Local Convergence Analysis of an Inexact Trust-Region Method for Nonsmooth Optimization

Robert J. Baraldi[1†] and Drew P. Kouri[1*†]

[1]Optimization and Uncertainty Quantification, Sandia National Laboratories, P.O. Box 5800, Albuquerque, 87125, NM, USA.

*Corresponding author(s). E-mail(s): dpkouri@sandia.gov;
Contributing authors: rjbaral@sandia.gov;
†These authors contributed equally to this work.

Abstract

In [R. J. Baraldi and D. P. Kouri, *Mathematical Programming*, (2022), pp. 1–40], we introduced an inexact trust-region algorithm for minimizing the sum of a smooth nonconvex function and a nonsmooth convex function in Hilbert space—a class of problems that is ubiquitous in data science, learning, optimal control, and inverse problems. This algorithm has demonstrated excellent performance and scalability with problem size. In this paper, we enrich the convergence analysis for this algorithm, proving strong convergence of the iterates with guaranteed rates. In particular, we demonstrate that the trust-region algorithm recovers superlinear, even quadratic, convergence rates when using a second-order Taylor approximation of the smooth objective function term.

Keywords: Nonsmooth Optimization, Nonconvex Optimization, Trust Region, Newton's Method, Superlinear Convergence, Quadratic Convergence

MSC Classification: 49M15 , 49M37 , 65K05 , 65K10 , 90C06 , 90C30

1 Introduction

We provide a comprehensive convergence analysis of the inexact trust-region method introduced in [1, Alg. 1] for solving the optimization problem

$$\min_{x \in X} \; f(x) + \phi(x), \tag{1}$$

where X is a real Hilbert space, $f : X \to \mathbb{R}$ is a smooth (possibly nonconvex) function, and $\phi : X \to [-\infty, +\infty]$ is a convex (possibly nonsmooth) function. Our analysis includes the strong convergence of the iterates and the derivation of convergence rates. As shown in [1, Th. 3], the trust-region method is guaranteed to converge even when the smooth objective function f and its gradient are computed inexactly. This feature is essential for the numerical solution of infinite-dimensional optimization problems, such as those governed by partial differential equations (PDEs), since approximations and discretizations are unavoidable [2–10]. As in [1], we assume that ϕ and its proximity operator can be computed exactly. Although this limits the problem setting covered by (1), relaxing this assumption raises fundamental complications with the analysis in [1]. However, when ϕ-related quantities are exactly computable, the trust-region method exhibits excellent numerical performance, especially on PDE-constrained optimization problems, where we observe mesh-independent behavior, cf. [1, Sect. 5.2].

Employing first-order necessary optimality conditions, one can reformulate (1) as the generalized equation: find $x \in X$ satisfying

$$0 \in f'(x) + \partial\phi(x), \tag{2}$$

where f' is the Fréchet derivative of f and $\partial\phi$ is the convex subdifferential of ϕ. In the seminal reports [11, 12], Josephy established convergence rates of Newton and quasi-Newton methods for solving the finite-dimensional generalized equation

$$0 \in h(x) + H(x), \tag{3}$$

where h is a smooth vector-valued function and H is the normal cone to a convex set. Josephy's Newton-type methods for (3) generate iterates by solving the auxiliary generalized equation

$$0 \in h(x_k) + B_k(x - x_k) + H(x), \tag{4}$$

where x_k is the current iterate and B_k is the Jacobian of h at x_k or an approximation thereof. Under the assumption of strong metric regularity, Josephy showed that the iterates generated by solving (4), with B_k chosen to be the Jacobian of h at x_k, converge q-quadratically, while the iterates of his quasi-Newton method converge q-superlinearly as long as B_k satisfies the Dennis-Moré condition [13, 14]. The authors in [15, 16], see also [17], extend Josephy's methods to infinite-dimensional Banach space, allowing for more general set-valued maps H.

Numerous authors have enhanced Josephy's pioneering work to handle inexact step computations and to weaken the differentiability and metric regularity assumptions, cf. [18–21]. For example, the authors of [21] propose an inexact Newton method for finite-dimensional generalized equations in which the iterates x_{k+1} satisfy

$$(h(x_k) + B_k(x_{k+1} - x_k) + H(x_{k+1})) \cap R_k(x_k, x_{k+1}) \neq \emptyset, \tag{5}$$

where R_k is a set-valued map that encapsulates the inexactness. They assume that $h + H$ is semistable at a reference solution—a property that is related to strong metric regularity as used in [11]—to prove q-superlinear, even q-quadratic, convergence under specific assumptions on R_k. In [19], the authors extend the inexact Newton method introduced in [21] to infinite dimensions and replace the semistability assumption with strong metric subregularity to achieve analogous convergence results. On the other hand, the authors of [20] study (5) in finite dimensions, where they assume that h is semismooth, but not necessarily differentiable. Under this assumption, along with strong metric regularity, the authors prove that the inexact Newton and quasi-Newton methods converge at a q-linear or q-superlinear rate. To achieve a q-quadratic rate, h is required to be strongly semismooth. These results were extended in [18] to the Banach space setting.

