



Convergence and Complexity Analysis of a Levenberg–Marquardt Algorithm for Inverse Problems

Item Type	Article
Authors	Bergou, El Houcine;Diouane, Youssef;Kungurtsev, Vyacheslav
Citation	Bergou, E. H., Diouane, Y., & Kungurtsev, V. (2020). Convergence and Complexity Analysis of a Levenberg–Marquardt Algorithm for Inverse Problems. <i>Journal of Optimization Theory and Applications</i> . doi:10.1007/s10957-020-01666-1
Eprint version	Post-print
DOI	10.1007/s10957-020-01666-1
Publisher	Springer Nature
Journal	<i>Journal of Optimization Theory and Applications</i>
Rights	Archived with thanks to <i>Journal of Optimization Theory and Applications</i>
Download date	2024-04-21 05:36:59
Link to Item	http://hdl.handle.net/10754/662924

Convergence and Complexity Analysis of a Levenberg-Marquardt Algorithm for Inverse Problems

E. Bergou ^{*} Y. Diouane[†] V. Kungurtsev [‡]

May 31, 2020

Abstract

The Levenberg-Marquardt algorithm is one of the most popular algorithms for finding the solution of nonlinear least squares problems. Across different modified variations of the basic procedure, the algorithm enjoys global convergence, a competitive worst case iteration complexity rate, and a guaranteed rate of local convergence for both zero and nonzero small residual problems, under suitable assumptions. We introduce a novel Levenberg-Marquardt method that matches, simultaneously, the state of the art in all of these convergence properties with a single seamless algorithm. Numerical experiments confirm the theoretical behavior of our proposed algorithm.

Keywords: Inverse problems; Levenberg-Marquardt method; worst-case complexity bound; global and local convergence.

1 Introduction

In this paper, we consider solving general nonlinear least squares problems where one may or may not have a solution with a zero residual. Problems of this nature arise in several important practical contexts, including inverse problems for ill-posed nonlinear continuous systems [1] with applications such as data assimilation [2]. Usually the resulting least squares problems do not necessarily have a zero residual at any point, although the minimum residual may be small.

Recall that the Gauss-Newton method is an iterative procedure for solving nonlinear least squares problems by iteratively solving a linearized least squares subproblem. This subproblem may not be well-posed in the case of rank deficiency of the residual Jacobian function. Furthermore, the Gauss-Newton method may not be globally convergent. The Levenberg-Marquardt (LM) method [3, 4, 5] was developed to deal with the rank deficiency of the Jacobian matrix and also to provide a globalization strategy for Gauss-Newton. In this paper, we will present and analyze the global and local convergence results of a novel LM method for solving general nonlinear least squares problems, that carefully balances the opposing objectives of ensuring global convergence and stabilizing a fast local convergence regime.

^{*}MaIAGE, INRAE, Université Paris-Saclay, 78350 Jouy-en-Josas, France. KAUST, Thuwal, Saudi Arabia. (elhoucine.bergou@inrae.fr).

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

[‡]Department of Computer Science, Faculty of Electrical Engineering, Czech Technical University in Prague (vyacheslav.kungurtsev@fel.cvut.cz).

In general, the goals of encouraging global and local convergence compete against each other. Namely, the regularization parameter appearing in the subproblem should be allowed to become arbitrarily large in order to encourage global convergence, by ensuring the local accuracy of the linearized subproblem, but the parameter must approach zero in order to function as a stabilizing regularization that encourages fast local convergence. In the original presentation of the LM method in [3, 4], the regularization parameter is not permitted to go to zero, and only global convergence is considered.

The strongest results for local convergence of LM are given in a series of papers beginning with [6] (followed by, e.g., [7] and [8]; see also [9]), wherein it is assumed that the residual function is zero at the solution. For the global convergence, the algorithm considered is a two-phase one, where quadratic decline in the residual is tested with each step that is otherwise globalized by a line-search procedure.

In the case of nonzero residuals, it has been found that a LM method converges locally at a linear rate, if the norm of the residual is sufficiently small and the regularization parameter goes to zero [10]. Our proof of linear convergence is simpler than in [10]. Worst-case iteration complexity bounds for LM methods applied to nonlinear least squares problems can be found in [13, 14, 15]. We show that our proposed algorithm has a complexity bound that matches these results, up to a logarithmic factor.

In this paper, we propose a method that successfully balances the multiple objectives in theoretical convergence properties, including (a) global convergence for exact and inexact solutions of the subproblem, (b) worst case iteration complexity, and (c) local convergence for both zero and nonzero residual problems. Table 1 summarizes the literature on this class of methods; our proposed algorithm uniquely matches the state of the art in all of these properties.

Table 1: Convergence Properties of Levenberg-Marquardt Algorithms.

	[3, 4, 5]	[13, 14, 15]	[6, 7, 9]	[10]	This work
Global convergence	Yes	Yes	Two-phase or No	Yes	Yes
Complexity analysis	No	Yes	No	No	Yes
Local zero residual	No	No	Quadratic	Superlinear	Quadratic
Local nonzero small residual	No	No	No	Linear	Linear

The outline of this paper is as follows. In Section 2, we present the proposed LM algorithm and address the inexact solution of the linearized least squares subproblem. Section 3 contains a worst-case complexity and global convergence analysis of the proposed method. In Section 4, we derive the local convergence theory. In Section 5, preliminary numerical experiments are presented to demonstrate the behavior of our algorithm. Finally, we conclude in Section 6.

2 A Novel Levenberg-Marquardt Algorithm

In this paper, we consider the following nonlinear least squares problem

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

where $F : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a (deterministic) vector-valued function, assumed twice continuously differentiable. Here and in the rest of the text, $\|\cdot\|$ denotes the vector or matrix l_2 -norm. At

each iteration j , the LM method computes (approximately) a step s_j of the form $-(J_j^\top J_j + \gamma_j I)^{-1} J_j^\top F_j$, corresponding to the unique solution of

$$\min_{s \in \mathbb{R}^n} m_j(s) := \frac{1}{2} \|F_j + J_j s\|^2 + \frac{1}{2} \gamma_j \|s\|^2, \quad (2)$$

where $\gamma_j > 0$ is an appropriately chosen regularization parameter, $F_j := F(x_j)$ and $J_j := J(x_j)$ denotes the Jacobian of F at x_j .

In deciding whether to accept a step s_j generated by the subproblem (2), the LM method can be seen as precursor of the trust-region method [12]. In fact, it seeks to determine when the Gauss-Newton step is applicable or when it should be replaced by a slower but safer steepest descent step. One considers the ratio ρ_j between the actual reduction $f(x_j) - f(x_j + s_j)$ attained in the objective function and the reduction $m_j(0) - m_j(s_j)$ predicted by the model. Then, if ρ_j is sufficiently greater than zero, the step is accepted and γ_j is possibly decreased. Otherwise the step is rejected and γ_j is increased.

In this paper, we use the standard choice of the regularization parameter

$$\gamma_j := \mu_j \|F(x_j)\|^2,$$

where μ_j is updated according to the ratio ρ_j . The considered LM algorithm using the above update strategy, as described in Algorithm 1, will be shown to be globally convergent with a complexity bound of order ϵ^{-2} and have strong local convergence properties.

