bar{l}^{(i)}(\mathbf{x}_k, \tilde{\mathbf{d}}_k)$, where

$$(2.24) \quad \tilde{\mathbf{d}}_k \approx \arg \min_{\|\mathbf{d}\| \leq 1} \bar{l}^{(i)}(\mathbf{x}_k, \mathbf{d}),$$

and

$$(2.25) \quad \bar{l}^{(i)}(\mathbf{x}_k, \mathbf{d}) = f(\mathbf{x}_k) + (\mathbf{g}_k^{(i)})^T \mathbf{d} + h(\mathbf{x}_k + \mathbf{d}).$$

Here $\tilde{\mathbf{d}}_k$ is calculated approximately so that the approximate estimate $\bar{\eta}_1^{(i)}(\mathbf{x}_k)$ satisfies

$$(2.26) \quad \eta_1^{(i)}(\mathbf{x}_k) - \bar{\eta}_1^{(i)}(\mathbf{x}_k) < e_2 \omega_C^{i-1} \Delta_k^{\text{init}}$$

- 6: **if** $\omega_C^{i-1} \Delta_k^{\text{init}} \leq \mu \bar{\eta}_1^{(i)}(\mathbf{x}_k)$ **then**
 - 7: **return** $Y_k^{(i)}$, $m_k^{(i)}$, $\Delta_k \leftarrow \omega_C^{i-1} \Delta_k^{\text{init}}$, $\bar{\eta}_1(\mathbf{x}_k) \leftarrow \bar{\eta}_1^{(i)}(\mathbf{x}_k)$.
 - 8: **end if**
 - 9: **end for**
-

Proof. Different versions of this result are applicable depending on the specific model construction used for p_k , e.g. [15, Theorems 2.11 & 2.12 or Theorem 5.4]. When we specialize to nonlinear least-squares problems in Section 4, this will be given by Lemma 4.3. \square

3.1. Global Convergence Analysis. The following lemma ensures that unless the current iterate is a critical point, Algorithm 2.2 for the criticality phase terminates in finite loops. It also provides a bound for the trust region radius Δ_k .

LEMMA 3.5. *Suppose Assumptions 3.1 and 3.2 hold and $\Psi_1(\mathbf{x}_k) \geq \epsilon > 0$. Then for any $\mu > 0$ and $\omega_C \in (0, 1)$, the criticality phase in Algorithm 2.2 terminates in finite time with Y_k Λ -poised in $B(\mathbf{x}_k, \Delta_k)$ and $\Delta_k \leq \mu \bar{\eta}_1(\mathbf{x}_k)$ for any $\mu > 0$ and $\eta_1(\mathbf{x}_k) - \bar{\eta}_1(\mathbf{x}_k) < e_2 \Delta_k$. We also have the bound*

$$(3.1) \quad \min \left(\Delta_k^{\text{init}}, \frac{\omega_C \epsilon}{\kappa_{\text{eg}} + 1/\mu + e_2} \right) \leq \Delta_k \leq \Delta_k^{\text{init}}.$$

Proof. This proof is similar to [9, Lemma B.1] except that we allow the inaccurate estimate of our criticality measure η_1 . First, suppose Algorithm 2.2 terminates on the first iteration. Then $\Delta_k = \Delta_k^{\text{init}}$, and the result holds. Otherwise, consider some iterations i where Algorithm 2.2 does not terminate, that is, where $\omega_C^{i-1} \Delta_k^{\text{init}} > \mu \bar{\eta}_1^{(i)}(\mathbf{x}_k)$. Then since $m_k^{(i)}$ is fully linear in $B(\mathbf{x}_k, \omega_C^{i-1} \Delta_k^{\text{init}})$, from Lemma 2.3 and (2.26), we have

$$(3.2) \quad \epsilon \leq \Psi_1(\mathbf{x}_k) = \Psi_1(\mathbf{x}_k) - \eta_1^{(i)}(\mathbf{x}_k) + \eta_1^{(i)}(\mathbf{x}_k)$$

$$(3.3) \quad \leq \kappa_{\text{eg}} \omega_C^{i-1} \Delta_k^{\text{init}} + \bar{\eta}_1^{(i)}(\mathbf{x}_k) + e_2 \omega_C^{i-1} \Delta_k^{\text{init}}$$

$$(3.4) \quad \leq \left(\kappa_{\text{eg}} + \frac{1}{\mu} + e_2 \right) \omega_C^{i-1} \Delta_k^{\text{init}},$$

or equivalently $\omega_C^{i-1} \geq \frac{\epsilon}{(\kappa_{\text{eg}} + 1/\mu + e_2)\Delta_k^{\text{init}}}$. That is, if termination does not occur on iteration i , we must have

$$(3.5) \quad i \leq 1 + \frac{1}{|\log \omega_C|} \log \left(\frac{(\kappa_{\text{eg}} + 1/\mu + e_2)\Delta_k^{\text{init}}}{\epsilon} \right),$$

so Algorithm 2.2 terminates finitely. We also have $\omega_C^{i-1}\Delta_k^{\text{init}} \geq \frac{\epsilon}{\kappa_{\text{eg}} + 1/\mu + e_2}$, which gives (3.1). \square

Next, we show that having a fully linear model p_k for f guarantees a condition similar to (2.4) for our full (nonsmooth) model m_k (2.17).

LEMMA 3.6. *Suppose Assumptions 3.1, 3.2 and 3.3 hold. If p_k is a fully linear model of f on the ball $B(\mathbf{x}_k, \Delta_k)$, then*

$$(3.6) \quad |\Phi(\mathbf{x}_k + \mathbf{s}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k)| \leq \kappa_{\text{ef}}\Delta_k^2.$$

Proof. From (2.23), \mathbf{s}_k must be calculated such that $\mathbf{x}_k + \mathbf{s}_k \in \text{dom } h$. Since $\mathbf{x}_0 \in \text{dom } h$ by definition, we must have $\mathbf{x}_k \in \text{dom } h$ for all k . Hence by definition of Φ and m_k we have

$$(3.7) \quad \Phi(\mathbf{x}_k + \mathbf{s}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k) = f(\mathbf{x}_k + \mathbf{s}_k) - p_k(\mathbf{x}_k + \mathbf{s}_k)$$

and the result follows from (2.4). \square

In Algorithm 2.1, if the safety phase is not called, we say an iteration is ‘successful’ if $R_k \geq \beta_1$ and ‘very successful’ if $R_k \geq \beta_2$.

LEMMA 3.7. *Suppose Assumptions 3.1, 3.2 and 3.3 hold. If p_k is a fully linear model of f on the ball $B(\mathbf{x}_k, \Delta_k)$, and*

$$(3.8) \quad \Delta_k \leq \min \left\{ \frac{1}{\kappa_H}, \frac{e_3 c_1 (1 - \beta_2)}{\kappa_{\text{ef}}} \right\} \bar{\eta}_1(\mathbf{x}_k),$$

then, the k -th iteration is either very successful or the safety phase is called.

Proof. This proof is based on [14, Lemma 5.3]. Since $\|H_k\| \leq \kappa_H$ and $\Delta_k \leq \frac{\bar{\eta}_1(\mathbf{x}_k)}{\kappa_H}$, it follows from Lemma 6.2 that

$$\begin{aligned} m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k) &\geq e_3 c_1 \bar{\eta}_1(\mathbf{x}_k) \min \left\{ \Delta_k, \frac{\bar{\eta}_1(\mathbf{x}_k)}{\max\{1, \|H_k\|\}} \right\} \\ &\geq e_3 c_1 \bar{\eta}_1(\mathbf{x}_k) \min \left\{ \Delta_k, \frac{\bar{\eta}_1(\mathbf{x}_k)}{\kappa_H} \right\} \\ &= e_3 c_1 \bar{\eta}_1(\mathbf{x}_k) \Delta_k. \end{aligned}$$

Then, applying Lemma 3.6, we obtain

$$\begin{aligned} |1 - R_k| &= \left| \frac{(m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k)) - (\Phi(\mathbf{x}_k) - \Phi(\mathbf{x}_k + \mathbf{s}_k))}{m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k)} \right| \\ &= \frac{|\Phi(\mathbf{x}_k + \mathbf{s}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k)|}{|m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k)|} \\ &\leq \frac{\kappa_{\text{ef}}\Delta_k}{e_3 c_1 \bar{\eta}_1(\mathbf{x}_k)} \\ &\leq 1 - \beta_2, \end{aligned}$$

where we use the fact that $\Delta_k \leq \frac{e_3 c_1 (1 - \beta_2) \bar{\eta}_1(\mathbf{x}_k)}{\kappa_{\text{ef}}}$. Hence, $R_k \geq \beta_2$, and by Algorithm 2.1, the safety phase is called if $\|\mathbf{s}_k\| \leq \tau_k \gamma_S \rho_k$, otherwise the iteration k is very successful. \square

The next lemma provides a lower bound for the trust region step size $\|\mathbf{s}_k\|$, which will be used later to determine that the safety phase is not called under appropriate assumptions. It is based on [9, Lemma 3.6] but with more work required to handle the non-standard assumption on the trust-region subproblem decrease (2.23).

LEMMA 3.8. *Suppose Assumptions 3.1 and 3.2 hold. Then the step \mathbf{s}_k satisfies*

$$(3.9) \quad \|\mathbf{s}_k\| \geq c_2 \tau_k \min \left\{ \Delta_k, \frac{\bar{\eta}_1(\mathbf{x}_k)}{\max\{1, \|H_k\|\}} \right\},$$

$$\text{where } c_2 := \frac{2e_3 c_1}{1 + \sqrt{1 + 2e_3 c_1}} < 1 \text{ and } \tau_k := \min \left\{ \frac{\bar{\eta}_1(\mathbf{x}_k)}{\|\mathbf{g}_k\| + L_h}, 1 \right\}.$$

Proof. Let $h_k := \max\{\|H_k\|, 1\} \geq 1$. Since $m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k) \geq 0$ by Algorithm 2.1, it follows that

$$(3.10) \quad m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k) = |m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k)|$$

$$(3.11) \quad = \left| \mathbf{g}_k^T \mathbf{s}_k + \frac{1}{2} \mathbf{s}_k^T H_k \mathbf{s}_k - h(\mathbf{x}_k) + h(\mathbf{x}_k + \mathbf{s}_k) \right|$$

$$(3.12) \quad \leq \left| \mathbf{g}_k^T \mathbf{s}_k + \frac{1}{2} \mathbf{s}_k^T H_k \mathbf{s}_k \right| + |h(\mathbf{x}_k) - h(\mathbf{x}_k + \mathbf{s}_k)|$$

$$(3.13) \quad \leq \|\mathbf{g}_k\| \cdot \|\mathbf{s}_k\| + \frac{h_k}{2} \|\mathbf{s}_k\|^2 + L_h \|\mathbf{s}_k\|.$$

Substituting this into (2.23), we have

$$(3.14) \quad \frac{1}{2} \|\mathbf{s}_k\|^2 + \frac{\|\mathbf{g}_k\| + L_h}{h_k} \|\mathbf{s}_k\| - e_3 c_1 \frac{\bar{\eta}_1(\mathbf{x}_k)}{h_k} \min \left\{ \Delta_k, \frac{\bar{\eta}_1(\mathbf{x}_k)}{h_k} \right\} \geq 0.$$

Define $C_k^g := (\|\mathbf{g}_k\| + L_h)/h_k$ and $C_k^\eta := \bar{\eta}_1(\mathbf{x}_k)/h_k$. For (3.14) to be satisfied, we require that $\|\mathbf{s}_k\|$ is no less than the positive root of the left-hand side of (3.14), which gives the inequality below

$$(3.15) \quad \|\mathbf{s}_k\| \geq \frac{2e_3 c_1 C_k^\eta}{C_k^g + \sqrt{(C_k^g)^2 + 2e_3 c_1 (C_k^\eta)^2}} \min \{ \Delta_k, C_k^\eta \}$$

Note that $\tau_k = \min\{C_k^\eta/C_k^g, 1\}$. If $\tau_k \geq 1$, i.e., $C_k^\eta \geq C_k^g$, from (3.15),

$$(3.16) \quad \begin{aligned} \|\mathbf{s}_k\| &\geq \frac{2e_3 c_1 C_k^\eta}{C_k^\eta + \sqrt{(C_k^\eta)^2 + 2e_3 c_1 (C_k^\eta)^2}} \min \{ \Delta_k, C_k^\eta \} \\ &= \frac{2e_3 c_1}{1 + \sqrt{1 + 2e_3 c_1}} \min \{ \Delta_k, C_k^\eta \} \end{aligned}$$

On the other hand, if $\tau_k < 1$, that is, $C_k^\eta < C_k^g$, it follows from (3.15) that

$$(3.17) \quad \begin{aligned} \|\mathbf{s}_k\| &\geq \frac{2e_3 c_1 C_k^g (C_k^\eta/C_k^g)}{C_k^g + \sqrt{(C_k^g)^2 + 2e_3 c_1 (C_k^g)^2}} \min \{ \Delta_k, C_k^\eta \} \\ &= \frac{2e_3 c_1 (C_k^\eta/C_k^g)}{1 + \sqrt{1 + 2e_3 c_1}} \min \{ \Delta_k, C_k^\eta \}. \end{aligned}$$

