On the inexact proximal Gauss-Newton methods for regularized nonlinear least squares problems

Federica Porta, Silvia Villa, Marco Viola, and Martin Zach

Abstract The Gauss-Newton method is one of the most common choices for solving nonlinear systems $F(x) = 0$. The idea is to minimize the corresponding least squares problem $\frac{1}{2}\|F(x)\|^2$ by solving a sequence of linearized problems. The method has been extended to account for non-smooth regularizers, leading to proximal Gauss-Newton algorithms. At each iteration, the algorithm avoids computation or storage of the Hessian, but requires two in principle costly operations: inversion of the matrix $F'(x)^*F'(x)$ and computation of the proximal point in a variable metric. To overcome this limitation, we propose an inexact version of the proximal Gauss-Newton algorithm based on an iterative approximation of the linearized sub-problems at each iteration. Numerical experiments on both convex ℓ_1 penalized nonlinear least squares problems arising in binary classification as well as non-convex bound-constrained nonlinear least squares problems show promising performance of the suggested approach.

1 Introduction

Given Hilbert spaces X and Y, we consider the following penalized nonlinear least squares problem

$$\min_{x \in X} \frac{1}{2}\|F(x)\|^2 + J(x) := \Phi(x) \tag{$\mathcal{P}$}$$

Federica Porta
University of Modena and Reggio Emilia, e-mail: federica.porta@unimore.it

Silvia Villa
University of Genoa, e-mail: silvia.villa@unige.it

Marco Viola
University College Dublin, e-mail: marco.viola@ucd.ie

Martin Zach
Graz University of Technology, e-mail: martin.zach@icg.tugraz.at

where $F : X \to Y$ is a Gâteaux differentiable nonlinear operator and $J : X \to \mathbb{R} \cup \{+\infty\}$ is a convex lower-semicontinuous penalty function. The optimization model ($\mathcal{P}$) allows one to formalize different problems arising in real-life applications (see for example [3, 6, 8, 13]), making it a considerably interesting area of investigation.

The objective function ($\mathcal{P}$) is composed of a differentiable possibly non-convex term and a possibly non-differentiable convex term. Forward-backward methods, also called proximal gradient methods, exploit such a structure and thus can be considered to solve ($\mathcal{P}$). They have been intensively investigated in the literature, leading to several different versions including inertial steps, inexact computation of the proximal operator, variable metric strategies and adaptive step size selection rules (see for example [1, 2, 5, 4, 11, 12] and references therein).

Particularly, in [11] the authors suggested a proximal Gauss-Newton method to solve ($\mathcal{P}$). The general iteration of this scheme can be written as

$$x_{n+1} = \mathrm{prox}_J^{H(x_n)}\big(x_n - H(x_n)^{-1}F'(x_n)^*F(x_n)\big), \tag{1}$$

where $H(x_n) = F'(x_n)^*F'(x_n)$ and $\mathrm{prox}_J^{H(x_n)}$ is the proximal operator of J with respect to the metric induced by the operator $H(x_n)$. For this proximal Gauss-Newton algorithm, convergence results of local type as well as an estimate of the radius of convergence can be obtained (see, e.g., [11, Theorem 1, Corollary 1, and Corollary 2]).

One drawback of this approach (and in general of the Gauss-Newton methods) is the inversion of $H(x_n)$ at each iteration. Another in principle costly operation due to the presence of the regularizer is the computation of the proximal map. The main aim of this paper is to overcome these two difficulties at once by proposing an inexact version of algorithm (1) based on a suitable approximation of the Gauss-Newton step. In this setting, we generalize the convergence results developed in [11] by detailing the assumptions required on the sequence of points which approximate the iterates x_n. The influence of computational errors on the proximal Gauss-Newton method has been investigated also in [9] for the particular case of J being the indicator function of a convex set. In [9], the inexactness is limited to the computation of the proximal operator (in their case an orthogonal projection), but the inversion of the Hessian is assumed to be exact. In this paper we consider a generic function J, which allows one to take into account a more general convex regularization term. In order to evaluate the effectiveness of the suggested inexact proximal Gauss-Newton method, we perform numerical experiments on penalized nonlinear least squares problems. In particular, we consider ℓ_1-regularized convex problems arising in binary classification applications and bound-constrained non-convex problems from standard repositories.

2 The algorithm and its convergence analysis

We start by setting the assumptions on F and J. Let Ω be an open convex subset of X, we assume that

$$\begin{cases} F : \Omega \subseteq X \to Y \text{ is Gâteaux differentiable and} \\ J : X \to \mathbb{R} \cup \{+\infty\} \text{ is proper, lower semicontinuous, and convex.} \end{cases} \tag{A1}$$

Note that the resulting optimization problem is non-convex, therefore we will focus our attention on the convergence towards local minimizers. Hereafter we will denote by x_* any local minimizers of the penalized nonlinear least squares problem. Now we recall first order conditions for local minimizers [11, Proposition 5]. If assumption (A1) is satisfied, any local minimizer x_* fulfills

$$-F'(x_*)^* F(x_*) \in \partial J(x_*),$$

where the subdifferential is the Rockafellar one [4]. To introduce another sufficient condition for local minimizers, we recall the definition of proximal operator induced by a linear operator.