Algorithm 1: A Levenberg-Marquardt algorithm.

Initialization

Choose the constants $\eta \in]0, 1[$, $\mu_{\min} > 0$ and $\lambda > 1$. Select x_0 and $\mu_0 \geq \mu_{\min}$. Set $\gamma_0 = \mu_0 \|F(x_0)\|^2$ and $\bar{\mu} = \mu_0$.

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

1. Solve (or approximately solve) (2), and let s_j denote such a solution.
 2. Compute $\rho_j := \frac{f(x_j) - f(x_j + s_j)}{m_j(0) - m_j(s_j)}$.
 3. If $\rho_j \geq \eta$, then set $x_{j+1} = x_j + s_j$ and $\mu_{j+1} \in [\max\{\mu_{\min}, \bar{\mu}/\lambda\}, \bar{\mu}]$ and $\bar{\mu} = \mu_j$.
Otherwise, set $x_{j+1} = x_j$ and $\mu_{j+1} = \lambda \mu_j$.
 4. Compute $\gamma_{j+1} = \mu_{j+1} \|F(x_{j+1})\|^2$.
-

This Algorithm has one particularly novel feature among LM methods: we have an auxiliary parameter $\bar{\mu}$ which represents the last parameter corresponding to a successful step, introduced to balance the requirements of global and local convergence. If the model is inaccurate, then μ_j is driven higher, however, when we reach a region associated with the local convergence regime, the residual $\|F(x_j)\|$ should ultimately dominate the behavior of γ_j for successful steps. Step 1 of Algorithm 1 requires the approximate solution of subproblem (2). As in trust-region methods, there are different techniques to approximate the solution of this subproblem that yield a globally convergent step. For that it suffices to compute a step s_j that provides a reduction

in the model at least as good as the one produced by the so-called Cauchy step (defined as the minimizer of the model along the negative gradient) which is given by

$$s_j^c := -\frac{\|\nabla f(x_j)\|^2}{\nabla f(x_j)^\top (J_j^\top J_j + \gamma_j I) \nabla f(x_j)} \nabla f(x_j).$$

The Cauchy step is cheap to calculate as it does not require any system solve. Moreover, the LM method will be globally convergent if it uses a step that attains a reduction in the model as good as a multiple of the Cauchy decrease. Thus we will impose the following assumption on the step calculation:

Assumption 2.1 *There exists $\theta_{fcd} > 0$ such that for every iteration j :*

$$m_j(0) - m_j(s_j) \geq \frac{\theta_{fcd} \|\nabla f(x_j)\|^2}{2 \|J_j\|^2 + \gamma_j}.$$

Despite providing a sufficient reduction in the model and being cheap to compute, the Cauchy step is scaled steepest descent. In practice, a version of Algorithm 1 based solely on the Cauchy step would suffer from the same drawbacks as the steepest descent algorithm on ill-conditioned problems.

Since the Cauchy step is the first step of the conjugate gradient method (CG) when applied to the minimization of the quadratic $s \rightarrow m_j(s)$, it is natural to propose running CG further and stopping only when the residual becomes sufficiently small. Since the CG generates iterates by minimizing the quadratic model over nested Krylov subspaces, and the first subspace is the one generated by $\nabla f(x_j)$ (see, e.g., [16, Theorem 5.2]), the decrease obtained at the first CG iteration (i.e., by the Cauchy step) is at least attained by the remaining iterations. Thus Assumption 2.1 holds for all the iterates s_j^{CG} generated by the truncated-CG whenever it is initialized by the null vector. The following lemma is similar to [17, Lemma 5.1] and will be useful for our worst-case complexity analysis.

Lemma 2.1 *For the three steps proposed (exact, Cauchy, and truncated-CG), one has that*

$$\|s_j\| \leq \frac{\|\nabla f(x_j)\|}{\gamma_j} \quad \text{and} \quad |s_j^\top (\gamma_j s_j + \nabla f(x_j))| \leq \frac{\|J_j\|^2 \|\nabla f(x_j)\|^2}{\gamma_j^2}.$$

In what comes next, we will call all iterations j for which $\rho_j \geq \eta$ *successful*, and we denote the set of their indices by the symbol \mathcal{S} , i.e., $\mathcal{S} := \{j \in \mathbb{N} \mid \rho_j \geq \eta\}$.

3 Worst-Case Iteration Complexity and Global Convergence

We now establish a worst-case complexity bound of Algorithm 1. Namely, given a tolerance $\epsilon \in]0, 1[$, we aim at deriving the number of iterations, in the worst case, needed to reach an iterate x_j such that

$$\|\nabla f(x_j)\| < \epsilon \quad \text{or} \quad \|F(x_j)\| < \max\{\epsilon, (1 + \epsilon)\|F(\bar{x}_j)\|\} \tag{3}$$

where $\bar{x}_j \in \operatorname{argmin}_{\{x \in \mathbb{R}^n : \nabla f(x) = 0\}} \|x_j - x\|$. Without loss of generality, we will assume that $\|F(\bar{x}_j)\|$ is unique and independent from x_j , we will denote it by \bar{f} . It can be seen that, if we drop this assumption, then the same arguments in this section show asymptotic global convergence. Then, $\bar{f} := \sqrt{2f(\bar{x}_j)} = \|F(\bar{x}_j)\|$ can just be taken to be the value at the limit point of the sequence. We start now by giving some classical assumptions.

Assumption 3.1 *The function f is continuously differentiable in an open set containing $L(x_0) := \{x \in \mathbb{R}^n : f(x) \leq f(x_0)\}$ with Lipschitz continuous gradient on $L(x_0)$ with the constant $\nu > 0$.*

Assumption 3.2 *The Jacobian J of F is uniformly bounded, i.e., there exists $\kappa_J > 0$ such that $\|J\| \leq \kappa_J$ for all x .*

We begin by obtaining a condition on the parameter μ_j that is sufficient for an iteration to be successful. We omit the proof as it is standard, see for instance Lemma 5.2 in [17].

Lemma 3.1 *Let Assumptions 2.1, 3.1, and 3.2 hold. Suppose that at the j -th iteration of Algorithm 1, one has*

$$\mu_j > \frac{\kappa}{\|F(x_j)\|^2} \quad (4)$$

where

$$\kappa := \frac{a + \sqrt{a^2 + 4a\kappa_J^2(1-\eta)}}{2(1-\eta)} \quad \text{and} \quad a := \frac{\frac{\nu}{2} + 2\kappa_J^2}{\theta_{fcd}}.$$

Then, the iteration is successful.

Our next result states that, when the gradient norm stays bounded away from zero, the parameter μ_j cannot grow indefinitely. Without loss of generality, we assume that $\epsilon \leq \sqrt{\frac{\lambda\kappa}{\mu_0}}$, where κ is the same as in the previous lemma.

