

A stochastic Levenberg-Marquardt method using random models with complexity results

El Houcine Bergou^{*}, Youssef Diouane[†], Vyacheslav Kungurtsev[‡], and Clément W. Royer[§]

Abstract. Globally convergent variants of the Gauss-Newton algorithm are often the methods of choice to tackle nonlinear least-squares problems. Among such frameworks, Levenberg-Marquardt and trust-region methods are two well-established, similar paradigms. Both schemes have been studied when the Gauss-Newton model is replaced by a random model that is only accurate with a given probability. Trust-region schemes have also been applied to problems where the objective value is subject to noise: this setting is of particular interest in fields such as data assimilation, where efficient methods that can adapt to noise are needed to account for the intrinsic uncertainty in the input data. In this paper, we describe a stochastic Levenberg-Marquardt algorithm that handles noisy objective function values and random models, provided sufficient accuracy is achieved in probability. Our method relies on a specific scaling of the regularization parameter, that allows us to leverage existing results for trust-region algorithms. Moreover, we exploit the structure of our objective through the use of a family of stationarity criteria tailored to least-squares problems. Provided the probability of accurate function estimates and models is sufficiently large, we bound the expected number of iterations needed to reach an approximate stationary point, which generalizes results based on using deterministic models or noiseless function values. We illustrate the link between our approach and several applications related to inverse problems and machine learning.

Key words. Levenberg-Marquardt method, nonlinear least squares, random models, noisy functions, worst-case complexity, data assimilation, machine learning.

AMS subject classifications. 49M05, 90C56, 90C60

1. Introduction. Minimizing a nonlinear least-squares function is one of the most classical problems in numerical optimization, arising in a variety of applications. In numerous cases, the objective function to be optimized can only be accessed through noisy estimates. Typical occurrences of such a formulation can be encountered while solving inverse problems [17, 34, 35] or while minimizing the error of a model in the context of machine learning [9]. In such cases, the presence of noise is often due to the estimation of the objective function via cheaper, less accurate calculations. For instance, in data-fitting problems, part of the data is often left aside while computing the function or derivative estimates, due to the cost of considering the entire available dataset.

Such concerns have motivated the development of optimization frameworks that cope with inexactness in the objective function or its derivatives. In particular, the field of derivative-free optimization [16], where it is assumed that the derivatives exist but are unavailable for use in an algorithm, has expanded in recent years with the introduction of random models [27]. In a seminal work, Bandeira et al. [1] applied arguments from compressed sensing to guarantee accuracy of quadratic models whenever the Hessian exhibits a certain (unknown) sparsity pattern. Trust-region methods based on

^{*} Mohammed VI Polytechnic University, Ben Guerir, Morocco (elhoucine.bergou@um6p.ma).

[†] ISAE-SUPAERO, Université de Toulouse, 31055 Toulouse Cedex 4, France (youssef.diouane@isae-supero.fr).

[‡] Department of Computer Science, Faculty of Electrical Engineering, Czech Technical University in Prague. Support for this author was provided by the OP VVV project CZ.02.1.01/0.0/0.0/16_019/0000765 “Research Center for Informatics” (vyacheslav.kungurtsev@fel.cvut.cz).

[§] LAMSADE, CNRS, Université Paris Dauphine-PSL, 75016 Paris, France. Support for this author was provided by Subcontract 3F-30222 from Argonne National Laboratory, by CNRS INS2I under the grant GASCON and by ANR through a PRAIRIE chair endowment (clement.royer@dauphine.psl.eu).

general probabilistic models were then proposed and convergence to first- and second-order stationary points was established under appropriate accuracy assumptions on the models [2]. Global convergence rates were derived for this approach in expectation and with high probability [22]. In a parallel line of work, trust-region methods with probabilistic models were extended so as to accommodate noisy function values by Blanchet et al [14]. To this end, the analysis accounts for two sources of randomness, arising from both the noisy function estimates and the random construction of the models. It can then be shown that the trust-region scheme requires at most $\mathcal{O}(\epsilon^{-2})$ iterations in expectation to drive the gradient norm below some threshold ϵ [8].

In the context of derivative-free least-squares problems with exact function values, various deterministic approaches based on globalization of the Gauss-Newton method have been studied. The algorithms developed in the derivative-free community are mostly of trust-region type, and rely on building models that satisfy the so-called fully linear property, which requires the introduction of a so-called criticality step to guarantee its satisfaction throughout the algorithmic process [12, 38, 39, 36]. The recent DFO-GN algorithm [12] was equipped with a complexity result with a bound of the same order as derivative-free trust-region methods for generic functions [20]. As for general problems, random models emerged as a way of relaxing the need for accuracy at every iteration. A Levenberg-Marquardt algorithm based in this idea was proposed by Bergou et al [7], motivated by problems from data assimilation: this method extends the classical Levenberg-Marquardt scheme by replacing the gradient of the objective function with a noisy estimate that is only accurate in probability. Using reasoning similar to the trust-region case [2], almost-sure global convergence to a first-order stationary point was established.

The case of noisy least squares has also been examined. A recent preprint [10] proposed an efficient approach for handling noisy values in practice, but did not provide theoretical guarantees. A Levenberg-Marquardt framework for noisy optimization without derivatives was proposed by Bellavia et al. [3]. This method assumes that function values can be estimated to a prescribed accuracy level and explicitly maintains a sequence of these levels throughout the iterations of the algorithm. Since the noise level must be small compared to the norm of the Levenberg-Marquardt step, one must be able to reduce the noise level when necessary (note that this idea resembles the criticality step of derivative-free model-based methods). In certain applications, this may be deemed as too expensive. By contrast, the use of random models and estimates with properties only guaranteed in probability allows for arbitrarily bad estimates, which seems more economical at the iteration level, and does not exclude the possibility of computing good steps from bad, cheap estimates. Probabilistic properties thus represent a valuable alternative to the above approach. Furthermore, the connections between Levenberg-Marquardt and trust-region methods [30] suggest that the analysis of the latter on noisy problems can help with studying the former.

In this paper, we propose a stochastic Levenberg-Marquardt framework that builds upon the algorithm of Bergou et al. [7] to handle both random models and noise in the function evaluations. Our setup allows for arbitrarily inaccurate models or function estimates: provided those occur at the same time with a small probability, we can equip our method with complexity guarantees. Our analysis adapts that of the stochastic trust-region framework using random models proposed in [8, 14], thanks to an appropriate definition of the Levenberg-Marquardt regularization parameter. In addition, we quantify convergence using a scaled stationarity criterion that accounts for the least-squares structure of our problem and covers standard as well as recently proposed metrics [11].

The remainder of the paper is organized as follows. Section 2 describes our Levenberg-Marquardt algorithm. Section 3 details accuracy requirements that we enforce for the noisy function values and the probabilistic models. Worst-case guarantees for our framework are provided in Section 4. Section 5 discusses several applications of our method to inverse problems and machine learning. Section 6 concludes our work.

2. A Levenberg-Marquardt algorithm based on estimated values. This paper is concerned with the following nonlinear least squares problem:

$$(2.1) \quad \min_{x \in \mathbb{R}^n} f(x) := \frac{1}{2} \|r(x)\|^2,$$

where $r : \mathbb{R}^n \rightarrow \mathbb{R}^\ell$ is a so-called residual, vector-valued function, which we assume to be continuously differentiable, and $\|\cdot\|$ is the Euclidean norm. We consider that r and its derivatives cannot be accessed directly for algorithmic purposes. Therefore, we will present an algorithm that relies on random approximations of these quantities that are of good quality with a certain probability.

In the rest of this section, we recall the main features of the Levenberg-Marquardt method, then describe our extension of this algorithm to handle inexact function and derivative values.

2.1. Deterministic Levenberg-Marquardt paradigm. Popular approaches for solving problem (2.1) are based on the Gauss-Newton model. Given a current iterate x_j , a step is computed as a solution of the linearized least-squares subproblem

$$\min_{s \in \mathbb{R}^n} \frac{1}{2} \|r(x_j) + J(x_j)s\|^2,$$

where $J(\cdot)$ denotes the Jacobian of r . This subproblem possesses a unique solution if $J(x_j)$ has full column rank, and in that case the step is a descent direction for f . When $J(x_j)$ is not of full column rank, the introduction of a regularization parameter can lead to similar properties. This is the underlying idea behind the Levenberg-Marquardt algorithm [28, 29, 31], a globally convergent method based upon the Gauss-Newton model. At each iteration, one considers a step of the form $-(J(x_j)^\top J(x_j) + \gamma_j I)^{-1} J(x_j)^\top r(x_j)$, corresponding to the unique solution of

$$(2.2) \quad \min_{s \in \mathbb{R}^n} \frac{1}{2} \|r(x_j) + J(x_j)s\|^2 + \frac{1}{2} \gamma_j \|s\|^2,$$

where $\gamma_j \geq 0$ is an appropriately chosen regularization parameter, typically updated in the spirit of the classical trust-region radius update strategy at each iteration. In our proposed scheme, we draw a closer connection between this parameter and the trust-region radius by scaling γ_j by the norm of the gradient of the Gauss-Newton model. This approach has been previously proposed in Levenberg-Marquardt-type methods and leads to complexity guarantees that match those of trust-region schemes [5, 40].

2.2. Algorithmic framework based on estimates. Our goal is to propose a method that applies to instances of (2.1) for which neither r nor J can be accessed directly. Consequently, we consider a variant of the Levenberg-Marquardt algorithm, described in Algorithm 2.2, in which both the function and gradient values are approximated. At every iteration, estimates of the values of f and its derivative at the current iterate are computed, and used to compute a Gauss-Newton type model (2.3). A regularized version of this model is then approximately minimized, yielding a trial step s_j . New estimates of the objective at the current and trial point are computed: this new point is accepted if the ratio ρ_j between the estimated function decrease and the model decrease is sufficiently large.

A key feature of our method is that the regularization parameter is defined using a specific scaling formula: namely, we set $\gamma_j = \mu_j \|J_{m_j}^\top r_{m_j}\|$ where $\mu_j \geq 0$. The parameter μ_j is updated depending on the value of ρ_j , and also on a condition involving the model gradient. Such updates are typical of derivative-free model-based methods based on random estimates [2, 7, 14, 22]. Note that we follow these earlier references in checking the condition $\|J_{m_j}^\top r_{m_j}\| \geq \frac{\eta_2}{\mu_j}$ at the end of the iteration, while a more practical application would evaluate it at Step 2 of our algorithm. Our theoretical analysis remains unchanged.

A Levenberg-Marquardt method using random models and estimates.

Initialization

Define $\eta_1 \in (0, 1)$, $\eta_2, \mu_{\min} > 0$, and $\lambda > 1$. Choose x_0 and $\mu_0 \geq \mu_{\min}$.

For $j = 0, 1, 2, \dots$

1. Compute an estimate $f_j^0 = \frac{1}{2}\|r_j^0\|^2$ of $f(x_j)$.
2. Compute r_{m_j} and J_{m_j} , the residual and the Jacobian estimate at x_j , set $\gamma_j = \mu_j \|J_{m_j}^\top r_{m_j}\|$, and define the model m_j of f around x_j by:

$$(2.3) \quad \forall s \in \mathbb{R}^n, \quad m_j(x_j + s) := \frac{1}{2}\|r_{m_j} + J_{m_j}s\|^2 \\ = \frac{1}{2}\|r_{m_j}\|^2 + (J_{m_j}^\top r_{m_j})^\top s + \frac{1}{2}s^\top J_{m_j}^\top J_{m_j}s.$$

3. Compute an approximate solution s_j of the subproblem

$$(2.4) \quad \min_{s \in \mathbb{R}^n} m_j(x_j + s) + \frac{\gamma_j}{2}\|s_j\|^2.$$

4. Compute an estimate $f_j^s = \frac{1}{2}\|r_j^s\|^2$ of $f(x_j + s_j)$, then compute

$$\rho_j := \frac{f_j^0 - f_j^s}{m_j(x_j) - m_j(x_j + s_j) - \frac{\gamma_j}{2}\|s_j\|^2}.$$

5. If $\rho_j \geq \eta_1$ and $\|J_{m_j}^\top r_{m_j}\| \geq \frac{\eta_2}{\mu_j}$, set $x_{j+1} = x_j + s_j$ and $\mu_{j+1} = \max\{\frac{\mu_j}{\lambda}, \mu_{\min}\}$.
Otherwise, set $x_{j+1} = x_j$ and $\mu_{j+1} = \lambda \mu_j$.
-

3. Probabilistic properties of models and function estimates. The framework of Algorithm 2.2 allows for approximations of the objective function and its derivatives to construct both the models and estimate the function values. In this section, we consider that the function values and the derivatives can only be accessed through noisy approximations, and we define accuracy formulas in a deterministic and probabilistic sense.

3.1. Deterministic accuracy. We begin by describing our accuracy requirements for models of the form given in (2.3). Following previous work on derivative-free Levenberg-Marquardt methods [7], we propose the following accuracy definition, and motivate its use further below.