Definition 1 Let $H\colon X \to X$ be a symmetric positive definite bounded operator and consider for all $x, z \in X$ the scalar product $\langle x, z\rangle_H = \langle x, Hz\rangle$ and the corresponding induced norm $\|\cdot\|_H$ on the space X. For any given $z \in X$, the proximal operator of J with respect to the metric induced by the linear operator H is defined as

$$\operatorname{prox}_J^H(z) = \arg\min_{x\in X} \left\{ J(x) + \frac{1}{2}\|x - z\|_H^2 \right\}.$$

Writing the first order conditions, this can be equivalently characterized as

$$\operatorname{prox}_J^H = (H + \partial J)^{-1} \circ H.$$

Relying on the previous definition, we state a different condition satisfied by any local minimizer x_* (see [11, Proposition 5]). If $F'(x_*)$ is injective and has closed range then x_* satisfies the fixed point equation

$$x_* = \operatorname{prox}_J^{H(x_*)}(x_* - F'(x_*)^\dagger F(x_*)), \tag{2}$$

where $H(x_*) = F'(x_*)^* F'(x_*)$ and $(\cdot)^\dagger$ denotes the Moore-Penrose inverse (see Appendix 4 for more details). In order to devise the convergence analysis we consider the following additional assumptions:

$$\begin{cases} \text{The function } \Phi \text{ has a local minimizer } x_*. \\ F'(x_*) \text{ is injective with closed range.} \\ F'\colon \Omega \subseteq X \to \mathbb{L}(X, Y) \text{ is Lipschitz continuous with constant } L \end{cases} \tag{A2}$$

In the assumption above, $\mathbb{L}(X, Y)$ denotes the space of bounded linear operators from X to Y. Assumptions (A2) are in line with the state-of-the-art of Gauss-Newton methods [10, 11]. Furthermore, the Lipschitz continuity of F' assumed in (A2) is a typical assumption in nonlinear inverse problems and roughly controls the nonlinearity of the problem. An example in which the above assumptions are satisfied is the one of linear inverse problems, namely $F(x) = Ax - y$ for some bounded linear operator $A\colon X \to Y$ with closed range and $y \in Y$. Another example is the one in which we need to solve a system of m quadratic equations in d variables. In the latter case the derivative of the function F is an affine operator and therefore it is Lipschitz continuous.

In the following we report the convergence analysis for an inexact version of the proximal Gauss-Newton method (1) based on the next definition.

Definition 2 Assume (A1)–(A2), let $x \in X$ and $\varepsilon = (\varepsilon_0, \varepsilon_1) \in \mathbb{R}_+^2$, and let x_* be a local minimizer of problem ($\mathcal{P}$). We say that $z \in X$ is an ε-approximate proximal Gauss-Newton step starting from x, and we write

$$z \approx_\varepsilon \operatorname{prox}_J^{H(x)}\bigl(x - F'(x)^\dagger F(x)\bigr),$$

if

$$\|z - \operatorname{prox}_J^{H(x)}(x - F'(x)^\dagger F(x))\| \le (\varepsilon_1 \|x - x_*\| + \varepsilon_0)\|x - x_*\|.$$

Next, we state the main result of the paper. The proof follows the same line as that of Theorem 1 in [11]. The main difference is that here we consider an inexact version of the Gauss-Newton step, while in [11] problem (1) is supposed to be exactly solved. Moreover, here we can simplify the proof, since we make a stronger Lipschitz assumption on $F'(x)$, which also allows us to obtain more explicit estimates. The analysis of inexact proximal Gauss-Newton method establishes convergence to a local optimum for suitable initialization, sufficiently close to the local minimizer.

Theorem 1 *Suppose that* (A1)–(A2) *are satisfied. Let $x_* \in \Omega$ be a local minimizer of Φ. Set $\alpha = \|F(x_*)\|$, $\beta = \left\|F'(x_*)^\dagger\right\|$, $\kappa = \left\|F'(x_*)^\dagger\right\| \|F'(x_*)\|$. Let $\varepsilon = (\varepsilon_0, \varepsilon_1) \in \mathbb{R}_+^2$. Suppose that $\dfrac{\beta^2 \alpha L(\kappa(1+\sqrt{2})+1)}{1-\varepsilon_0} < 1$ and $R = \frac{1}{L\beta}\left(1 - \beta\sqrt{\frac{L\alpha(\kappa(1+\sqrt{2})+1)}{1-\varepsilon_0}}\right)$ and let $r \in \left]0, R\right[$ be such that $B_r(x_*) \subseteq \Omega$. Define*

$$q_0 = \frac{(\kappa(1+\sqrt{2})+1)\beta^2\alpha L}{(1-\beta L r)^2},$$

$$q_1 = \frac{\beta L}{2\,(1-\beta L r)^2}\left(\beta L r + 2(1+\sqrt{2})L\beta^2\alpha + \kappa\right),$$

and assume that $r(q_1 + \varepsilon_1) + q_0 + \varepsilon_0 < 1$. Consider the approximate Gauss-Newton sequence

$$x_0 \in B_r(x_*),$$

$$x_{n+1} = G(x_n) \approx_\varepsilon \operatorname{prox}_J^{H(x_n)}\left(x_n - F'(x_n)^\dagger F(x_n)\right)$$

with $H(x_n) = F'(x_n)^* F'(x_n)$ and $G(x)$ an ε-approximation of the Gauss-Newton step for every $x \in \Omega$.