Lemma 3.2 *Under Assumptions 2.1, 3.1, and 3.2, let j be a given iteration index such that for every $l \leq j$ it holds that $\|F(x_l)\| \geq \max\{\epsilon, (1+\epsilon)\bar{f}\}$ where $\epsilon \in]0, 1[$. Then, for every $l \leq j$, one also has*

$$\mu_l \leq \mu_{\max} := \frac{\lambda\kappa}{\max\{\epsilon^2, (1+\epsilon)^2\bar{f}^2\}}.$$

Proof. We prove this result by contradiction. Suppose that $l \geq 1$ is the first index such that

$$\mu_l > \frac{\lambda\kappa}{\max\{\epsilon^2, (1+\epsilon)^2\bar{f}^2\}}. \quad (5)$$

By the updating rules on μ_l , either the iteration $l-1$ is successful, in which case $\mu_l \leq \mu_{l-1} \leq \frac{\lambda\kappa}{\max\{\epsilon^2, (1+\epsilon)^2\bar{f}^2\}}$ which contradicts (5), or the iteration $l-1$ is unsuccessful and thus

$$\mu_l = \lambda\mu_{l-1} \quad \Rightarrow \quad \mu_{l-1} = \frac{\mu_l}{\lambda} > \frac{\kappa}{\max\{\epsilon^2, (1+\epsilon)^2\bar{f}^2\}} > \frac{\kappa}{\|F(x_l)\|^2},$$

therefore using Lemma 3.1 this implies that the $(l-1)$ -th iteration is successful which leads to a contradiction again.

■ Thanks to Lemma 3.2, we can now bound the number of successful iterations needed to drive the gradient norm below a given threshold.

Proposition 3.1 *Under Assumptions 2.1, 3.1, and 3.2. Let $\epsilon \in]0, 1[$ and j_ϵ be the first iteration index such that (3) holds. Then, if \mathcal{S}_ϵ is the set of indexes of successful iterations prior to j_ϵ , one has:*

$$|\mathcal{S}_\epsilon| \leq C\epsilon^{-2} \quad \text{with} \quad C := \frac{2(\kappa_J^2 + \mu_{\max}\|F(x_0)\|^2)}{\eta\theta_{fcd}} f(x_0).$$

Proof. For any $j \in \mathcal{S}_\epsilon$, one has

$$f(x_j) - f(x_{j+1}) \geq \eta(m_j(0) - m_j(s_j)) \geq \eta \frac{\theta_{fcd}}{2} \frac{\|\nabla f(x_j)\|^2}{\kappa_j^2 + \mu_j \|F(x_j)\|^2}.$$

Hence, using the fact that $\|\nabla f(x_j)\| \geq \epsilon$, $\|F(x_j)\| \leq \|F(x_0)\|$ and $\mu_j \leq \mu_{\max}$, we arrive at

$$f(x_j) - f(x_{j+1}) \geq \eta \frac{\theta_{fcd}}{2} \frac{\epsilon^2}{\kappa_j^2 + \mu_{\max} \|F(x_0)\|^2}.$$

Consequently, by summing on all iteration indices within \mathcal{S}_ϵ , we obtain

$$f(x_0) - 0 \geq \sum_{j \in \mathcal{S}_\epsilon} f(x_j) - f(x_{j+1}) \geq |\mathcal{S}_\epsilon| \frac{\eta \theta_{fcd}}{2 (\kappa_j^2 + \mu_{\max} \|F(x_0)\|^2)} \epsilon^2,$$

hence the result. \blacksquare

Lemma 3.3 *Under Assumptions 2.1, 3.1, and 3.2. Let \mathcal{U}_ϵ denote the set of unsuccessful iterations of index less than or equal to j_ϵ . Then,*

$$|\mathcal{U}_\epsilon| \leq \log_\lambda \left(\frac{\kappa}{\mu_{\min} \epsilon^2} \right) |\mathcal{S}_\epsilon|. \quad (6)$$

Proof. Note that we necessarily have $j_\epsilon \in \mathcal{S}_\epsilon$ (otherwise it would contradict the definition of j_ϵ). Our objective is to bound the number of unsuccessful iterations between two successful ones. Let thus $\{j_0, \dots, j_t = j_\epsilon\}$ be an ordering of \mathcal{S}_ϵ , and $i \in \{0, \dots, t-1\}$.

Due to the updating formulas for μ_j on successful iterations, we have:

$$\mu_{j_{i+1}} \geq \max\{\mu_{\min}, \bar{\mu}/\lambda\} \geq \mu_{\min}.$$

Moreover, we have $\|F(x_{j_{i+1}})\| \geq \epsilon$ by assumption. By Lemma 3.1, for any unsuccessful iteration $j \in \{j_i + 1, \dots, j_{i+1} - 1\}$, we must have: $\mu_j \leq \frac{\kappa}{\epsilon^2}$, since otherwise $\mu_j > \frac{\kappa}{\epsilon^2} \geq \frac{\kappa}{\|F(x_j)\|^2}$ and the iteration would be successful.

Using the updating rules for μ_j on unsuccessful iterations, we obtain:

$$\forall j = j_i + 1, \dots, j_{i+1} - 1, \quad \mu_j = \lambda^{j-j_i-1} \mu_{j_{i+1}} \geq \lambda^{j-j_i-1} \mu_{\min}.$$

Therefore, the number of unsuccessful iterations between j_i and j_{i+1} , equal to $j_{i+1} - j_i - 1$, satisfies:

$$j_{i+1} - j_i - 1 \leq \log_\lambda \left(\frac{\kappa}{\mu_{\min} \epsilon^2} \right). \quad (7)$$

By considering (7) for $i = 0, \dots, t-1$, we arrive at

$$\sum_{i=0}^{t-1} (j_{i+1} - j_i - 1) \leq \log_\lambda \left(\frac{\kappa}{\mu_{\min} \epsilon^2} \right) [|\mathcal{S}_\epsilon| - 1]. \quad (8)$$

What is left to bound is the number of possible unsuccessful iterations between the iteration of index 0 and the first successful iteration j_0 . Since $\mu_0 \geq \mu_{\min}$, a similar reasoning as the one used to obtain (7) leads to

$$j_0 - 1 \leq \log_\lambda \left(\frac{\kappa}{\mu_{\min} \epsilon^2} \right). \quad (9)$$

Putting (8) and (9) together yields the expected result.

■ By combining the results from Proposition 3.1 and Lemma 3.3, we thus get the following complexity estimate.

Theorem 3.1 *Under Assumptions 2.1, 3.1, and 3.2. Let $\epsilon \in]0, 1[$. Then, the first index j_ϵ for which $\|\nabla f(x_{j_\epsilon+1})\| < \epsilon$ or $\|F(x_{j_\epsilon+1})\| < \max\{\epsilon, (1+\epsilon)\bar{f}\}$, is bounded above by*

$$\mathcal{C} \left(1 + \log_\lambda \left[\frac{\kappa}{\mu_{\min} \epsilon^2} \right] \right) \epsilon^{-2}, \quad (10)$$

where \mathcal{C} is the constant defined in Proposition 3.1.

For the LM method proposed in this paper, we thus obtain an iteration complexity bound of $\tilde{\mathcal{O}}(\epsilon^{-2})$, where the notation $\tilde{\mathcal{O}}(\cdot)$ indicates the presence of logarithmic factors in ϵ . Note that the evaluation complexity bounds are of the same order.

We note that by the definition of \bar{f} it holds that: $\|F(x_j)\| = \bar{f}$ implies that $\nabla f(x_j) = 0$. Thus, by letting $\epsilon \rightarrow 0$, Theorem 3.1 implies that

$$\liminf_{j \rightarrow \infty} \|\nabla f(x_j)\| = 0.$$

In order to derive the global convergence result, we need to extend this to a limit result.