The well-known proximal (quasi-)Newton method [22, 23] is a specific realization of Josephy's (quasi-)Newton method applied to (2). The authors in [23] provide a thorough convergence analysis for the proximal (quasi-)Newton method in finite dimensions, globalized with a linesearch. In particular, they demonstrate that proximal (quasi)-Newton methods converge q-superlinearly, and that inexact proximal Newton methods also converge q-superlinearly for specific subproblem stopping tolerances. However, these results require that both f and ϕ are convex—an assumption that was later relaxed in [22]. In contrast with [22], the authors in [24] analyze a proximal Newton method for (1) with nonconvex f and ϕ restricted to the L^1 penalty.

We establish similar convergence rates for the trust-region algorithm introduced in [1, Alg. 1]. Like [22–24], we show q-linear, q-superlinear and q-quadratic rates when using second-order Taylor approximations of the smooth objective function f for the trust-region subproblem model. In Section 2, we describe the assumptions on the problem data in (1) as well as on the trust-region subproblem model. In Section 3, we review the trust-region algorithm from [1] and in Section 4, we provide assumptions on the trial step computation that are required to obtain convergence rates. We present the convergence analysis in Section 5, proving strong convergence of the trust-region iterates along with convergence rates.

2 Notation and Problem Assumptions

We now introduce a list of assumptions on the problem data f and ϕ in (1) as well as on the approximations of f used in the trust-region subproblems.

Throughout, X denotes a real Hilbert space with inner product $\langle \cdot, \cdot \rangle$ and norm $\| \cdot \|$. We associate X with its dual space X^* and denote the Riesz representer of the Fréchet derivative of a function $w : X \to \mathbb{R}$ (i.e., the gradient of w) by ∇w. To simplify the presentation, we employ the following notation. The effective domain of ϕ is $\operatorname{dom} \phi := \{x \in X \,|\, \phi(x) < +\infty\}$. Moreover, we denote the proximal gradient operator associated with ϕ and f by

$$G(x,t) := \frac{1}{t}(x - \operatorname{Prox}_{t\phi}(x - t\nabla f(x)))$$

and its norm by $h(x,t) := \|G(x,t)\|$ for all $x \in X$ and $t > 0$. Here, $\operatorname{Prox}_{t\phi}(\cdot)$ denotes the usual proximity operator,

$$\operatorname{Prox}_{t\phi}(x) := \arg\min_{y \in X} \left\{ \frac{1}{2t}\|y - x\|^2 + \phi(y) \right\}.$$

Recall that if $h(x,t) = 0$ for some $t > 0$, then $h(x,t) = 0$ for all $t > 0$ and x is a critical point of (1) [1, Lem. 5]. We further denote by $\mathcal{L}(X)$, the space of continuous linear operators that map X into itself and for $A \in \mathcal{L}(X)$, A^* denotes the adjoint of A.

The results in the subsequent sections do not depend on all of the following assumptions. As such, we explicitly state the required assumptions for each result.

P1 *The function $\phi : X \to [-\infty, +\infty]$ is proper, closed and convex.*

P2 *The function $f : X \to \mathbb{R}$ is L_1-smooth on* $\operatorname{dom} \phi$. *That is, there exists an open set $V \subseteq X$, containing* $\operatorname{dom} \phi$, *on which f is Fréchet differentiable and its gradient ∇f is Lipschitz continuous with modulus $L_1 > 0$.*

P3 *The objective function $F := f + \phi$ is bounded below.*

P4 *The function f is locally strongly convex. In particular, there exist a positive constant $\lambda > 0$ and a nonempty convex set $U \subseteq V$ satisfying*

$$\langle \nabla f(x) - \nabla f(y), x - y \rangle \geq \lambda \|x - y\|^2 \quad \forall x, y \in U.$$

P5 *Assumption P4 holds and there exists $\bar{x} \in U$ at which $h(\bar{x}, t) = 0$ for all $t > 0$.*

P6 *The function f is twice Fréchet differentiable on V and its Hessian $\nabla^2 f$ is Lipschitz continuous with modulus $L_2 > 0$.*

At each iteration of the trust-region algorithm, we compute a trial iterate $x_k^+ \in X$ that approximately solves the trust-region subproblem

$$\min_{x \in X} \{m_k(x) := f_k(x) + \phi(x)\} \qquad \text{subject to} \qquad \|x - x_k\| \le \Delta_k, \tag{6}$$

where $x_k \in \operatorname{dom} \phi$ is the current iterate, f_k is a smooth local model of f around x_k, and $\Delta_k > 0$ is the trust-region radius. Analogous to G, we define $G_k(x,t)$ with f replaced by f_k and similarly define $h_k := \|G_k(x_k, r_0)\|$, where $r_0 > 0$ is fixed. As with the problem data assumptions, the following model assumptions are not required for all results.