Then the approximate Gauss-Newton sequence is well-defined, i.e., $x_n \in B_r(x_*)$ and $F'(x_n)$ is injective with closed range and it holds

$$\|x_{n+1} - x_*\| \leq q^n \|x_0 - x_*\|,$$

where $q < 1$. *More precisely,*

$$\|x_{n+1} - x_*\| \leq (q_1 + \varepsilon_1) \|x_n - x_*\|^2 + (q_0 + \varepsilon_0) \|x_n - x_*\|.$$

The proof of Theorem 1 is postponed to the end of the section, since we first need some auxiliary results. The following fixed point theorem will be crucial to prove convergence of the inexact proximal Gauss Newton method. Let us comment here that r in the statement can be chosen in such a way that the condition $r(q_1+\varepsilon_1)+q_0+\varepsilon_0 < 1$ is satisfied since the function $r \mapsto r(q_1 + \varepsilon_1) + q_0 + \varepsilon_0$ is continuous and is strictly less than 1 at $r = 0$. Indeed the hypothesis $\frac{\beta^2 \alpha L(\kappa(1+\sqrt{2})+1)}{1-\varepsilon_0} < 1$ and the definition of R ensure that $q_0 + \varepsilon_0 < 1$.

Proposition 1 *Let D be an open and convex subset of X. Let $\bar{G} : D \subseteq X \to X$ be a mapping and $x_* \in D$ a fixed point of $\bar{G}$. Assume that there exist positive constants q_0 and q_1 such that $\bar{G}$ satisfies the inequality*

$$\left\|\bar{G}(x) - x_*\right\| \leq (q_1\|x - x_*\| + q_0)\|x - x_*\|, \qquad \text{for all } x \in D. \tag{3}$$

Let $G\colon D \subseteq X \to X$ and $\varepsilon_0 > 0$ and $\varepsilon_1 > 0$ be such that

$$\left\|G(x) - \bar{G}(x)\right\| \leq (\varepsilon_1\|x - x_*\| + \varepsilon_0)\|x - x_*\|, \qquad \text{for all } x \in D. \tag{4}$$

Suppose that $q_0 + \varepsilon_0 < 1$ and define

$$\bar{r} = \frac{1 - q_0 - \varepsilon_0}{q_1 + \varepsilon_1}.$$

Then given $r \in \mathbb{R}$ with $0 < r \leq \bar{r}$ and $B_r(x_*) \subseteq D$, it follows $G(B_r(x_*)) \subseteq B_r(x_*)$, thus, given $x_0 \in B_r(x_*)$ the sequence defined by setting $x_{n+1} = G(x_n)$ is well-defined. Moreover, denoting $q = (q_1 + \varepsilon_1)\|x_0 - x_*\| + q_0 + \varepsilon_0$, it holds $q < 1$,

$$\|x_{n+1} - x_*\| \leq [(q_1 + \varepsilon_1)\|x_n - x_*\| + q_0 + \varepsilon_0] \|x_n - x_*\| \tag{5}$$

and

$$\|x_{n+1} - x_*\| \leq q^n\|x_0 - x_*\|. \tag{6}$$

Proof First note that $\bar{r} > 0$ being $q_0 + \varepsilon_0 < 1$. Fix $r \in \mathbb{R}$, $0 < r \leq \bar{r}$ such that $B_r(x_*) \subseteq D$ and let $x \in B_r(x_*)$. Then, it follows from $\|x - x_*\| < \bar{r}$ that $(q_1 + \varepsilon_1)\|x - x_*\| + q_0 + \varepsilon_0 < 1$. Therefore (3) and (4) imply that

$$\|G(x) - x_*\| \leq \left\|G(x) - \bar{G}(x)\right\| + \left\|\bar{G}(x) - x_*\right\| \leq \|x - x_*\| < r,$$

i.e., $G(x) \in B_r(x_*)$. Thus $G(B_r(x_*)) \subseteq B_r(x_*)$ and a sequence can be defined in $B_r(x_*)$ by choosing $x_0 \in B_r(x_*)$ and setting $x_{n+1} = G(x_n)$. Being $\|x_n - x_*\| < \bar{r}$, again from the definition of $\bar{r}$ we have $[(q_1 + \varepsilon_1)\|x_n - x_*\| + q_0 + \varepsilon_0] < 1$, and from (3) we get

$$
\begin{aligned}
\|x_{n+1} - x_*\| &= \|G(x_n) - x_*\| \\
&\le [(q_1 + \varepsilon_1)\|x_n - x_*\| + q_0 + \varepsilon_0]\, \|x_n - x_*\| \qquad (7)\\
&< \|x_n - x_*\|.
\end{aligned}
$$

This implies (5) and that $(q_1+\varepsilon_1)\|x_{n+1}-x_*\|+q_0+\varepsilon_0 \le (q_1+\varepsilon_1)\|x_n-x_*\|+q_0+\varepsilon_0$. Therefore, recalling the definition of q, we get $(q_1 + \varepsilon_1)\|x_n - x_*\| + q_0 + \varepsilon_0 \le q < 1$ for all $n \in \mathbb{N}$ and equation (7) yields

$$\|x_{n+1} - x_*\| \le q\|x_n - x_*\| \le \ldots \le q^n \|x_0 - x_*\|.$$

□

Given the following definitions

$$M\colon X \to X, \quad x \mapsto x - F'(x)^\dagger F(x),$$

and

$$\bar{G}\colon D \subseteq X \to X, \qquad \bar{G}(x) = \operatorname{prox}_J^{H(x)}(M(x)), \tag{8}$$