Theorem 3.2 *Under Assumptions 2.1, 3.1, and 3.2, the sequence $\{x_j\}$ generated by Algorithm 1 satisfies*

$$\lim_{j \rightarrow \infty} \|\nabla f(x_j)\| = 0.$$

Proof. Consider the case that \mathcal{S} , the set of successful iterations, is finite. Then $\exists j_0$ such that for all $j \geq j_0$, $\|\nabla f(x_j)\| = \|\nabla f(x_{j_0})\|$ therefore from the previous theorem we conclude that in this case

$$\|\nabla f(x_{j_0})\| = \lim_{j \rightarrow \infty} \|\nabla f(x_j)\| = \liminf_{j \rightarrow \infty} \|\nabla f(x_j)\| = 0.$$

Alternatively, assume \mathcal{S} is infinite and let $\epsilon > 0$. Since $(f(x_j))$ is monotonically decreasing and bounded from below, one has $\lim_{j \rightarrow \infty} f(x_j) - f(x_{j+1}) = 0$. Since $\frac{\eta\theta_{fcd}\epsilon^2}{2(\kappa_j^2 + \mu_0\|F(x_0)\|^2)} > 0$, we conclude, for j sufficiently large, that

$$0 \leq f(x_j) - f(x_{j+1}) \leq \frac{\eta\theta_{fcd}\epsilon^2}{2(\kappa_j^2 + \mu_0\|F(x_0)\|^2)}.$$

Thus, for a sufficiently large j , consider the case that $j \in \mathcal{S}$, then $\rho_j \geq \eta$ and by rearranging the terms and using Assumption 2.1 we conclude that

$$\begin{aligned} \|\nabla f(x_j)\| &\leq \sqrt{\frac{2(m_j(0) - m_j(s_j))(\kappa_j^2 + \mu_0\|F(x_0)\|^2)}{\theta_{fcd}}} \\ &\leq \sqrt{\frac{2(\kappa_j^2 + \mu_0\|F(x_0)\|^2)(f(x_j) - f(x_{j+1}))}{\eta\theta_{fcd}}} \leq \epsilon. \end{aligned}$$

If $j \notin \mathcal{S}$ than $\|\nabla f(x_j)\| = \|\nabla f(x_{\hat{j}})\| \leq \epsilon$, where $\hat{j} \in \mathcal{S}$ is the last successful iteration before j . Hence, it must hold that $\lim_{j \rightarrow \infty} \|\nabla f(x_j)\| = 0$.

■

4 Local Convergence

In this section, we prove local convergence for the Algorithm, showing a quadratic rate for zero residual problems and explicit linear rate for nonzero residuals. Since the problem is generally nonconvex and with possibly nonzero residuals, there can be multiple sets of stationary points with varying objective values. We consider a particular subset with a constant value of the objective.

Assumption 4.1 *There exists a connected isolated set X^* composed of stationary points to (1) each with the same value f^* of $f(\cdot) = \frac{1}{2}\|F(\cdot)\|^2$ and Algorithm 1 generates a sequence with an accumulation point $x^* \in X^*$.*

We note, that under Assumption 4.1, the value $\|F(\bar{x})\|$ is unique for all $\bar{x} \in X^*$ which may not be the case for the residual vector $F(\bar{x})$. Thus we define $\bar{f} = \sqrt{2f(\bar{x})} = \|F(\bar{x})\|$ for $\bar{x} \in X^*$. Henceforth, from the global convergence analysis, we can assume, without loss of generality, that there exists a subsequence of iterates approaching this X^* . This subsequence does not need to be unique, i.e., there may be more than one subsequence converging to separate connected sets of stationary points. We shall see that eventually, one of these sets shall “catch” the subsequence and result in direct convergence to the set of stationary points at a quadratic or linear rate, depending on f .

In the sequel, $N(x, \delta)$ denotes the closed ball with center x (a given vector) and radius $\delta > 0$ and $\text{dist}(x, X^*)$ denotes the distance between the vector x and the set X^* , i.e., $\text{dist}(x, X^*) = \min_{y \in X^*} \|x - y\|$, and $\bar{x} \in \arg\min_{y \in X^*} \|x - y\|$.

Assumption 4.2 *It holds that $F(x)$ is twice continuously differentiable around $x^* \in X^*$ with x^* satisfying Assumption 4.1. In particular this implies, that there exists $\delta_1 > 0$ such that for all $x, y \in N(x^*, \delta_1)$,*

$$\|\nabla f(x)\|^2 = \|J(x)^\top F(x) - J(\bar{x})^\top F(\bar{x})\|^2 \leq L_1 \|x - \bar{x}\|^2, \quad (11)$$

$$\|F(x) - F(y)\| \leq L_2 \|x - y\|, \quad (12)$$

$$\|F(y) - F(x) - J(x)(y - x)\| \leq L_3 \|x - y\|^2, \quad (13)$$

where L_1, L_2 , and L_3 are positive constants.

From the triangle inequality and assuming (12), we get

$$\|F(x)\| - \|F(y)\| \leq \|F(x) - F(y)\| \leq L_2 \|x - y\|. \quad (14)$$

We introduce the following additional assumption.

Assumption 4.3 *There exist a $\delta_3 > 0$ and $M > 0$ such that*

$$\forall x \in N(x^*, \delta_3), \quad \text{dist}(x, X^*) \leq M \|F(x) - F(\bar{x})\|.$$

As the function $x \rightarrow F(x) - F(\bar{x})$ is zero residual, the proposed error bound assumption can be seen as a generalization of the zero residual case [6, 7, 8, 9]. Thus any ill-posed zero residual problem, as considered in this line of work on quadratic local convergence for LM methods, satisfies the assumptions. Our assumptions are also covered by a range of nonzero residual problems, for instance, standard data assimilation problems [2] as given by Example 4.1.

Example 4.1 Consider the following data assimilation problem $F : \mathbb{R}^n \rightarrow \mathbb{R}^m$ defined, for a given $x \in \mathbb{R}^n$, by $F(x) = ((x - x_b)^\top, (H(x) - y)^\top)^\top$, where $x_b \in \mathbb{R}^n$ is a background vector, $y \in \mathbb{R}^{m-n}$ the vector of observations and $H : \mathbb{R}^n \rightarrow \mathbb{R}^{m-n}$ is a smooth operator modeling the observations. For such problems, the set of stationary points $X^* = \{\{\bar{x}\}\} \subset \mathbb{R}^n$ is a finite disjoint set, $\bar{F} = F(\bar{x})$ and $\text{dist}(x, X^*) = \|x - \bar{x}\|$ for \bar{x} closest to x . Clearly, one has

$$\|F(x) - F(\bar{x})\|^2 = \|x - \bar{x}\|^2 + \|H(x) - H(\bar{x})\|^2 \geq \|x - \bar{x}\|^2 = \text{dist}(x, X^*)^2.$$

Thus, for these typical problems arising in data assimilation, Assumptions 4.2, 4.1 and 4.3 are satisfied.