Definition 3.1. *Consider a realization of Algorithm 2.2, and the model m_j of f defined around the iterate x_j of the form (2.3), and let $\kappa_{ef}, \kappa_{eg} > 0$. Then, the model m_j is called $(\kappa_{ef}, \kappa_{eg})$ -first-order accurate with respect to (x_j, μ_j) if the following properties hold:*

$$(3.1) \quad \|J_{m_j}^\top r_{m_j} - J(x_j)^\top r(x_j)\| \leq \frac{\kappa_{eg}}{\mu_j},$$

$$(3.2) \quad \left| \|r(x_j)\|^2 - \|r_{m_j}\|^2 \right| \leq \frac{2\kappa_{ef}}{\mu_j^2}.$$

Remark 3.2. Definition 3.1 resembles that of fully linear models in derivative-free optimization [16], thanks to a specific scaling of the regularization parameter. In particular, the accuracy requirement for the model gradient (3.1) differs from the first-order accuracy property introduced by Bergou, Gratton

and Vicente [7]. With our choice of notation, the latter corresponds to:

$$\left\| J_{m_j}^\top r_{m_j} - J(x_j)^\top r(x_j) \right\| \leq \frac{\kappa \epsilon_g}{\gamma_j},$$

One thus sees that this property uses $\gamma_j = \mu_j \|J_{m_j}^\top r_{m_j}\|$ in the right-hand side, while ours (3.1) uses μ_j . The purpose of our new property is twofold. First, it allows us to measure the accuracy in formulas (3.1) and (3.2) through a parameter that is updated in an explicit fashion throughout the algorithmic run: this is a key property for performing a probabilistic analysis of optimization methods. Secondly, we believe this choice to be a better reflection of the relationship between the Levenberg-Marquardt and the trust-region parameter. Indeed, the global solution of the subproblem (2.4) is given by $d_j = -(J(x_j)^\top J(x_j) + \gamma_j I)^{-1} J(x_j)^\top r(x_j)$, which is also the solution of the trust-region subproblem

$$(3.3) \quad \begin{cases} \min_d & \frac{1}{2} \|r(x_j) + J(x_j)d\|^2 \\ \text{s.t.} & \|d\| \leq \delta_j = \|d_j\|. \end{cases}$$

As a result, we see that for a large value of γ_j , one would have $\delta_j = \mathcal{O}\left(\frac{\|J(x_j)^\top r(x_j)\|}{\gamma_j}\right)$, which suggests that γ_j is not exactly equivalent to the inverse of the trust-region radius (as used in earlier work [7]), but rather is an equivalent to $\frac{\|J(x_j)^\top r(x_j)\|}{\delta_j}$. As a result, the parameter μ_j can be thought as equivalent to $\frac{1}{\delta_j}$: our property (3.1) thus matches the gradient accuracy condition in fully linear models [16].

Note that Definition 3.1 contains two conditions related to the model Jacobian and the value at the current point. The latter property, described by (3.2), is necessary because our method relies on inexact residual values. For the same reasons, we define accuracy conditions for the estimates computed at every iteration of our method.

Definition 3.3. *Consider a realization of Algorithm 2.2, and the residual estimates r_j^0 and r_j^s computed at iteration j . Given $\epsilon_f > 0$, we say that r_j^0 and r_j^s are ϵ_f -accurate estimates of $f(x_j)$ and $f(x_j + s_j)$ with respect to (x_j, μ_j) if*

$$(3.4) \quad \left| \|r_j^0\|^2 - \|r(x_j)\|^2 \right| \leq \frac{2\epsilon_f}{\mu_j^2} \quad \text{and} \quad \left| \|r_j^s\|^2 - \|r(x_j + s_j)\|^2 \right| \leq \frac{2\epsilon_f}{\mu_j^2}.$$

Here again, we point out that the parameter μ_j plays the role of a reciprocal of the trust-region radius. In that sense, the previous definitions are consistent with the definitions of sufficient accuracy presented in the case of stochastic trust-region methods [14].

3.2. Probabilistic properties. The deterministic properties described in the previous section allow for a deterministic inexact analysis of Algorithm 2.2. We are further interested in the case where the models and the estimates are computed in a stochastic fashion. This introduction of randomness implies that the iterates, regularization parameters and trial steps become stochastic processes. We will thus denote by X_j, Γ_j, Π_j and S_j the random quantities at iteration j ; the notations $x_j = X_j(\omega)$, $\gamma_j = \Gamma_j(\omega)$, $\mu_j = \Pi_j(\omega)$ and $s_j = S_j(\omega)$ correspond to realizations of these processes. The random model at iteration j of Algorithm 2.2 will be denoted by M_j , and we use $m_j = M_j(\omega)$ for a realization of that model (corresponding to a realization of the algorithm). Similarly, we let r_{M_j} and J_{M_j} denote the estimates of the residual $r(X_j)$ and the Jacobian $J(X_j)$ at iteration j , with their realizations denoted by $r_{m_j} = r_{M_j}(\omega)$, and $J_{m_j} = J_{M_j}(\omega)$. We also define R_j^0 and R_j^s as the random estimates of $r(X_j)$ and $r(X_j + S_j)$. The realizations of R_j^0 and R_j^s will be denoted by r_j^0 and r_j^s .

Finally, we give the probabilistic equivalents of Definitions 3.1 and 3.3 below.

Definition 3.4. Let $p \in (0, 1]$, $\kappa_{ef} > 0$ and $\kappa_{eg} > 0$. A sequence of random models $\{M_j\}$ is said to be p -probabilistically $\{\kappa_{ef}, \kappa_{eg}\}$ -first-order accurate with respect to the sequence $\{X_j, \Gamma_j\}_j$ if the events

$$U_j := \left\{ \left\| J_{M_j}^\top r_{M_j} - J(X_j)^\top r(X_j) \right\| \leq \frac{\kappa_{eg}}{\Gamma_j} \ \& \ \left| \|r(X_j)\|^2 - \|r_{M_j}\|^2 \right| \leq \frac{2\kappa_{ef}}{\Gamma_j^2} \right\}$$

satisfy the following condition

$$(3.5) \quad p_j^* := P(U_j | \mathcal{F}_{j-1}^{M \cdot R}) \geq p,$$

where $\mathcal{F}_j^{M \cdot R} = \sigma(M_0, \dots, M_{j-1}, R_0^0, R_0^s, \dots, R_{j-1}^0, R_{j-1}^s)$ is the σ -algebra generated by M_0, \dots, M_{j-1} and $R_0^0, R_0^s, \dots, R_{j-1}^0, R_{j-1}^s$.

Definition 3.5. Given constants $\varepsilon_f > 0$, and $q \in (0, 1]$, the sequences of random quantities R_j^0 and R_j^s is called q -probabilistically ε_f -accurate, for corresponding sequence $\{X_j, \Gamma_j\}_j$, if the events

$$V_j := \left\{ \left| \|R_j^0\|^2 - \|R(X_j)\|^2 \right| \leq \frac{2\varepsilon_f}{\Gamma_j^2} \quad \text{and} \quad \left| \|R_j^s\| - \|R(X_j + S_j)\| \right| \leq \frac{2\varepsilon_f}{\Gamma_j^2} \right\}$$

satisfy the following condition

$$(3.6) \quad q_j^* := P(V_j | \mathcal{F}_{j-1/2}^{M \cdot R}) \geq q,$$

where $\mathcal{F}_{j-1/2}^{M \cdot R}$ is the σ -algebra generated by $M_0, \dots, M_j, R_0^0, R_0^s, \dots, R_{j-1}^0, R_{j-1}^s$.

4. Convergence rate analysis. In this section, we provide a theoretical study of our algorithm using stochastic process theory. Our methodology follows the approach by Blanchet et al. [8] for trust-region methods.¹ However, our setup introduces a number of variations that require us to make some modifications in key components of the analysis. In particular, we consider a measure of stationarity that exploits the least-squares form of the problem. Indeed, in order to take advantage of the least-squares structure of our problem, we focus on a scaled optimality criterion inspired by previous proposals for least-squares problems [11]. Rather than considering $\|\nabla f(x)\| \leq \epsilon$ as our stationarity condition, we introduce the criterion:

$$(4.1) \quad \|r(x)\| \leq \epsilon_p \quad \text{or} \quad \|g_r^i(x)\| \leq \epsilon_d,$$

where g_r^i is the so-called scaled gradient defined for a fixed integer $i \in \mathbb{N} \cup \{-1\}$ by

$$(4.2) \quad g_r^i(x) := \begin{cases} \frac{\|J(x)^\top r(x)\|}{\|r(x)\|^{\frac{2^{i+1}-1}{2^i}}} & \text{if } \|r(x)\| \neq 0, \\ 0 & \text{otherwise.} \end{cases}$$

In our framework, the tolerance ϵ_p corresponds to a tolerance after which the noise from the estimated values would dominate the actual residual value in the case of small or zero residuals. When the residuals at the optimum are non-zero, however, we consider a scaled version of the optimality conditions, captured by the scaled gradient g_r^i . In that case, the tolerance ϵ_d can be seen as a scaled version of the classical gradient tolerance. Note that our definition of the scaled gradient (4.2) matches previous

¹When the original, unpublished version of this paper [6] was released, our complexity results differed from those that Blanchet et al. [8] had obtained at that time for their algorithm. The analysis was then improved in the final, published version [8] and matches the the results of this earlier, unpublished version [6].

proposals [11, 21] for $i = 0$, while it corresponds to the classical gradient for $i = -1$. As $i \rightarrow \infty$, we have $g_r^i(x) \rightarrow \frac{\|J(x)^T r(x)\|}{\|r(x)\|^2}$ (when $\|r(x)\| \neq 0$), and the condition $\frac{\|J(x)^T r(x)\|}{\|r(x)\|^2} \geq \epsilon_d$ resembles that of gradient dominance of degree 1 (see [11]).

In order to connect the variation in the objective value with the scaled gradient criterion, we will rely on the lemma below, that is a direct corollary of Gould et al. [21, Lemma 3.11] with fixed i (we refer the reader to this earlier result [21, Lemma 3.11] for a full proof).

Lemma 4.1. *For any $a > b \geq 0$ and $c \in \mathbb{R}$,*

$$(4.3) \quad a^2 - b^2 \leq c \quad \Rightarrow \quad a^{1/2^i} - b^{1/2^i} \leq \frac{c}{a \frac{2^{i+1}-1}{2^i}}.$$

4.1. Assumptions and deterministic results. The objective function will be required to satisfy the following assumptions.

Assumption 4.2. *f is continuously differentiable on an open set containing the level set $\mathcal{L}(x_0) = \{x \in \mathbb{R}^n | f(x) \leq f(x_0)\}$, with Lipschitz continuous gradient, of Lipschitz constant ν .*

We also require that the Jacobian model is uniformly bounded over the sequence of iterates, for every realization of the algorithm.

Assumption 4.3. *There exists $\kappa_{J_m} > 0$ such that for all j and all realizations J_{m_j} of the j -th model Jacobian, one has:*

$$\|J_{m_j}\| \leq \kappa_{J_m}.$$

Finally, we state the assumptions on the approximate solve of the subproblem (2.2). The first assumption states that the trial step achieves a fraction of Cauchy decrease for the regularized model.

Assumption 4.4. *There exists $\theta_{fcd} > 0$ such that for every iteration j of any realization of the algorithm,*

$$(4.4) \quad m_j(x_j) - m_j(x_j + s_j) - \frac{\gamma_j}{2} \|s_j\|^2 \geq \frac{\theta_{fcd}}{2} \frac{\|J_{m_j}^\top r_{m_j}\|^2}{\|J_{m_j}\|^2 + \gamma_j}.$$

The second assumption states that the trial step satisfies desirable bounds on its norm and part of the model decrease, two terms that arise naturally in the theoretical analysis (see the proof of Lemma 4.9).

Assumption 4.5. *At each iteration j and for every realization of the algorithm, the step size satisfies*

$$(4.5) \quad \|s_j\| \leq \frac{2\|J_{m_j}^\top r_{m_j}\|}{\gamma_j} = \frac{2}{\mu_j},$$

and there exists $\theta_{in} > 0$ such that

$$(4.6) \quad |s_j^\top (\gamma_j s_j + J_{m_j}^\top r_{m_j})| \leq \frac{4\|J_{m_j}\|^2 \|J_{m_j}^\top r_{m_j}\|^2 + 2\theta_{in} \|J_{m_j}^\top r_{m_j}\|^2}{\gamma_j^2} = \frac{4\|J_{m_j}\|^2 + 2\theta_{in}}{\mu_j^2}.$$

Several choices for the approximate minimization of $m_j(x_j + s)$ verify relations (4.4), (4.5) and (4.6). In particular, Assumptions 4.4 and 4.5 hold for the exact minimizer of the quadratic subproblem (for any $\theta_{in} > 0$), but they are also valid for inexact steps (for some $\theta_{in} > 0$) such as the Cauchy step or a step computed by the truncated Conjugate Gradient algorithm [7, Lemma 5.1].

As shown by the lemma below, our assumptions guarantee that an accurate model also provides an accurate estimate for the trial step.

Lemma 4.6. *Let Assumptions 4.2, 4.3, and 4.5 hold for a realization of Algorithm 2.2. Consider the j -th iteration of that realization, and suppose that m_j is $(\kappa_{ef}, \kappa_{eg})$ -first-order accurate. Then,*

$$(4.7) \quad |f(x_j + s_j) - m_j(x_j + s_j)| \leq \frac{\kappa_{efs}}{\mu_j^2},$$

where $\kappa_{efs} := \frac{2\kappa_{ef} + 4\kappa_{eg} + 4(\nu + \kappa_{J_m}^2)}{2}$.

Proof. Using a Taylor expansion of the function f around x_j and the definition of m_j , we have:

$$\begin{aligned} |f(x_j + s_j) - m_j(x_j + s_j)| &\leq |f(x_j) + (J(x_j)^\top r(x_j))^\top s_j - m_j(x_j + s_j)| + \frac{\nu}{2} \|s_j\|^2 \\ &\leq |f(x_j) - m_j(x_j)| + \left| \left(J(x_j)^\top r(x_j) - J_{m_j}^\top r_{m_j} \right)^\top s_j \right| + \frac{\|J_{m_j}^\top J_{m_j}\| + \nu}{2} \|s_j\|^2. \end{aligned}$$

By Assumptions 4.3 and 4.5, this leads to

$$|f(x_j + s_j) - m_j(x_j + s_j)| \leq \frac{\kappa_{ef}}{\mu_j^2} + \frac{\kappa_{eg}}{\mu_j} \|s_j\| + \frac{\kappa_{J_m}^2 + \nu}{2} \|s_j\|^2 \leq \frac{\kappa_{ef}}{\mu_j^2} + \frac{2\kappa_{eg}}{\mu_j^2} + \frac{2(\kappa_{J_m}^2 + \nu)}{\mu_j^2},$$

hence the result. ■

The next lemmas describe useful results that hold for any realization of Algorithm 2.2: they will be instrumental in studying the behavior of the method in a probabilistic setting (see Section 4.2).