the domain of $\bar{G}$ is the subset D of Ω defined as

$$D = \{x \in \Omega\colon F'(x) \text{ is injective and range}(F'(x)) \text{ is closed}\}, \tag{9}$$

where Proposition 2 shows that the operator $\bar{G}$ in (8) satisfies requirement (3). Moreover, note that, if $x_* \in D$ is a local minimizer of $(\mathcal{P})$, the fixed point equation (2) can be restated by saying that x_* is a fixed point for $\bar{G}$, namely

$$x_* = \bar{G}(x_*). \tag{10}$$

Proposition 2 *Assume* (A1)–(A2)*, and suppose that* $x_* \in \Omega$ *is a local minimizer of* Φ*. Set* $\alpha = \|F(x_*)\|$, $\beta = \left\|F'(x_*)^\dagger\right\|$ *and* $\kappa = \left\|F'(x_*)^\dagger\right\| \|F'(x_*)\|$*. Then, for all* $r \in \mathbb{R}$ *with* $0 < r < 1/(L\beta)$ *and* $B_r(x_*) \subseteq \Omega$*, it follows that, for every* $x \in B_r(x_*)$

$$\|\bar{G}(x) - x_*\| \le (q_1\|x - x_*\| + q_0)\|x - x_*\|,$$

with $q_0 = \dfrac{(\kappa(1+\sqrt{2})+1)\beta^2\alpha L}{(1-\beta L r)^2}$ *and* $q_1 = \dfrac{\beta L}{2\,(1-\beta L r)^2}(\beta L r + 2(1+\sqrt{2})L\beta^2\alpha + \kappa)$.

Proof Since x_* is a local minimizer of Φ and $x_* \in D$ by (A2), with D defined as in (9), x_* is a fixed point of $\bar{G}$, see (10). Therefore

$$M(x_*) - \operatorname{prox}_J^{H(x_*)}(M(x_*)) = M(x_*) - x_* = -F'(x_*)^\dagger F(x_*). \tag{11}$$

Given $x \in D$, since F' is L-Lipschitz

$$\|F'(x) - F'(x_*)\| \left\|F'(x_*)^\dagger\right\| \le L\beta\, \|x - x_*\| < 1,$$

and thus by (22), $F'(x)$ is injective, with closed range, and

$$\left\|F'(x)^{\dagger}\right\| \leq \frac{\beta}{1-\beta L r}. \tag{12}$$

It follows from [11, Remark 4] with $H_1 = H(x)$, $H_2 = H(x_*)$, $z_1 = M(x)$ and $z_2 = M(x_*)$, and taking into account (11) we get

$$\begin{aligned}
\left\|\bar{G}(x) - x_*\right\| &= \left\|\mathrm{prox}_J^{H(x)}(M(x)) - \mathrm{prox}_J^{H(x_*)}(M(x_*))\right\| \\
&\leq \left(\|H(x)\| \left\|H(x)^{-1}\right\|\right)^{1/2} \|M(x) - M(x_*)\| \\
&\quad + \left\|H(x)^{-1}\right\| \left\|(H(x) - H(x_*))\big(M(x_*) - \mathrm{prox}_J^{H(x_*)}(M(x_*))\big)\right\| \\
&= \left(\|H(x)\| \left\|H(x)^{-1}\right\|\right)^{1/2} \|M(x) - M(x_*)\| \\
&\quad + \left\|H(x)^{-1}\right\| \left\|(H(x) - H(x_*))F'(x_*)^{\dagger} F(x_*)\right\|.
\end{aligned} \tag{13}$$

Moreover

$$\|H(x)\| = \|F'(x)^* F'(x)\| = \|F'(x)\|^2$$

$$\left\|H(x)^{-1}\right\| = \left\|[F'(x)^* F'(x)]^{-1}\right\| = \left\|F'(x)^{\dagger}\right\|^2.$$

Recalling the properties of the Moore-Penrose generalized inverse in (21) we get

$$\begin{aligned}
(H(x) - H(x_*))F'(x_*)^{\dagger} &= (F'(x)^* F'(x) - F'(x_*)^* F'(x_*))F'(x_*)^{\dagger} \\
&= F'(x)^* F'(x) F'(x_*)^{\dagger} - F'(x_*)^* F'(x_*) F'(x_*)^{\dagger} \\
&= F'(x)^* F'(x) F'(x_*)^{\dagger} - F'(x)^* P_{R(F'(x_*))} \\
&\quad + F'(x)^* P_{R(F'(x_*))} - F'(x_*)^* P_{R(F'(x_*))} \\
&= F'(x)^* (F'(x) - F'(x_*)) F'(x_*)^{\dagger} \\
&\quad + (F'(x) - F'(x_*))^* P_{R(F'(x_*))},
\end{aligned}$$

therefore,

$$\left\|(H(x) - H(x_*))F'(x_*)^{\dagger}\right\| \leq \left(\|F'(x)\| \left\|F'(x_*)^{\dagger}\right\| + 1\right) \|F'(x) - F'(x_*)\|. \tag{14}$$

Hence, substituting in (13) the bound derived in (14) we obtain

$$\begin{aligned}
\left\|\bar{G}(x) - x_*\right\| &\leq \|F'(x)\| \left\|F'(x)^{\dagger}\right\| \|M(x) - M(x_*)\| \\
&\quad + \left\|F'(x)^{\dagger}\right\|^2 \left(\|F'(x)\| \left\|F'(x_*)^{\dagger}\right\| + 1\right) \|F'(x) - F'(x_*)\| \|F(x_*)\|.
\end{aligned} \tag{15}$$