Example 4.2 Consider $F : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ defined, for a given $x \in \mathbb{R}^3$, by

$$F(x) = (\exp(x_1 - x_2) - 1, x_3 - 1, x_3 + 1)^\top.$$

Clearly, F is Lipschitz smooth, hence Assumption 4.2 is satisfied. We note that for all $x, \bar{x} \in X^* = \{x \in \mathbb{R}^3 : x_1 = x_2 \text{ and } x_3 = 0\}$ satisfies $F(\bar{x}) = (0, -1, 1)^\top$ and, for a given x , one has $\text{dist}(x, X^*) = \sqrt{\frac{(x_1 - x_2)^2}{2} + x_3^2}$. Hence, for all $x \in \mathbb{R}^3$ such that $|x_1 - x_2| \leq \frac{1}{2}$, one concludes that

$$\|F(x) - F(\bar{x})\| = \sqrt{(\exp(x_1 - x_2) - 1)^2 + 2x_3^2} \geq \text{dist}(x, X^*).$$

Thus, for this example, Assumptions 4.1 and 4.3 are also satisfied.

From the global convergence results, we have established that there is a subsequence of successful iterations converging to the set of stationary points X^* . In this section, we begin by considering the subsequence of iterations that succeed the successful iterations, i.e., we consider the subsequence $\mathcal{K} = \{j + 1 : j \in \mathcal{S}\}$. We shall present the results with a slight abuse of notation that simplifies the presentation without sacrificing accuracy or generality: in particular every time we denote a quantity a_j , the index j corresponds to an element of this subsequence \mathcal{K} denoted above, thus when we say a particular statement holds eventually, this means that it holds for all $j \in \mathcal{S} + 1$ with j sufficiently large. Let $\hat{\mu}$ be an upper bound for μ_j , note that this exists by the formulation of Algorithm 1. We shall denote also δ as $\delta = \min(\delta_1, \delta_2, \delta_3)$, with $\{\delta_i\}_{i=1,2,3}$. In the proof, we follow the structure of the local convergence proof in [6], with the addition that the step is accepted by the globalization procedure. We use \bar{F}_j to denote $F(\bar{x}_j)$. The first lemma is similar to [6, Lemma 2.1].

Lemma 4.1 Suppose that Assumptions 4.1 and 4.2 are satisfied. If $x_j \in N(x^*, \frac{\delta}{2})$, then the solution s_j to (2) satisfies

$$\|J_j s_j + F_j\| - \bar{f} \leq C_1 \text{dist}(x_j, X^*)^2,$$

where C_1 is a positive constant independent of j .

Proof. Let us assume that $\bar{f} > 0$. Otherwise the proof is the same as in [6, Lemma 2.1]. Without loss of generality, since $f(x_j)$ is monotonically decreasing, we can consider that j is

sufficiently large such that $\|F_j\| \leq 2\bar{f}$. Hence, we get

$$\begin{aligned}
\|J_j s_j + F_j\|^2 &\leq 2m_j(\bar{x}_j - x_j) = \|J_j(\bar{x}_j - x_j) + F_j\|^2 + \mu_j \|F_j\|^2 \|x_j - \bar{x}_j\|^2 \\
&\leq (\|\bar{F}_j\| + \|J_j(\bar{x}_j - x_j) + F_j - \bar{F}_j\|)^2 + \mu_j \|F_j\|^2 \|x_j - \bar{x}_j\|^2 \\
&\leq (\|\bar{F}_j\| + L_3 \|x_j - \bar{x}_j\|^2)^2 + 4\mu_j \|\bar{F}_j\|^2 \|x_j - \bar{x}_j\|^2 \\
&= L_3^2 \|x_j - \bar{x}_j\|^4 + 2\|\bar{F}_j\| (2\mu_j \|\bar{F}_j\| + L_3) \|x_j - \bar{x}_j\|^2 + \|\bar{F}_j\|^2 \\
&\leq ((2\hat{\mu}\bar{f} + L_3) \|x_j - \bar{x}_j\|^2 + \bar{f})^2,
\end{aligned}$$

which concludes the proof. \blacksquare

Lemma 4.2 *Suppose that Assumptions 4.1 and 4.2 are satisfied.*

If $x_j \in N(x^, \frac{\delta}{2})$, then the solution s_j to (2) satisfies*

$$\|s_j\| \leq C_2 \text{dist}(x_j, X^*), \quad (15)$$

where C_2 is a positive constant independent of j .

Proof. We assume that $\|\bar{F}_j\| = \bar{f} > 0$, otherwise the proof is the same as in [6, Lemma 2.1]. In this case, using the fact that $\|F_j\| \geq \bar{f}$ and $\mu_j \geq \mu_{\min}$, one has

$$\|s_j\| \leq \frac{\|\nabla f(x_j)\|}{\gamma_j} = \frac{\|\nabla f(x_j)\|}{\mu_j \|F_j\|^2} \leq \frac{\sqrt{L_1}}{\mu_{\min} \bar{f}^2} \text{dist}(x_j, X^*),$$

which concludes the proof. \blacksquare

Lemma 4.3 *Suppose that Assumptions 4.1, 4.2 and 4.3 are satisfied. Consider the case where $\bar{f} = 0$, then for j sufficiently large, one has $\rho_j \geq \eta$.*

Proof. In fact, for j sufficiently large, using Lemmas 4.1 and 4.2, one gets

$$\begin{aligned}
2(m_j(0) - m_j(s_j)) &= \|F_j\|^2 - \|F_j + J_j s_j\|^2 - \gamma_j \|s_j\|^2 \\
&\geq (\|F_j\| + \|F_j + J_j s_j\|) \left(\frac{1}{M} \|x_j - \bar{x}_j\| - C_1 \|x_j - \bar{x}_j\|^2 \right) - C_2^2 \gamma_j \|x_j - \bar{x}_j\|^2 \\
&\geq \|F_j\| \left(\frac{1}{M} \|x_j - \bar{x}_j\| - C_1 \|x_j - \bar{x}_j\|^2 \right) - C_2^2 \mu_j \|F_j\|^2 \|x_j - \bar{x}_j\|^2 \\
&\geq \frac{1}{M} \|x_j - \bar{x}_j\| \left(\frac{1}{M} \|x_j - \bar{x}_j\| - C_1 \|x_j - \bar{x}_j\|^2 \right) - L_2^2 C_2^2 \hat{\mu} \|x_j - \bar{x}_j\|^4 \\
&= \frac{1}{M^2} \|x_j - \bar{x}_j\|^2 - \frac{C_1}{M} \|x_j - \bar{x}_j\|^3 - L_2^2 C_2^2 \hat{\mu} \|x_j - \bar{x}_j\|^4 > 0.
\end{aligned}$$