M1 *The function $f_k : X \to \mathbb{R}$ is M_k-smooth on $\operatorname{dom} \phi$ for all k. We denote the associated open set on which f_k is Lipschitz continuously differentiable by $V_k \subseteq X$.*

M2 *There exists a constant $\kappa_{\text{grad}} \ge 0$ such that $g_k := \nabla f_k(x_k)$ satisfies*

$$\|g_k - \nabla f(x_k)\| \le \kappa_{\text{grad}} \min\{h_k, \Delta_k\} \quad \forall\, k.$$

M3 *There exists a constant $\kappa_{\text{curv}} > 0$ such that $M_k \le \kappa_{\text{curv}}$ for all k.*

M4 *The function f_k is locally strongly convex for k sufficiently large. In particular, there exist a positive integer K_0, a positive constant $m > 0$, and nonempty convex sets $U_k \subseteq V_k$ satisfying*

$$\langle \nabla f_k(x) - \nabla f_k(y), x - y\rangle \ge m\|x - y\|^2 \quad \forall\, x,\, y \in U_k, \quad \forall\, k \ge K_0.$$

M5 *The gradient of f_k satisfies $\nabla f_k(x_k) = \nabla f(x_k)$ for $k = 1, 2, \ldots$ and for any sequence $\{v_k\} \in V$ with $\|v_k - x_k\| \to 0$, we have that*

$$\lim_{k \to \infty} \frac{\|\nabla f_k(v_k) - \nabla f(v_k)\|}{\|v_k - x_k\|} = 0.$$

Assumptions P1–P3 and M1–M2 are the basic conditions required to prove global convergence of [1, Alg. 1], cf. [1, Th. 3]. Although many of these assumptions can be relaxed to hold on some open, convex set containing the iterates $\{x_k\}$ and trial iterates $\{x_k^+\}$, we will work with the stated assumptions to maintain consistency with [1]. In the forthcoming analysis, we use Assumptions P4–P5 and M3 to prove strong convergence of the trust-region iterates $\{x_k\}$. Moreover, we use P6 and M4–M5 to obtain convergence rates. Note that assumption M5 requires that the model gradient is exact at x_k, i.e., $g_k = \nabla f(x_k)$ and so M2 is automatically satisfied. Additionally, when the model f_k is the common quadratic model

$$f_k(x) = \frac{1}{2}\langle B_k(x - x_k), x - x_k\rangle + \langle g_k, x - x_k\rangle, \tag{7}$$

where $B_k = B_k^* \in \mathcal{L}(X)$, Assumption M5 is closely related to the Dennis-Moré condition [13, 14]

$$\lim_{k\to\infty} \frac{\|(B_k - \nabla^2 f(\bar{x}))(x_{k+1} - x_k)\|}{\|x_{k+1} - x_k\|} = 0. \tag{8}$$

The following result extends [23, Lem. 3.8] to infinite dimensions when Assumption P6 holds and $B_k = \nabla^2 f(x_k)$.

Lemma 1 *Suppose assumptions P1 and P6 hold, f_k is given by the quadratic model (7), $g_k = \nabla f(x_k)$, and $B_k = \nabla^2 f(x_k)$. Then, M5 holds and*

$$\|G(x,t) - G_k(x,t)\| \le \frac{L_2}{2}\|x - x_k\|^2 \quad \forall\, x \in V, \quad \forall\, t > 0.$$

Proof Fix $x \in V$ and $v \in X$. Applying the fundamental theorem of calculus to $\psi(t) = \langle \nabla f(x_k + t(x - x_k)), v\rangle$ yields

$$\psi(1) - \psi(0) = \langle \nabla f(x) - \nabla f(x_k), v\rangle = \int_0^1 \langle \nabla^2 f(x_k + t(x - x_k))(x - x_k), v\rangle \,\mathrm{d}t$$

and consequently

$$\begin{aligned} |\langle \nabla f(x) - \nabla f_k(x), v\rangle| &= |\langle \nabla f(x) - \nabla f(x_k) - \nabla^2 f(x_k)(x - x_k), v\rangle| \\ &\le \int_0^1 |\langle (\nabla^2 f(x_k + t(x - x_k)) - \nabla^2 f(x_k))(x - x_k), v\rangle| \,\mathrm{d}t \\ &\le \frac{L_2}{2}\|v\|\|x - x_k\|^2. \end{aligned} \tag{9}$$

Maximizing both sides of the above inequality with respect to $v \in X$ with $\|v\| = 1$ yields the inequality

$$\|\nabla f(x) - \nabla f_k(x)\| = \|\nabla f(x) - \nabla f(x_k) - \nabla^2 f(x_k)(x - x_k)\| \le \frac{L_2}{2}\|x - x_k\|^2. \tag{10}$$