On the other hand, thanks to the properties of the Moore-Penrose pseudoinverse reported in (21) and the injectivity of $F'(x)$ we derive

$$
\begin{aligned}
M(x) - M(x_*) &= x - x_* - F'(x)^\dagger F(x) + F'(x_*)^\dagger F(x_*) \\
&= F'(x)^\dagger [F'(x)(x - x_*) - F(x) + F(x_*)] \\
&\quad + (F'(x_*)^\dagger - F'(x)^\dagger) F(x_*)
\end{aligned}
$$

and thus, by (23)

$$
\begin{aligned}
\|M(x) - M(x_*)\| &\le \left\|F'(x)^\dagger\right\| \|F(x_*) - F(x) - F'(x)(x_* - x)\| \\
&\quad + \left\|F'(x_*)^\dagger - F'(x)^\dagger\right\| \|F(x_*)\| \\
&\le \left\|F'(x)^\dagger\right\| \|F(x_*) - F(x) - F'(x)(x_* - x)\| \\
&\quad + \sqrt{2}\left\|F'(x)^\dagger\right\| \left\|F'(x_*)^\dagger\right\| \|F'(x) - F'(x_*)\| \|F(x_*)\| \\
&= \left\|F'(x)^\dagger\right\| \big\{ \|F(x_*) - F(x) - F'(x)(x_* - x)\| \\
&\quad + \sqrt{2}\left\|F'(x_*)^\dagger\right\| \|F(x_*)\| \|F'(x) - F'(x_*)\| \big\}.
\end{aligned} \tag{16}
$$

Substituting (16) in (15)

$$
\begin{aligned}
\left\|\bar{G}(x) - x_*\right\| \le \left\|F'(x)^\dagger\right\|^2 \Big\{ &\|F'(x)\| \big(\|F(x_*) - F(x) - F'(x)(x_* - x)\| \\
&+ (1 + \sqrt{2}) \left\|F'(x_*)^\dagger\right\| \|F(x_*)\| \|F'(x) - F'(x_*)\| \big) \\
&+ \|F'(x) - F'(x_*)\| \|F(x_*)\| \Big\}.
\end{aligned} \tag{17}
$$

Lipschitz continuity of F' and (12) imply

$$
\begin{aligned}
\left\|\bar{G}(x) - x_*\right\| \le \frac{\beta^2}{(1 - \beta L r)^2} \Big\{ &\|F'(x)\| \, [(L/2)\|x - x_*\|^2 \\
&+ (1 + \sqrt{2})\beta\alpha L \|x - x_*\|] + \alpha L \|x - x_*\| \Big\}.
\end{aligned}
$$

By the triangle inequality, we get

$$
\|F'(x)\| \le \|F'(x) - F'(x_*)\| + \|F'(x_*)\| \le L\|x - x_*\| + \frac{\kappa}{\beta}.
$$

Thus, we finally obtain

$$
\begin{aligned}
\left\|\bar{G}(x) - x_*\right\| \le \frac{\beta^2}{(1 - \beta L r)^2} \Big\{ &(L\|x - x_*\| + \kappa/\beta) [(L/2)\|x - x_*\|^2 \\
&+ (1 + \sqrt{2})\beta\alpha L \|x - x_*\|] + \alpha L \|x - x_*\| \Big\},
\end{aligned}
$$

or, equivalently

$$\left\|\bar{G}(x) - x_*\right\| \leq \frac{\beta L}{2\,(1-\beta L r)^2}(\beta L r + 2(1+\sqrt{2})L\beta^2\alpha + \kappa)\|x - x_*\|^2 + \frac{(\kappa(1+\sqrt{2})+1)\beta^2\alpha L}{(1-\beta L r)^2}\|x - x_*\|.$$

□

We can finally detail the proof of Theorem 1.

Proof (Theorem 1) Let $\bar{G}$ be defined as in (8) and G be the approximation of $\bar{G}$ in the sense of Definition 2 for every $x \in D$. Thanks to the assumption on r and the fact that the hypotheses (3) and (4) are ensured by Proposition 1 and Definition 2 respectively, implies that $\bar{G}$ satisfies Proposition 2. So the statement follows from (5) and (6) in Proposition 1.

3 Numerical experiments

We performed numerical experiments on realistic problems to assess the performance of inexact proximal Gauss-Newton methods. Before detailing the numerical experiments, we clarify the implementation of the inexact proximal Gauss-Newton algorithm. Observe that x_{n+1} in (1) can be reformulated by linearizing the operator F at each iterate x_n. In more detail, if $F'(x_n)$ is injective with closed range, x_{n+1} in (1) can be equivalently written as [11, Proposition 6]

$$x_{n+1} = \arg\min_{x\in X} \frac{1}{2}\left\|F(x_n) + F'(x_n)(x - x_n)\right\|^2 + J(x). \tag{18}$$

As a consequence, an inexact proximal Gauss-Newton approach can be realized by inexactly solving the minimization problem in (18) by means of efficient iterative first order methods (or other approaches). As an example, when using proximal gradient to solve (18), the resulting method can be outlined as

$$\begin{cases} x_0 \in X \\ \text{For } n = 0, \dots \\ \quad z_n^0 = x_n \\ \quad \gamma = \lambda_{\max}(H(x_n))^{-1} \\ \quad \text{For } k = 0, \dots, M \\ \quad\quad z_n^{k+1} = \operatorname{prox}_J\left(z_n^k - \gamma F'(x_n)^*\left(F'(x_n)(z_n^k - x_n) + F(x_n)\right)\right) \\ \quad x_{n+1} = z_n^M \end{cases} \tag{19}$$

We remark that Definition 2 introduces an inexact criterion of impractical use and procedure (19) is employed to obtain an approximation of the Gauss-Newton step.