Combining (3.16) and (3.17),

$$(3.18) \quad \|\mathbf{s}_k\| \geq \frac{2e_3c_1\tau_k}{1 + \sqrt{1 + 2e_3c_1}} \min\{\Delta_k, C_k^\eta\},$$

Then we can recover (3.9) from the definition of h_k , C_k^η and c_2 . We conclude by noting $c_2 < 1$ since $e_3 < 1$ and $c_1 < 1/2$. \square

LEMMA 3.9. *In all iterations, $\eta_1(\mathbf{x}_k) \geq \bar{\eta}_1(\mathbf{x}_k) \geq \min\left\{e_1\epsilon_C, \frac{\Delta_k}{\mu}\right\}$. Also, if $\Psi_1(\mathbf{x}_k) \geq \epsilon > 0$, then*

$$(3.19) \quad \eta_1(\mathbf{x}_k) \geq \bar{\eta}_1(\mathbf{x}_k) \geq \epsilon_g := \min\left\{e_1\epsilon_C, \frac{\epsilon}{1 + (\kappa_{\text{eg}} + e_2)\mu}\right\} > 0.$$

Proof. Firstly, since $\eta_1(\mathbf{x}_k) \geq \bar{\eta}_1(\mathbf{x}_k)$, if the criticality phase is not called, then we must have $\eta_1(\mathbf{x}_k) \geq \bar{\eta}_1(\mathbf{x}_k) = \bar{\eta}_1^{\text{init}}(\mathbf{x}_k) > e_1\epsilon_C$. Otherwise, we have $\Delta_k \leq \mu\bar{\eta}_1(\mathbf{x}_k) \leq \mu\eta_1(\mathbf{x}_k)$. Hence, $\eta_1(\mathbf{x}_k) \geq \bar{\eta}_1(\mathbf{x}_k) \geq \min\left\{e_1\epsilon_C, \frac{\Delta_k}{\mu}\right\}$.

The proof of (3.19) follows from [15, Lemma 10.11] except that we must take into account the inaccuracy when computing $\bar{\eta}_1(\mathbf{x}_k)$. We first suppose that the criticality phase is not called. Then $\bar{\eta}_1(\mathbf{x}_k) = \bar{\eta}_1^{\text{init}}(\mathbf{x}_k) > e_1\epsilon_C$ and hence (3.19) holds. Otherwise, the criticality phase is called and m_k is fully linear in $B(\mathbf{x}_k, \Delta_k)$ with $\Delta_k \leq \mu\bar{\eta}_1(\mathbf{x}_k)$. In this case, applying Lemma 2.3, we obtain

$$(3.20) \quad \Psi_1(\mathbf{x}_k) = \Psi_1(\mathbf{x}_k) - \eta_1(\mathbf{x}_k) + \eta_1(\mathbf{x}_k)$$

$$(3.21) \quad \leq \kappa_{\text{eg}}\Delta_k + \bar{\eta}_1(\mathbf{x}_k) + e_2\Delta_k$$

$$(3.22) \quad \leq (\kappa_{\text{eg}} + e_2)\mu\bar{\eta}_1(\mathbf{x}_k) + \bar{\eta}_1(\mathbf{x}_k).$$

Since $\Psi_1(\mathbf{x}_k) \geq \epsilon$, then $\bar{\eta}_1(\mathbf{x}_k) \geq \frac{\epsilon}{1 + (\kappa_{\text{eg}} + e_2)\mu}$ and (3.19) holds. \square

LEMMA 3.10. *Suppose that Assumption 3.1, Assumption 3.2 and Assumption 3.3 hold. If $\Psi_1(\mathbf{x}_k) \geq \epsilon > 0$ for all k , then $\rho_k \geq \rho_{\min} > 0$ for all k , where*

$$(3.23) \quad \rho_{\min} := \min\left\{\Delta_0^{\text{init}}, \frac{\omega_C\epsilon}{\kappa_{\text{eg}} + 1/\mu + e_2}, \frac{c_2\gamma_{\text{dec}}\alpha_1\epsilon_g}{2\kappa_H}, \frac{c_2\gamma_{\text{dec}}\alpha_1}{2} \left(\kappa_{\text{eg}} + \frac{\kappa_{\text{ef}}}{e_3c_1(1 - \beta_2)} + e_2\right)^{-1} \epsilon\right\}.$$

Proof. This proof is similar to [9, Lemma 3.8]. To find a contradiction, suppose that $k(0)$ is the first k such that $\rho_k < \rho_{\min}$. That is, we have

$$(3.24) \quad \rho_0^{\text{init}} \geq \rho_0 \geq \rho_1^{\text{init}} \geq \rho_1 \geq \dots \geq \rho_{k(0)-1}^{\text{init}} \geq \rho_{k(0)-1} \geq \rho_{\min} \quad \text{and} \quad \rho_{k(0)} < \rho_{\min}.$$

We first claim that

$$(3.25) \quad \rho_{k(0)} = \rho_{k(0)}^{\text{init}} < \rho_{\min}.$$

To see this, note that we have either $\rho_{k(0)} = \rho_{k(0)}^{\text{init}}$ or $\rho_{k(0)} = \Delta_{k(0)}$ from Algorithm 2.1. In the former case, (3.25) holds trivially. In the latter, using Lemma 3.5 and (3.24), we have that

$$(3.26) \quad \rho_{\min} > \Delta_{k(0)} \geq \min\left(\Delta_{k(0)}^{\text{init}}, \frac{\omega_C\epsilon}{\kappa_{\text{eg}} + 1/\mu + e_2}\right) \geq \min\left(\rho_{k(0)}^{\text{init}}, \frac{\omega_C\epsilon}{\kappa_{\text{eg}} + 1/\mu + e_2}\right).$$

Since $\rho_{\min} \leq \frac{\omega_C \epsilon}{\kappa_{\text{eg}} + 1/\mu + e_2}$, then we conclude that (3.25) holds.

Since $\rho_{\min} \leq \Delta_0^{\text{init}} = \rho_0^{\text{init}}$ by Algorithm 2.1, we must have $k(0) > 0$ and $\rho_{k(0)-1} \geq \rho_{\min} > \rho_{k(0)}^{\text{init}}$. This reduction in ρ can only happen from a safety step or an unsuccessful step, and in both cases, we have $\rho_{k(0)}^{\text{init}} = \alpha_1 \rho_{k(0)-1}$. Therefore, $\rho_{k(0)-1} \leq \rho_{\min}/\alpha_1$. Also, by Lemma 3.8, Lemma 3.9 and Assumption 3.3, we have

$$(3.27) \quad \|\mathbf{s}_{k(0)-1}\| \geq c_2 \tau_{k(0)-1} \min \left\{ \Delta_{k(0)-1}, \frac{\epsilon_g}{\kappa_H} \right\}.$$

If we have a safety step, we know

$$(3.28) \quad \|\mathbf{s}_{k(0)-1}\| \leq \tau_{k(0)-1} \gamma_S \rho_{k(0)-1} \leq \frac{\tau_{k(0)-1} \gamma_S}{\alpha_1} \rho_{\min}.$$

Combining (3.27) and (3.28), it follows from the assumption $\gamma_S < c_2$ in Algorithm 2.1 that

$$(3.29) \quad \min \left\{ \Delta_{k(0)-1}, \frac{\epsilon_g}{\kappa_H} \right\} \leq \frac{\gamma_S \rho_{\min}}{c_2 \alpha_1} < \frac{\rho_{\min}}{\alpha_1}.$$

Since $\rho_{\min}/\alpha_1 \leq c_2 \gamma_{\text{dec}} \epsilon_g / (2\kappa_H) \leq \epsilon_g / \kappa_H$, then $\min \{ \Delta_{k(0)-1}, \epsilon_g / \kappa_H \} = \Delta_{k(0)-1}$. Hence in (3.27), $\|\mathbf{s}_{k(0)-1}\| \geq c_2 \tau_{k(0)-1} \Delta_{k(0)-1} > \tau_{k(0)-1} \gamma_S \rho_{k(0)-1}$, which contradicts our assumption that the safety step is called. Therefore, the iteration $k(0) - 1$ must be an unsuccessful step. To obtain reduction in ρ , the update rule of $\Delta_{k(0)}^{\text{init}}$ in Algorithm 2.1 implies that

$$(3.30) \quad \gamma_{\text{dec}} \|\mathbf{s}_{k(0)-1}\| \leq \min \{ \gamma_{\text{dec}} \Delta_{k(0)-1}, \|\mathbf{s}_{k(0)-1}\| \} \leq \tau_{k(0)-1} \rho_{k(0)-1}.$$

Therefore, $\|\mathbf{s}_{k(0)-1}\| \leq \tau_{k(0)-1} \gamma_{\text{dec}}^{-1} \rho_{k(0)-1} \leq \tau_{k(0)-1} \gamma_{\text{dec}}^{-1} \alpha_1^{-1} \rho_{\min}$. From this, (3.27) together with (3.23), we have

$$(3.31) \quad \begin{aligned} \min \left\{ \Delta_{k(0)-1}, \frac{\epsilon_g}{\kappa_H} \right\} &\leq \frac{\rho_{\min}}{c_2 \gamma_{\text{dec}} \alpha_1} \\ &< \frac{2\rho_{\min}}{c_2 \gamma_{\text{dec}} \alpha_1} \\ &\leq \min \left\{ \frac{\epsilon_g}{\kappa_H}, \left(\kappa_{\text{eg}} + \frac{\kappa_{\text{ef}}}{e_3 c_1 (1 - \beta_2)} + e_2 \right)^{-1} \epsilon \right\}. \end{aligned}$$

Then $\min \{ \Delta_{k(0)-1}, \epsilon_g / \kappa_H \} = \Delta_{k(0)-1}$. From this and Lemma 3.9, we obtain

$$(3.32) \quad \Delta_{k(0)-1} \leq \left(\kappa_{\text{eg}} + \frac{\kappa_{\text{ef}}}{e_3 c_1 (1 - \beta_2)} + e_2 \right)^{-1} \epsilon \quad \text{and} \quad \Delta_{k(0)-1} \leq \frac{\epsilon_g}{\kappa_H} \leq \frac{\bar{\eta}_1(\mathbf{x}_{k(0)-1})}{\kappa_H}.$$

Now suppose that

$$(3.33) \quad \Delta_{k(0)-1} > \frac{e_3 c_1 (1 - \beta_2)}{\kappa_{\text{ef}}} \bar{\eta}_1(\mathbf{x}_{k(0)-1}).$$

From Algorithm 2.1, regardless of the call of the criticality phase, we have

$$(3.34) \quad \eta_1(\mathbf{x}_{k(0)-1}) - \bar{\eta}_1(\mathbf{x}_{k(0)-1}) < e_2 \Delta_{k(0)-1}.$$

Then using Lemma 2.3 and (3.33) we have

$$(3.35) \quad \epsilon \leq \Psi_1(\mathbf{x}_{k(0)-1}) = \Psi_1(\mathbf{x}_{k(0)-1}) - \eta_1(\mathbf{x}_{k(0)-1}) + \eta_1(\mathbf{x}_{k(0)-1})$$

$$(3.36) \quad \leq \kappa_{\text{eg}} \Delta_{k(0)-1} + \eta_1(\mathbf{x}_{k(0)-1})$$

$$(3.37) \quad \leq (\kappa_{\text{eg}} + e_2) \Delta_{k(0)-1} + \bar{\eta}_1(\mathbf{x}_{k(0)-1})$$

$$(3.38) \quad < \left(\kappa_{\text{eg}} + \frac{\kappa_{\text{ef}}}{e_3 c_1 (1 - \beta_2)} + e_2 \right) \Delta_{k(0)-1},$$

contradicting (3.32). That is, (3.33) is false and so together with (3.32), we have (3.8). Note that $p_{k(0)-1}$ is a fully linear model of f in the unsuccessful step, then Lemma 3.7 implies that iteration $k(0) - 1$ is not an unsuccessful step, which leads to a contradiction. \square

Remark 3.11. The proof of Lemma 3.10 fixes a small error in the proof of [9, Lemma 3.8]: When deriving a bound for the step size $\|\mathbf{s}_{k(0)-1}\|$ at iteration $k(0) - 1$ which is either a safety step or an unsuccessful step, the first inequality in (3.14) of [9, Lemma 3.8] should be $\max(\gamma_S, \gamma_{\text{dec}}^{-1})\rho_{k(0)-1}$ instead of $\min(\gamma_S, \gamma_{\text{dec}}^{-1})\rho_{k(0)-1}$. This affects the bound of $\Delta_{k(0)-1}$ in (3.16) of [9, Lemma 3.8]. Setting $\tau_{k(0)-1} = 1$ in our proof, we follow the same approach as [9] to conclude that the safety step is not called. To deal with the case of an unsuccessful step, we adjust the constant in the expression of ρ_{\min} so that we can bound $\Delta_{k(0)-1}$ in (3.31) similar to (3.16) in [9, Lemma 3.8]. After these corrections, the result in [9, Lemma 3.8] remains unchanged except that ρ_{\min} has a different expression

$$\rho_{\min} := \min \left(\Delta_0^{\text{init}}, \frac{\omega_C \epsilon}{\kappa_{\text{eg}} + 1/\mu}, \frac{c_2 \gamma_{\text{dec}} \alpha_1 \epsilon_g}{2\kappa_H}, \frac{c_2 \gamma_{\text{dec}} \alpha_1}{2} \left(\kappa_{\text{eg}} + \frac{2\kappa_{\text{ef}}}{c_1(1 - \eta_2)} \right)^{-1} \epsilon \right),$$

which does not affect any of the results proved later in [9].