Substituting x with the elements of a sequence $\{v_k\} \subset X$ satisfying $\|v_k - x_k\| \to 0$ in (10) proves that M5 is satisfied. For the final result, recall that the proximity operator is nonexpansive [25, Cor. 23.11(i)] (i.e., Lipschitz continuous with unit modulus). Therefore,

$$\|G(x,t) - G_k(x,t)\| \le \|\nabla f(x) - \nabla f(x_k) - \nabla^2 f(x_k)(x - x_k)\|$$

and the result follows from (10). □

We conclude this section with a technical lemma regarding strong monotonicity of the maps $x \mapsto G(x,t)$ and $x \mapsto G_k(x,t)$ for sufficiently small $t > 0$. This lemma is closely related to [23, Lem. 3.9] and [24, Lem. 4.1], which are restricted to finite dimensions. In Section 5, we use this result to help prove convergence rates for the trust-region iterates $\{x_k\}$.

Lemma 2 *Let assumptions P1, P2 and P4 hold. Then for fixed $t \in (0, 2\lambda L_1^{-2})$, the map $x \mapsto G(x,t)$ is strongly monotone on U with parameter $(\lambda - \frac{1}{2}tL_1^2)$, i.e.,*

$$\langle G(x,t) - G(y,t), x - y\rangle \geq (\lambda - \tfrac{1}{2}tL_1^{-2})\|x - y\|^2 \qquad \forall\, x,\, y \in U.$$

On the other hand, if P1, M1 and M4 hold, then $x \mapsto G_k(x,t)$ is strongly monotone on U_k with parameter $(m - \frac{1}{2}tM_k^2)$ for any $t \in (0, 2mM_k^{-2})$.

Proof Applying Moreau's decomposition [25, Th. 14.3(ii)] to G yields

$$G(x,t) = \nabla f(x) + \frac{1}{t}\mathrm{Prox}_{(t\phi)^*}(x - t\nabla f(x))$$

and thus

$$\begin{aligned} G(x,t) - G(y,t) = {}& (\nabla f(x) - \nabla f(y)) \\ &+ \frac{1}{t}(\mathrm{Prox}_{(t\phi)^*}(x - t\nabla f(x)) - \mathrm{Prox}_{(t\phi)^*}(y - t\nabla f(y))) \end{aligned}$$

for arbitrary, fixed $x,\, y \in U$. In the subsequent arguments, we use the notation

$$d \coloneqq (x - t\nabla f(x)) - (y - t\nabla f(y)) = (x - y) - t(\nabla f(x) - \nabla f(y))$$

and

$$w \coloneqq \mathrm{Prox}_{(t\phi)^*}(x - t\nabla f(x)) - \mathrm{Prox}_{(t\phi)^*}(y - t\nabla f(y)).$$

Note that if $w = 0$, then P4 ensures that $\langle G(x,t) - G(y,t), x - y\rangle \geq \lambda\|x - y\|^2$ for all $t > 0$. As such, we assume that $w \neq 0$ and employ the self-adjoint linear operator $W \in \mathcal{L}(X)$ defined by

$$Wv = \frac{\langle w, v\rangle}{\langle w, d\rangle} w \qquad \forall\, v \in X.$$

Since $\mathrm{Prox}_{(t\phi)^*}(\cdot)$ is nonexpansive, we have that $\|w\|^2 \leq \langle w, d\rangle$ and the denominator in the definition of W is positive. These definitions allow use to write

$$G(x,t) - G(y,t) = (\nabla f(x) - \nabla f(y)) + \frac{1}{t}Wd.$$

Computing the inner product of this quantity with $(x - y)$ yields

$$\begin{aligned} \langle G(x,t) - G(y,t), x - y\rangle &= \langle \nabla f(x) - \nabla f(y), x - y\rangle + \frac{1}{t}\langle Wd, x - y\rangle \\ &= \langle (I - W)(\nabla f(x) - \nabla f(y)), x - y\rangle \\ &\quad + \frac{1}{t}\langle W(x - y), x - y\rangle. \end{aligned}$$

By the Fenchel-Young inequality, we have that

$$\langle W(\nabla f(x) - \nabla f(y)), x - y\rangle \leq \frac{t}{2}\|\nabla f(x) - \nabla f(y)\|^2 + \frac{1}{2t}\|W(x - y)\|^2,$$

which yields the bound

$$\begin{aligned} \langle G(x,t) - G(y,t), x - y\rangle \geq {}& \left(\langle \nabla f(x) - \nabla f(y), x - y\rangle - \frac{t}{2}\|\nabla f(x) - \nabla f(y)\|^2\right) \\ &+ \frac{1}{t}\left(\langle W(x - y), x - y\rangle - \frac{1}{2}\|W(x - y)\|^2\right). \end{aligned}$$