3.1 Convex problems

For the first experiment we considered 5 datasets from the LIBSVM[1] binary classification repository and built convex regularized nonlinear least squares problems of the form

$$\min_{x \in \mathbb{R}^n} \frac{1}{2N} \sum_{i=1}^{N} \phi_i(x)^2 + \frac{1}{N} \|x\|_1, \tag{20}$$

where $\phi_i(z) = \sqrt{\log\left(1 + e^{-b_i\, a_i^\top z}\right)}$, with (a_i, b_i) representing an instance in the training set. These problems are meant to reproduce ℓ_1-regularized logistic regression models which are typically used in the context of data classification. In Table 1 we report the number of training points N and the problem dimensionality n for each of the 5 considered instances.

	N	n
a9a	32 561	123
cod-rna	59 535	8
covtype	406 709	54
ijcnn1	34 994	22
phishing	11 055	68

Table 1 Number of training points N and problem dimensionality n for regularized logistic regression problems.

For this test we implemented an inexact proximal Gauss-Newton method, where we solve the inner regularized linear least squares problem by means of FISTA[2]. To realize the inexactness, we limited the number of inner FISTA iterations M to 25, 10, and 5. We compare the resulting three algorithms (named respectively GN_25, GN_10, and GN_5) to running FISTA directly on the original problem (20). In Figures 1 and 2, we plot the value of the objective function value against the iteration count (left column) and elapsed time in seconds (right column). Here, in the case of our proposed algorithms, iteration count refers to the outer iterations. For FISTA, it is simply the number of iterations in the algorithm.

The figure indicates that GN_25 outperforms all the other algorithms when looking at the iterations count and that there is a progressive decrease in the performance when the iterations are reduced to 10 and 5. In terms of time, GN_10 performs best on the cod-rna and covtype datasets. For a9a, we observe that FISTA performs best when little time is elapsed, but GN_10 is competitive afterwards.

[1] https://www.csie.ntu.edu.tw/ cjlin/libsvmtools/datasets/binary.html

[2] https://www.tau.ac.il/ becka/solvers/fista

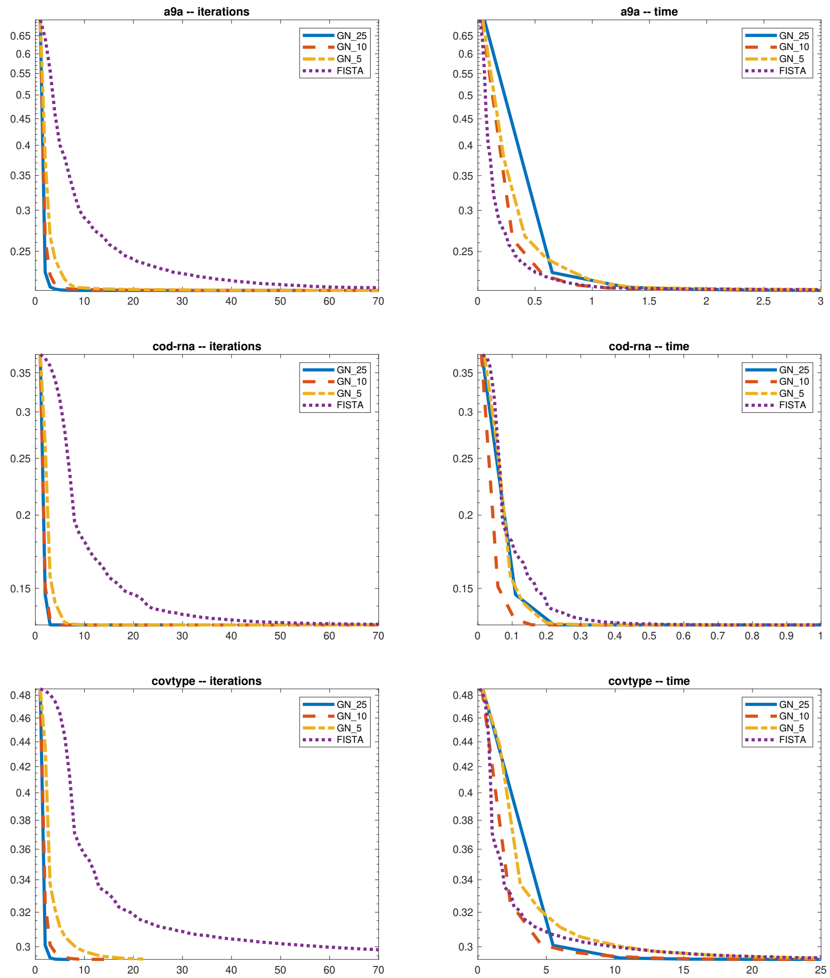

Fig. 1 Results in terms of iterations (*left*) and time (*right*) for datasets A9A (*top*), COD-RNA (*center*), and COVTYPE (*bottom*).