Our convergence results below are based on [15, Chapter 10]. Moreover, we consider the possibility of safety phases and maintenance on ρ_k as presented in [9]. However, we use a different criticality measure Ψ_1 and its inaccurate estimation $\bar{\eta}_1$ in our proofs. The following lemma proves the Lipschitz continuity of Ψ_1 , which allows us to follow the convergence analysis in [15, Chapter 10]. Note that when h is chosen to be an indicator function of a convex set, this lemma is equivalent to [8, Theorem 3.4].

LEMMA 3.12. *Suppose that Assumptions 3.1 and 3.2 hold. Then Ψ_1 is Lipschitz continuous in $\text{dom } h$. That is, for all $x, y \in \text{dom } h$,*

$$(3.39) \quad |\Psi_1(x) - \Psi_1(y)| \leq L_\Psi \|x - y\|,$$

where constant $L_\Psi := L_{\nabla f} + 2L_h$.

Proof. Take any $x, y \in \text{dom } h$. From (2.2), we have

$$(3.40) \quad \Psi_1(\mathbf{x}) - \Psi_1(\mathbf{y}) = h(\mathbf{x}) - h(\mathbf{y}) + \min_{\|\mathbf{s}\| \leq 1} [\nabla f(\mathbf{y})^T \mathbf{s} + h(\mathbf{y} + \mathbf{s})]$$

$$(3.41) \quad - \min_{\|\mathbf{s}\| \leq 1} [\nabla f(\mathbf{x})^T \mathbf{s} + h(\mathbf{x} + \mathbf{s})].$$

Since ∇f and h are continuous by Assumptions 3.1 and 3.2, applying the Weierstrass

Theorem, there exists $\mathbf{s}_x, \mathbf{s}_y \in B(\mathbf{0}, 1) \cap \text{dom}h$ such that

$$(3.42) \quad \min_{\|\mathbf{s}\| \leq 1} \nabla f(\mathbf{x})^T \mathbf{s} + h(\mathbf{x} + \mathbf{s}) = \nabla f(\mathbf{x})^T \mathbf{s}_x + h(\mathbf{x} + \mathbf{s}_x),$$

$$(3.43) \quad \min_{\|\mathbf{s}\| \leq 1} \nabla f(\mathbf{y})^T \mathbf{s} + h(\mathbf{y} + \mathbf{s}) = \nabla f(\mathbf{y})^T \mathbf{s}_y + h(\mathbf{y} + \mathbf{s}_y).$$

Plugging them into (3.40), we obtain

$$\begin{aligned} \Psi_1(\mathbf{x}) - \Psi_1(\mathbf{y}) &= h(\mathbf{x}) - h(\mathbf{y}) + (\nabla f(\mathbf{y})^T \mathbf{s}_y + h(\mathbf{y} + \mathbf{s}_y)) \\ &\quad - (\nabla f(\mathbf{x})^T \mathbf{s}_x + h(\mathbf{x} + \mathbf{s}_x)) \\ &= h(\mathbf{x}) - h(\mathbf{y}) + \nabla f(\mathbf{y})^T \mathbf{s}_y + h(\mathbf{y} + \mathbf{s}_y) - \nabla f(\mathbf{y})^T \mathbf{s}_x - h(\mathbf{y} + \mathbf{s}_x) \\ &\quad + (\nabla f(\mathbf{y}) - \nabla f(\mathbf{x}))^T \mathbf{s}_x + h(\mathbf{y} + \mathbf{s}_x) - h(\mathbf{x} + \mathbf{s}_x) \\ &\leq |h(\mathbf{x}) - h(\mathbf{y})| + \left| (\nabla f(\mathbf{y}) - \nabla f(\mathbf{x}))^T \mathbf{s}_x \right| + |h(\mathbf{y} + \mathbf{s}_x) - h(\mathbf{x} + \mathbf{s}_x)| \\ &\leq (L_{\nabla f} + 2L_h) \|\mathbf{x} - \mathbf{y}\|. \end{aligned}$$

In the first inequality, we use the triangle inequality and the fact that \mathbf{s}_y is a minimizer of (3.43). The last line follows from Assumptions 3.1 and 3.2, and $\|\mathbf{s}_x\|, \|\mathbf{s}_y\| \leq 1$. Since we can interchange the role of \mathbf{x} and \mathbf{y} in the above argument, Ψ_1 is Lipschitz continuous in $\text{dom}h$ with Lipschitz constant $L_{\nabla f} + 2L_h$. \square

The following convergence results consider the case when the number of successful iterations is finite.

LEMMA 3.13. *Suppose that Assumptions 3.1, 3.2 and 3.3 hold. If there are finitely many successful iterations, then $\lim_{k \rightarrow \infty} \Delta_k = 0$ and $\lim_{k \rightarrow \infty} \Psi_1(\mathbf{x}_k) = 0$.*

Proof. Let us consider iterations that come after the last successful iteration k_0 . From Algorithm 2.1 and Lemma 3.5, we know that we can have either one model-improving phase or finite many (uniformly bounded by an integer N) iterations of Algorithm 2.2 in the criticality phase before the model is modified to be fully linear. Therefore, we have an infinite number of iterations where the model is fully linear. That is, there are infinitely many iterations that are either safety or unsuccessful steps. In either case, the trust region radius is reduced by a factor $\max\{\alpha_2, \omega_S, \gamma_{\text{dec}}\} < 1$. Since Δ_k is increased only in successful iterations, then Δ_k converges to zero. For each iteration $k > k_0$, let i_k denote the index of the first iteration after k_0 for which p_k is a fully linear model for f . Since Δ_k is reduced in Algorithm 2.2, then as $k \rightarrow \infty$,

$$(3.44) \quad \|\mathbf{x}_k - \mathbf{x}_{i_k}\| \leq N\Delta_k \rightarrow 0.$$

Note that we can rewrite Ψ_1 and apply the triangle inequality to obtain

$$(3.45) \quad \begin{aligned} \Psi_1(\mathbf{x}_k) &\leq |\Psi_1(\mathbf{x}_k) - \Psi_1(\mathbf{x}_{i_k})| + |\Psi_1(\mathbf{x}_{i_k}) - \eta_1(\mathbf{x}_{i_k})| + |\eta_1(\mathbf{x}_{i_k}) - \bar{\eta}_1(\mathbf{x}_{i_k})| + \bar{\eta}_1(\mathbf{x}_{i_k}) \\ &\leq L_{\Psi} \|\mathbf{x}_k - \mathbf{x}_{i_k}\| + \kappa_{\text{eg}} \Delta_{i_k} + e_2 \Delta_{i_k} + \bar{\eta}_1(\mathbf{x}_{i_k}), \end{aligned}$$

which follows from Lemma 3.12, Lemma 2.3 and the approximation error of η_1 given by Algorithm 2.1. We observe that $\bar{\eta}_1(\mathbf{x}_{i_k}) \rightarrow 0$ as $k \rightarrow \infty$. Otherwise, (3.8) holds for sufficiently large k and then by Lemma 3.7, iteration k must be a safety phase. However, when k is large enough, Lemma 3.8 implies that $\|\mathbf{s}_k\| \geq c_2 \tau_k \Delta_k \geq \tau_k \gamma_S \Delta_k$, \square

LEMMA 3.14. *Suppose that Assumptions 3.1, 3.2 and 3.3 hold. Then we have $\lim_{k \rightarrow \infty} \Delta_k = 0$ and hence $\lim_{k \rightarrow \infty} \rho_k = 0$.*

Proof. This proof is similar to [15, Lemma 10.9]. Let \mathcal{S} denote the set of successful iterations. When $|\mathcal{S}| < \infty$, this follows from Lemma 3.13. Now we consider the case when \mathcal{S} is infinite. For any $k \in \mathcal{S}$, we have

$$(3.46) \quad \Phi(\mathbf{x}_k) - \Phi(\mathbf{x}_{k+1}) \geq \beta_1 [m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k)].$$

Applying Lemma 6.2, Lemma 3.9 and Assumption 3.3, we obtain that

$$(3.47) \quad \Phi(\mathbf{x}_k) - \Phi(\mathbf{x}_{k+1}) \geq \beta_1 e_3 c_1 \min \{e_1 \epsilon_C, \Delta_k / \mu\} \min \left\{ \Delta_k, \frac{\min \{e_1 \epsilon_C, \Delta_k / \mu\}}{\kappa_H} \right\}.$$

Since \mathcal{S} is infinite and Φ is bounded below by Assumption 3.2, then the left hand side of (3.47) converges to 0. Thus, $\lim_{k \in \mathcal{S}} \Delta_k = 0$. Let $k \notin \mathcal{S}$ be the index of an iteration after the first successful iteration, and let s_k denote the index of the last successful iteration before k . Since the trust region radius can be increased only during a successful iteration, and only by a factor of at most $\tilde{\gamma}_{\text{inc}}$, then $\Delta_k \leq \tilde{\gamma}_{\text{inc}} \Delta_{s_k}$. Since $s_k \in \mathcal{S}$, then $\Delta_{s_k} \rightarrow 0$ as $k \rightarrow \infty$ and hence $\Delta_k \rightarrow 0$ for $k \notin \mathcal{S}$. Therefore, $\lim_{k \rightarrow \infty} \Delta_k = 0$. It follows immediately from the fact $\rho_k \leq \Delta_k$ that $\lim_{k \rightarrow \infty} \rho_k = 0$. \square

LEMMA 3.15. *Suppose that Assumptions 3.1, 3.2 and 3.3 hold. Then*

$$(3.48) \quad \liminf_{k \rightarrow \infty} \Psi_1(\mathbf{x}_k) = 0.$$

Proof. Assume for contradiction that there is a constant $\epsilon > 0$ such that $\Psi_1(\mathbf{x}_k) \geq \epsilon$ for all k . In this case, by Lemma 3.10, we have that $\Delta_k \geq \rho_k \geq \rho_{\min} > 0$ for all k , contradicting Lemma 3.14. \square

THEOREM 3.16. *Suppose that Assumptions 3.1, 3.2 and 3.3 hold. Then*

$$(3.49) \quad \lim_{k \rightarrow \infty} \Psi_1(\mathbf{x}_k) = 0.$$

Proof. This proof is based on [15, Theorem 10.13] and [9, Theorem 3.12]. Let \mathcal{S} and \mathcal{M} denote the set of successful and model-improving iterations, respectively. Our theorem holds when $|\mathcal{S}| < \infty$ by Lemma 3.13. Assume, for the purpose of establishing a contradiction, that there exists a subsequence $\{t_i\} \in \mathcal{S}$ of successful iterations such that

$$(3.50) \quad \Psi_1(\mathbf{x}_{t_i}) \geq 2\epsilon > 0,$$

for some $\epsilon > 0$ and for all i . From Lemma 3.15, we know that for each t_i , there exists a first successful iteration $l_i > t_i$ such that $\Psi_1(\mathbf{x}_{l_i}) < \epsilon$. Now we consider iterations whose indices are in the set \mathcal{K} defined by

$$(3.51) \quad \mathcal{K} := \bigcup_{i \in \mathbb{N}} \{k \in \mathbb{N} : t_i \leq k < l_i\},$$

where t_i and l_i belong to the two subsequences defined above. Therefore, for every $k \in \mathcal{K}$, we have $\Psi_1(\mathbf{x}_k) \geq \epsilon$. By Lemma 3.9, $\bar{\eta}_1(\mathbf{x}_k) \geq \epsilon_g > 0$ for $k \in \mathcal{K}$. From Lemma 3.14, $\lim_{k \rightarrow \infty} \Delta_k = 0$. Therefore, Lemma 3.8 implies that $\|\mathbf{s}_k\| \geq c_2 \tau_k \Delta_k \geq \tau_k \gamma_S \Delta_k$ if k is sufficiently large.

$$(3.52) \quad \|\mathbf{x}_{t_i} - \mathbf{x}_{l_i}\| \leq \sum_{\substack{k=t_i \\ k \in \mathcal{K}}}^{l_i-1} \|\mathbf{x}_k - \mathbf{x}_{k+1}\| = \sum_{\substack{k=t_i \\ k \in \mathcal{K} \cap \mathcal{S}}}^{l_i-1} \|\mathbf{x}_k - \mathbf{x}_{k+1}\| \leq \sum_{\substack{k=t_i \\ k \in \mathcal{K} \cap \mathcal{S}}}^{l_i-1} \Delta_k.$$