Lemma 4.7. *Let Assumptions 4.2, 4.3, 4.4, and 4.5 hold for a realization of Algorithm 2.2, and consider its j -th iteration. If the model is $(\kappa_{ef}, \kappa_{eg})$ -first-order accurate and*

$$(4.8) \quad \mu_j \geq \max \left\{ \kappa_{J_m}^2, \frac{8(\kappa_{ef} + \kappa_{efs})}{\eta_1 \theta_{fcd}} \right\} \frac{1}{\|J_{m_j}^\top r_{m_j}\|},$$

then the trial step s_j satisfies

$$(4.9) \quad \|r(x_j + s_j)\|^2 - \|r(x_j)\|^2 \leq -\frac{\eta_1 \theta_{fcd}}{4} \frac{\|J_{m_j}^\top r_{m_j}\|}{\mu_j}.$$

Proof. Since the model is $(\kappa_{ef}, \kappa_{eg})$ -first-order accurate, we have:

$$\begin{aligned} f(x_j + s_j) - f(x_j) &= f(x_j + s_j) - m(x_j + s_j) + m(x_j + s_j) - m_j(x_j) + m_j(x_j) - f(x_j) \\ &\leq \frac{\kappa_{efs}}{\mu_j^2} + m(x_j + s_j) - m_j(x_j) + \frac{\kappa_{ef}}{\mu_j^2} \\ &\leq \frac{\kappa_{ef} + \kappa_{efs}}{\mu_j^2} - \frac{\eta_1 \theta_{fcd}}{2} \frac{\|J_{m_j}^\top r_{m_j}\|^2}{\kappa_{J_m}^2 + \gamma_j} = \frac{\kappa_{ef} + \kappa_{efs}}{\mu_j^2} - \frac{\eta_1 \theta_{fcd}}{2} \frac{\|J_{m_j}^\top r_{m_j}\|^2}{\kappa_{J_m}^2 + \mu_j \|J_{m_j}^\top r_{m_j}\|}, \end{aligned}$$

where we used the result of Lemma 4.6 (where κ_{efs} is defined) and Assumption 4.4. Using the first part of (4.8), we then have $\mu_j \|J_{m_j}^\top r_{m_j}\| \geq \kappa_{J_m}^2$, and therefore

$$\begin{aligned} f(x_{j+1}) - f(x_j) &\leq \frac{\kappa_{ef} + \kappa_{efs}}{\mu_j^2} - \frac{\eta_1 \theta_{fcd}}{2} \frac{\|J_{m_j}^\top r_{m_j}\|^2}{2\mu_j \|J_{m_j}^\top r_{m_j}\|} \\ &= \frac{1}{\mu_j} \left[\frac{\kappa_{ef} + \kappa_{efs}}{\mu_j} - \frac{\eta_1 \theta_{fcd}}{4} \|J_{m_j}^\top r_{m_j}\| \right] \leq \frac{1}{\mu_j} \left[-\frac{\eta_1 \theta_{fcd}}{8} \|J_{m_j}^\top r_{m_j}\| \right], \end{aligned}$$

where the second part of the maximum in (4.8) was used in the last line, yielding (4.9). ■

The next result is a consequence of Lemma 4.7.

Lemma 4.8. *Let the assumptions of Lemma 4.7 hold. If m_j is $(\kappa_{ef}, \kappa_{eg})$ -first-order accurate and*

$$(4.10) \quad \mu_j \geq \left(\kappa_{eg} + \max \left\{ \kappa_{J_m}^2, \frac{8(\kappa_{ef} + \kappa_{efs})}{\eta_1 \theta_{fcd}} \right\} \right) \frac{1}{\|J(x_j)^\top r(x_j)\|},$$

then the trial step s_j satisfies

$$(4.11) \quad \|r(x_j + s_j)\|^2 - \|r(x_j)\|^2 \leq -C_1 \frac{\|J(x_j)^\top r(x_j)\|}{\mu_j},$$

where $C_1 := \frac{\eta_1 \theta_{fcd}}{4} \frac{\max \left\{ \kappa_{J_m}^2, \frac{4(\kappa_{ef} + \kappa_{efs})}{\eta_1 \theta_{fcd}} \right\}}{\kappa_{eg} + \max \left\{ \kappa_{J_m}^2, \frac{4(\kappa_{ef} + \kappa_{efs})}{\eta_1 \theta_{fcd}} \right\}}$. Moreover,

$$(4.12) \quad \|r(x_j + s_j)\|^{1/2^i} - \|r(x_j)\|^{1/2^i} \leq -\frac{C_1}{2^{i+1}} \frac{\|J(x_j)^\top r(x_j)\|}{\|r(x_j)\|^{2^i-1} \mu_j},$$

Proof. Since the model is $(\kappa_{ef}, \kappa_{eg})$ -first-order accurate, we have

$$(4.13) \quad \|J(x_j)^\top r(x_j)\| \leq \|J(x_j)^\top r(x_j) - J_{m_j}^\top r_{m_j}\| + \|J_{m_j}^\top r_{m_j}\| \leq \frac{\kappa_{eg}}{\mu_j} + \|J_{m_j}^\top r_{m_j}\|.$$

Using (4.10) to bound the left-hand side, we obtain:

$$\frac{\kappa_{eg} + \max \left\{ \kappa_{J_m}^2, \frac{8(\kappa_{ef} + \kappa_{efs})}{\eta_1 \theta_{fcd}} \right\}}{\mu_j} \leq \frac{\kappa_{eg}}{\mu_j} + \|J_{m_j}^\top r_{m_j}\|$$

which gives $\mu_j \geq \max \left\{ \kappa_{J_m}^2, \frac{8(\kappa_{ef} + \kappa_{efs})}{\eta_1 \theta_{fcd}} \right\} \frac{1}{\|J_{m_j}^\top r_{m_j}\|}$. We are thus in the assumptions of Lemma 4.7, and (4.9) holds. Using the fact that the model is $(\kappa_{ef}, \kappa_{eg})$ -first-order accurate together with (4.10) and (4.13), we have:

$$\|J(x_j)^\top r(x_j)\| \leq \frac{\kappa_{eg}}{\mu_j} + \|J_{m_j}^\top r_{m_j}\| \leq \frac{\kappa_{eg}}{\kappa_{eg} + \max \left\{ \kappa_{J_m}^2, \frac{8(\kappa_{ef} + \kappa_{efs})}{\eta_1 \theta_{fcd}} \right\}} \|J(x_j)^\top r(x_j)\| + \|J_{m_j}^\top r_{m_j}\|,$$

leading to

$$\|J_{m_j}^\top r_{m_j}\| \geq \frac{\max \left\{ \kappa_{J_m}^2, \frac{8(\kappa_{ef} + \kappa_{efs})}{\eta_1 \theta_{fcd}} \right\}}{\kappa_{eg} + \max \left\{ \kappa_{J_m}^2, \frac{8(\kappa_{ef} + \kappa_{efs})}{\eta_1 \theta_{fcd}} \right\}} \|J(x_j)^\top r(x_j)\|.$$

Combining this relation with (4.9) finally gives (4.11). To obtain (4.12), we simply invoke Lemma 4.1 starting from (4.11). \blacksquare

Lemma 4.9. *Let Assumptions 4.2, 4.3, 4.4 and 4.5 hold. Consider the j -th iteration of a realization of Algorithm 2.2. Suppose that m_j is $(\kappa_{ef}, \kappa_{eg})$ -first-order accurate, (r_j^0, r_j^s) is ε_f -accurate, and that*

$$(4.14) \quad \mu_j \geq \max \left\{ \frac{\alpha + \sqrt{\alpha^2 + 4\alpha\kappa_{J_m}^2(1-\eta_1)}}{2(1-\eta_1)}, \eta_2 \right\} \frac{1}{\|J_{m_j}^\top r_{m_j}\|} := \frac{\kappa_{\mu g}}{\|J_{m_j}^\top r_{m_j}\|},$$

holds, where $\alpha := \varepsilon_f + \kappa_{eg} + \nu + 3\kappa_{J_m}^2 + \theta_{in}$. Then, the j -th iteration is successful (i.e. $\rho_j \geq \eta_1$ and $\|J_{m_j}^\top r_{m_j}\| \geq \frac{\eta_2}{\mu_j}$).

Proof. To simplify the notations, we will omit the indices j in the proof. By definition of the ratio ρ and the model m , we have:

$$\begin{aligned} \left|1 - \frac{\rho}{2}\right| &= \left|1 - \frac{1}{2} \frac{f^0 - f^s}{m(x) - m(x+s) - \frac{\gamma}{2}\|s\|^2}\right| = \frac{|m(x) - m(x+s) - \frac{\gamma}{2}\|s\|^2 - \frac{1}{2}f^0 + \frac{1}{2}f^s|}{|m(x) - m(x+s) - \frac{\gamma}{2}\|s\|^2|} \\ &= \frac{|-g_m^\top s - \frac{1}{2}s^\top(J_m^\top J_m + \gamma I)s - \frac{1}{2}f^0 + \frac{1}{2}f^s|}{|m(x) - m(x+s) - \frac{\gamma}{2}\|s\|^2|} \\ &= \frac{|\frac{1}{2}(f^s - f^0 - g_m^\top s - s^\top J_m^\top J_m s) - \frac{1}{2}s^\top(g_m + \gamma s)|}{|m(x) - m(x+s) - \frac{\gamma}{2}\|s\|^2|} \\ &\leq \frac{\frac{1}{2}|f^s - f^0 - g_m^\top s - s^\top J_m^\top J_m s| + \frac{1}{2}|s^\top(g_m + \gamma s)|}{|m(x) - m(x+s) - \frac{\gamma}{2}\|s\|^2|}. \end{aligned}$$

The first term in the numerator can be bounded using a Taylor expansion of $f(x+s)$. Indeed,

$$\begin{aligned} |f^s - f^0 - g_m^\top s - s^\top J_m^\top J_m s| &= |f^s - f(x+s) + f(x) - f^0 + f(x+s) - f(x) - g_m^\top s - \frac{1}{2}s^\top J_m^\top J_m s| \\ &\leq |f^s - f(x+s)| + |f(x) - f^0| + |f(x+s) - f(x) - g_m^\top s - \frac{1}{2}s^\top J_m^\top J_m s| \\ &\leq |f^s - f(x+s)| + |f(x) - f^0| + |[\nabla f(x) - g_m]^\top s| + \frac{\nu + \|J_m^\top J_m\|}{2}\|s\|^2. \end{aligned}$$

Recalling that $\nabla f(x) = J(x)^\top r(x)$, $g_m = J_m^\top r_m$ and using the accuracy properties of the model and function estimates as well as Assumption 4.3, we obtain:

$$|f^s - f^0 - g_m^\top s - s^\top J_m^\top J_m s| \leq \frac{2\varepsilon_f}{\mu^2} + \frac{\kappa_{eg}\|s\|}{\mu} + \frac{\nu + \kappa_{J_m}^2}{2}\|s\|^2.$$

Thus, we have

$$\left|1 - \frac{\rho}{2}\right| \leq \frac{\frac{\varepsilon_f}{\mu^2} + \frac{\kappa_{eg}\|s\|}{2\mu} + \frac{\nu + \kappa_{J_m}^2}{4}\|s\|^2 + \frac{1}{2}|s^\top(g_m + \gamma s)|}{m(x) - m(x+s) - \frac{\gamma}{2}\|s\|^2}.$$

Using Assumption 4.5 on the numerator and Assumption 4.4 on the denominator, we arrive at

$$\begin{aligned} \left|1 - \frac{\rho}{2}\right| &\leq \frac{\frac{\varepsilon_f}{\mu^2} + \frac{\kappa_{eg}}{\mu^2} + \frac{\nu + \kappa_{J_m}^2}{\mu^2} + \frac{2\kappa_{J_m}^2 + \theta_{in}}{\mu^2}}{m(x) - m(x+s) - \frac{\gamma}{2}\|s\|^2} \leq \frac{(\varepsilon_f + \kappa_{eg} + \nu + 3\kappa_{J_m}^2 + \theta_{in}) \frac{1}{\mu^2}}{\frac{\theta_{fcd}}{2} \frac{\|J_m^\top r_m\|^2}{\|J_m\|^2 + \gamma}} \\ &= \frac{(\varepsilon_f + \kappa_{eg} + \nu + 3\kappa_{J_m}^2 + \theta_{in}) \frac{\|J_m^\top r_m\|^2}{\gamma^2}}{\frac{\theta_{fcd}}{2} \frac{\|J_m^\top r_m\|^2}{\kappa_{J_m}^2 + \gamma}} \\ &= (\varepsilon_f + \kappa_{eg} + \nu + 3\kappa_{J_m}^2 + \theta_{in}) \frac{\kappa_{J_m}^2 + \gamma}{\gamma^2} = \alpha \frac{\kappa_{J_m}^2 + \gamma}{\gamma^2}, \end{aligned}$$

where the last equality uses the definition of α from the lemma's statement.