On the other hand, using the fact that $\|J_j\|$ is bounded (by $\kappa_J > 0$), Lemma 4.2 and (since $\bar{F}_j = 0$) $\frac{1}{M} \|\bar{x}_j - x_j\| \leq \|F_j\| \leq L_2 \|\bar{x}_j - x_j\|$, one gets

$$\begin{aligned}
2|f(x_j + s_j) - m_j(s_j)| &= \left| \|F(x_j + s_j)\|^2 - \|F_j + J_j s_j\|^2 - \gamma_j \|s_j\|^2 \right| \\
&= \left| (\|F(x_j + s_j)\| - \|F_j + J_j s_j\|) (\|F(x_j + s_j)\| + \|F_j + J_j s_j\|) - \gamma_j \|s_j\|^2 \right| \\
&\leq L_3 \|s_j\|^2 (\|F(x_j + s_j)\| + \|F_j\| + \|J_j\| \|s_j\|) + \gamma_j \|s_j\|^2 \\
&\leq L_3 \|s_j\|^2 (L_2 \|x_j + s_j - \bar{x}_j\| + L_2 \|x_j - \bar{x}_j\| + \|J_j\| \|s_j\|) + \gamma_j \|s_j\|^2 \\
&\leq (C_2 L_2 + 2L_2 + C_2 \kappa_J) C_2^2 L_3 \|x_j - \bar{x}_j\|^3 + L_2^2 C_2^2 \hat{\mu} \|x_j - \bar{x}_j\|^4.
\end{aligned}$$

Hence, for j sufficiently large

$$\begin{aligned}
|1 - \rho_j| &= \left| \frac{m_j(0) - f(x_j) + f(x_j + s_j) - m_j(s_j)}{m_j(0) - m_j(s_j)} \right| \\
&\leq \frac{(C_2 L_2 + 2L_2 + C_2 \kappa_J) C_2^2 L_3 \|x_j - \bar{x}_j\| + L_2^2 C_2^2 \hat{\mu} \|x_j - \bar{x}_j\|^2}{\frac{1}{M^2} - \frac{C_1}{M} \|x_j - \bar{x}_j\| - L_2^2 C_2^2 \hat{\mu} \|x_j - \bar{x}_j\|^2}.
\end{aligned}$$

Thus, $|1 - \rho_j| \rightarrow 0$ as j goes to $+\infty$.

■ For the nonzero residual case, we must consider a specific instance of the Algorithm. In particular, we specify Step 3 of Algorithm 1 to be,

$$\text{If } \rho_j \geq \eta, \text{ then set } x_{j+1} = x_j + s_j, \mu_{j+1} = \bar{\mu} \text{ and } \bar{\mu} = \mu_j.$$

Note that this is still consistent with the presentation of the Algorithm.

Lemma 4.4 *Suppose that Assumptions 4.1, 4.2 and 4.3 are satisfied. Consider that $\bar{f} > 0$, then for j sufficiently large, one has $\rho_j \geq \eta$.*

Proof. Indeed, by the new updating mechanism, the parameter $\bar{\mu}$ is monotonically nondecreasing. In particular, if there is an infinite set of unsuccessful steps, then $\mu_j \rightarrow \infty$. This implies that for some j_0 it holds that for $j \geq j_0$, $\mu_j > \frac{\kappa}{\bar{f}^2} > \frac{\kappa}{\|F(x_j)\|^2}$, which together with Lemma 3.1 reach a contradiction. Thus, there is a finite number of unsuccessful steps, and every step is accepted for j sufficiently large.

■

Proposition 4.1 *Suppose that Assumptions 4.1, 4.2, and 4.3 are satisfied. Let $x_j, x_{j+1} \in N(x^*, \delta/2)$. One has*

$$(1 - \sqrt{m}L_1M^2\bar{f}) \text{dist}(x_{j+1}, X^*)^2 \leq C_3^2 \text{dist}(x_j, X^*)^4 + \hat{C}_3^2 \bar{f} \text{dist}(x_j, X^*)^2,$$

where $C_3 := M\sqrt{C_1^2 + 2L_3C_1C_2^2 + L_3^2C_2^4}$ and $\hat{C}_3 := M\sqrt{2C_1 + 2L_3C_2^2}$.

Proof. Indeed, using Assumption 4.3, Lemma 4.1, the fact that the step is accepted for j sufficiently large, and $\bar{f} = \|\bar{F}_{j+1}\|$, one has

$$\begin{aligned} \|x_{j+1} - \bar{x}_{j+1}\|^2 &\leq M^2 \|F(x_j + s_j) - \bar{F}_{j+1}\|^2 \\ &\leq M^2 (\|F(x_j + s_j)\|^2 - 2F(x_j + s_j)^\top \bar{F}_{j+1} + \bar{f}^2) \\ &\leq M^2 \left((\|J(x_j)s_j + F_j\| + L_3\|s_j\|)^2 - 2F(x_j + s_j)^\top \bar{F}_{j+1} + \bar{f}^2 \right) \\ &\leq M^2 \left(\|J(x_j)s_j + F_j\|^2 + 2L_3\|J(x_j)s_j + F_j\|\|s_j\|^2 + L_3^2\|s_j\|^4 \right. \\ &\quad \left. - 2F(x_j + s_j)^\top \bar{F}_{j+1} + \bar{f}^2 \right) \\ &\leq M^2 \left(C_1^2 \|x_j - \bar{x}_j\|^4 + 2C_1\|x_j - \bar{x}_j\|^2 \bar{f} + \bar{f}^2 + 2L_3C_1\|x_j - \bar{x}_j\|^2 \|s_j\|^2 \right. \\ &\quad \left. + 2L_3\bar{f}\|s_j\|^2 + L_3^2\|s_j\|^4 - 2F(x_j + s_j)^\top \bar{F}_{j+1} + \bar{f}^2 \right). \end{aligned}$$

Therefore, using Lemma 4.2, one gets

$$\|x_{j+1} - \bar{x}_{j+1}\|^2 \leq C_3^2 \|x_j - \bar{x}_j\|^4 + \hat{C}_3^2 \bar{f} \|x_j - \bar{x}_j\|^2 + 2M^2 |F(x_j + s_j)^\top \bar{F}_{j+1} - \bar{f}^2|, \quad (16)$$

where $C_3 := M\sqrt{C_1^2 + 2L_3C_1C_2^2 + L_3^2C_2^4}$ and $\hat{C}_3 := M\sqrt{2C_1 + 2L_3C_2^2}$ are positive constants. Moreover, by applying Taylor expansion to $x \rightarrow F(x)^\top \bar{F}_{j+1}$ at the point $x_{j+1} = x_j + s_j$ around \bar{x}_{j+1} , there exists $R > 0$ such that

$$\begin{aligned} |F_{j+1}^\top \bar{F}_{j+1} - \bar{f}^2| &\leq |(J(\bar{x}_{j+1})^\top \bar{F}_{j+1})^\top (x_{j+1} - \bar{x}_{j+1})| + R\|x_{j+1} - \bar{x}_{j+1}\|^2 \\ &= |\nabla f(\bar{x}_{j+1})^\top (x_{j+1} - \bar{x}_{j+1})| + R\|x_{j+1} - \bar{x}_{j+1}\|^2 \\ &= R\|x_{j+1} - \bar{x}_{j+1}\|^2 \quad \text{by using } \bar{x}_{j+1} \in X^*. \end{aligned}$$

Note that the Hessian of $x \rightarrow F(x)^\top \bar{F}_{j+1}$ is equal to $\sum_{i=1}^m \bar{F}(x_{j+1})_i \nabla^2 F_i(x)$, and from Assumption 4.2 we have $\nabla^2 F_i(x)$ are bounded. Hence, the constant R is bounded as follows $R \leq \frac{L_1}{2} \sum_{i=1}^m |(\bar{F}_{j+1})_i| \leq \frac{\sqrt{m}L_1}{2} \|\bar{F}_{j+1}\|$. Combining the obtained Taylor expansion and (16) gives

$$\|x_{j+1} - \bar{x}_{j+1}\|^2 \leq C_3^2 \|x_j - \bar{x}_j\|^4 + \hat{C}_3^2 \bar{f} \|x_j - \bar{x}_j\|^2 + \sqrt{m}L_1 M^2 \bar{f} \|x_{j+1} - \bar{x}_{j+1}\|^2,$$

which completes this proof.