3.2 Non-convex problems

To test the performance of the inexact proximal Gauss-Newton method on non-convex problems, we considered modified versions of nonlinear non-convex optimization problems from [7]. In detail, as done in [11] we took 3 unconstrained optimization problems (ROSENBROCK, OSBORNE1, and KOWALIK) and equipped them with bound constraints, thus obtaining problems of the form ($\mathcal{P}$) where J models the bound constraints (see [11, Section 8] for further details).

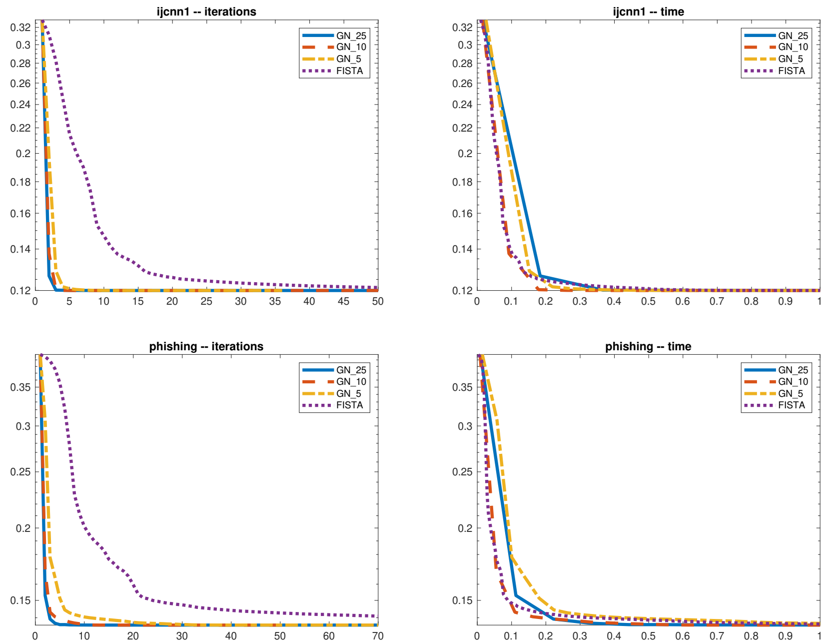

Fig. 2 Results in terms of iterations (*left*) and time (*right*) for datasets IJCNN1 (*top*) and PHISHING (*bottom*).

Again we considered inexact proximal Gauss-Newton approaches by setting a prefixed number of inner iterations of the solver for the minimization subproblems. In this case they are bound-constrained least squares problems which we solved by a proximal gradient method. In detail we considered a maximum number of inner solver iterations M equal to 100, 50, 10, and 5 and compared the four resulting algorithms (named respectively GN_100, GN_50, GN_10, and GN_5) with the proximal gradient (indicated as PG) method applied directly to the original problem. It is worth mentioning that for the inner problems in the inexact Gauss-Newton implementations we used a fixed steplength related to the subproblem Lipschitz constant. For the proximal gradient method to guarantee convergence we adopted a line-search along the projection arc with Armijo stopping conditions, using an Adaptive Barzilai-Borwein steplength [14] as starting value for the line search. The results obtained are reported in Figure 3 in terms of objective function value versus iterations (left column) and elapsed time in seconds (right column) for each of the 3 considered instances.

From the plots it is clear that GN_5 outperforms the other algorithms in terms of elapsed time for all the three problems. It is interesting to observe that in terms of iteration count all the GN algorithms perform equivalently on OSBORNE1, while GN_5 is the best performing GN also in terms of iterations on KOWALIK and the worst performing on ROSENBROCK.

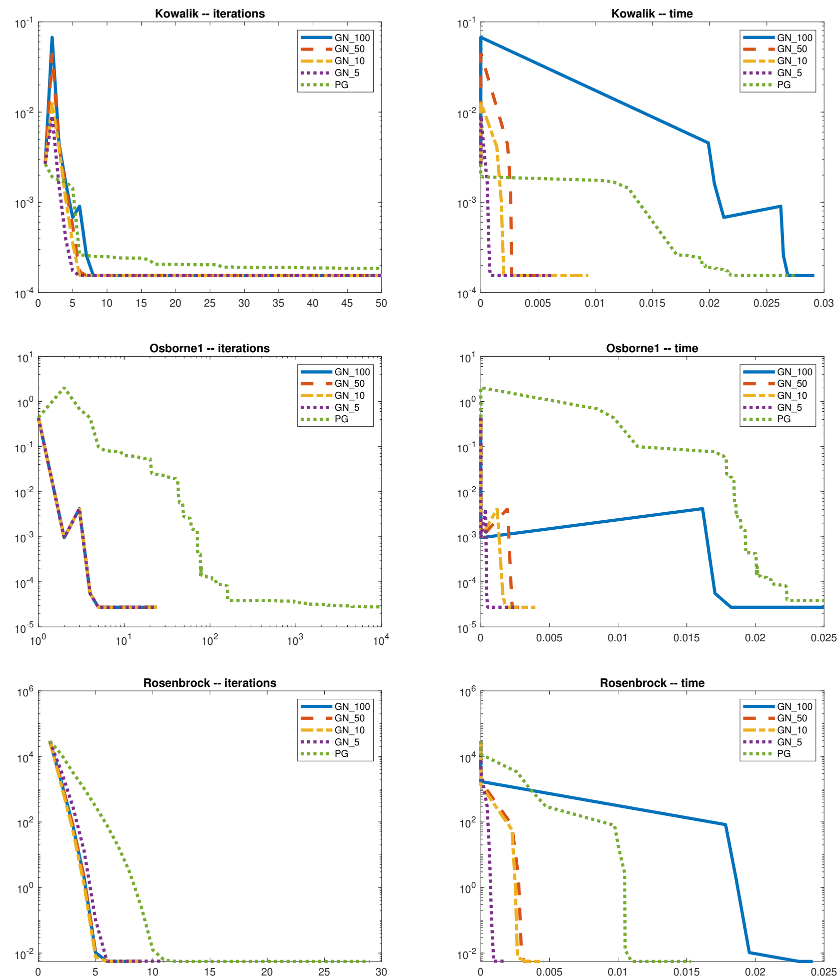

Fig. 3 Results in terms of iterations (*left*) and time (*right*) for problems Kowalik (*top*), Osborne1 (*center*), and Rosenbrock (*bottom*).