As a result, we have

$$\begin{aligned} \left|1 - \frac{\rho}{2}\right| \geq 1 - \eta_1 &\Rightarrow \alpha \frac{\kappa_{J_m}^2 + \gamma}{\gamma^2} \geq 1 - \eta_1 \\ \Leftrightarrow 0 &\geq (1 - \eta_1)\gamma^2 - \alpha\gamma - \alpha\kappa_{J_m}^2. \end{aligned}$$

Since the right-hand side is a second-order polynomial in γ , this gives

$$\gamma \leq \frac{\alpha + \sqrt{\alpha^2 + 4\alpha\kappa_{J_m}^2(1-\eta_1)}}{2(1-\eta_1)} \Leftrightarrow \mu \leq \frac{\alpha + \sqrt{\alpha^2 + 4\alpha\kappa_{J_m}^2(1-\eta_1)}}{2(1-\eta_1)} \frac{1}{\|J_m^\top r_m\|}.$$

But this contradicts (4.14), from which we conclude that we necessarily have $|1 - \frac{\rho}{2}| < 1 - \eta_1$, and thus $\rho > \eta_1$. Since $\|J_m^\top r_m\| \geq \frac{\eta_2}{\mu}$ as a direct consequence of (4.14), the iteration is a successful one, and the parameter μ is not increased. \blacksquare

The next lemma shows that having accurate function estimates but inaccurate models still leads to a decrease in the residual on successful iterations.

Lemma 4.10. *Let Assumptions 4.2, 4.3, 4.4, and 4.5 hold. Given a realization of Algorithm 2.2, let j be the index of an unsuccessful iteration. Suppose that (r_j^0, r_j^s) is ε_f -accurate, and suppose*

$$(4.15) \quad \eta_2 \geq \max \left\{ \kappa_{J_m}^2, \frac{8\varepsilon_f}{\eta_1\theta_{fcd}} \right\}$$

Then, one has:

$$(4.16) \quad \|r(x_j + s_j)\|^2 - \|r(x_j)\|^2 \leq -\frac{C_2}{\mu_j^2}.$$

where $C_2 := \frac{\eta_1\eta_2\theta_{fcd}}{2} - 4\varepsilon_f > 0$. Moreover, if $\|r(x_j)\| > 0$, we have

$$(4.17) \quad \|r(x_j + s_j)\|^{1/2^i} - \|r(x_j)\|^{1/2^i} \leq -\frac{C_2}{2^{i+1}} \frac{1}{\mu_j^2 \|r(x_j)\|^{\frac{2^{i+1}-1}{2^i}}}.$$

Proof. By definition of a successful iteration and using the accuracy properties of the models and the estimates, we have

$$\begin{aligned} f(x_{j+1}) - f(x_j) &= f(x_j + s_j) - f(x_j) = \frac{1}{2} [\|r(x_j + s_j)\|^2 - \|r_j^s\|^2 + \|r_j^s\|^2 - \|r_j^0\|^2 + \|r_j^0\|^2 - \|r(x_j)\|^2] \\ &\leq 2\frac{\varepsilon_f}{\mu_j^2} + \frac{1}{2}\|r_j^s\|^2 - \frac{1}{2}\|r_j^0\|^2 \\ &\leq 2\frac{\varepsilon_f}{\mu_j^2} + \eta_1(m(x_j + s_j) - m(x_j)) \\ &\leq 2\frac{\varepsilon_f}{\mu_j^2} - \frac{\eta_1\theta_{fcd}}{2} \frac{\|J_{m_j}^\top r_{m_j}\|^2}{\kappa_{J_m}^2 + \mu_j \|J_{m_j}^\top r_{m_j}\|} \\ &\leq 2\frac{\varepsilon_f}{\mu_j^2} - \frac{\eta_1\theta_{fcd}}{2} \frac{\|J_{m_j}^\top r_{m_j}\|^2}{\eta_2 + \mu_j \|J_{m_j}^\top r_{m_j}\|} \quad \text{since } \eta_2 \geq \kappa_{J_m}^2. \end{aligned}$$

Since the iteration is successful, we have $\mu_j \|J_{m_j}^\top r_{m_j}\| \geq \eta_2$, leading to

$$\|r(x_{j+1})\|^2 - \|r(x_j)\|^2 \leq 4\frac{\varepsilon_f}{\mu_j^2} - \frac{\eta_1\theta_{fcd}}{2} \frac{\|J_{m_j}^\top r_{m_j}\|}{\mu_j} \leq 4\frac{\varepsilon_f}{\mu_j^2} - \frac{\eta_1\eta_2\theta_{fcd}}{2} \frac{1}{\mu_j^2} = -\frac{C_2}{\mu_j^2}.$$

(Note that the positivity of C_2 comes from (4.15)). This establishes the desired result (4.16); by applying Lemma 4.1, we then arrive at (4.17). \blacksquare

To end this section, we formalize our assumptions regarding the probabilistic properties satisfied by our method.

Assumption 4.11. *The random model sequence $\{M_j\}$ is p -probabilistically $\{\kappa_{ef}, \kappa_{eg}\}$ -first-order accurate for some $p \in (0, 1]$, $\kappa_{ef} > 0$, and $\kappa_{eg} > 0$.*

Assumption 4.12. *The sequence of random function estimates $\{(R_j^0, R_j^s)\}_j$ is q -probabilistically ε_f -accurate for some $q \in (0, 1]$ and $\varepsilon_f > 0$.*

Assumption 4.13. *The constant η_2 is chosen such as*

$$(4.18) \quad \eta_2 \geq \max \left\{ \kappa_{J_m}^2, \frac{6(\kappa_{ef} + \kappa_{efs})}{\theta_{fcd}}, \frac{8\varepsilon_f}{\eta_1 \theta_{fcd}} \right\}.$$

In the rest of the paper, we will assume that $pq \neq 1$, since if $pq = 1$, then for every j , we have $p_j^* = P(U_j | \mathcal{F}_{j-1}^M) = p = q_j^* = P(V_j | \mathcal{F}_{j-1}^{M,R}) = q = 1$. and the behavior of the algorithm reduces to that of a deterministic method. Note that we still allow p or q to be equal to 1.

4.2. A key property. Similarly to existing analyzes, the main task in deriving our complexity result consists in proving the following theorem.

Theorem 4.14. *Let Assumptions 4.2, 4.3, 4.4, 4.5 and 4.13 hold. Suppose that Assumptions 4.11 and 4.12 are also satisfied, with the probabilities p and q chosen in a way specified later on. Then,*

$$(4.19) \quad \mathbb{P} \left(\sum_{j=0}^{\infty} \frac{1}{\Pi_j^2 \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}} < \infty \mid \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}} > 0 \forall j \right) = 1.$$

Our proof technique (available in Appendix) is adapted from that in the trust-region setting [14, Theorem 4.11] [26, Lemma 4]. It relies on a Lyapunov function combining the parameter Π_j and a measure of stationarity. Previous analyzes, including an earlier, unpublished version of this paper [6], considered the sequence

$$(4.20) \quad \tau f(X_j) + \frac{1 - \tau}{\Pi_j^2}$$

for an appropriately chosen $\tau \in (0, 1)$. In order to employ our scaled stationarity criterion, we fix an integer $i \in \mathbb{N} \cup \{-1\}$ and define

$$(4.21) \quad \Phi_j := \tau \|r(X_j)\|^{1/2^i} + \frac{1 - \tau}{\Pi_j^2 \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}}.$$

When $i = -1$, we recover the choice (4.20) up to a constant factor. As we will see, the more generic definition (4.21) is well suited to our use of the scaled gradient (4.2).

The proof requires $\tau \in (0, 1)$ to be chosen such that

$$(4.22) \quad \frac{\tau}{1 - \tau} > \max \left\{ \frac{2^{i+1} \left(\lambda^2 - \frac{1}{\lambda^2} \right)}{C_1 \zeta}, \frac{2^{i+1} \left(\lambda^2 - \frac{1}{\lambda^2} \right)}{C_2}, \frac{2^{i+2} \left(\lambda^2 - \frac{1}{\lambda^2} \right)}{\kappa_{ef} + \kappa_{efs}} \right\}$$

where ζ is a parameter such that

$$(4.23) \quad \zeta \geq \left(\kappa_{eg} + \max \left\{ \kappa_{\mu g}, \frac{8(\kappa_{ef} + \kappa_{efs})}{\eta_1 \theta_{fcd}}, \kappa_{J_m}^2, \eta_2 \right\} \right).$$

In addition, the probabilities p and q are required to satisfy:

$$(4.24) \quad \frac{pq - 1/2}{(1-p)(1-q)} \geq \frac{C_3}{C_1}, \quad \text{where } C_3 := 4 \left(1 + \frac{\nu}{\zeta}\right).$$

The bound (4.24) implies that $pq \geq 1/2$, as well as

$$(4.25) \quad (1-p)(1-q) \leq \frac{(1-\tau) \left(1 - \frac{1}{\lambda^2}\right)}{2(\tau C_3 \zeta + (1-\tau)(\lambda^2 - 1))}.$$

We remark that conditions (4.24) and (4.25) will be satisfied for p and q sufficiently close to 1; we also point out that when $q = 1$, these conditions would essentially reduce to $p \geq 1/2$ [7].

Provided conditions (4.22)–(4.25) hold, Theorem 4.14 is obtained by proving that there exists $\sigma > 0$ such that, at every iteration j ,

$$(4.26) \quad \mathbb{E} [\Phi_{j+1} - \Phi_j | \mathcal{F}_{j-1}^{M,R} \cap E_j^0] \leq -\frac{\sigma}{\Pi_j^2 \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}},$$

where $\mathcal{F}_{j-1}^{M,R} \cap E_j^0$ is the trace σ -algebra² produced by the event $E_j^0 = \{\|r(X_k)\|^{\frac{2^{i+1}-1}{2^i}} > 0 \forall k = 0, \dots, j\}$, and the expectation is taken over the product trace σ -algebra generated by all models and function value estimates. We point out that the right-hand side is measurable with respect to $\mathcal{F}_{j-1}^{M,R}$. Since in our case, both $\|r(x_j)\|$ and $f(x_j)$ are bounded from below by 0, we have that $\Phi_j \geq 0$ and $\Pi_j > 0$, (4.26) guarantees that the series converges almost surely (see, e.g., [15, Proposition 4.24]).

The proof focuses on a realization of the process Φ_k , and divides the iterations into two subsets, depending on whether the following condition holds:

$$(4.27) \quad \|J(x_j)^\top r(x_j)\| \geq \frac{\zeta}{\mu_j}.$$

This condition is strongly related to the requirements on μ_j in the lemmas of Section 4.1.

Note that the result of Theorem 4.14 holds for any $i \in \mathbb{N} \cup -1$. In particular, when $i = -1$, the conditioning event $\|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}} = 1 > 0$ is true for all realizations of the method and we simply have:

$$\mathbb{P} \left(\sum_{j=0}^{\infty} \frac{1}{\Pi_j^2} < \infty \right) = 1.$$

As a result, one can follow the lines of the proof of [14, Theorem 4.16] and establish that Algorithm 2.2 to a stationary point, i.e. the sequence of iterates satisfies $\liminf_j \|J(X_j)^\top r(X_j)\| = 0$, with probability 1. In this work, we focus is on complexity results, thus we refer to the first, unpublished version of this work for a full convergence proof [6].

4.3. Complexity bound. We now introduce the necessary probability tools to derive our complexity results. Given a stochastic process $\{X_j\}$, T is said to be a *stopping time* for $\{X_j\}$, if, for all $j \geq 1$, the event $\{T \leq j\}$ belongs to the σ -algebra associated with X_1, X_2, \dots, X_j . For a given $\epsilon > 0$, we define a random time T_ϵ by

$$T_\epsilon := \inf \{j \geq 0, \|r(X_j)\| \leq \epsilon_p \text{ or } \|g_r^i(X_j)\| \leq \epsilon_d\}.$$

We also define $\Pi_\epsilon = \zeta / (\epsilon_p^{\frac{2^{i+1}-1}{2^i}} \epsilon_d)$, where ζ satisfies (4.23). Based on the above analysis, T_ϵ is a stopping time for the stochastic process defined by Algorithm 2.2 and hence for $\{\Phi_j, \Pi_j\}$ where Φ_j is given by (4.21).

²Given a σ -algebra \mathcal{F} and an event E , the trace σ -algebra is defined as $\{F \cap E | F \in \mathcal{F}\}$.

Assumption 4.15. *There exists a positive constant $\kappa_r > 0$ such for all j , w.p. 1, $\|r(X_j)\| \leq \kappa_r$.*

By Assumption 4.15, we have that for any j such that $\|r(X_j)\| \geq \epsilon_p$:

$$\forall j, \quad \Phi_j = \tau \|r(X_j)\|^{1/2^i} + (1 - \tau) \frac{1}{\square_j^2 \|r(X_j)\| \frac{2^{i+1}-1}{2^i}} \leq \kappa_r^{1/2^i} + \frac{1}{\mu_{\min}^2 \epsilon_p \frac{2^{i+1}-1}{2^i}} \leq \Phi_{\max} \epsilon_p^{\frac{1-2^{i+1}}{2^i}},$$

where we define $\Phi_{\max} = \kappa_r^{1/2^i} + \frac{1}{\mu_{\min}^2}$. The last inequality comes from $\frac{2^{i+1}-1}{2^i} \geq 0$ for $i \geq -1$ together with $\epsilon_p \in (0, 1)$.

For simplicity reasons, we will assume that $\mu_0 = \frac{\square_\epsilon}{\lambda^s}$ and $\mu_{\min} = \frac{\square_\epsilon}{\lambda^t}$ for some integers $s, t > 0$, hence for all j , one has $\square_j = \frac{\square_\epsilon}{\lambda^k}$ for some integer k . We note that, in this case, whenever $\square_j < \square_\epsilon$, one has $\square_j \leq \frac{\square_\epsilon}{\lambda}$, and hence $\square_{j+1} \leq \square_\epsilon$. This assumption can be made without loss of generality, for instance, provided $\mu_{\min} = \mu_0 \lambda^{s-t}$ (one can choose μ_{\min} so that this is true) and $\zeta = \mu_0 \lambda^s \epsilon$, where s is the smallest integer such that ζ satisfies (4.23).