Expanding the second parenthetical term in the above lower bound and recalling that $\|w\|^2 \leq \langle w, d\rangle$ yields

$$\begin{aligned}\langle W(x-y), x-y\rangle - \frac{1}{2}\|W(x-y)\|^2 &= \frac{\langle w, x-y\rangle^2}{\langle w, d\rangle^2}\left(\langle w, d\rangle - \frac{1}{2}\|w\|^2\right) \\ &\geq \frac{1}{2}\frac{\langle w, x-y\rangle^2}{\langle w, d\rangle^2}\|w\|^2 \geq 0\end{aligned}$$

and consequently,

$$\langle G(x,t) - G(y,t), x-y\rangle \geq \langle \nabla f(x) - \nabla f(y), x-y\rangle - \frac{t}{2}\|\nabla f(x) - \nabla f(y)\|^2.$$

By P2 and P4, we have that

$$\langle G(x,t) - G(y,t), x-y\rangle \geq \lambda\|x-y\|^2 - \frac{tL_1^2}{2}\|x-y\|^2 = \left(\lambda - \frac{tL_1^2}{2}\right)\|x-y\|^2.$$

The lower bound is positive for $t \in (0, 2\lambda L_1^{-2})$ and hence $G(\cdot, t)$ is strongly monotone for $t \in (0, 2\lambda L_1^{-2})$. We omit the proof for G_k as it is identical to the proof for G. □

3 Trust-Region Algorithm

As discussed in Section 2, the trial iterate x_k^+ is only required to be an approximate solution to the trust-region subproblem (6). In particular, we require that x_k^+ satisfies the following two assumptions.

A1 *There exists a constant* $\kappa_{\text{rad}} > 0$ *such that*

$$x_k^+ \in \mathcal{B}_k := \{x \in X \mid \|x - x_k\| \leq \kappa_{\text{rad}}\Delta_k\} \quad \forall\, k.$$

A2 *There exists a constant* $\kappa_{\text{fcd}} > 0$ *such that*

$$m_k(x_k) - m_k(x_k^+) \geq \kappa_{\text{fcd}} h_k \min\left\{\frac{h_k}{1+\omega_k}, \Delta_k\right\} \quad \forall\, k,$$

where ω_k *measures the curvature of* f_k *over the trust region and is defined by*

$$\omega_k := \sup\left\{ \frac{2}{\|s\|^2}|f_k(x_k+s) - f_k(x_k) - \langle g_k, s\rangle| \;\middle|\; 0 < \|s\| \leq \kappa_{\text{rad}}\Delta_k \right\}.$$

Given a trial iterate x_k^+ that satisfies A1 and A2, traditional trust-region algorithms decide whether or not to accept x_k^+ based on the ratio of actual and predicted reduction

$$\rho_k^* := \frac{\text{ared}_k}{\text{pred}_k},$$

where

$$\text{ared}_k := F(x_k) - F(x_k^+) \qquad \text{and} \qquad \text{pred}_k := m_k(x_k) - m_k(x_k^+).$$

Here, ared_k is the reduction of the objective function F achieved by x_k^+ relative to x_k and pred_k is the reduction of the model m_k. In many practical applications, the objective function F cannot be computed accurately [3, 5, 6, 10, 26]. Instead, ared_k is replaced by an approximation denoted cred_k—the *computed reduction*. The algorithm then decides whether or not to accept x_k^+ based on the ratio of computed and predicted reduction

$$\rho_k := \frac{\text{cred}_k}{\text{pred}_k}. \tag{11}$$

If $\rho_k \geq \eta_1$, we set $x_{k+1} = x_k^+$. Otherwise, we set $x_{k+1} = x_k$. The trust-region algorithm then increases the radius Δ_k if $\rho_k \geq \eta_2$ and reduces Δ_k if $\rho_k < \eta_1$. The algorithmic parameters $0 < \eta_1 < \eta_2 < 1$ are user-specified with common values $\eta_1 = 10^{-4}$ and $\eta_2 = 0.75$. To ensure that cred_k is a sufficiently accurate approximation of ared_k, we require the following assumption.

A3 *There exists a constant $\kappa_{\text{obj}} \geq 0$ such that*

$$|\text{ared}_k - \text{cred}_k| \leq \kappa_{\text{obj}}[\eta \min\{\text{pred}_k, \theta_k\}]^\zeta \quad \forall\, k,$$

where ζ, η, and θ_k are (user-specified) positive real numbers that satisfy

$$\zeta > 1, \qquad 0 < \eta < \min\{\eta_1, (1 - \eta_2)\}, \qquad \textit{and} \qquad \lim_{k\to\infty} \theta_k = 0.$$