■ In next lemma, we show that, once the iterates $\{x_j\}_j$ lie sufficiently near their limit point x^* , the sequence $\{\text{dist}(x_j, X^*)\}_j$ converges to zero quadratically if the problem has a zero residual, or linearly when the residual is small.

Lemma 4.5 *Suppose that Assumptions 4.1, 4.2, and 4.3 are satisfied. Let $\{x_j\}_j$ be a sequence generated by the proposed Algorithm. Suppose that both x_j and x_{j+1} belong to $N(x^*, \delta/2)$. If the problem has a zero residual, i.e., $\bar{f} = 0$, then*

$$\text{dist}(x_{j+1}, X^*) \leq C_3 \text{dist}(x_j, X^*)^2, \quad (17)$$

where C_3 is a constant defined according to Proposition 4.1.

Otherwise, if $\bar{f} < \min \left\{ \frac{1}{\sqrt{m}L_1 M^2}, \frac{1 - C_3^2 \delta^2}{\hat{C}_3^2 + \sqrt{m}L_1 M^2} \right\}$, then

$$\text{dist}(x_{j+1}, X^*) \leq C_4 \text{dist}(x_j, X^*), \quad (18)$$

where $C_4 \in]0, 1[$ is a positive constant independent of j .

Proof. Indeed, under the zero residual case, i.e., $\bar{f} = 0$, then Proposition 4.1 is equivalent to $\text{dist}(x_{j+1}, X^*) \leq C_3 \text{dist}(x_j, X^*)^2$.

If the problem has a small residual, i.e., $\bar{f} < \min \left\{ \frac{1}{\sqrt{m}L_1 M^2}, \frac{1 - C_3^2 \delta^2}{\hat{C}_3^2 + \sqrt{m}L_1 M^2} \right\}$, then Proposition 4.1 will be equivalent to

$$\text{dist}(x_{j+1}, X^*)^2 \leq \frac{C_3^2 \delta^2 + \hat{C}_3^2 \bar{f}}{1 - \sqrt{m}L_1 M^2 \bar{f}} \text{dist}(x_j, X^*)^2 = C_4^2 \text{dist}(x_j, X^*)^2,$$

where $C_4 := \sqrt{\frac{C_3^2 \delta^2 + \hat{C}_3^2 \bar{f}}{1 - \sqrt{m}L_1 M^2 \bar{f}}} \in]0, 1[$.

■ The final theorem is standard (see, e.g., [6, Lemma 2.3]).

Theorem 4.1 *Suppose that Assumptions 4.1, 4.2, and 4.3 are satisfied. If $\bar{f} = 0$ then Algorithm 1 converges locally quadratically to X^* . Otherwise, if the problem has a small nonzero residual as in Lemma 4.5, Algorithm 1 converges locally at a linear rate to X^* .*

5 Numerical Results

In this section, we report the results of some preliminary experiments performed to test the practical behavior of Algorithm 1. All procedures were implemented in Matlab and run using Matlab 2019a on a MacBook Pro 2,4 GHz Intel Core i5, 4 GB RAM; the machine precision is $\epsilon_m \sim 2 \cdot 10^{-16}$.

We will compare our proposed algorithm with the LM method proposed in [15]. In fact, the latter algorithm can be seen to be similar to Algorithm 1 except that $\gamma_j := \mu_j \|\nabla f(x_j)\|$ where the parameter μ_j is updated in the following way. Given some constants $c_0 > 1$, $c_1 = c_0(c_0 + 1)^{-1}$ and $0 < \eta < \eta_1 < \eta_2 < 1$, if the iteration is unsuccessful then μ_j is increased (i.e., $\mu_{j+1} := c_0\mu_j$). Otherwise, if $\|\nabla f(x_j)\|\mu_j < \eta_1$ then $\mu_{j+1} = c_0\mu_j$, if $\|\nabla f(x_j)\|\mu_j > \eta_2$ then $\mu_{j+1} = \max\{c_1\mu_j, \mu_{\min}\}$, and μ_j is kept unchanged otherwise. The LM method proposed in [15] was shown to be globally convergent with a favorable complexity bound but its local behavior was not investigated. In our comparison, we will refer to the implementation of this method as **LM-(global)**, while Algorithm 1 will be referred to as **LM-(global and local)** (since it theoretically guarantees both global and local convergence properties). Both algorithms were written in Matlab and the subproblem was solved using the backslash operator. For the **LM-(global and local)** method, two variants were tested. In the first one, named **LM-(global and local)-V1**, we set the parameter μ_{j+1} equal to $\max\{\bar{\mu}/\lambda, \mu_{\min}\}$ if the iteration is declared successful. In the second variant, named **LM-(global and local)-V2**, the parameter is set $\mu_{j+1} = \bar{\mu}$ if the iteration j is successful. The initial parameters defining the implemented algorithms were set as: $\eta = 10^{-2}$, $\eta_1 = 0.1$, $\eta_2 = 0.9$, $\lambda = c_0 = 5$, $\mu_0 = 1$ and $\mu_{\min} = 10^{-16}$. As a set of problems \mathcal{P} , we used the well known 33 Moré/Garbow/Hillstom problems [20]. All the tested problems are smooth and have a least-squares structure. The residual function F and the Jacobian matrix for all the test problems [20] are implemented in Matlab. Some of these problems have a nonzero value at the optimum and thus are consistent with the scope of the paper. To obtain a larger test set, we created a set of additional 14 optimization problems by varying the problem dimension n when this was possible. For all the tested problems, we used the proposed starting points x_0 given in the problems' original presentation [20]. All algorithms are stopped when $\|\nabla f(x_j)\| \leq \epsilon$ where ϵ is the regarded accuracy level. If they did not converge within a maximum number of iterations $j_{\max} := 10000$, then they were considered to have failed.

For our test comparison, we used the performance profiles proposed by Dolan and Moré [21] over the set of problems \mathcal{P} (of cardinality $|\mathcal{P}|$). For a set of algorithms \mathcal{S} , the performance profile $\rho_s(\tau)$ of an algorithm s is defined as the fraction of problems where the performance ratio $r_{p,s}$ is at most τ , $\rho_s(\tau) = \frac{1}{|\mathcal{P}|} \text{size}\{p \in \mathcal{P} : r_{p,s} \leq \tau\}$. The performance ratio $r_{p,s}$ is in turn defined by $r_{p,s} = \frac{t_{p,s}}{\min\{t_{p,s} : s \in \mathcal{S}\}}$, where $t_{p,s} > 0$ measures the performance of the algorithm s when solving problem p , seen here as the number of iterations. Better performance of the algorithm s , relatively to the other algorithms on the set of problems, is indicated by higher values of $\rho_s(\tau)$. In particular, efficiency is measured by $\rho_s(1)$ (the fraction of problems for which algorithm s performs the best) and robustness is measured by $\rho_s(\tau)$ for τ sufficiently large (the fraction of problems solved by s). For a better visualization, we plot the performance profiles in a \log_2 -scale.