For each $k \in \mathcal{K} \cap \mathcal{S}$, using $\bar{\eta}_1(\mathbf{x}_k) \geq \epsilon_g$ and Assumption 3.3, we obtain

$$(3.53) \quad \Phi(\mathbf{x}_k) - \Phi(\mathbf{x}_{k+1}) \geq \beta_1 [m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k)]$$

$$(3.54) \quad \geq \beta_1 e_3 c_1 \bar{\eta}_1(\mathbf{x}_k) \min \left\{ \Delta_k, \frac{\bar{\eta}_1(\mathbf{x}_k)}{\max\{1, \|H_k\|\}} \right\}$$

$$(3.55) \quad \geq \beta_1 e_3 c_1 \epsilon_g \min \left\{ \Delta_k, \frac{\epsilon_g}{\kappa_H} \right\}$$

Since $\lim_{k \rightarrow \infty} \Delta_k = 0$ by Lemma 3.14, then $\min\{\Delta_k, \epsilon_g/\kappa_H\} = \Delta_k$ for k large enough. Thus,

$$(3.56) \quad \Delta_k \leq \frac{1}{\beta_1 e_3 c_1 \epsilon_g} [\Phi(\mathbf{x}_k) - \Phi(\mathbf{x}_{k+1})]$$

Combining (3.52) and (3.56), we have

$$(3.57) \quad \begin{aligned} \|\mathbf{x}_{t_i} - \mathbf{x}_{l_i}\| &\leq \sum_{\substack{k=t_i \\ k \in \mathcal{K} \cap \mathcal{S}}}^{l_i-1} \frac{1}{\beta_1 e_3 c_1 \epsilon_g} [\Phi(\mathbf{x}_k) - \Phi(\mathbf{x}_{k+1})] \\ &= \frac{1}{\beta_1 e_3 c_1 \epsilon_g} \sum_{\substack{k=t_i \\ k \in \mathcal{K}}}^{l_i-1} [\Phi(\mathbf{x}_k) - \Phi(\mathbf{x}_{k+1})] \\ &= \frac{1}{\beta_1 e_3 c_1 \epsilon_g} [\Phi(\mathbf{x}_{t_i}) - \Phi(\mathbf{x}_{l_i})]. \end{aligned}$$

Since Φ is bounded below by Assumption 3.2 and the sequence $\{\Phi(\mathbf{x}_k)\}$ is monotone decreasing, then the right hand side of (3.57) converges to zero. Therefore,

$$(3.58) \quad \lim_{i \rightarrow \infty} \|\mathbf{x}_{t_i} - \mathbf{x}_{l_i}\| = 0.$$

However, from Lemma 3.12, we have that $|\Psi_1(\mathbf{x}_{t_i}) - \Psi_1(\mathbf{x}_{l_i})| \leq L_\Psi \|\mathbf{x}_{t_i} - \mathbf{x}_{l_i}\| \rightarrow 0$ as $i \rightarrow \infty$. This contradicts our construction of subsequences $\{t_i\}$ and $\{l_i\}$ that $\Psi_1(\mathbf{x}_{t_i}) - \Psi_1(\mathbf{x}_{l_i}) \geq \epsilon > 0$. By the principle of contradiction, we conclude that $\lim_{k \rightarrow \infty} \Psi_1(\mathbf{x}_k) = 0$. \square

3.2. Worst-Case Complexity. Now we study the worst-case complexity of Algorithm 2.1 following the approach of [9]. We bound the number of iterations and objective evaluations until $\Psi_1(\mathbf{x}_k) < \epsilon$, where the existence of bounds is guaranteed by Lemma 3.15 for each optimality level ϵ . Let i_ϵ be the last iteration before $\Psi_1(\mathbf{x}_{i_\epsilon+1}) < \epsilon$ for the first time. We classify all iterations until iteration i_ϵ (inclusive) into five types with descriptions and symbols below:

1. **Criticality Iteration (\mathcal{C}_{i_ϵ}):** the set of criticality iterations $k \leq i_\epsilon$. Moreover, we collect all iterations in \mathcal{C}_{i_ϵ} where Δ_k is not reduced into a set $\mathcal{C}_{i_\epsilon}^M$. That is, $\mathcal{C}_{i_\epsilon}^M$ is the set of the first iteration of every call of Algorithm 2.2. We denote the remaining iterations in \mathcal{C}_{i_ϵ} by $\mathcal{C}_{i_\epsilon}^U := \mathcal{C}_{i_\epsilon} - \mathcal{C}_{i_\epsilon}^M$.
2. **Safety Iteration (\mathcal{S}_{i_ϵ}):** the set of iterations $k \leq i_\epsilon$ where the safety phase is called.
3. **Successful Iteration (\mathcal{S}):** the set of successful iterations $k \leq i_\epsilon$.
4. **Model Improvement Iteration (\mathcal{M}_{i_ϵ}):** the set of iterations $k \leq i_\epsilon$ where the model improvement phase is called.

5. Unsuccessful Iteration (\mathcal{U}_{i_ϵ}): the set of unsuccessful iterations $k \leq i_\epsilon$. We first count the number of iterations up to iteration i_ϵ that are successful in the following lemma.

LEMMA 3.17. *Suppose that Assumptions 3.1, 3.2 and 3.3 hold. Then*

$$(3.59) \quad |\mathcal{S}_{i_\epsilon}| \leq \frac{\Phi(\mathbf{x}_0) - f_{\text{low}} - h_{\text{low}}}{\beta_1 e_3 c_1} \max \{ \kappa_H \epsilon_g^{-2}, \epsilon_g^{-1} \rho_{\min}^{-1} \},$$

where ϵ_g is defined in (3.19) and ρ_{\min} in (3.23).

Proof. For all $k \in \mathcal{S}_{i_\epsilon}$, by Lemma 6.2 and Assumption 3.3, we have that

$$(3.60) \quad \Phi(\mathbf{x}_k) - \Phi(\mathbf{x}_{k+1}) \geq \beta_1 [m_k(\mathbf{0}) - m_k(\mathbf{s}_k)] \geq \beta_1 e_3 c_1 \bar{\eta}_1(\mathbf{x}_k) \min \left\{ \Delta_k, \frac{\bar{\eta}_1(\mathbf{x}_k)}{\kappa_H} \right\}.$$

From Lemma 3.9 and Lemma 3.10, $\bar{\eta}_1(\mathbf{x}_k) \geq \epsilon_g$ and $\Delta_k \geq \rho_k \geq \rho_{\min}$ for any $k \geq i_\epsilon$. This means

$$(3.61) \quad \Phi(\mathbf{x}_k) - \Phi(\mathbf{x}_{k+1}) \geq \beta_1 e_3 c_1 \epsilon_g \min \left\{ \rho_{\min}, \frac{\epsilon_g}{\kappa_H} \right\}.$$

Summing (3.61) over all $k \in \mathcal{S}_{i_\epsilon}$, and noting that $f_{\text{low}} + h_{\text{low}} \leq \Phi(\mathbf{x}_k) \leq \Phi(\mathbf{x}_0)$, we obtain

$$(3.62) \quad \Phi(\mathbf{x}_0) - f_{\text{low}} - h_{\text{low}} \geq |\mathcal{S}_{i_\epsilon}| \beta_1 e_3 c_1 \epsilon_g \min \left\{ \rho_{\min}, \frac{\epsilon_g}{\kappa_H} \right\},$$

from which (3.59) follows. \square

LEMMA 3.18. *Suppose that Assumption 3.1, Assumption 3.2 and Assumption 3.3 hold. Then we have the bounds*

$$(3.63) \quad |\mathcal{C}_{i_\epsilon}^U| + |\mathcal{F}_{i_\epsilon}| + |\mathcal{U}_{i_\epsilon}| \leq |\mathcal{S}_{i_\epsilon}| \cdot \frac{\log \bar{\gamma}_{\text{inc}}}{|\log \alpha_3|} + \frac{1}{|\log \alpha_3|} \log \left(\frac{\Delta_0^{\text{init}}}{\rho_{\min}} \right),$$

$$(3.64) \quad |\mathcal{C}_{i_\epsilon}^M| \leq |\mathcal{F}_{i_\epsilon}| + |\mathcal{S}_{i_\epsilon}| + |\mathcal{U}_{i_\epsilon}|,$$

$$(3.65) \quad |\mathcal{M}_{i_\epsilon}| \leq |\mathcal{C}_{i_\epsilon}^M| + |\mathcal{C}_{i_\epsilon}^U| + |\mathcal{F}_{i_\epsilon}| + |\mathcal{S}_{i_\epsilon}| + |\mathcal{U}_{i_\epsilon}|,$$

where $\alpha_3 := \max\{\omega_C, \omega_S, \gamma_{\text{dec}}, \alpha_2\} < 1$ and ρ_{\min} is defined in (3.23).

Proof. See the proof of [9, Lemma 3.14]. \square

Following [9], we make an additional assumption below, which can be easily satisfied by appropriate choices of parameters in Algorithm 2.1.

ASSUMPTION 3.19. *The algorithm parameter ϵ_C satisfies $e_1 \epsilon_C = c_3 \epsilon$ for some constant $c_3 > 0$.*

THEOREM 3.20. *Suppose Assumptions 3.1, 3.2, 3.3 and 3.19 hold. Then the number of iterations i_ϵ until $\Psi_1(\mathbf{x}_{i_\epsilon+1}) < \epsilon$ is at most*

$$(3.66) \quad \left\lfloor \frac{4(\Phi(\mathbf{x}_0) - f_{\text{low}} - h_{\text{low}})}{\beta_1 e_3 c_1} \left(1 + \frac{\log \bar{\gamma}_{\text{inc}}}{|\log \alpha_3|} \right) \max \left\{ \frac{\kappa_H}{c_4^2 \epsilon^2}, \frac{1}{c_4 c_5 \epsilon^2}, \frac{1}{c_4 \Delta_0^{\text{init}} \epsilon} \right\} \right\rfloor$$

$$(3.67) \quad + \left\lfloor \frac{4}{|\log \alpha_3|} \max \{ 0, \log(\Delta_0^{\text{init}} c_5^{-1} \epsilon^{-1}) \} \right\rfloor,$$

where $c_4 := \min\{e_1 c_3, 1/(1 + (\kappa_{\text{eg}} + e_2)\mu)\}$ and

(3.68)

$$c_5 := \min \left\{ \frac{\omega_C}{\kappa_{\text{eg}} + 1/\mu + e_2}, \frac{c_2 c_4 \gamma_{\text{dec}} \alpha_1}{2\kappa_H}, \frac{c_2 \gamma_{\text{dec}} \alpha_1}{2} \left(\kappa_{\text{eg}} + \frac{\kappa_{\text{ef}}}{e_3 c_1 (1 - \beta_2)} + e_2 \right)^{-1} \right\}.$$

Proof. From Assumption 3.19 and the definition of ϵ_g in (3.19), we have $\epsilon_g = c_4 \epsilon$. Then ρ_{\min} defined in (3.23) can be rewritten as $\rho_{\min} = \min\{\Delta_0^{\text{init}}, c_5 \epsilon\}$. Using Lemma 3.18, the total number of iterations can be bounded by

$$(3.69) \quad i_\epsilon = |\mathcal{C}_{i_\epsilon}^M| + |\mathcal{C}_{i_\epsilon}^U| + |\mathcal{F}_{i_\epsilon}| + |\mathcal{S}_{i_\epsilon}| + |\mathcal{M}_{i_\epsilon}| + |\mathcal{U}_{i_\epsilon}|$$

$$(3.70) \quad \leq 2 (|\mathcal{C}_{i_\epsilon}^M| + |\mathcal{C}_{i_\epsilon}^U| + |\mathcal{F}_{i_\epsilon}| + |\mathcal{S}_{i_\epsilon}| + |\mathcal{U}_{i_\epsilon}|)$$

$$(3.71) \quad \leq 2 |\mathcal{C}_{i_\epsilon}^U| + 4 (|\mathcal{F}_{i_\epsilon}| + |\mathcal{S}_{i_\epsilon}| + |\mathcal{U}_{i_\epsilon}|)$$

$$(3.72) \quad \leq 4 (|\mathcal{C}_{i_\epsilon}^U| + |\mathcal{F}_{i_\epsilon}| + |\mathcal{U}_{i_\epsilon}|) + 4 |\mathcal{S}_{i_\epsilon}|$$

$$(3.73) \quad \leq 4 |\mathcal{S}_{i_\epsilon}| \left(1 + \frac{\log \bar{\gamma}_{\text{inc}}}{|\log \alpha_3|} \right) + \frac{4}{|\log \alpha_3|} \log \left(\frac{\Delta_0^{\text{init}}}{\rho_{\min}} \right).$$