Assumption A3 was originally used in [5], where it was motivated by [9, Sect. 5.3.3]. The forcing sequence $\{\theta_k\}$ in A3 enables the inclusion of the arbitrary constant κ_{obj} within the error bound—a feature that is indispensable when using error indicators that depend on uncomputable quantities. In particular, the inclusion of $\{\theta_k\}$ in Assumption A3 ensures that there exists a positive integer K_η for which $|\text{ared}_k - \text{cred}_k| \leq \eta\, \text{pred}_k$ for all $k \geq K_\eta$ (cf. [1, Lem. 6]), which is closely related to the inexactness condition described in [26, Ch. 10.6]. We state the trust-region algorithm as Algorithm 1. As shown in [1], if A1–A3, M1–M2 and P1–P3 hold and if

$$\sum_{k=1}^{\infty} (1 + \max\{\omega_i \,|\, i = 1, \ldots, k\})^{-1} = +\infty, \tag{12}$$

then the iterates generated by Algorithm 1 satisfy

$$\liminf_{k\to\infty} h_k = 0 \qquad \text{and} \qquad \liminf_{k\to\infty} h(x_k, r_0) = 0.$$

Note that if M3 holds, then (12) holds. Our task now is to employ assumptions M3–M5 and P4–P6 to arrive at stronger convergence results.

Algorithm 1 Nonsmooth Trust-Region Algorithm

Require: Initial guess $x_1 \in \operatorname{dom}\phi$, initial radius $\Delta_1 > 0$, $0 < \eta_1 < \eta_2 < 1$, and $0 < \gamma_1 \le \gamma_2 < 1 \le \gamma_3$

for $k = 1, 2, \ldots$ **do**
  Model Selection: Choose a model f_k that satisfies M1 and M2
  Step Computation: Compute $x_k^+ \in X$ that satisfies A1 and A2
  Computed Reduction: Compute cred_k that satisfies A3
  Step Acceptance and Radius Update: Compute ρ_k as in (11)
  if $\rho_k < \eta_1$ **then**
    $x_{k+1} \leftarrow x_k$
    $\Delta_{k+1} \in [\gamma_1 \Delta_k, \gamma_2 \Delta_k]$
  else
    $x_{k+1} \leftarrow x_k^+$
    if $\rho_k \in [\eta_1, \eta_2)$ **then**
      $\Delta_{k+1} \in [\gamma_2 \Delta_k, \Delta_k]$
    else
      $\Delta_{k+1} \in [\Delta_k, \gamma_3 \Delta_k]$
    end if
  end if
end for

4 Trial Step Computation

The numerical implementation of Algorithm 1 from [1] approximately solves the trust-region subproblem by first computing a Cauchy point [1, Sect. 3.1] and then improving upon the Cauchy point using the spectral proximal gradient method with a monotonic line search [1, Alg. 5]. With this as motivation, we assume that each trial step x_k^+ is computed by a procedure that generates finitely many iterates $\{x_{k,0}, x_{k,1}, \ldots, x_{k,n_k}\}$ with $x_{k,0} = x_k$ and $x_k^+ = x_{k,n_k}$, where $n_k \le$ `maxit` for all k and `maxit` is a (user-specified) positive integer. Moreover, the iterates satisfy the following assumptions.

S1 *There exists a constant $\mu \in (0,1)$, independent of k, such that*

$$m_k(x_{k,j+1}) \le m_k(x_{k,j}) + \mu Q_{k,j} \qquad \text{and} \qquad x_{k,j+1} \in \mathcal{B}_k$$

for $j = 0, \ldots, n_k - 1$, where

$$Q_{k,j} := \langle \nabla f_k(x_{k,j}), x_{k,j+1} - x_{k,j} \rangle + \phi(x_{k,j+1}) - \phi(x_{k,j}).$$

S2 *The final iterate $x_k^+ = x_{k,n_k}$ satisfies at least one of the stopping conditions:*

$$\|x_k^+ - x_k\| = \kappa_{\mathrm{rad}} \Delta_k \tag{S2.1}$$

$$\|G_k(x_k^+, r_0)\| \le \tau_k h_k, \tag{S2.2}$$

where $\{\tau_k\} \subset [0, \infty)$ is a bounded sequence of relative stopping tolerances.

S3 *If there exists $x_k^{\mathrm{n}} \in \mathcal{B}_k$ that satisfies (S2.2), then x_k^+ also satisfies (S2.2).*