The next lemma defines a geometric random walk based on successful iterations. The final complexity result heavily depends upon the behavior of this random walk. Note that this reasoning departs from the existing analysis of stochastic trust-region schemes [8].

Lemma 4.16. *Let Assumptions 4.2, 4.3, 4.4, 4.5 and 4.13 hold. For all $j < T_\epsilon$, whenever $\square_j \geq \square_\epsilon$, one has $\square_{j+1} = \frac{\square_j}{\lambda} \mathbf{1}_{\Omega_j} + \lambda \square_j (1 - \mathbf{1}_{\Omega_j})$, or, equivalently, letting $\gamma = \log(\lambda)$, one has*

$$(4.28) \quad \square_{j+1} = \square_j e^{\gamma \Lambda_j},$$

where $\mathbf{1}_{\Omega_j}$ is equal to 1 if the iteration j is successful and 0 otherwise, and $\Lambda_j = 2\mathbf{1}_{\Omega_j} - 1$ defines a birth-and-death process. That is, $\{\Lambda_j\}$ satisfies

$$\mathbb{P}(\Lambda_j = 1 | \mathcal{F}_{j-1}^{M,R}, \square_j \geq \square_\epsilon) = 1 - \mathbb{P}(\Lambda_j = -1 | \mathcal{F}_{j-1}^{M,R}, \square_j \geq \square_\epsilon) = \omega_j, \quad \text{with } \omega_j \geq pq.$$

Proof. By the mechanism of the algorithm one has $\square_{j+1} = \frac{\square_j}{\lambda} \mathbf{1}_{\Omega_j} + \lambda \square_j (1 - \mathbf{1}_{\Omega_j})$. Moreover, if $\mu_j \geq \square_\epsilon$ for a given $j < T_\epsilon$, one has

$$\|\nabla f(x_j)\| = \|g_r(x_j)\| \|r(x_j)\| \frac{2^{i+1}-1}{2^i} \geq \epsilon_d \epsilon_p \frac{2^{i+1}-1}{2^i} = \frac{\zeta}{\square_\epsilon} \geq \frac{\zeta}{\mu_j}.$$

Assume $\mathbf{1}_{U_j} = 1$ and $\mathbf{1}_{V_j} = 1$ (i.e. both m_j and (r_j^0, r_j^s) are accurate). Because the model is $(\kappa_{ef}, \kappa_{eg})$ -first-order accurate, this implies

$$\|g_{m_j}\| \geq \|\nabla f(x_j)\| - \frac{\kappa_{eg}}{\mu_j} \geq (\zeta - \kappa_{eg}) \frac{1}{\mu_j} \geq \frac{\kappa_{\mu g}}{\mu_j},$$

and since the estimates are also accurate, the iteration is must be successful per Lemma 4.9. Hence, one gets $\omega_j = \mathbb{P}(\Lambda_j = 1 | \mathcal{F}_{j-1}^{M,R}, \square_j \geq \square_\epsilon) \geq pq$. \blacksquare

Lemma 4.16 is analogous to [8, Lemma 7], however, in our case, the birth-and-process $\{\Lambda_j\}$ is based on successful iterations, whereas [8] considered the iterations where both the function estimates and the model were accurate.

For convenience, conditioned on the fact that $T_\epsilon > j$, the following proposition recalls the main argument in proving Theorem 4.14 in the case where $\|J(X_j)^\top r(X_j)\| \geq \frac{\zeta}{\square_j}$ (see ‘‘Case 1’’ in the proof of Theorem 4.14 in Appendix A).

Proposition 4.17. *Let Assumptions 4.2, 4.3, 4.4, 4.5 and 4.13 hold. Suppose that Assumptions 4.11 and 4.12 are also satisfied. Moreover, assume that the probabilities p and q satisfy:*

$$(4.29) \quad \frac{pq - 1/2}{(1-p)(1-q)} \geq \frac{C_3}{C_1},$$

and

$$(4.30) \quad (1-p)(1-q) \leq \frac{(1-\tau)(1-\frac{1}{\lambda^2})}{2(\tau C_3 \zeta + (1-\tau)(\lambda^2 - 1))}.$$

Then, there exists a constant $\sigma > 0$ such that, conditioned on $T_\epsilon > j$, one has

$$(4.31) \quad \mathbb{E} [\Phi_{j+1} - \Phi_j | \mathcal{F}_{j-1}^{M.R}] = \mathbb{E} [\Phi_{j+1} | \mathcal{F}_{j-1}^{M.R}] - \Phi_j < -\frac{\sigma}{\square_j^2 \|r(X_j)\|^{\frac{2^{j+1}-1}{2^j}}},$$

where $\sigma := \frac{1}{4}(1-\tau)(1-\frac{1}{\lambda^2})$.

We now define a renewal process $\{A_i\}$ by $A_0 = 0$ and $A_i = \min\{k > A_{i-1} : \square_k \leq \square_\epsilon\}$, that tracks the iterations indices for which \square_j returns to a prescribed level \square_ϵ . For all $j \geq 1$, we let $\tau_j = A_j - A_{j-1}$. The next result provides a bound on the expected value of τ_j .

Lemma 4.18. *Let Assumptions 4.2, 4.3, 4.4, 4.5 and 4.13 hold. Assuming that $pq > \frac{1}{2}$, one has for all j*

$$(4.32) \quad \mathbb{E} [\tau_j] \leq \frac{pq}{2pq - 1}.$$

Proof. One has

$$(4.33) \quad \begin{aligned} \mathbb{E} [\tau_j] &= \mathbb{E} [\tau_j | \square_{A_{j-1}} < \square_\epsilon] \mathbb{P}(\square_{A_{j-1}} < \square_\epsilon) + \mathbb{E} [\tau_j | \square_{A_{j-1}} = \square_\epsilon] \mathbb{P}(\square_{A_{j-1}} = \square_\epsilon) \\ &\leq \max\{\mathbb{E} [\tau_j | \square_{A_{j-1}} < \square_\epsilon], \mathbb{E} [\tau_j | \square_{A_{j-1}} = \square_\epsilon]\}. \end{aligned}$$

First we note that whenever $\square_j < \square_\epsilon$, one has $\square_j \leq \frac{\square_\epsilon}{\lambda}$, and hence $\square_{j+1} \leq \square_\epsilon$. Thus, if $\square_{A_{j-1}} < \square_\epsilon$, one deduces that $A_j = A_{j-1} + 1$ and then

$$(4.34) \quad \mathbb{E} [\tau_j | \square_{A_{j-1}} < \square_\epsilon] = 1.$$

Assuming now that $A_j > A_{j-1} + 1$ (if not, meaning that $A_j = A_{j-1} + 1$, the proof is straightforward), then conditioned on $\square_{A_{j-1}} = \square_\epsilon$, one has $\square_{A_j} = \square_\epsilon$ as well. We note also that for all $k_j \in [A_{j-1}, A_j]$, one has $\square_{k_j} \geq \square_\epsilon$. Hence, using Lemma 4.16, one has

$$\square_{k_j+1} = \square_{k_j} e^{\gamma \Lambda_{k_j}},$$

where $\gamma = \log(\lambda)$ and $\mathbb{P}(\Lambda_{k_j} = 1 | \mathcal{F}_{k_j-1}^{M.R}, \square_{k_j} \geq \square_\epsilon) = \omega_{k_j}$ and $\mathbb{P}(\Lambda_{k_j} = -1 | \mathcal{F}_{k_j-1}^{M.R}, \square_{k_j} \geq \square_\epsilon) = 1 - \omega_{k_j}$. Moreover, one has $\omega_{k_j} \geq pq$.

The process $\{\square_{A_{j-1}}, \square_{A_{j-1}+1}, \dots, \square_{A_j}\}$ then defines a geometric random walk between two returns to the same state (i.e., \square_ϵ) and τ_j represents the number of iterations until a return to the initial state. For such a geometric random walk, one can define the state probability

vector $\pi = (\pi_k)_k$ corresponding to the limiting stationary distribution [32]. Using the local balance equation between the two states k and $k + 1$, see [32, Theorem 12.13], one has

$$(1 - \omega_k)\pi_k = \omega_k\pi_{k+1}.$$

Since $\omega_k \geq pq$, one deduces that $(1 - pq)\pi_k \geq pq\pi_{k+1}$. Hence,

$$\pi_k \leq \varpi^k \pi_0 \quad \text{where } \varpi = \frac{1 - pq}{pq}.$$

Using the assumption $\varpi < 1$ (i.e. $pq > \frac{1}{2}$) and the definition of the state probability $\sum_k^\infty \pi_k = 1$, one has $\pi_0 \geq 1 - \varpi$ (this is a classical result for geometric random walk, see for instance [32, Example 12.26]).

Applying the properties of ergodic Markov chains, one deduces that the expected number of iterations until a return to the initial state (the state 0) is given by $\frac{1}{\pi_0}$. Hence

$$(4.35) \quad \mathbb{E}[\tau_j | \square_{A_{j-1}} = \square_\epsilon] = \frac{1}{\pi_0} \leq \frac{1}{1 - \varpi} = \frac{pq}{2pq - 1}.$$

By substituting (4.34) and (4.35) into (4.33), one deduces $\mathbb{E}[\tau_j] \leq \frac{pq}{2pq - 1}$ and hence the proof is completed. \blacksquare

We now introduce a counting process $N(j)$ given by the number of renewals that occur before time j :

$$N(j) := \max\{i : A_i \leq j\}.$$

We also consider the sequence of random variables defined by $Y_0 = \Phi_0$ and

$$Y_j = \Phi_{\min(j, T_\epsilon)} + \sigma \sum_{k=0}^{\min(j, T_\epsilon) - 1} \left(\frac{1}{\square_k^2 \|r(X_k)\| \frac{2^{i+1} - 1}{2^i}} \right) \quad \text{for all } j \geq 1.$$

The definition of $\{Y_j\}$ is our second and main distinction from the analysis of stochastic trust region (see [8, Lemma 2.2]).

Lemma 4.19. *Under the assumptions of Proposition 4.17. Let Assumption 4.15 hold. One has,*

$$\mathbb{E}[N(T_\epsilon)] \leq \frac{\Phi_0 \kappa r \frac{2^{i+1} - 1}{2^i}}{\sigma} \square_\epsilon^2.$$

Proof. Note that Y_j defines a supermartingale with respect to $\mathcal{F}_{j-1}^{M.R}$. Indeed, if $j < T_\epsilon$, then using Proposition 4.17 one has,

$$\begin{aligned} \mathbb{E}[Y_{j+1} | \mathcal{F}_{j-1}^{M.R}] &= \mathbb{E}[\Phi_{j+1} | \mathcal{F}_{j-1}^{M.R}] + \mathbb{E}\left[\sigma \sum_{k=0}^j \frac{1}{\square_k^2 \|r(X_k)\| \frac{2^{i+1} - 1}{2^i}} \middle| \mathcal{F}_{j-1}^{M.R}\right] \\ &\leq \Phi_j - \sigma \left(\frac{1}{\square_j^2 \|r(X_j)\| \frac{2^{i+1} - 1}{2^i}} \right) + \sigma \sum_{k=0}^j \left(\frac{1}{\square_k^2 \|r(X_k)\| \frac{2^{i+1} - 1}{2^i}} \right) \\ &= \Phi_j + \sigma \sum_{k=0}^{j-1} \left(\frac{1}{\square_k^2 \|r(X_k)\| \frac{2^{i+1} - 1}{2^i}} \right) = Y_j. \end{aligned}$$

If $j \geq T_\epsilon$, one has $Y_{j+1} = Y_j$ and thus $\mathbb{E} \left[Y_{j+1} | \mathcal{F}_{j-1}^{M,R} \right] = Y_j$. Using Assumption 4.15, one has for all $j \geq T_\epsilon$, $|Y_j| = |Y_{T_\epsilon}| \leq \left(\Phi_{\max} + \frac{(T_\epsilon+1)\sigma}{\mu_{\min}^2} \right) \epsilon_p \frac{1-2^{i+1}}{2^i}$. Hence, since T_ϵ is bounded, Y_j is also bounded. Because T_ϵ is a stopping time, the optimal stopping theorem [33, Theorem 6.4.1] for supermartingales applies, and we have

$$\mathbb{E} [Y_{T_\epsilon}] \leq \mathbb{E} [Y_0].$$

Hence,

$$(4.36) \quad \sigma \mathbb{E} \left[\sum_{k=0}^{T_\epsilon} \frac{1}{\Gamma_k^2 \|r(X_k)\| \frac{2^{i+1}-1}{2^i}} \right] \leq \mathbb{E} [Y_{T_\epsilon}] \leq \mathbb{E} [Y_0] = \Phi_0.$$

By the definition of the counting process $N(T_\epsilon)$, since the renewal times A_i (which satisfy $\Gamma_{A_i} \leq \Gamma_\epsilon$) are a subset of the iterations $0, 1, \dots, T_\epsilon$, one has

$$\sum_{k=0}^{T_\epsilon} \frac{1}{\Gamma_k^2 \|r(X_k)\| \frac{2^{i+1}-1}{2^i}} \geq \sum_{k=0}^{T_\epsilon} \frac{1}{\Gamma_k^2 \kappa_r \frac{2^{i+1}-1}{2^i}} \geq \frac{N(T_\epsilon)}{\Gamma_\epsilon^2 \kappa_r \frac{2^{i+1}-1}{2^i}}$$

Inserting the latter inequality in (4.36), one gets

$$\mathbb{E} [N(T_\epsilon)] \leq \frac{\Phi_0 \kappa_r \frac{2^{i+1}-1}{2^i}}{\sigma} \Gamma_\epsilon^2,$$

which concludes the proof. ■

Using Wald's equation [33, Corollary 6.2.3], we can finally obtain a bound on the expected value of T_ϵ .