Therefore, (3.66) follows from this and Lemma 3.17. \square

COROLLARY 3.21. *Suppose Assumptions 3.1, 3.2, 3.3 and 3.19 hold. For $\epsilon \in (0, 1]$, the number of iterations i_ϵ until $\Psi_1(\mathbf{x}_{i_\epsilon+1}) < \epsilon$ for the first time is at most $\mathcal{O}(\kappa_H \kappa_d^2 \epsilon^{-2})$, where $\kappa_d := \max(\kappa_{\text{ef}}, \kappa_{\text{eg}})$.*

Proof. We note that $c_4^{-1} = \mathcal{O}(\kappa_{\text{eg}}) = \mathcal{O}(\kappa_d)$ from Theorem 3.20.

$$(3.74) \quad c_5^{-1} = \mathcal{O}(\max\{\kappa_{\text{eg}}, \kappa_H c_4^{-1}, \kappa_{\text{ef}} + \kappa_{\text{eg}}\}) = \mathcal{O}(\kappa_H \kappa_d).$$

To leading order in (3.66), the number of iterations is

$$(3.75) \quad \mathcal{O}(\max\{\kappa_H c_4^{-2}, c_4^{-1} c_5^{-1}\} \epsilon^{-2}) = \mathcal{O}(\kappa_H \kappa_d^2 \epsilon^{-2}),$$

as required. \square

4. Adaptation to Regularized Nonlinear Least-Squares. We now consider the specific case of regularized nonlinear least-squares problems, where (1.1) becomes

$$(4.1) \quad \min_{\mathbf{x} \in \mathbb{R}^n} \Phi(\mathbf{x}) := f(\mathbf{x}) + h(\mathbf{x}) = \frac{1}{2} \|\mathbf{r}(\mathbf{x})\|^2 + h(\mathbf{x}),$$

for some function $\mathbf{r} : \mathbb{R}^n \rightarrow \mathbb{R}^m$. As above, we assume that \mathbf{r} is continuously differentiable with Jacobian $[J(\mathbf{x})]_{i,j} = \frac{\partial r_i(\mathbf{x})}{\partial x_j}$ but these derivatives are not accessible.

ASSUMPTION 4.1. *$\mathbf{r}(\mathbf{x})$ is continuously differentiable and its Jacobian is Lipschitz continuous with Lipschitz constant L_J in the convex hull of $\mathcal{B} := \cup_k B(\mathbf{x}_k, \Delta_{\max})$. Furthermore, $\mathbf{r}(\mathbf{x})$ and $J(\mathbf{x})$ are uniformly bounded in \mathcal{B} , i.e. $\|\mathbf{r}(\mathbf{x})\| \leq r_{\max}$ and $\|J(\mathbf{x})\| \leq J_{\max}$ for all $\mathbf{x} \in \mathcal{B}$.*

Note that Assumption 4.1 (taken from [9, Assumption 3.1]) automatically gives us the smoothness requirements from Section 3.

LEMMA 4.2 (Lemma 3.2, [9]). *If Assumption 4.1 holds, then Assumption 3.1 holds with*

$$(4.2) \quad L_{\nabla f} := r_{\max} L_J + J_{\max}^2.$$

and $f_{\text{low}} = 0$.

We are now in a position to precisely specify how we construct the model p_k (2.16) to be used in m_k (2.17). Our construction is based on a linear interpolating model for \mathbf{r} , a derivative-free analog of the Gauss-Newton method [9]. At k -th iteration, the linear Taylor series of the residual $\mathbf{r}(\mathbf{x})$ around \mathbf{x}_k is given by

$$(4.3) \quad \mathbf{r}(\mathbf{x}_k + \mathbf{s}) \approx \mathbf{r}(\mathbf{x}_k) + J(\mathbf{x}_k)\mathbf{s},$$

Since $J(\mathbf{x})$ is not accessible, we maintain a set of $n+1$ interpolation points containing \mathbf{x}_k , denoted by $Y_k := \{\mathbf{y}_0, \dots, \mathbf{y}_n\}$ with $\mathbf{y}_0 := \mathbf{x}_k$, and approximate $J(\mathbf{x}_k)$ by a matrix $J_k \in \mathbb{R}^{m \times n}$. That is, we define $\mathbf{m}_k(\mathbf{x}_k + \mathbf{s}) := \mathbf{r}(\mathbf{x}_k) + J_k\mathbf{s}$, where J_k satisfies the following interpolation conditions

$$(4.4) \quad \mathbf{r}(\mathbf{y}_t) = \mathbf{m}_k(\mathbf{y}_t), \quad \forall t = 1, \dots, n.$$

The uniqueness of J_k is guaranteed once the interpolation directions $\{\mathbf{y}_1 - \mathbf{x}_k, \dots, \mathbf{y}_n - \mathbf{x}_k\}$ are linearly independent, and we say that Y_k is *poised* for linear interpolation [15, Section 2.3]. Our resulting quadratic model p_k for f is

$$(4.5) \quad f(\mathbf{x}_k + \mathbf{s}) \approx p_k(\mathbf{x}_k + \mathbf{s}) := \frac{1}{2} \|\mathbf{m}_k(\mathbf{x}_k + \mathbf{s})\|^2,$$

and so $\mathbf{g}_k := J_k^T \mathbf{r}(\mathbf{x}_k)$ and $H_k := J_k^T J_k$ in (2.16). Therefore, our full model m_k (2.17) for Φ is

$$(4.6) \quad \Phi(\mathbf{x}_k + \mathbf{s}) \approx m_k(\mathbf{x}_k + \mathbf{s}) := p_k(\mathbf{x}_k + \mathbf{s}) + h(\mathbf{x}_k + \mathbf{s}).$$

In this setting, we now have a specific version of Lemma 3.4 for our model construction (4.5).

LEMMA 4.3 (Lemma 3.3, [9]). *Suppose Assumption 4.1 holds and Y_k is Λ -poised in $B(\mathbf{x}_k, \Delta_k)$. Then p_k (4.5) is a fully linear model for f in $B(\mathbf{x}_k, \Delta_k)$ with*

$$(4.7) \quad \kappa_{\text{ef}} = \kappa_{\text{eg}} + \frac{L_{\nabla f} + (\kappa_{\text{eg}}^r \Delta_{\text{max}} + J_{\text{max}})^2}{2} \quad \text{and}$$

$$(4.8) \quad \kappa_{\text{eg}} = L_{\nabla f} + \kappa_{\text{eg}}^r r_{\text{max}} + (\kappa_{\text{eg}}^r \Delta_{\text{max}} + J_{\text{max}})^2,$$

where $\kappa_{\text{eg}}^r := \frac{1}{2} L_J (\sqrt{n}C + 2)$ and $C = \mathcal{O}(\sqrt{n}\Lambda)$. We also have the bound

$$(4.9) \quad \|H_k\| \leq (J_{\text{max}}^k)^2 = (\kappa_{\text{eg}}^r \Delta_{\text{max}} + J_{\text{max}})^2.$$

We are now in a position to report the convergence and complexity results for Algorithm 2.1 applied to (4.1).

COROLLARY 4.4. *Suppose Assumptions 4.1, 3.2, 3.3 and 3.19 hold, and we run Algorithm 2.1 with the model m_k from (4.6). Then $\lim_{k \rightarrow \infty} \Psi_1(\mathbf{x}_k) = 0$. For $\epsilon \in (0, 1]$, the number of iterations i_ϵ until $\Psi_1(\mathbf{x}_{i_\epsilon+1}) < \epsilon$ for the first time is at most $\mathcal{O}(\kappa_H \kappa_d^2 \epsilon^{-2})$, where $\kappa_d = \mathcal{O}(n^2 L_J^2 \Lambda^2)$, and the number of evaluations of \mathbf{r} until i_ϵ is at most $\mathcal{O}(\kappa_H \kappa_d^2 n \epsilon^{-2})$.*

Proof. Lemma 4.2 shows Assumption 3.1 holds. Otherwise, we replace Lemma 3.4 with Lemma 4.3 and apply Theorem 3.16 and Corollary 3.21. Our interpolation set contains $n+1$ points, so we require no more than $n+1$ objective evaluations for each iteration. Thus, the total number of objective evaluations up to iteration i_ϵ is $\mathcal{O}(ni_\epsilon)$, as expected. The result $\kappa_d = \max(\kappa_{\text{ef}}, \kappa_{\text{eg}}) = \mathcal{O}(n^2 L_J^2 \Lambda^2)$ follows from Lemma 4.3. \square

The bound (4.9) suggests that $\kappa_H = \mathcal{O}(\kappa_d)$ is a reasonable estimate (even though it only holds in iterations where p_k is fully linear). This would give a final complexity bound of $\mathcal{O}(n^6\epsilon^{-2})$ iterations or $\mathcal{O}(n^7\epsilon^{-2})$ evaluations of \mathbf{r} .

Remark 4.5. Regarding the constant C in Lemma 4.3, we point the reader to [9, Remark 3.20] for a detailed discussion of the dependency of C on n . In particular, the results from [18] effectively assume $C = \mathcal{O}(\Lambda)$ and so report an improved dependency on n in their final complexity bound. In the case of Corollary 4.4, we would instead get $\kappa_d = \mathcal{O}(n)$ and a final complexity bound of $\mathcal{O}(n^3\epsilon^{-2})$ iterations or $\mathcal{O}(n^4\epsilon^{-2})$ evaluations of \mathbf{r} .

Subproblem solutions. Although Algorithm 2.1 is for the general problem (1.1), the trust-region subproblem (2.14) (where we require a solution with accuracy (2.23)) is now minimizing the sum of a nonconvex quadratic function and nonsmooth convex regularizer subject to a ball constraint. However, this problem becomes convex in the nonlinear least-squares case, where m_k (4.6) is convex (since $H_k = J_k^T J_k$ is positive semidefinite).

We also have the convex subproblem of computing the criticality measure (2.20) to accuracy (2.22) (or (2.24) to accuracy (2.26) in the case of Algorithm 2.2). This is also convex, requiring the minimization of a linear function plus a nonsmooth convex regularizer subject to a ball constraint.

We will discuss how both (convex) subproblems can be solved to the desired accuracy in Section 6.

5. Smoothing-Based Algorithm. We now introduce our second approach for solving (1.1), based on the smoothing technique from [18]. Here, we introduce a smoothing parameter $\gamma > 0$ and build an approximation $\Phi^\gamma(\mathbf{x}) \approx \Phi(\mathbf{x})$ with Φ^γ smooth and $\Phi^\gamma \rightarrow \Phi$ as $\gamma \rightarrow 0$. For decreasing values of γ , we approximately minimize Φ^γ using a DFO method suitable for smooth objectives. This algorithm is summarized in Algorithm 5.1.

Algorithm 5.1 Smoothing model-based DFO method for (1.1) [18, Algorithm 4.1]

Input: Starting point $\mathbf{x}_0 \in \text{dom } h$, initial smoothing parameter $\gamma_0 > 0$.

Parameters: trust-region termination function $d : (0, \infty) \rightarrow (0, \infty)$, smoothing update parameter $\sigma \in (0, 1)$.

- 1: **for** $j = 0, 1, 2, \dots$ **do**
 - 2: Approximately minimize $\Phi^{\gamma_j}(\mathbf{x})$ starting from \mathbf{x}_j using a globally convergent DFO trust-region method, terminating when it reaches minimum trust-region radius $\Delta_{j,k} < d(\gamma_j)$ and returning approximate minimizer \mathbf{x}_{j+1} .
 - 3: Set $\gamma_{j+1} = \sigma\gamma_j$.
 - 4: **end for**
-

For Algorithm 5.1 to converge, we require the following properties about the smoothed function Φ^γ .

ASSUMPTION 5.1. *For any $\gamma > 0$, the function $\Phi^\gamma : \mathbb{R}^n \rightarrow \mathbb{R}$ satisfies:*

- (a) Φ^γ has $L_\Phi(\gamma)$ -Lipschitz continuous gradient for some $L_\Phi(\gamma) > 0$ and is bounded below on \mathbb{R}^n .
- (b) For any $\mathbf{x} \in \mathbb{R}^n$, $\lim_{\gamma \rightarrow 0^+} \Phi^\gamma(\mathbf{x}) = \Phi(\mathbf{x})$.

In [18], the sufficient decrease condition in the inner loop (line 2 of Algorithm 5.1) is required to be based on the ratio

$$(5.1) \quad R_k := \frac{\Phi^{\gamma_j}(\mathbf{x}_k) - \Phi^{\gamma_j}(\mathbf{x}_k + \mathbf{s}_k) - c\Delta_k^p}{m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k)}$$

for some $c > 0$ and $p > 1$, which is a slight modification of (2.15).

The global convergence for Algorithm 5.1 is given by the following result.