The quantity $Q_{k,j}$ in S1 is a commonly used measure of sufficient decrease for problems of the form (1) [22, 27, 28] and provides an upper bound on the directional derivative of m_k. The subproblem stopping condition (S2.2) is closely related to stopping conditions used for inexact Newton methods, cf., [22–24, 29]. Assumptions S1 and S2 are satisfied by [1, Alg. 5] and provide guidance for developing algorithms to compute trial iterates that satisfy A1 and A2. Assumption S3 ensures that that any point generated by the subproblem solver eventually behaves like an inexact Newton step. Similar assumptions exist in the trust-region literature. For smooth linearly-constrained problems, [30] assumes that if $\|x_k^+ - x_k\| \le \kappa_{\mathrm{rad}}^* \Delta_k$ for fixed $\kappa_{\mathrm{rad}}^* \in (0, \kappa_{\mathrm{rad}})$, then x_k^+ satisfies (S2.2). An analogous assumption is used in [31, Th. 4.9] for smooth unconstrained problems. Although concrete subproblem solvers are beyond the scope of this paper, one can satisfy S3 using a dogleg method [32] by requiring that the Newton step used to construct the dogleg path satisfies (S2.2). Assumptions S1-S3 are not required to prove global convergence of Algorithm 1, but are required to prove that the trial iterates x_k^+ are eventually always accepted.

5 Local Convergence Analysis

Our first result is common for trust-region methods (e.g., [26, Th. 6.4.6] is a similar result for smooth, unconstrained optimization) and strengthens the result in [1, Th. 3] under the additional assumption M3. In particular, we show that the limit, not just the lower limit, of h_k is zero.

Theorem 1 *Suppose assumptions P1–P3, M1–M3, and A1–A3 are satisfied. Let $\{x_k\} \subset X$ denote the iterates generated by Algorithm 1. Then,*

$$\lim_{k\to\infty} h_k = 0 \qquad \textit{and} \qquad \lim_{k\to\infty} h(x_k, t) = 0,$$

for any fixed $t > 0$.

Proof The proof of this result mirrors that of [26, Th. 6.4.6] with modifications to account for inexact objective function values A3 and gradients M2. Let $\mathcal{S} := \{k \in \mathbb{N} \,|\, \rho_k \ge \eta_1\}$ denote the indices of the successful iterations. We assume that $\mathcal{S}$ is infinite since otherwise the result would follow from [1, Cor. 4]. To arrive at a contradiction, we assume that there exists $\epsilon > 0$ and a subsequence $\{k_i\} \subseteq \mathcal{S}$ satisfying

$$h_{k_i} \ge 2\epsilon > 0 \tag{14}$$

for all i. By [1, Th. 3], for each i there exists an iteration index $\ell > k_i$ for which $h_\ell < \epsilon$. Let $\ell_i > k_i$ denote the first such index and define

$$\mathcal{K} := \{k \in \mathcal{S} \,|\, k_i \le k < \ell_i \;\; \forall\, i\}.$$

Note that for all $k \in \mathcal{K}$, we have that $h_k \ge \epsilon$. By [1, Lem. 6], there exists a positive integer K_η such that $|\rho_k^* - \rho_k| \le \eta$ for all $k \ge K_\eta$. This and assumptions A2 and M3 ensure that

$$F(x_k) - F(x_{k+1}) \ge (\eta_1 - \eta)\kappa_{\mathrm{fcd}}\epsilon \min\left\{\frac{\epsilon}{1+\kappa_{\mathrm{curv}}}, \Delta_k\right\}$$

for all $k \in \mathcal{K}$ with $k \geq K_\eta$. Since $\{F(x_k)\}$ is monotonically decreasing and bounded below by P3, we have that the left-hand side converges to zero and hence so does Δ_k. Consequently, for sufficiently large $k \in \mathcal{K}$, we obtain that

$$\Delta_k \leq \frac{1}{\kappa_{\mathrm{fcd}}\epsilon(\eta_1 - \eta)}[F(x_{k+1}) - F(x_k)].$$

For i sufficiently large, this and the triangle inequality ensure that

$$\|x_{k_i} - x_{\ell_i}\| \leq \sum_{\substack{j=k_i\\ j\in\mathcal{K}}}^{\ell_i - 1} \|x_j - x_{j+1}\| \leq \kappa_{\mathrm{rad}} \sum_{\substack{j=k_i\\ j\in\mathcal{K}}}^{\ell_i - 1} \Delta_j \leq \frac{\kappa_{\mathrm{rad}}}{\kappa_{\mathrm{fcd}}\epsilon(\eta_1 - \eta)}[F(x_{k_i}) - F(x_{\ell_i})].$$

Again, owing to the monotonicity of $\{F(x_k)\}$ and P3, we have that the right-hand side above tends to zero and therefore so does $\|x_{k_i} - x_{\ell_i}\|$ as i goes to infinity. By P2, we also have that $\|\nabla f(x_{k_i}) - \nabla f(x_{\ell_i})\|$ tends to zero as i goes to infinity. Now, P2, M2, and the nonexpansivity of the proximity operator ensure that

$$\begin{aligned}
\epsilon &\leq |h_{k_i} - h_{\ell_i}| \\
&\leq \frac{2}{r_0}\|x_{k_i} - x_{\ell_i}\| + \|g_{k_i} - \nabla f(x_{k_i})\| + \|\nabla f(x_{k_i}) - \nabla f(x_{\ell_i})\| + \|\nabla f(x_{\ell_i}) - g_{\ell_i}\| \\
&\leq \left(\frac{2}{r_0} + L_1\right)\|x_{k_i} - x_{\ell_i}\| + \kappa_{\mathrm{grad}}(\Delta_{k_i} + \gamma_3\Delta_{\ell_i - 1}).
\end{aligned}$$