Theorem 4.20. *Under the assumptions of Proposition 4.17 as well as Assumption 4.15, one has*

$$\mathbb{E} [T_\epsilon] \leq \frac{pq}{2pq-1} \left(\frac{\Phi_0 \kappa_r \frac{2^{i+1}-1}{2^i}}{\sigma} \Gamma_\epsilon^2 + 1 \right) - 1 \leq \frac{pq}{2pq-1} \left(\kappa_s \kappa_r \frac{2^{i+1}-1}{2^i} \epsilon_d^{-2} \epsilon_p^{-\frac{2^{i+1}-1}{2^{i-1}}} + 1 \right) - 1.$$

where $\kappa_s = \frac{\tau f(x_0) + (1-\tau)\mu_0^{-2}}{\frac{1}{4}(1-\tau)(1-\frac{1}{\lambda^2})} \zeta^2$, $\tau \in (0, 1)$ satisfies (4.22) for a value ζ such that (4.23) holds.

Proof. First note that the renewal process $A_{N(T_\epsilon)+1} = \sum_{i=0}^{N(T_\epsilon)} \tau_i$ where τ_i defines independent inter-arrival times. Moreover, since the probabilities p and q satisfy (4.24), one has $pq > 1/2$ and hence, by applying Lemma 4.18, for all $i = 1, \dots, N(T_\epsilon)$ one has $\mathbb{E} [\tau_i] \leq \frac{pq}{2pq-1} < +\infty$. Thus, by Wald's equation [33, Corollary 6.2.3], one gets,

$$\mathbb{E} [A_{N(T_\epsilon)+1}] = \mathbb{E} [\tau_1] \mathbb{E} [N(T_\epsilon) + 1] \leq \frac{pq}{2pq-1} \mathbb{E} [N(T_\epsilon) + 1].$$

By the definition of $N(T_\epsilon)$ one has $A_{N(T_\epsilon)+1} > T_\epsilon$, hence using Lemma 4.19 one gets

$$\mathbb{E}[T_\epsilon] \leq \mathbb{E}[A_{N(T_\epsilon)+1}] - 1 \leq \frac{pq}{2pq-1} \left(\frac{\Phi_0 \kappa_r^{\frac{2^{i+1}-1}{2^i}}}{\sigma} \Gamma_\epsilon^2 + 1 \right) - 1,$$

which establishes the result by definition of κ_s , Φ_0 and Γ_ϵ . \blacksquare

We now comment on the complexity orders appearing in Theorem 4.20. When $i = -1$, the result of the theorem matches the bounds derived for stochastic trust-region methods [8]. For $i = 0$, the order becomes $\epsilon_d^{-2} \epsilon_p^{-2}$ in expectation for the classical scaled optimality criterion. In addition, as $i \rightarrow \infty$, the bound of Theorem 4.20 asymptotically becomes of order $\epsilon_d^{-2} \epsilon_p^{-4}$ for the scaled criterion $\frac{\|J(x)^\top r(x)\|}{\|r(x)\|^2} \leq \epsilon_d$ or $\|r(x)\| \leq \epsilon_p$. Note that if we set $\epsilon_p = \epsilon \epsilon_d^{-1}$ for some $\epsilon \in (0, \epsilon_d)$, we recover a bound in ϵ^{-2} for the case $i = 0$, matching the order obtained asymptotically by Gould et al. [21].

5. Applications. In this section, we discuss the relevance of our approach for several formulations arising in inverse problems, data assimilation and machine learning.

5.1. Ensemble methods for inverse problems and data assimilation. We describe the problem of interest using the terminology of inverse problems. Let x^* be a ground truth vector and y be a vector of observations such that

$$(5.1) \quad y = \mathcal{H}(x^*) + \zeta,$$

where \mathcal{H} is the observational operator mapping the unknowns from the state space to the observation space and $\zeta \sim \mathcal{N}(0, R)$ is a Gaussian noise vector with given covariance matrix R . The maximum likelihood estimator for the vector x^* can then be obtained by solving:

$$(5.2) \quad \min_x \frac{1}{2} (\|y - \mathcal{H}(x)\|_{R^{-1}}^2).$$

In inverse problems, the formulation (5.2) is typically ill-posed, and a regularized version of this problem is used instead. The formulation of Iglesias et al. [24] incorporates prior knowledge of x in the form of a finite dimensional space where problem (5.2) is solved. This approach bears strong connection with data assimilation, in which a prior knowledge on the state of the form $x = b + \xi$ is typically assumed, with b being a fixed prior (or a background) on x and $\xi \sim \mathcal{N}(0, B^\infty)$ being a Gaussian noise vector with unknown covariance matrix B^∞ (see for instance [23, 25, 13]). With such a prior, problem (5.2) is then replaced by

$$(5.3) \quad \min_x \frac{1}{2} \left(\|y - \mathcal{H}(x)\|_{R^{-1}}^2 + \|x - b\|_{(B^\infty)^{-1}}^2 \right),$$

which is a maximum a posteriori estimation problem. The latter formulation matches the structure of problem (2.1) with $r(x) := \begin{pmatrix} (B^\infty)^{-1/2}(x - x_b) \\ R^{-1/2}(y - \mathcal{H}(x)) \end{pmatrix}$.

In both data assimilation and inverse problems, an important issue with formulation (5.3) is the estimation of the covariance matrix B^∞ . One possible approach to address this issue

consists in estimating B^∞ using a random approximation via an ensemble. Assuming that an initial ensemble of \hat{n} elements $\{\hat{x}^k\}_{k=1}^{\hat{n}}$ is available, the matrix B^∞ is approximated by the empirical covariance matrix of the ensemble, given by

$$(5.4) \quad B^{\hat{n}} := \frac{1}{\hat{n} - 1} \sum_{k=1}^{\hat{n}} (\hat{x}^k - b) (\hat{x}^k - b)^\top.$$

By assuming that the initial ensemble is sampled from the Gaussian distribution with mean b and unknown covariance matrix B^∞ , then the matrix $B^{\hat{n}}$ follows a Wishart distribution [37]. Thus, if $\hat{n} \geq n + 1$, then $B^{\hat{n}}$ is nonsingular with probability one, $(B^{\hat{n}})^{-1}$ follows an inverse Wishart distribution, and

$$(5.5) \quad \mathbb{E} [B^{\hat{n}}] = B^\infty \text{ and } \mathbb{E} [(B^{\hat{n}})^{-1}] = \frac{\hat{n} - 1}{\hat{n} - 1 - n} (B^\infty)^{-1}.$$

Provided \hat{n} is large enough relative to n , the empirical covariance matrix $B^{\hat{n}}$ can be assumed to be non-singular and $\mathbb{E} [(B^{\hat{n}})^{-1}]$ is very close to $(B^\infty)^{-1}$. Using $B^{\hat{n}}$ to approximate B^∞ , a noisy approximation of the objective function in (5.3) is given by

$$(5.6) \quad \frac{1}{2} \left(\|y - \mathcal{H}(x)\|_{R^{-1}}^2 + \|x - b\|_{(B^{\hat{n}})^{-1}}^2 \right).$$

Iglesias et al. [24, 23] proposed approximately solving (5.6) using the Ensemble Kalman Inversion (EKI) method [34], an adaptation of the ensemble Kalman filter (EnKF) used in data assimilation [18, 19] to inverse problems. The TEKI (Tikhonov-EKI) approach of Chada et al. [13] adds an extra regularization of the form $\frac{1}{2}\gamma\|x\|^2$ to the formulation (5.6), note the resemblance to the Levenberg-Marquardt approach. More recently, Iglesias and Yang [25] introduced a new adaptive regularization strategy for EKI that comes even closer to our work, with the most significant distinction being that the subproblems considered in the EKI framework are typically nonlinear. By linearizing the operator \mathcal{H} as in Gauss-Newton techniques, one would obtain an algorithm quite close in spirit to Algorithm 2.2, and our analysis could help in endowing that method with global convergence and complexity guarantees even in the presence of noise in the objective value.

Despite this connection, we acknowledge that our setup (and in particular, satisfaction of the properties (5.5)) requires significantly more samples than what is usually used in practical ensemble methods. Furthermore, our analysis relies on conditional independence between the iterations, which would amount to generate new ensembles at every iteration. Such approaches would likely be impractical in data assimilation and inverse problems, yet we believe that the links between our results could be informative in developing new variants of the EKI approach with global guarantees.

5.2. Subsampling approaches. Least-squares problems play a prominent role in data analysis, as the least-squares loss is a common way to assess the accuracy of a machine learning model. In such a setup, the residual function can be written as $r(x) = [y_i - h(z_i, x)]_{i=1}^m$, where $\{(z_i, y_i)\}_{i=1}^m$ are data points and h is a learning model parameterized by x , such as the output

of a neural network. In modern data science applications, the number of data points is usually quite large, but the data is assumed to follow a known distribution form. As a result, one can rely on using subsamples of the data points to estimate the function values and the gradient: such a paradigm is at the essence of stochastic gradient methods [9].

In this context, it is possible to guarantee the probabilistic properties defined in Section 3 with a sufficiently large number of samples. Following recent work in optimization and subsampling [4], we know that $\mathcal{O}\left(\ln\left(\frac{1}{\delta}\right)\lceil\lceil_j^4\rceil\right)$ samples are needed to get accurate function estimates with probability $1 - \delta$, while $\mathcal{O}\left(\ln\left(\frac{1}{\delta}\right)\lceil\lceil_j^2\rceil\right)$ samples can be used to build a gradient estimate that is accurate with probability $1 - \delta$. Therefore, in order to satisfy Assumptions 4.11 and 4.12, we require

$$\mathcal{O}\left(\ln\left(\frac{1}{1-p}\right)\max\{\lceil\lceil_j^4\rceil, \lceil\lceil_j^2\rceil\}\right) + \mathcal{O}\left(\ln\left(\frac{1}{1-q}\right)\lceil\lceil_j^4\rceil\right)$$

samples per iteration. Note that these rates naturally extend to the stochastic case [26].

6. Conclusion. We proposed a stochastic Levenberg-Marquardt method to solve nonlinear least-squares problems wherein the objective function and its gradient are subject to noise and can only be computed accurately within a certain probability. By employing a scaling formula for the Levenberg-Marquardt parameter, we leveraged the link between our approach and a trust-region-type framework to obtain complexity bounds in expectation for our method. Our guarantees are based upon a scaled gradient criterion that exploits the least-squares structures of our problem, and generalizes previously proposed metrics. We have illustrated how our working assumptions can hold in the context of inverse problems and machine learning where the noise in the objective arises from empirical covariance estimators and subsampling techniques, respectively.

The study of the performance of our approach when applied to large-scale data assimilation problems is a natural continuation of the present work, that poses additional challenges. Indeed, in practical situations, a single ensemble may be used at each iteration, which would introduce correlations between the model and the estimates. In addition, the ensemble size might be significantly smaller than the dimension of the state space, which can jeopardize the quality of the ensemble approximations. Extending our analysis to these settings raises a number of theoretical issues, which we plan to investigate so as to hew our method closer to standard practice.

Acknowledgments. The authors would like to thank Matt Menickelly and Katya Scheinberg for useful discussions regarding the STORM algorithm. The authors are also grateful to the associate editor and two anonymous referees for the careful reading of the paper and their suggestions.

References.

- [1] A. S. BANDEIRA, K. SCHEINBERG, AND L. N. VICENTE, *Computation of sparse low degree interpolating polynomials and their application to derivative-free optimization*, Math. Program., 134 (2012), pp. 223–257.
- [2] A. S. BANDEIRA, K. SCHEINBERG, AND L. N. VICENTE, *Convergence of trust-region methods based on probabilistic models*, SIAM J. Optim., 24 (2014), pp. 1238–1264.

- [3] S. BELLAVIA, S. GRATTON, AND E. RICCIETTI, *A Levenberg-Marquardt method for large nonlinear least-squares problems with dynamic accuracy in functions and gradients*, Numer. Math., 140 (2018), pp. 791–824.
- [4] E. BERGOU, Y. DIOUANE, V. KUNC, V. KUNGURTSEV, AND C. W. ROYER, *A subsampling line-search method with second-order results*, (June 2021). arXiv:1810.07211v4.
- [5] E. BERGOU, Y. DIOUANE, AND V. KUNGURTSEV, *Convergence and iteration complexity analysis of a Levenberg-Marquardt algorithm for inverse problems*, J. Optim. Theory Appl., 185 (2020), pp. 927–944.
- [6] E. BERGOU, Y. DIOUANE, V. KUNGURTSEV, AND C. W. ROYER, *A stochastic Levenberg-Marquardt method using random models with application to data assimilation*. arXiv:1807.02176v1, 2018.
- [7] E. BERGOU, S. GRATTON, AND L. N. VICENTE, *Levenberg-Marquardt methods based on probabilistic gradient models and inexact subproblem solution, with application to data assimilation*, SIAM/ASA J. Uncertain. Quantif., 4 (2016), pp. 924–951.
- [8] J. BLANCHET, C. CARTIS, M. MENICKELLY, AND K. SCHEINBERG, *Convergence rate analysis of a stochastic trust region method for nonconvex optimization*, INFORMS Journal on Optimization, 1 (2019), pp. 92–119.
- [9] L. BOTTOU, F. E. CURTIS, AND J. NOCEDAL, *Optimization Methods for Large-Scale Machine Learning*, SIAM Rev., 60 (2018), pp. 223–311.
- [10] C. CARTIS, J. FIALA, B. MARTEAU, AND L. ROBERTS, *Improving the flexibility and robustness of model-based derivative-free optimization solvers*, ACM Trans. Math. Software, (2019 (to appear)).
- [11] C. CARTIS, N. I. M. GOULD, AND P. L. TOINT, *On the evaluation complexity of cubic regularization methods for potentially rank-deficient nonlinear least-squares problems and its relevance to constrained nonlinear optimization*, SIAM J. Optim., 23 (2013), pp. 1553–1574.
- [12] C. CARTIS AND L. ROBERTS, *A derivative-free Gauss-Newton method*, Math. Program. Comput., (2019 (available online)).
- [13] N. K. CHADA, A. M. STUART, AND X. T. TONG, *Adaptive regularisation for ensemble kalman inversion*, sinum, 29 (2020).
- [14] R. CHEN, M. MENICKELLY, AND K. SCHEINBERG, *Stochastic optimization using a trust-region method and random models*, Math. Program., 169 (2018), pp. 447–487.
- [15] E. ÇINLAR, *Probability and Stochastics*, Graduate Texts in Mathematics, Springer, New York, 2011.
- [16] A. R. CONN, K. SCHEINBERG, AND L. N. VICENTE, *Introduction to Derivative-Free Optimization*, MPS-SIAM Series on Optimization, SIAM, Philadelphia, 2009.
- [17] P. COURTIER, J. N. THYÉPAUT, AND A. HOLLINGSWORTH, *A strategy for operational implementation of 4d-var, using an incremental approach*, Quarterly Journal of the Royal Meteorological Society, 120 (1994), pp. 1367–1387.
- [18] G. EVENSEN, *Sequential data assimilation with a nonlinear quasi-geostrophic model using monte carlo methods to forecast error statistics*, Journal of Geophysical Research: Oceans, 99 (1994), pp. 10143–10162.
- [19] G. EVENSEN, *Data Assimilation: The Ensemble Kalman Filter*, Springer, second ed., 2009.
- [20] R. GARMANJANI, D. JÚDICE, AND L. N. VICENTE, *Trust-region methods without using derivatives: Worst case complexity and the non-smooth case*, SIAM J. Optim., 26 (2016), pp. 1987–2011.
- [21] N. I. M. GOULD, T. REES, AND J. A. SCOTT, *Convergence and evaluation-complexity analysis of a regularized tensor-Newton method for solving nonlinear least-squares problems*, Comput. Optim. Appl., (2019 (available online)).