THEOREM 5.2 (Theorem 4.3, [18]). *Suppose Assumption 5.1 holds, and all model Hessians used in line 2 of Algorithm 5.1 are uniformly bounded. Then if $\lim_{\gamma \rightarrow 0^+} d(\gamma) = 0$ and $\lim_{\gamma \rightarrow 0^+} L_\Phi(\gamma)d(\gamma) = 0$, then $\lim_{j \rightarrow \infty} \|\nabla \Phi^{\gamma_j}(\mathbf{x}_j)\| = 0$.*

We also note that [18, Corollary 4.1] provides a worst-case complexity bound for Algorithm 5.1 of $\mathcal{O}(|\log(\epsilon)|\epsilon^{-3})$ function evaluations to reach an outer iteration j with $\|\nabla \Phi^{\gamma_j}(\mathbf{x}_j)\| < \mathcal{O}(\epsilon)$ (with constants depending on the dimension n), provided we choose $p = 3/2$ in (5.1) and trust-region termination function $d(\gamma) = \gamma^2$, and $L_\Phi(\gamma) = \mathcal{O}(1/\gamma)$ in Assumption 5.1.

5.1. Adaptation to Regularized Least-Squares. We now outline how we adapt Algorithm 5.1 to the regularized least-squares setting (4.1). Our approach uses DFO-GN [9, Algorithm 1]—essentially Algorithm 2.1 with $h = 0$ —with modified interpolation models similar to (4.6).

For the objective Φ (4.1), our smoothed approximation will be

$$(5.2) \quad \Phi^\gamma(\mathbf{x}) := f(\mathbf{x}) + M_h^{\mu(\gamma)}(\mathbf{x}) = \frac{1}{2} \|\mathbf{r}(\mathbf{x})\|^2 + M_h^{\mu(\gamma)}(\mathbf{x}),$$

where

$$(5.3) \quad M_h^\mu(\mathbf{x}) := \min_{\mathbf{z} \in \text{dom } h} \left\{ h(\mathbf{z}) + \frac{1}{2\mu} \|\mathbf{z} - \mathbf{x}\|^2 \right\},$$

is the Moreau envelope of h for the parameter $\mu > 0$. Note that for a given γ , we will choose a value $\mu(\gamma)$ for the Moreau envelope parameter. We will use the following standard properties of M_h^μ .

LEMMA 5.3. *Suppose h satisfies Assumption 3.2. Then for any $\mu > 0$, we have:*

- (a) $M_h^\mu(\mathbf{x}) \leq h(\mathbf{x}) \leq M_h^\mu(\mathbf{x}) + \frac{L_h \mu}{2}$ for all $\mathbf{x} \in \text{dom } h$.
- (b) M_h^μ is a convex and differentiable function on $\text{dom } h$, and ∇M_h^μ is Lipschitz continuous with Lipschitz constant $\frac{1}{\mu}$.

Proof. See for example [5, Theorems 6.55, 6.60 & 10.51]. \square

From this, we see that our smoothed approximation (5.2) is a suitable choice for Algorithm 5.1.

LEMMA 5.4. *Suppose \mathbf{r} satisfies Assumption 4.1 and h satisfies Assumption 3.2. If $\mu(\gamma) = \Theta(\gamma)$ as $\gamma \rightarrow 0^+$ then Φ^γ (5.2) satisfies Assumption 5.1 with $L_\Phi(\gamma) = \mathcal{O}(1/\gamma)$.*

Proof. This follows immediately from Assumptions 4.1 and 3.2, together with Lemmas 4.2 and 5.3. \square

In DFO-GN (called in line 2 of Algorithm 5.1), as in Section 4, we build linear interpolation models of $\mathbf{r}(\mathbf{x})$ of the form

$$(5.4) \quad \mathbf{r}(\mathbf{x}_k + \mathbf{s}) \approx \mathbf{m}_k(\mathbf{x}_k + \mathbf{s}) := \mathbf{r}(\mathbf{x}_k) + J_k \mathbf{s},$$

for some J_k derived from the interpolation conditions (4.4), yielding our model

$$(5.5) \quad \Phi^\gamma(\mathbf{x}_k + \mathbf{s}) \approx m_k(\mathbf{x}_k + \mathbf{s}) := \frac{1}{2} \|\mathbf{m}_k(\mathbf{x}_k + \mathbf{s})\|^2 + M_h^{\mu(\gamma)}(\mathbf{x}_k + \mathbf{s}).$$

From Lemma 4.3, this model is a fully linear approximation for Φ^γ whenever the interpolation set used for (4.4) is Λ -poised, and so all geometry-improving procedures from DFO-GN can be used here.

It only remains to describe how the trust-region subproblem

$$(5.6) \quad \min_{\|\mathbf{s}\| \leq \Delta_k} m_k(\mathbf{x}_k + \mathbf{s}) = \frac{1}{2} \|\mathbf{m}_k(\mathbf{x}_k + \mathbf{s})\|^2 + M_h^{\mu(\gamma)}(\mathbf{x}_k + \mathbf{s})$$

is solved, and the choice of $\mu(\gamma)$. These are discussed in Section 6.

5.2. Global Convergence. Theorem 5.2 establishes the global convergence of Algorithm 5.1, but it does not provide any standard first-order optimality results. We now show that any accumulation point of the \mathbf{x}_j generated by Algorithm 5.1 is a Clarke-stationary point of Φ .

THEOREM 5.5. *Suppose the assumptions of Theorem 5.2 and Lemma 5.4 hold. If \mathbf{x}^* is any accumulation point of the iterates \mathbf{x}_j of Algorithm 5.1, then $\mathbf{0} \in \partial_C \Phi(\mathbf{x}^*)$.*

Proof. From Theorem 5.2, we have $\nabla f(\mathbf{x}_j) + \nabla M_h^{\mu(\gamma_j)}(\mathbf{x}_j) \rightarrow \mathbf{0}$, and so since $\mathbf{x}_{j_k} \rightarrow \mathbf{x}^*$ for some subsequence of iterates and f is continuously differentiable, we have $\nabla M_h^{\mu(\gamma_{j_k})}(\mathbf{x}_{j_k}) \rightarrow -\nabla f(\mathbf{x}^*)$. From [23, Theorem 5.5] and $\mu(\gamma_{j_k}) \rightarrow 0$, we conclude that $-\nabla f(\mathbf{x}^*) \in \partial h(\mathbf{x}^*)$. Finally, we have $\partial_C \Phi(\mathbf{x}) = \nabla f(\mathbf{x}) + \partial h(\mathbf{x})$ from [11, Corollary 1, p. 39] and the result follows. \square

6. Calculating Subproblem Solutions. There are three places where we need an algorithm to calculate subproblem solutions: the trust-region subproblem (2.14) in Algorithm 2.1, calculating $\bar{\eta}_1(\mathbf{x}_k)$ in Algorithms 2.1 and 2.2 (in (2.20) and (2.24) respectively), and the trust-region subproblem (5.6) in the smoothed DFO method Algorithm 5.1. In the case where H_k is positive semidefinite, as in the case solving regularized nonlinear least-squares problems (4.1) (with models (4.6) and (5.5)), this becomes practical.

In all three cases, our subproblems can be written in the form

$$(6.1) \quad \min_{\mathbf{d}} G(\mathbf{d}) := \mathbf{g}^T \mathbf{d} + \frac{1}{2} \mathbf{d}^T H \mathbf{d} + h(\mathbf{x} + \mathbf{d}) + I_C(\mathbf{d}),$$

where H is positive semidefinite and I_C is an indicator function for the Euclidean ball $C := B(\mathbf{0}, r)$. This problem is the sum of a smooth convex function and two nonsmooth convex functions. To solve this to arbitrary accuracy, we will use S-FISTA [5, Chapter 10.8], a smoothed version of FISTA.

Using Lemma 5.3, we approximate h by M_h^μ to obtain a smoothed formulation of (6.1):

$$(6.2) \quad \min_{\mathbf{d}} G_\mu(\mathbf{d}) := f(\mathbf{x}) + \mathbf{g}^T \mathbf{d} + \frac{1}{2} \mathbf{d}^T H \mathbf{d} + M_h^\mu(\mathbf{x} + \mathbf{d}) + I_C(\mathbf{d}),$$

for some smoothing parameter $\mu > 0$. Defining $F_\mu(\mathbf{d}) := f(\mathbf{x}) + \mathbf{g}^T \mathbf{d} + \frac{1}{2} \mathbf{d}^T H \mathbf{d} + M_h^\mu(\mathbf{x} + \mathbf{d})$, from [5, Theorem 6.60] we have

$$(6.3) \quad \nabla F_\mu(\mathbf{d}) = \mathbf{g} + H \mathbf{d} + \frac{1}{\mu} (\mathbf{x} + \mathbf{d} - \text{prox}_{\mu h}(\mathbf{x} + \mathbf{d})),$$

where the proximal operator $\text{prox}_{\mu h}(\cdot)$ is defined by

$$(6.4) \quad \text{prox}_{\mu h}(\mathbf{y}) := \arg \min_{\mathbf{z} \in \text{dom } h} \left\{ h(\mathbf{z}) + \frac{1}{2\mu} \|\mathbf{z} - \mathbf{y}\|^2 \right\}.$$

Note that the proximal operator of I_C is the (well-defined) Euclidean projection P_C onto the set C . The S-FISTA algorithm [5, Chapter 10.8.4], specialized to the problem (6.1) is given in Algorithm 6.1.

Algorithm 6.1 S-FISTA [5, Chapter 10.8.4] for solving (6.1)

Input: smoothing parameter $\mu > 0$.

- 1: Set $\mathbf{d}^0 = \mathbf{y}^0 = \mathbf{0}$, $t_0 = 1$, and step size $L = \|H\| + \frac{1}{\mu}$.
 - 2: **for** $k=0, 1, 2, \dots$ **do**
 - 3: set $\mathbf{d}^{k+1} = P_C(\mathbf{y}^k - \frac{1}{L}\nabla F_\mu(\mathbf{y}^k))$;
 - 4: set $t_{k+1} = \frac{1+\sqrt{1+4t_k^2}}{2}$;
 - 5: compute $\mathbf{y}^{k+1} = \mathbf{d}^{k+1} + \left(\frac{t_k-1}{t_{k+1}}\right)(\mathbf{d}^{k+1} - \mathbf{d}^k)$.
 - 6: **end for**
-

The following theorem shows that for a given accuracy level ϵ , the smoothing parameter μ and the number of iterations K in Algorithm 6.1 can be chosen carefully to guarantee optimality ϵ in $\mathcal{O}(\frac{1}{\epsilon})$ iterations.

THEOREM 6.1. *Suppose Assumption 3.1, Assumption 3.2 and Assumption 3.3 hold. Consider the set $C := \{\mathbf{d} : \|\mathbf{d}\| \leq r\}$. Let $\{\mathbf{d}^k\}_{k \geq 0}$ be the sequence generated by S-FISTA. For an accuracy level $\epsilon > 0$, if the smoothing parameter μ and the number of iterations K are set as*

$$(6.5) \quad \mu = \frac{2\epsilon}{L_h(L_h + \sqrt{L_h^2 + 2\|H\|\epsilon})} \quad \text{and} \quad K = \frac{r(2L_h + \sqrt{2\|H\|\epsilon})}{\epsilon},$$

then for any $k \geq K$, it holds that $G(\mathbf{d}^k) - G(\mathbf{d}^*) \leq \epsilon$, where \mathbf{d}^* is in the optimal set of problem in (6.1).

Proof. Substitute $(\alpha, \beta, \Gamma, L_f) = (1, \frac{L_h^2}{2}, r^2, \|H\|)$ into [5, Theorem 10.57]. \square

6.1. Trust-region subproblem. We are now able to show that Algorithm 6.1 can be used to solve (2.14) for Algorithm 2.1 in the case of nonlinear least-squares problems (where m_k is given by (4.6)). Within Algorithm 6.1, we set $r := \Delta_k$, $\mathbf{g} := \mathbf{g}_k$, $H := H_k$ and $\epsilon := (1 - e_3)c_1\bar{\eta}_1(\mathbf{x}_k) \min\left\{\Delta_k, \frac{\bar{\eta}_1(\mathbf{x}_k)}{\max\{1, \|H_k\|\}}\right\}$. We use K iterations of Algorithm 6.1 with smoothing parameter μ as per Theorem 6.1. Lemma 6.2 below shows that this gives us the required decrease (2.23).

LEMMA 6.2. *Suppose Assumption 3.2 holds and the step \mathbf{s}_k is calculated such that*

$$(6.6) \quad m_k(\mathbf{x}_k + \mathbf{s}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k^*) \leq (1 - e_3)c_1\bar{\eta}_1(\mathbf{x}_k) \min\left\{\Delta_k, \frac{\bar{\eta}_1(\mathbf{x}_k)}{\max\{1, \|H_k\|\}}\right\},$$

where \mathbf{s}_k^* is an exact minimizer of (2.14). Then the step \mathbf{s}_k satisfies (2.23) with $c_1 := \frac{1}{2} \min\left\{1, \frac{1}{\Delta_{\max}^2}\right\}$.

Proof. This proof follows [19, Lemma 11], with the main difference being the use of $\bar{\eta}_1(\mathbf{x}_k)$ instead of $\Psi_1(\mathbf{x}_k)$.