4 Auxiliary results

In this section we collect some well-known facts useful in our analysis. Given to Hilbert spaces X and Y and a linear operator $A\colon X \to Y$ with closed range, the pseudoinverse $A^\dagger\colon Y \to X$ of A is the linear operator satisfying the four Moore-Penrose equations. In particular, if we denote by $P_{N(A)}$ and $P_{R(A)}$ the projection onto the kernel and the range of A respectively, the following identities are satisfied

$$A^\dagger A = I - P_{N(A)}, \qquad AA^\dagger = P_{R(A)}. \tag{21}$$

If A is injective and has a closed range then an explicit formula for the pseudoinverse is given by $A^\dagger = (A^*A)^{-1}A^*$. In addition, if B is another linear operator such that $\|(B-A)A^\dagger\| < 1$, then B is injective, the range of B is closed and

$$\|B^\dagger\| \le \frac{\|A^\dagger\|}{1 - \|(B-A)A^\dagger\|}, \tag{22}$$

$$\|A^\dagger - B^\dagger\| \le \sqrt{2}\,\|A^\dagger\|\|B^\dagger\|\|A-B\|. \tag{23}$$

Acknowledgements S.V. acknowledges the financial support of the European Research Council (grant SLING 819789), the H2020-MSCA-ITN Project Trade-OPT 2019,funded by the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 861137, the AFOSR projects FA9550-18-1-7009 and FA8655-22-1-7034 (European Office of Aerospace Research and Development). The research by S.V. was partially supported by the MUR Excellence Department Project awarded to Dipartimento di Matematica, Università di Genova, CUP D33C23001110001. S. V. is part of the INdAM group "Gruppo Nazionale per l'Analisi Matematica, la Probabilità e le loro applicazioni". F. P. and M. V. are part of the INdAM group "Gruppo Nazionale per il Calcolo Scientifico".

References

1. H. Attouch, J. Bolte, and B. F Svaiter. Convergence of descent methods for semi–algebraic and tame problems: proximal algorithms, forward–backward splitting, and regularized Gauss–Seidel methods. *Math. Program.*, 137:91–129, 2013.
2. S. Bonettini, F. Porta, M. Prato, S. Rebegoldi, V. Ruggiero, and L. Zanni. Recent advances in variable metric first-order methods. In *Computational Methods for Inverse Problems in Imaging*, volume 36, pages 1–31. Springer INDAM Series, 2019.
3. J. Chung and J. G. Nagy. Nonlinear least squares and super resolution. In *Journal of Physics: Conference Series*, volume 124, page 012019. IOP Publishing, 2008.
4. P. L. Combettes and J.-C. Pesquet. Proximal splitting methods in signal processing. In *Fixed-point algorithms for inverse problems in science and engineering*, Springer Optimization and Its Applications, pages 185–212. Springer, New York, NY, 2011.
5. P. L. Combettes and V. R. Wajs. Signal recovery by proximal forward-backward splitting. *SIAM Multiscale Model. Simul.*, 4:1168–1200, 2005.
6. H. W. Engl, M. Hanke, and A. Neubauer. *Regularization of Inverse Problems*, volume 375 of *Mathematics and Its Applications*. Springer Dordrecht, 1996.
7. R. Fletcher. *Practical Methods of Optimization*. John Wiley & Sons, Ltd, 2000.
8. C. A. Floudas and P. M. Pardalos. *Collection of Test Problems for Constrained Global Optimization Algorithms*, volume 455 of *Lecture Notes in Computer Science*. Springer Berlin, Heidelberg, 1990.
9. M. L. N. Gonçalves and T. C. Menezes. Gauss–newton methods with approximate projections for solving constrained nonlinear least squares problems. *Journal of Complexity*, 58:101459, 2020.
10. C. Li, W. Zhang, and X. Jin. Convergence and uniqueness properties of Gauss-Newton's method. *Comput. Optim. Appl.*, 47:1057–1067, 2004.
11. S. Salzo and S. Villa. Convergence analysis of a proximal Gauss-Newton method. *Comput. Optim. Appl.*, 53:557–589, 2012.
12. S. Salzo and S. Villa. Inexact and accelerated proximal point algorithms. *J. Convex Anal.*, 19:1167–1192, 2012.

13. O. Scherzer, M. Grasmair, H. Grossauer, M. Haltmeier, and F. Lenzen. *Variational Methods in Imaging*, volume 167 of *Applied Mathematical Sciences*. Springer, New York, 2009.
14. Bin Zhou, Li Gao, and Yu-Hong Dai. Gradient methods with adaptive step-sizes. *Computational Optimization and Applications*, 35(1):69–86, 2006.