- [22] S. GRATTON, C. W. ROYER, L. N. VICENTE, AND Z. ZHANG, *Complexity and global rates of trust-region methods based on probabilistic models*, IMA J. Numer. Anal., 38 (2018), pp. 1579–1597.
- [23] M. IGLESIAS, *A regularizing iterative ensemble kalman method for pde-constrained inverse problems*, Inverse Problems, 32 (2016).
- [24] M. IGLESIAS, K. J. H. LAW, AND A. M. STUART, *Ensemble kalman methods for inverse problems*, Inverse Problems, 29 (2013).
- [25] M. IGLESIAS AND Y. YANG, *Adaptive regularisation for ensemble kalman inversion*, Inverse Problems, 37 (2020).
- [26] J. LARSON AND S. C. BILLUPS, *Stochastic derivative-free optimization using a trust region framework*, Comput. Optim. Appl., 64 (2016), pp. 619–645.
- [27] J. LARSON, M. MENICKELLY, AND S. M. WILD, *Derivative-free optimization methods*, Acta Numer., 28 (2019), pp. 287–404.
- [28] K. LEVENBERG, *A method for the solution of certain problems in least squares*, Quart. Appl. Math., 2 (1944), pp. 164–168.
- [29] D. MARQUARDT, *An algorithm for least-squares estimation of nonlinear parameters*, SIAM J. Appl. Math., 11 (1963), pp. 431–441.
- [30] J. J. MORÉ, *The Levenberg-Marquardt algorithm: implementations and theory*, in Lecture Notes in Math., G. A. Watson, ed., vol. 360, Springer-Verlag, Berlin, 1977, pp. 105–116.
- [31] M. R. OSBORNE, *Nonlinear least squares—the Levenberg algorithm revisited*, J. Austral. Math. Soc. Ser. B, 19 (1976), pp. 343–357.
- [32] D. J. G. R. D. YATES, *Probability and Stochastic Processes: A Friendly Introduction for Electrical and Computer Engineers*, Engineering mathematics, John Wiley & Sons, Inc., New York, 2005.
- [33] S. M. ROSS, *Stochastic Processes*, Mathematics, Wiley, New York, 1983.
- [34] A. TARANTOLA, *Inverse Problem Theory and Methods for Model Parameter Estimation*, SIAM, Philadelphia, 2005.
- [35] Y. TRÉMOLET, *Model-error estimation in 4D-Var*, Quarterly Journal of the Royal Meteorological Society, 133 (2007), pp. 1267–1280.
- [36] S. M. WILD, *POUNDERS in TAO: Solving Derivative-Free Nonlinear Least-Squares Problems with POUNDERS*, in Advances and Trends in Optimization with Engineering Applications, T. Terlaky, M. F. Anjos, and S. Ahmed, eds., MOS-SIAM Book Series on Optimization, Society for Industrial and Applied Mathematics, Philadelphia, 2017, ch. 40.
- [37] J. WISHART, *The generalised product moment distribution in samples from a normal multivariate population*, Biometrika, 20A (1928), pp. 32–52.
- [38] H. ZHANG AND A. R. CONN, *On the local convergence of a derivative-free algorithm for least-squares minimization*, Comput. Optim. Appl., 51 (2012).
- [39] H. ZHANG, A. R. CONN, AND K. SCHEINBERG, *A derivative-free algorithm for the least-square minimization*, SIAM J. Optim., 20 (2010), pp. 3555–3576.
- [40] R. ZHAO AND J. FAN, *Global complexity bound of the Levenberg-Marquardt method*, Optim. Methods Softw., 31 (2016), pp. 805–814.

A. Proof of Theorem 4.14.

Proof of Theorem 4.14. Consider a realization of Algorithm 2.2, and let ϕ_j be the corresponding realization of Φ_j . If j is the index of a successful iteration, then $x_{j+1} = x_j + s_j$, and $\mu_{j+1} \geq \frac{\mu_j}{\lambda}$. One thus has:

$$(A.1) \quad \phi_{j+1} - \phi_j \leq \tau \left(\|r(x_{j+1})\|^{1/2^i} - \|r(x_j)\|^{1/2^i} \right) + (1 - \tau) \frac{\lambda^2 - 1}{\mu_j^2 \|r_j\|^{2^{i+1}-1}}.$$

If j is the index of an unsuccessful iteration, $x_{j+1} = x_j$ and $\mu_{j+1} = \lambda\mu_j$, leading to

$$(A.2) \quad \phi_{j+1} - \phi_j = (1 - \tau) \left(\frac{1}{\lambda^2} - 1 \right) \frac{1}{\mu_j^2 \|r(x_j)\|^{2^{i+1}-1}} < 0.$$

For both types of iterations, we will consider four possible outcomes, involving the quality of the model and that of the estimates.

Case 1: $\|J(x_j)^\top r(x_j)\| \geq \frac{\zeta}{\mu_j}$.

1. **Both m_j and (r_j^0, r_j^s) are accurate.** Since we are in Case 1,

$$\left\| J(x_j)^\top r(x_j) \right\| \geq (\kappa_{eg} + \kappa_{\mu g}) \frac{1}{\mu_j}.$$

Because the model is $(\kappa_{ef}, \kappa_{eg})$ -first-order accurate, this implies

$$\|J_{m_j}^\top r_{m_j}\| \geq \|J(x_j)^\top r(x_j)\| - \frac{\kappa_{eg}}{\mu_j} \geq (\zeta - \kappa_{eg}) \frac{1}{\mu_j} \geq \frac{\kappa_{\mu g}}{\mu_j}$$

so (4.14) holds; since the estimates are also accurate, the iteration is necessarily successful by Lemma 4.9. Moreover,

$$\|J(x_j)^\top r(x_j)\| \geq \frac{\zeta}{\mu_j} \geq \left(\kappa_{eg} + \max \left\{ \kappa_{J_m}^2, \frac{8(\kappa_{ef} + \kappa_{efs})}{\eta_1 \theta_{fcd}} \right\} \right) \frac{1}{\mu_j},$$

so the condition (4.10) is satisfied, and by Lemma 4.8, we can guarantee a decrease on the function value. More precisely,

$$\|r(x_{j+1})\|^{1/2^i} - \|r(x_j)\|^{1/2^i} \leq -\frac{C_1}{2^{i+1} \frac{\|J(x_j)^\top r(x_j)\|}{\|r(x_j)\|^{2^{i+1}-1}} \mu_j},$$

leading to

$$(A.3) \quad \varphi_{j+1} - \varphi_j \leq -\tau \frac{C_1}{2^{i+1}} \frac{\|J(x_j)^\top r(x_j)\|}{\|r(x_j)\|^{2^{i+1}-1} \mu_j} + (1 - \tau)(\lambda^2 - 1) \frac{1}{\mu_j^2 \|r(x_j)\|^{2^{i+1}-1}}.$$

2. **Only m_j is accurate.** The decrease formula of Lemma 4.8 is valid in this case: therefore, if the iteration is successful, then (A.3) holds, and we have

$$\begin{aligned}
\varphi_{j+1} - \varphi_j &\leq -\tau \frac{C_1}{2^{i+1}} \frac{\|J(x_j)^\top r(x_j)\|}{\|r(x_j)\|^{\frac{2^{i+1}-1}{2^i}} \mu_j} + (1-\tau)(\lambda^2 - 1) \frac{1}{\mu_j^2 \|r(x_j)\|^{\frac{2^{i+1}-1}{2^i}}} \\
&\leq \frac{-\tau \frac{C_1}{2^{i+1}} \zeta + (1-\tau)(\lambda^2 - 1)}{\mu_j^2 \|r(x_j)\|^{\frac{2^{i+1}-1}{2^i}}} \\
&\leq \frac{-\tau \frac{C_1}{2^{i+2}} \zeta + (1-\tau)(\lambda^2 - 1)}{\mu_j^2 \|r(x_j)\|^{\frac{2^{i+1}-1}{2^i}}} \\
&< (1-\tau) \left(\frac{1}{\lambda^2} - 1 \right) \frac{1}{\mu_j^2 \|r(x_j)\|^{\frac{2^{i+1}-1}{2^i}}},
\end{aligned}$$

using (4.22) to obtain the last inequality. Therefore, (A.2) holds when the iteration is successful. From the beginning of the proof, we know that (A.2) also holds if the iteration is unsuccessful.

3. **Only (r_j^0, r_j^s) is accurate.** If the iteration is unsuccessful, then (A.2) is satisfied. Otherwise, we can apply Lemma 4.10 and have a guarantee of decrease in the case of a successful iteration, namely,

$$\|r(x_{j+1})\|^{1/2^i} - \|r(x_j)\|^{1/2^i} \leq -\frac{C_2}{2^{i+1}} \frac{1}{\mu_j^2 \|r(x_j)\|^{\frac{2^{i+1}-1}{2^i}}},$$

hence

$$\begin{aligned}
\varphi_{j+1} - \varphi_j &\leq \left[-\tau \frac{C_2}{2^{i+1}} + (1-\tau)(\lambda^2 - 1) \right] \frac{1}{\mu_j^2 \|r(x_j)\|^{\frac{2^{i+1}-1}{2^i}}} \\
&\leq (1-\tau) \left(\frac{1}{\lambda^2} - 1 \right) \frac{1}{\mu_j^2 \|r(x_j)\|^{\frac{2^{i+1}-1}{2^i}}},
\end{aligned}$$

where we used (4.22) to obtain the last inequality. Thus (A.2) also holds if the iteration is successful.

4. **Both m_j and (r_j^0, r_j^s) are inaccurate.** If the iteration is unsuccessful, then (A.2) holds. Suppose now that the iteration is successful: by considering a Taylor expansion of $f(x_j + s_j)$ around x_j , we know that the possible increase in the step is bounded

above by:

$$\begin{aligned}
f(x_j + s_j) - f(x_j) &\leq \nabla f(x_j)^\top s_j + \frac{\nu}{2} \|s_j\|^2 \\
&= \left[J(x_j)^\top r(x_j) \right]^\top s_j + \frac{\nu}{2} \|s_j\|^2 \\
&\leq \|J(x_j)^\top r(x_j)\| \|s_j\| + \frac{\nu}{2} \|s_j\|^2 \\
&\leq \|J(x_j)^\top r(x_j)\| \|s_j\| + \nu \frac{\|s_j\|}{\mu_j} \\
&\leq \left(1 + \frac{\nu}{\zeta} \right) \|J(x_j)^\top r(x_j)\| \|s_j\| \\
&\leq 2 \left(1 + \frac{\nu}{\zeta} \right) \frac{\|J(x_j)^\top r(x_j)\|}{\mu_j}.
\end{aligned}$$

We thus have

$$\|r(x_j + s_j)\|^2 - \|r(x_j)\|^2 \leq 4 \left(1 + \frac{\nu}{\zeta} \right) \frac{\|J(x_j)^\top r(x_j)\|}{\mu_j}.$$

Suppose that $\|r(x_j + s_j)\| > \|r(x_j)\|$. Applying Lemma 4.1 to the above equation yields:

$$\|r(x_j + s_j)\|^{1/2^i} - \|r(x_j)\|^{1/2^i} \leq \frac{4(1 + \frac{\nu}{\zeta})}{\mu_j} \frac{\|J(x_j)^\top r(x_j)\|}{\|r(x_j + s_j)\|^{2^{i+1}-1}} \leq \frac{4(1 + \frac{\nu}{\zeta})}{\mu_j} \frac{\|J(x_j)^\top r(x_j)\|}{\|r(x_j)\|^{2^{i+1}-1}},$$

and note that this relation still holds when $\|r(x_j + s_j)\| \leq \|r(x_j)\|$ as the left-hand side is negative in that case. We thus obtain the following bound on the change in ϕ :

$$(A.4) \quad \phi_{j+1} - \phi_j \leq \tau \frac{C_3}{2^{i+1}} \frac{\|J(x_j)^\top r(x_j)\|}{\|r(x_j)\|^{2^{i+1}-1}} \frac{1}{\mu_j} + (1 - \tau)(\lambda^2 - 1) \frac{1}{\mu_j^2},$$

with $C_3 = 4 \left(1 + \frac{\nu}{\zeta} \right)$. The right-hand side of (A.4) is nonnegative and larger than that of (A.2), thus the bound (A.4) holds when the iteration is successful and when it is unsuccessful.