We begin by defining \mathbf{s}_k^* as a global minimizer of the subproblem (2.14). Let us also define $\tilde{\mathbf{s}}_k$ to be a global minimizer of $\tilde{l}(\mathbf{x}_k, \mathbf{s})$ for $\mathbf{s} \in B(\mathbf{0}, \Delta_k)$ (which exists since h is continuous by Assumption 3.2), and so $\eta_{\Delta_k}(\mathbf{x}_k) = \tilde{l}(\mathbf{x}_k, \mathbf{0}) - \tilde{l}(\mathbf{x}_k, \tilde{\mathbf{s}}_k)$. Note also

that since h is convex, we have, for any $\theta \in [0, 1]$,

$$(6.7) \quad h(\mathbf{x} + \theta \tilde{\mathbf{s}}_k) \leq (1 - \theta)h(\mathbf{x}) + \theta h(\mathbf{x} + \tilde{\mathbf{s}}_k).$$

Hence, for any $\theta \in [0, 1]$, we have

$$(6.8) \quad \begin{aligned} m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k^*) \\ \geq m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \theta \tilde{\mathbf{s}}_k) \end{aligned}$$

$$(6.9) \quad = h(\mathbf{x}_k) - \theta \nabla p_k(\mathbf{x}_k)^T \tilde{\mathbf{s}}_k - \frac{1}{2} \theta^2 \tilde{\mathbf{s}}_k^T H_k \tilde{\mathbf{s}}_k - h(\mathbf{x}_k + \theta \tilde{\mathbf{s}}_k)$$

$$(6.10) \quad \geq \theta h(\mathbf{x}_k) - \theta \nabla p_k(\mathbf{x}_k)^T \tilde{\mathbf{s}}_k - \frac{1}{2} \max\{1, \|H_k\|\} \Delta_k^2 \theta^2 - \theta h(\mathbf{x} + \tilde{\mathbf{s}}_k)$$

$$(6.11) \quad = \theta \eta_{\Delta_k}(\mathbf{x}_k) - \frac{1}{2} \max\{1, \|H_k\|\} \Delta_k^2 \theta^2.$$

Since $\theta \in [0, 1]$ was arbitrary, (6.11) is equivalent to

$$(6.12) \quad m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k^*) \geq \max_{0 \leq \theta \leq 1} \left\{ \theta \eta_{\Delta_k}(\mathbf{x}_k) - \frac{1}{2} \max\{1, \|H_k\|\} \Delta_k^2 \theta^2 \right\}.$$

This is a concave quadratic in θ , with constrained maximizer

$$(6.13) \quad \theta^* = \min \left(1, \frac{\eta_1(\mathbf{x}_k)}{\max\{1, \|H_k\|\} \Delta_k^2} \right).$$

Using this value of θ^* , we get

$$(6.14) \quad m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k^*) \geq \frac{1}{2} \min \left\{ \eta_{\Delta_k}(\mathbf{x}_k), \frac{[\eta_{\Delta_k}(\mathbf{x}_k)]^2}{\max\{1, \|H_k\|\} \Delta_k^2} \right\}.$$

Then [7, Lemma 2.1]³ gives

$$(6.15) \quad m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k^*) \geq \frac{1}{2} \min\{1, \Delta_k\} \eta_1(\mathbf{x}_k) \min \left\{ 1, \frac{\min\{1, \Delta_k\} \eta_1(\mathbf{x}_k)}{\max\{1, \|H_k\|\} \Delta_k^2} \right\}.$$

If $\Delta_k \leq 1$, then (6.15) reduces to

$$(6.16) \quad m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k^*) \geq \frac{1}{2} \eta_1(\mathbf{x}_k) \min \left\{ \Delta_k, \frac{\eta_1(\mathbf{x}_k)}{\max\{1, \|H_k\|\}} \right\}.$$

If $\Delta_k \geq 1$, and noting that $\Delta_{\max} > 1$ (by assumption in Algorithm 2.1), it follows from (6.15) and $\Delta_k \leq \Delta_{\max} \leq \Delta_{\max}^2$ that

$$(6.17) \quad \begin{aligned} m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k^*) &\geq \frac{1}{2} \eta_1(\mathbf{x}_k) \min \left\{ 1, \frac{\eta_1(\mathbf{x}_k)}{\max\{1, \|H_k\|\} \Delta_k^2} \right\} \\ &\geq \frac{1}{2 \Delta_{\max}^2} \eta_1(\mathbf{x}_k) \min \left\{ \Delta_{\max}^2, \frac{\eta_1(\mathbf{x}_k)}{\max\{1, \|H_k\|\}} \right\} \\ &\geq \frac{1}{2 \Delta_{\max}^2} \eta_1(\mathbf{x}_k) \min \left\{ \Delta_k, \frac{\eta_1(\mathbf{x}_k)}{\max\{1, \|H_k\|\}} \right\} \end{aligned}$$

Therefore, combining (6.16) and (6.17), since $\bar{\eta}_1(\mathbf{x}_k) \leq \eta_1(\mathbf{x}_k)$, we obtain:

$$(6.18) \quad m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k^*) \geq c_1 \bar{\eta}_1(\mathbf{x}_k) \min \left\{ \Delta_k, \frac{\bar{\eta}_1(\mathbf{x}_k)}{\max\{1, \|H_k\|\}} \right\}.$$

Subtracting (6.6) from (6.18) gives (2.23) as required. \square

³Stating that $\eta_r(\mathbf{x}) \geq \min\{1, r\} \eta_1(\mathbf{x})$ for any $r > 0$ and $\mathbf{x} \in \mathbb{R}^n$.

6.2. Computing criticality measures. When we implement S-FISTA for the approximate criticality measure $\bar{\eta}_1$ in Algorithm 2.1 (i.e. in (2.20) and (2.24)), we choose $r := 1$, $\mathbf{g} := \nabla p_k(\mathbf{x}_k) = \mathbf{g}_k$ and $H := 0$ in (6.5). The accuracy level is set to be $\epsilon := \min\{(1 - e_1)\epsilon_C, e_2\Delta_k^{\text{init}}\}$ in (2.20) and $\epsilon := e_2\omega_C^{i-1}\Delta_k^{\text{init}}$ in (2.24), matching our requirements (2.22) and (2.26) respectively.

6.3. Smoothed trust-region subproblem. In the case of solving the trust-region problem (5.6) for Algorithm 5.1, we follow (6.5) from Theorem 6.1 and at each iteration use the value

$$(6.19) \quad \mu(\gamma) := \frac{2\gamma}{L_h(L_h + \sqrt{L_h^2 + 2\|H_k\|\gamma})},$$

where $\|H_k\| = \|J_k^T J_k\|$ from (5.5). The total number of iterations K is also given by substituting $\epsilon \leftarrow \gamma$ in (6.5). This ensures that our trust-region subproblem is solved to global optimality level γ , which decreases to zero as Algorithm 5.1 progresses.

7. Numerical Experiments.

7.1. Implementation. We investigate the performance of both approaches for solving the regularized least-squares problem (4.1) by implementing two modifications of the original DFO-LS software [6]:

- DFO-LSR (DFO-LS with Regularization): Algorithm 2.1 with directly solves (4.1).⁴
- DFO-LSSR (DFO-LS with Smoothed Regularization): Algorithm 5.1, which iteratively smoothing (4.1) to (5.2).⁵ We used parameters $\sigma = 0.1$, $\gamma_0 = 0.01$ and trust-region termination function $d(\gamma_j) = \mu(\gamma_j)^2$ for $\mu(\gamma)$ given by (6.19).

All other algorithm parameters are set to their default values in DFO-LS. As described in Section 6, the different subproblems are solved using S-FISTA (Algorithm 6.1) with parameters as described above. In practice, we additionally terminate S-FISTA after 500 iterations, which produces comparable numerical results to using the expected number of iterations, but improves the overall algorithm runtime.

7.2. Testing Setup. The computational performance of DFO-LSR is compared with PyNOMAD, which is a Python interface for NOMAD [26], a solver for black-box optimization based on the Mesh Adaptive Direct Search (MADS) algorithm and is capable of solving nonsmooth problems. We note that NOMAD is not aware of the problem structure in (4.1), it simply can evaluate $\Phi(\mathbf{x})$ as a black-box, whereas our algorithms receive more problem information, namely $\mathbf{r}(\mathbf{x})$ and h (plus its proximal operator). We build our test suite based on [27], a collection of 53 unconstrained nonlinear least squares problems with dimension $2 \leq n \leq 12$ and $2 \leq m \leq 65$. For each problem, we set $f(\mathbf{x})$ to be the nonlinear least squares from [27] and $h(\mathbf{x}) = \|\mathbf{x}\|_1$. Both solvers are tested additionally on our test suite where stochastic noise is introduced in the evaluations of the residuals r_i . Specifically, we introduce i.i.d. $\epsilon \sim N(0, \sigma^2)$ for each i and \mathbf{x} and implement two unbiased noise models as follows:

- Multiplicative Gaussian noise: we evaluate residual $\tilde{r}_i(\mathbf{x}) = r_i(\mathbf{x})(1 + \epsilon)$; and
- Additive Gaussian noise: we evaluate residual $\tilde{r}_i(\mathbf{x}) = r_i(\mathbf{x}) + \epsilon$.

For our testing, we took the noise level to be $\sigma = 10^{-2}$. Each solver was run 10 times on every problem with noise, and a maximum budget of $100(n + 1)$ evaluations was given (for an n -dimensional problem).

⁴Implementation available from <https://github.com/yanjunliu-regina/dfols>.

⁵Implementation available from <https://github.com/khflam/dfols/>.




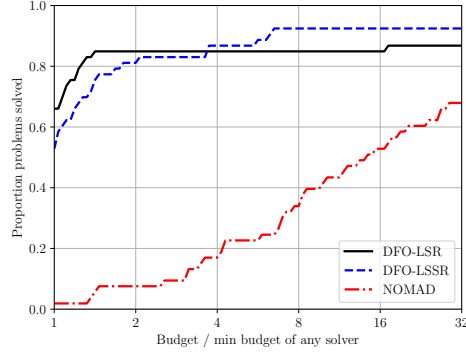
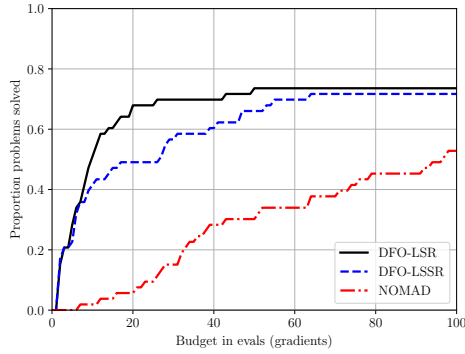


 (a) Data profile, $\tau = 10^{-3}$

 (b) Performance profile, $\tau = 10^{-3}$

 (c) Data profile, $\tau = 10^{-5}$

 (d) Performance profile, $\tau = 10^{-5}$

 (e) Data profile, $\tau = 10^{-7}$

 (f) Performance profile, $\tau = 10^{-7}$

Fig. 1. Comparison of DFO-LSR, DFO-LSSR and NOMAD on smooth/noiseless test problems for increasing accuracy levels $\tau = 10^{-3}$, $\tau = 10^{-5}$ and $\tau = 10^{-7}$.

To compare solvers, we use the data and performance profiles [16]. For every solver \mathcal{S} , each problem p , we determine the number of function evaluations $N_p(\mathcal{S}; \tau)$ required for a problem to be solved up to a given accuracy level $\tau \in (0, 1)$:

$$(7.1) \quad N_p(\mathcal{S}; \tau) := \# \text{ objective evals before } \Phi(\mathbf{x}_k) \leq \Phi^* + \tau(\Phi(\mathbf{x}_0) - \Phi^*),$$

where Φ^* is an approximation to the true minimum of problem p . In our numerical experiments, we take Φ^* to be the smallest objective value generated by either of the two solvers. We set $N_p(\mathcal{S}; \tau) = \infty$ if the inequality in (7.1) is not achieved for a corresponding \mathcal{S} , p and r in the maximum computational budget allowed, here $100(n_p + 1)$.