We present the obtained performance profiles using two levels of accuracy in Figure 1. For a level of accuracy of 10^{-3} , **LM-(global and local)** variants present a better efficiency compared to **LM-(global)** (in more than 60% of the tested problems **LM-(global and local)-V1** performed best, and **LM-(global and local)-V2** performed better on 40% while **LM-(global)** was better on less than 30%). When it comes to robustness, all the solvers exhibit good performance. Using a higher accuracy, the two variants of **LM-(global and local)** outperform **LM-(global)**. The **LM-(global and local)-V1** variant shows the best performance both in terms of efficiency and robustness.




Figure 1: Obtained performance profiles considering the two levels of accuracy, 10^{-3} and 10^{-5} .

In order to estimate the local convergence rate, we estimated the order of convergence by

$$\text{EOC} := \log \left(\frac{\|\nabla f(x_{j_f})\|}{\max\{1, \|\nabla f(x_0)\|\}} \right) / \log \left(\frac{\|\nabla f(x_{j_f-1})\|}{\max\{1, \|\nabla f(x_0)\|\}} \right),$$

where j_f is the index of the final computed iterate. When $\text{EOC} \geq 1.8$, the algorithm will be said quadratically convergent. If $1.8 > \text{EOC} \geq 1.1$, then the algorithm will be seen as super-linearly convergent. Otherwise, the algorithm is linearly convergent or worse. The estimation

Table 2: Order of convergence for problems in \mathcal{P} with the accuracy level $\epsilon = 10^{-5}$.

\mathcal{P}	Method	Number of problems to converge		
		Linear or worse	Superlinear	Quadratic
zero residual	LM-(global and local)-V1	2	8	18
	LM-(global and local)-V2	4	9	15
	LM-(global)	18	9	1
nonzero residual	LM-(global and local)-V1	7	7	5
	LM-(global and local)-V2	10	8	1
	LM-(global)	12	5	2

of the order of convergence (see Table 2) shows the good local behavior of the LM-(global and local) variants compared to LM-(global). In fact, LM-(global and local) variants converged quadratically or superlinearly on 38 (v1) and 33 (v2) problems respectively, while LM-(global) showed quadratic or superlinear convergence for only 17 problems.

6 Conclusions

In this paper, we presented and analyzed a novel LM method for solving nonlinear least-squares problems. We were able to formulate a globally convergent LM method with strong worst-case iteration complexity bounds. The proposed method is locally convergent at quadratic rate for

zero residual problems and at a linear rate for small residuals. Preliminary numerical results confirmed the theoretical behavior. Future research can include problems with constraints as well as those with noisy data.

Acknowledgements

We would like to thank Clément Royer and the referees for their careful readings and corrections that helped us to improve our manuscript significantly. Support for Vyacheslav Kungurtsev was provided by the OP VVV project CZ.02.1.01/0.0/0.0/16_019/0000765 “Research Center for Informatics”.

References

- [1] Tarantola, A.: Inverse Problem Theory and Methods for Model Parameter Estimation. SIAM, Philadelphia (2005)
- [2] Trémolet, Y.: Model Error Estimation in 4D-Var. *Q. J. R. Meteorol. Soc.* **133**, 1267–1280 (2007)
- [3] Levenberg, K.: A Method for the Solution of Certain Problems in Least Squares. *Quart. Appl. Math.* **2**, 164–168 (1944)
- [4] Marquardt, D.: An Algorithm for Least-Squares Estimation of Nonlinear Parameters. *SIAM J. Appl. Math.* **11**, 431–441 (1963)
- [5] Osborne, M.R.: Nonlinear Least Squares – The Levenberg Algorithm Revisited. *J. Austral. Math. Soc. Ser. B* **19**, 343–357 (1976)
- [6] Yamashita, N., Fukushima, M.: On the Rate of Convergence of the Levenberg-Marquardt Method. In: *Topics in numerical analysis*, pp. 239–249. Springer (2001)
- [7] Fan, J., Yuan, Y.: On the Quadratic Convergence of the Levenberg-Marquardt Method without Nonsingularity Assumption. *Computing* **74**, 23–39 (2005)
- [8] Dan, H., Yamashita, N., Fukushima, M.: Convergence Properties of the Inexact Levenberg-Marquardt Method under Local Error Bound Conditions. *Optim. Methods Softw.* **17**, 605–626 (2002)
- [9] Facchinei, F., Fischer, A., Herrich, M.: A Family of Newton Methods for Nonsmooth Constrained Systems with Nonisolated Solutions. *Math. Methods Operations Research* **77**, 433–443 (2013)
- [10] Ipsen, I.C.F., Kelley, C.T., Pope, S.R.: Rank-deficient Nonlinear Least Squares Problems and Subset Selection. *SIAM J. Numer. Anal.* **49**, 1244–1266 (2011)
- [11] Fan, J.: Convergence Rate of the Trust Region Method for Nonlinear Equations under Local Error Bound Condition. *Comput. Optim. Appl.* **34**, 215–227 (2006)
- [12] Conn, A.R., Gould, N.I.M., Toint, Ph.L.: *Trust-Region Methods*. SIAM, Philadelphia, PA, USA (2000)
- [13] Ueda, K., Yamashita, N.: On a global complexity bound of the Levenberg-Marquardt method. *J. Optim. Theory Appl.* **147**, 443–453 (2010)
- [14] Ueda, K., Yamashita, N.: Global Complexity Bound Analysis of the Levenberg-Marquardt Method for Nonsmooth Equations and Its Application to the Nonlinear Complementarity Problem. *J. Optim. Theory Appl.* **152**, 450–467 (2012)
- [15] Zhao, R., Fan, J.: Global complexity bound of the Levenberg-Marquardt method. *Optim. Methods Softw.* **31**, 805–814 (2016)

- [16] Nocedal, J., Wright, S.J.: Numerical Optimization, second edn. Springer, Berlin (2006)
- [17] Bergou, E., Gratton, S., Vicente, L.N.: Levenberg-Marquardt Methods Based on Probabilistic Gradient Models and Inexact Subproblem Solution, with Application to Data Assimilation. *SIAM/ASA J. Uncertainty Quantification* **4**, 924–951 (2016)
- [18] Bonnans, J.F., Shapiro, A.: Perturbation Analysis of Optimization Problems. Springer Science & Business Media (2013)
- [19] Fischer, A., Shukla, P., Wang, M.: On the inexactness level of robust Levenberg-Marquardt methods. *Optimization* **59**, 273–287 (2010)
- [20] Moré, J.J., Garbow, B.S., Hillstom, K.E.: Testing Unconstrained Optimization Software. *ACM Trans. Math. Softw.* **7**, 17–41 (1981)
- [21] Dolan, E. D. and Moré, J. J.: Benchmarking optimization software with performance profiles. *Math. Prog.* **91(2)**, 201–213 (2002)