Summarizing the four cases, we have that the bound (A.3) on $\Phi_{j+1} - \Phi_j$ holds in case 1-1, the bound (A.2) holds in both cases 1-2 and 1-3, and the bound (A.4) holds in case 1-4. Putting

those together with their associated probability of occurrence, we obtain:

$$\begin{aligned}
& \mathbb{E} \left[\Phi_{j+1} - \Phi_j | \mathcal{F}_{j-1}^{M,R} \cap E_j^0 \cap \left\{ \|J(X_j)^\top r(X_j)\| \geq \frac{\zeta}{\Gamma_j} \right\} \right] \\
& \leq pq \left[-\tau \frac{C_1}{2^{i+1}} \frac{\|J(X_j)^\top r(X_j)\|}{\|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}} \frac{1}{\Gamma_j} + (1-\tau)(\lambda^2-1) \frac{1}{\Gamma_j^2 \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}} \right] \\
& \quad + [p(1-q) + (1-p)q] \left[(1-\tau) \left(\frac{1}{\lambda^2} - 1 \right) \frac{1}{\Gamma_j^2} \right] \\
& \quad + (1-p)(1-q) \left[\tau \frac{C_3}{2^{i+1}} \frac{\|J(X_j)^\top r(X_j)\|}{\|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}} \frac{1}{\Gamma_j} + (1-\tau)(\lambda^2-1) \frac{1}{\Gamma_j^2 \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}} \right] \\
& = \frac{-C_1 pq + (1-p)(1-q)C_3}{2^{i+1}} \tau \frac{\|J(X_j)^\top r(X_j)\|}{\|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}} \frac{1}{\Gamma_j^2} \\
& \quad + [pq - \frac{1}{\lambda^2}(p(1-q) + (1-p)q) + (1-p)(1-q)] (1-\tau)(\lambda^2-1) \frac{1}{\Gamma_j^2 \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}} \\
& \leq \frac{-C_1 pq + (1-p)(1-q)C_3}{2^{i+1}} \tau \frac{\|J(X_j)^\top r(X_j)\|}{\|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}} \frac{1}{\Gamma_j^2} + (1-\tau)(\lambda^2-1) \frac{1}{\Gamma_j^2 \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}},
\end{aligned}$$

where the last line uses

$$pq - \frac{1}{\lambda^2}(p(1-q) + (1-p)q) + (1-p)(1-q) \leq (p + (1-p))(q + (1-q)) = 1.$$

Suppose p and q are chosen such that

$$(A.5) \quad \frac{pq - 1/2}{(1-p)(1-q)} \geq \frac{C_3}{C_1},$$

holds. Then, one has by combining (A.5) and (4.22):

$$(A.6) \quad \frac{-C_1 pq + (1-p)(1-q)C_3}{2^{i+1}} \leq -\frac{1}{2^{i+2}} C_1 \leq -2 \frac{(1-\tau)(\lambda^2-1)}{\tau \zeta}.$$

On the other hand, since $\|J(X_j)^\top r(X_j)\| \geq \zeta/\Gamma_j$, we have:

$$(1-\tau)(\lambda^2-1) \frac{1}{\Gamma_j^2 \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}} \leq -\frac{1}{2} [-C_1 pq + (1-p)(1-q)C_3] \tau \frac{\|J(X_j)^\top r(X_j)\|}{\|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}} \frac{1}{\Gamma_j}.$$

This leads to

$$\begin{aligned}
\mathbb{E} \left[\Phi_{j+1} - \Phi_j | \mathcal{F}_{j-1}^{M,R} \cap E_j^0 \cap \left\{ \|J(X_j)^\top r(X_j)\| \geq \frac{\zeta}{\Gamma_j} \right\} \right] & \leq -\frac{1}{4} C_1 \tau \frac{\|J(X_j)^\top r(X_j)\|}{\|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}} \frac{1}{\Gamma_j} \\
& \leq -\frac{1}{4} C_1 \tau \zeta \frac{1}{\Gamma_j^2 \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}},
\end{aligned}$$

which, using (A.6), finally gives:

$$(A.7) \quad \mathbb{E} \left[\Phi_{j+1} - \Phi_j | \mathcal{F}_{j-1}^{M,R} \cap E_j^0 \cap \left\{ \|J(X_j)^\top r(X_j)\| \geq \frac{\zeta}{\Gamma_j} \right\} \right] \leq -\frac{(1-\tau)(\lambda^2-1)}{4} \frac{1}{\Gamma_j^2 \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}} \\ \leq -\frac{(1-\tau)(1-\frac{1}{\lambda^2})}{4} \frac{1}{\Gamma_j^2 \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}}.$$

Case 2: $\|J(x_j)^\top r(x_j)\| < \frac{\zeta}{\mu_j}$.

Whenever $\|J_{m_j}^\top r_{m_j}\| < \frac{\eta_2}{\mu_j}$, the iteration is necessarily unsuccessful and (A.2) holds. We thus assume in what follows that $\|J_{m_j}^\top r_{m_j}\| \geq \frac{\eta_2}{\mu_j}$, and consider again four cases.

1. **Both m_j and (r_j^0, r_j^s) are accurate.** Unlike case 1-a), it is now possible for the iteration to be unsuccessful: in that case, we have (A.2). Otherwise, if the iteration is successful, then we can use the result from Lemma 4.10, and we have:

$$\|r(x_{j+1})\|^{1/2^i} - \|r(x_j)\|^{1/2^i} \leq -\frac{C_2}{2^{i+1}} \frac{1}{\mu_j^2 \|r(x_j)\|^{\frac{2^{i+1}-1}{2^i}}}.$$

We can thus apply the same reasoning than in case 1-3, which implies that (A.2) also holds when the iteration is successful.

2. **Only m_j is accurate.** If the iteration is unsuccessful, it is clear that (A.2) holds. Otherwise, using $\eta_2 \geq \kappa_{J_m}^2$ that arises from (4.18) and applying the same argument as in the proof of Lemma 4.10, we have $\|J_{m_j}^\top r_{m_j}\| \geq \frac{\eta_2}{\mu_j} \geq \frac{\kappa_{J_m}^2}{\mu_j}$, leading to

$$m_j(x_j) - m_j(x_j + s_j) \geq \frac{\theta_{fcd}}{2} \frac{\|J_{m_j}^\top r_{m_j}\|^2}{\kappa_{J_m}^2 + \mu_j \|J_{m_j}^\top r_{m_j}\|} \\ \geq \frac{\theta_{fcd}}{2} \frac{\|J_{m_j}^\top r_{m_j}\|^2}{2\mu_j \|J_{m_j}^\top r_{m_j}\|} = \frac{\theta_{fcd}}{4} \frac{\|J_{m_j}^\top r_{m_j}\|}{\mu_j} \geq \frac{\eta_2 \theta_{fcd}}{4} \frac{1}{\mu_j^2}.$$

Since the model is $(\kappa_{ef}, \kappa_{eg})$ -first-order accurate, the function variation satisfies:

$$f(x_j) - f(x_j + s_j) = f(x_j) - m_j(x_j) + m_j(x_j) - m_j(x_j + s_j) + m_j(x_j + s_j) - f(x_j + s_j) \\ \geq -\frac{\kappa_{ef}}{\mu_j^2} + \frac{\eta_2 \theta_{fcd}}{4} \frac{1}{\mu_j^2} - \frac{\kappa_{efs}}{\mu_j^2} \\ \geq \left(\frac{\eta_2 \theta_{fcd}}{4} - (\kappa_{ef} + \kappa_{efs}) \right) \frac{1}{\mu_j^2} \geq \frac{\kappa_{ef} + \kappa_{efs}}{2\mu_j^2},$$

where the last line comes from (4.18).

As a result, we have $\|r(x_j + s_j)\|^2 - \|r(x_j)\|^2 \leq -\frac{\kappa_{ef} + \kappa_{efs}}{\mu_j^2}$. Applying Lemma 4.1 then gives:

$$\|r(x_j + s_j)\|^{1/2^i} - \|r(x_j)\|^{1/2^i} \leq -\frac{\kappa_{ef} + \kappa_{efs}}{2^{i+1}} \frac{1}{\mu_j^2 \|r(x_j)\|^{\frac{2^{i+1}-1}{2^i}}},$$

and this leads to

$$(A.8) \quad \phi_{j+1} - \phi_j \leq \left[-\tau \frac{\kappa_{ef} + \kappa_{efs}}{2^{i+1}} \frac{1}{\mu_j^2 \|r(x_j)\|^{\frac{2^{i+1}-1}{2^i}}} + (1-\tau)(\lambda^2 - 1) \right] \frac{1}{\mu_j^2}$$

$$(A.8) \quad \phi_{j+1} - \phi_j \leq -(1-\tau) \left(1 - \frac{1}{\lambda^2}\right) \frac{1}{\mu_j^2}$$

by (4.22).

3. **Only (r_j^0, r_j^s) is accurate.** This case can be analyzed the same way as Case 2.1 to show that (A.2) holds regardless of whether the iteration is successful or unsuccessful.
4. **Both m_j and (r_j^0, r_j^s) are inaccurate.** As in Case 1.4, we have

$$f(x_j + s_j) - f(x_j) \leq \|J(x_j)^\top r(x_j)\| \|s_j\| + \frac{\nu}{2} \|s_j\|^2.$$

Using $\|J(x_j)^\top r(x_j)\| < \frac{\zeta}{\mu_j}$ and (4.5), we obtain:

$$\begin{aligned} f(x_j + s_j) - f(x_j) &\leq \zeta \frac{\|s_j\|}{\mu_j} + \frac{\nu}{2} \|s_j\|^2 \\ &\leq 2(\zeta + \nu) \frac{1}{\mu_j^2}. \end{aligned}$$

Then, applying Lemma 4.1 as in Case 1.4 yields

$$\|r(x_j + s_j)\|^{1/2^i} - \|r(x_j)\|^{1/2^i} \leq \frac{4(\zeta + \nu) \|J(x_j)^\top r(x_j)\|}{\mu_j \|r(x_j)\|^{\frac{2^{i+1}-1}{2^i}}},$$

thus

$$\|r(x_j + s_j)\|^{1/2^i} - \|r(x_j)\|^{1/2^i} \leq \frac{4\zeta \left(1 + \frac{\nu}{\zeta}\right)}{\mu_j^2 \|r(x_j)\|^{\frac{2^{i+1}-1}{2^i}}} = \zeta C_3 \frac{1}{\mu_j^2 \|r(x_j)\|^{\frac{2^{i+1}-1}{2^i}}}.$$

and

$$(A.9) \quad \phi_{j+1} - \phi_j \leq [\tau C_3 \zeta + (1-\tau)(\lambda^2 - 1)] \frac{1}{\mu_j^2 \|r(x_j)\|^{\frac{2^{i+1}-1}{2^i}}}.$$

Combining all the subcases for Case 2, we can bound all of those by (A.2) save for Case 2.4,

for which (A.9) applies. Thus, we obtain:

$$\begin{aligned}
& \mathbb{E} \left[\Phi_{j+1} - \Phi_j | \mathcal{F}_{j-1}^{M,R} \cap E_j^0 \cap \{ \|\nabla f(X_j)\| < \frac{\zeta}{\Gamma_j} \} \right] \\
& \leq [pq + p(1-q) + q(1-p)](1-\tau) \left(\frac{1}{\lambda^2} - 1 \right) \frac{1}{\Gamma_j^2 \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}} \\
& \quad + (1-p)(1-q) [\tau C_3 \zeta + (1-\tau)(\lambda^2 - 1)] \frac{1}{\Gamma_j^2 \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}} \\
& \leq -pq(1-\tau) \left(1 - \frac{1}{\lambda^2} \right) \frac{1}{\Gamma_j^2 \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}} \\
& \quad + (1-p)(1-q) [\tau C_3 \zeta + (1-\tau)(\lambda^2 - 1)] \frac{1}{\Gamma_j^2 \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}}.
\end{aligned}$$

We now assume that p and q have been chosen such that $pq \geq \frac{1}{2}$ and

$$(A.10) \quad (1-p)(1-q) \leq \frac{1}{4} \frac{(1-\tau) \left(1 - \frac{1}{\lambda^2} \right)}{\tau C_3 \zeta + (1-\tau)(\lambda^2 - 1)}$$

holds. Using (A.10), we obtain

$$(A.11) \mathbb{E} \left[\Phi_{j+1} - \Phi_j | \mathcal{F}_{j-1}^{M,R} \cap E_j^0 \cap \{ \|\nabla f(X_j)\| < \frac{\zeta}{\Gamma_j} \} \right] \leq -\frac{1}{4}(1-\tau) \left(1 - \frac{1}{\lambda^2} \right) \frac{1}{\Gamma_j^2 \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}},$$

which is the same amount of decrease as in (A.7). Letting $\sigma := \frac{1}{4}(1-\tau) \left(1 - \frac{1}{\lambda^2} \right)$, we have then established that for every iteration j ,

$$\mathbb{E} \left[\Phi_{j+1} - \Phi_j | \mathcal{F}_{j-1}^{M,R} \cap E_j^0 \right] < -\frac{\sigma}{\Gamma_j^2 \|r(X_j)\|^{\frac{2^{i+1}-1}{2^i}}}.$$

As a result, the statement of the theorem holds. ■