We compare solvers by calculating the proportion of problems solved in our problem suite \mathcal{P} for a given computational budget. Consider solver \mathcal{S} , problem p and accuracy level $\tau \in (0, 1)$, for *data profiles*, we normalize $N_p(\mathcal{S}; \tau)$ by the dimension for problem p and plot

$$(7.2) \quad d_{\mathcal{S}, \tau}(\alpha) := \frac{|\{p \in \mathcal{P} : N_p(\mathcal{S}; \tau) \leq \alpha(n_p + 1)\}|}{|\mathcal{P}|}, \quad \alpha \in [0, N_g],$$

where N_g is the maximum computational budget, measured in simplex gradients. For *performance profiles*, we plot $N_p(\mathcal{S}; \tau)$ normalized by the minimum number of objective evaluations $N_p^*(\tau) := \min_{\mathcal{S}} N_p(\mathcal{S}; \tau)$:

$$(7.3) \quad \pi_{\mathcal{S}, \tau}(\alpha) := \frac{|\{p \in \mathcal{P} : N_p(\mathcal{S}; \tau) \leq \alpha N_p^*(\tau)\}|}{|\mathcal{P}|}, \quad \alpha \geq 1.$$

7.3. Test Results. In Figure 1, we compare our methods DFO-LSR and DFO-LSSR with NOMAD for problems with noiseless objective evaluations, for accuracy levels $\tau \in \{10^{-3}, 10^{-5}, 10^{-7}\}$. As we would expect, both DFO-LSR and DFO-LSSR significantly outperform NOMAD at all accuracy levels, as they have much more information about the problem structure, including the full residual vector $\mathbf{r}(\mathbf{x})$ and the proximal operator of h . Overall, the performance of DFO-LSR and DFO-LSSR are similar, but we observe that DFO-LSSR outperforms DFO-LSR for lower accuracy levels, but DFO-LSR is better able to reach high accuracy solutions, $\tau = 10^{-7}$.

For problems with noisy objective evaluations, we show the full numerical results at accuracy levels $\tau \in \{10^{-3}, 10^{-5}\}$ in Appendix A. We do not include the highest accuracy level $\tau = 10^{-7}$ here, since such reaching high accuracy requirements is usually impractical for noisy problems without further algorithmic improvements (e.g. sample averaging). For noisy problems, again both DFO-LSR and DFO-LSSR outperform NOMAD. The performance difference between the DFO-LSR and DFO-LSSR is quite small, and so it is not clear from these results which variant is to be preferred in this setting.

8. Conclusions and Future Work. We introduce two model-based derivative-free approaches for solving (1.1), and in particular its specialization to nonlinear least-squares (4.1). First, Algorithm 2.1 extends the approach from [20] to handle inexact stationarity measures, a simpler sufficient decrease condition for the trust-region subproblem, and a more complicated algorithmic framework inherited from [9], without sacrificing global convergence or worst-case complexity results. Secondly, Algorithm 5.1 adapts the smoothing approach from [18] to our setting, and we add to its convergence theory by showing that all accumulation points of the algorithm are Clarke stationary. Our numerical results indicate that both approaches perform well, although there is some evidence to suggest that the smoothing approach is slightly better at achieving low accuracy solutions, while Algorithm 2.1 is better at finding high accuracy solutions.

The most promising direction for future research on this topic is to extend our approaches to the case where f does not have a nonlinear least-squares structure. The

extra difficulty in this setting is to appropriately handle the nonconvex, nonsmooth subproblems that would arise in this setting.

REFERENCES

- [1] A. Y. ARAVKIN, R. BARALDI, AND D. ORBAN, *A proximal quasi-newton trust-region method for nonsmooth regularized optimization*, SIAM Journal on Optimization, 32 (2022), pp. 900–929, <https://doi.org/10.1137/21M1409536>.
- [2] H. ATTOUCH AND J. PEYPOUQUET, *The rate of convergence of nesterov’s accelerated forward-backward method is actually faster than $1/k^2$* , SIAM Journal on Optimization, 26 (2016), pp. 1824–1834.
- [3] C. AUDET AND J. E. DENNIS, *Mesh adaptive direct search algorithms for constrained optimization*, SIAM Journal on Optimization, 17 (2006), pp. 188–217, <https://doi.org/10.1137/040603371>.
- [4] C. AUDET AND W. HARE, *Model-based methods in derivative-free nonsmooth optimization*, in Numerical Nonsmooth Optimization: State of the Art Algorithms, A. M. Bagirov, M. Gaudio, N. Karimtsa, M. M. Mäkelä, and S. Taheri, eds., Springer International Publishing, Cham, 2020, <https://doi.org/10.1007/978-3-030-34910-3>.
- [5] A. BECK, *First-order methods in optimization*, SIAM, 2017.
- [6] C. CARTIS, J. FIALA, B. MARTEAU, AND L. ROBERTS, *Improving the flexibility and robustness of model-based derivative-free optimization solvers*, ACM Transactions on Mathematical Software (TOMS), 45 (2019), pp. 1–41.
- [7] C. CARTIS, N. I. GOULD, AND P. L. TOINT, *On the evaluation complexity of composite function minimization with applications to nonconvex nonlinear programming*, SIAM Journal on Optimization, 21 (2011), pp. 1721–1739.
- [8] C. CARTIS, N. I. GOULD, AND P. L. TOINT, *An adaptive cubic regularization algorithm for nonconvex optimization with convex constraints and its function-evaluation complexity*, IMA Journal of Numerical Analysis, 32 (2012), pp. 1662–1695.
- [9] C. CARTIS AND L. ROBERTS, *A derivative-free Gauss–Newton method*, Mathematical Programming Computation, 11 (2019), pp. 631–674.
- [10] X. CHEN, L. NIU, AND Y. YUAN, *Optimality conditions and a smoothing trust region newton method for nonlipschitz optimization*, SIAM Journal on Optimization, 23 (2013), pp. 1528–1552.
- [11] F. H. CLARKE, *Optimization and nonsmooth analysis*, SIAM, 1990.
- [12] P. CONEJO, E. W. KARAS, L. G. PEDROSO, A. A. RIBEIRO, AND M. SACHINE, *Global convergence of trust-region algorithms for convex constrained minimization without derivatives*, Applied Mathematics and Computation, 220 (2013), pp. 324–330.
- [13] A. R. CONN, N. I. GOULD, AND P. L. TOINT, *Trust region methods*, SIAM, 2000.
- [14] A. R. CONN, K. SCHEINBERG, AND L. N. VICENTE, *Global convergence of general derivative-free trust-region algorithms to first-and second-order critical points*, SIAM Journal on Optimization, 20 (2009), pp. 387–415.
- [15] A. R. CONN, K. SCHEINBERG, AND L. N. VICENTE, *Introduction to derivative-free optimization*, SIAM, 2009.
- [16] E. D. DOLAN AND J. J. MORÉ, *Benchmarking optimization software with performance profiles*, Mathematical programming, 91 (2002), pp. 201–213.
- [17] M. J. EHRHARDT AND L. ROBERTS, *Inexact derivative-free optimization for bilevel learning*, Journal of Mathematical Imaging and Vision, 63 (2021), pp. 580–600, <https://doi.org/10.1007/s10851-021-01020-8>.
- [18] R. GARMANJANI, D. JÚDICE, AND L. N. VICENTE, *Trust-region methods without using derivatives: worst case complexity and the nonsmooth case*, SIAM Journal on Optimization, 26 (2016), pp. 1987–2011.
- [19] G. N. GRAPIGLIA, J. YUAN, AND Y.-X. YUAN, *On the convergence and worst-case complexity of trust-region and regularization methods for unconstrained optimization*, Mathematical Programming, 152 (2015), pp. 491–520.
- [20] G. N. GRAPIGLIA, J. YUAN, AND Y.-X. YUAN, *A derivative-free trust-region algorithm for composite nonsmooth optimization*, Computational and Applied Mathematics, 35 (2016), pp. 475–499.
- [21] W. HARE, C. PLANIDEN, AND C. SAGASTIZÁBAL, *A derivative-free VU-algorithm for convex finite-max problems*, Optimization Methods and Software, 35 (2020), pp. 521–559, <https://doi.org/10.1080/10556788.2019.1668944>.
- [22] M. HOUGH AND L. ROBERTS, *Model-based derivative-free methods for convex-constrained*

- optimization*, SIAM Journal on Optimization, 32 (2022), pp. 2552–2579.
- [23] A. JOURANI, *Limit superior of subdifferentials of uniformly convergent functions*, Positivity, 3 (1999), pp. 33–47.
- [24] J. LARSON AND M. MENICKELLY, *Structure-aware methods for expensive derivative-free nonsmooth composite optimization*, Mathematical Programming Computation, (2023), <https://doi.org/10.1007/s12532-023-00245-5>.
- [25] J. LARSON, M. MENICKELLY, AND S. M. WILD, *Derivative-free optimization methods*, Acta Numerica, 28 (2019), pp. 287–404.
- [26] S. LE DIGABEL, *Algorithm 909: Nomad: Nonlinear optimization with the mads algorithm*, ACM Transactions on Mathematical Software (TOMS), 37 (2011), pp. 1–15.
- [27] J. J. MORÉ AND S. M. WILD, *Benchmarking derivative-free optimization algorithms*, SIAM Journal on Optimization, 20 (2009), pp. 172–191.
- [28] N. PARIKH, S. BOYD, ET AL., *Proximal algorithms*, Foundations and trends® in Optimization, 1 (2014), pp. 127–239.
- [29] M. J. POWELL, *The bobyqa algorithm for bound constrained optimization without derivatives*, Cambridge NA Report NA2009/06, University of Cambridge, Cambridge, 26 (2009).
- [30] M. J. D. POWELL, *The BOBYQA algorithm for bound constrained optimization without derivatives*, Tech. Report DAMTP 2009/NA06, University of Cambridge, 2009.
- [31] S. J. WRIGHT AND B. RECHT, *Optimization for Data Analysis*, Cambridge University Press, 2022.
- [32] Y.-X. YUAN, *Conditions for convergence of trust region algorithms for nonsmooth optimization*, Mathematical Programming, 31 (1985), pp. 220–228.
- [33] Y.-X. YUAN, *Recent advances in trust region algorithms*, Mathematical Programming, 151 (2015), pp. 249–281.
- [34] H. ZHANG, A. R. CONN, AND K. SCHEINBERG, *A derivative-free algorithm for least-squares minimization*, SIAM Journal on Optimization, 20 (2010), pp. 3555–3576.

Appendix A. Supplementary Numerical Results. Below we show our numerical results for problems with noisy objectives, described in Section 7.2. Results for problems with multiplicative Gaussian noise are shown in Figure 2 and problems with additive Gaussian noise in Figure 3.

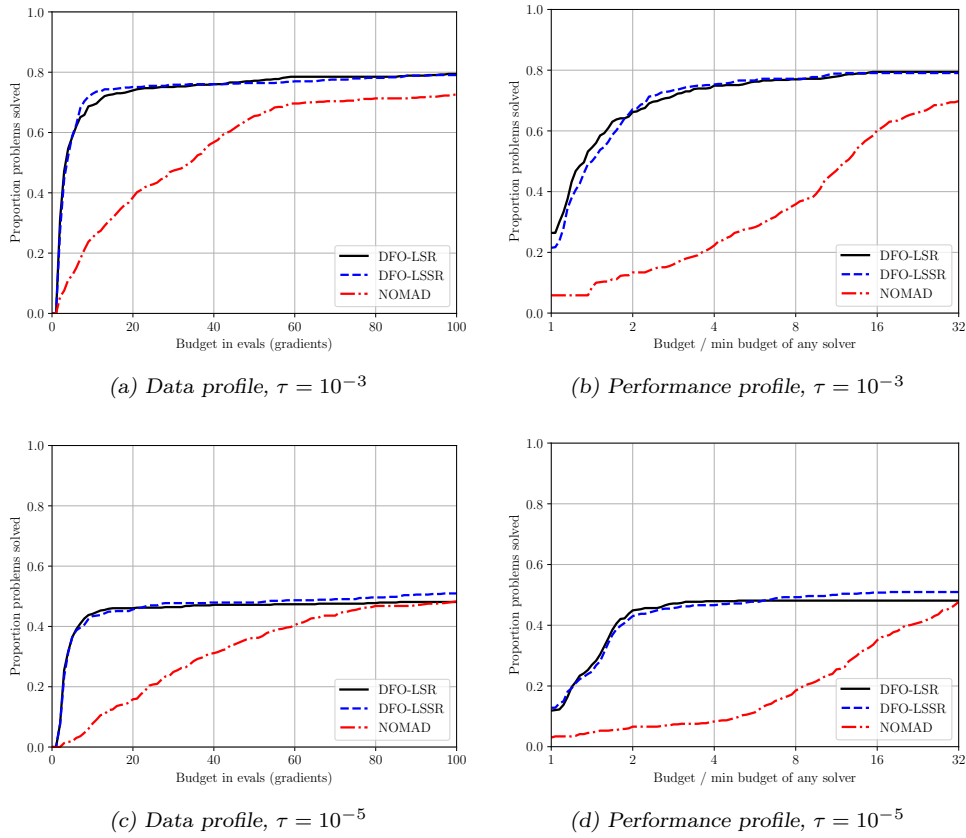


Fig. 2. Comparison of DFO-LSR, DFO-LSSR and NOMAD on test problems with multiplicative Gaussian noise for accuracy levels $\tau = 10^{-3}$ and $\tau = 10^{-5}$.





(a) Data profile, $\tau = 10^{-3}$ (b) Performance profile, $\tau = 10^{-3}$ (c) Data profile, $\tau = 10^{-5}$ (d) Performance profile, $\tau = 10^{-5}$

Fig. 3. Comparison of DFO-LSR, DFO-LSSR and NOMAD on test problems with additive Gaussian noise for accuracy levels $\tau = 10^{-3}$ and $\tau = 10^{-5}$.