Since $\{\Delta_k\}_{k\in\mathcal{K}}$ converges to zero, so does the above upper bound, leading to a contradiction. Hence, no subsequence satisfying (14) exists and consequently, $h_k \to 0$. Now, owing to M2, we have that

$$\lim_{k\to\infty} h(x_k, r_0) \leq \lim_{k\to\infty} \|G(x_k, r_0) - G_k(x_k, r_0)\| + h_k \leq \lim_{k\to\infty} (1 + \kappa_{\mathrm{grad}})h_k = 0.$$

Since $h(x, r_0) \geq h(x, t)$ for all $t \geq r_0$ [1, Lem. 2, part 2], we have that $h(x_k, t) \to 0$ for any fixed $t \geq r_0$. On the other hand, since $th(x, t) \leq r_0 h(x, r_0)$ for all $t \in (0, r_0]$ [1, Lem. 2, part 1], we have that $h(x_k, t) \to 0$ for any fixed $t > 0$. □

When compared with [1, Th. 3], Theorem 1 only requires the additional assumption of bounded curvature, M3, which again implies (12). Under additional assumptions on f, namely P4 and P5, we can prove that the sequence of iterates $\{x_k\}$ generated by Algorithm 1 strongly converge to a critical point of (1). This improves upon [1, Cor. 3], which proves that strong accumulation points of $\{x_k\}$ are critical points of (1).

Corollary 1 *Let assumptions P1–P5, M1–M3 and A1–A3 hold, and suppose that there exists a positive integer K_U such that $x_k \in U$ for all $k \geq K_U$. Here, U is defined in P4. Then, $x_k \to \bar{x}$.*

Proof By Theorem 1, we have that $h(x_k, t) \to 0$ for fixed $t \in (0, 2\lambda L_1^{-2})$. Hence, the strong monotonicity of G from Lemma 2 and the criticality of $\bar{x}$ ensure that

$$\left(\lambda - \frac{tL_1^2}{2}\right)\lim_{k\to\infty}\|x_k - \bar{x}\| \leq \lim_{k\to\infty}\|G(x_k, t) - G(\bar{x}, t)\| = \lim_{k\to\infty} h(x_k, t) = 0,$$

proving the desired result. □

Our final three results require the strongest assumptions on the model to ultimately prove convergence rates for $\{x_k\}$. Much like Theorem 1, the next two results are common in the trust-region literature and ultimately demonstrate that Algorithm 1 eventually accepts every trial iterate. See, for example, [31, Th. 4.9] for unconstrained optimization and [30, Th. 5.3] for linearly-constrained optimization.

Theorem 2 *Let P1–P3, M1–M5, A1–A3 and S1 hold. Then, Δ_k is bounded away from zero for all k.*

Proof We first bound $|\rho_k^* - 1|$ and then show that this bound converges to zero if $\|x_k^+ - x_k\| \to 0$. To bound $|\rho_k^* - 1|$, we note that $|\rho_k^* - 1| = |\mathrm{ared}_k - \mathrm{pred}_k|/\mathrm{pred}_k$. The ϕ terms cancel in $|\mathrm{ared}_k - \mathrm{pred}_k|$ and we can bound the numerator as

$$\begin{aligned}|\mathrm{ared}_k - \mathrm{pred}_k| &= |f(x_k) - f(x_k^+) - f_k(x_k) + f_k(x_k^+)| \\ &= \left|\int_0^1 \left\langle \nabla f(x_k + t(x_k^+ - x_k)) - \nabla f_k(x_k + t(x_k^+ - x_k)), x_k^+ - x_k \right\rangle \mathrm{d}t\right| \\ &\le \int_0^1 \|\nabla f(x_k + t(x_k^+ - x_k)) - \nabla f_k(x_k + t(x_k^+ - x_k))\| \,\mathrm{d}t \|x_k^+ - x_k\|,\end{aligned}$$

where the second equality follows from the fundamental theorem of calculus and the inequality follows from Cauchy-Schwarz. By assumption M5, $\nabla f_k(x_k) = \nabla f(x_k)$ and therefore, P2, M1 and M3 ensure that

$$\frac{\|\nabla f(x_k + t(x_k^+ - x_k)) - \nabla f_k(x_k + t(x_k^+ - x_k))\|}{\|x_k^+ - x_k\|} \le t(L_1 + \kappa_{\mathrm{curv}}).$$

Consequently, M5 and the Lebesgue dominated convergence theorem yield

$$\lim_{k\to\infty} \int_0^1 \frac{\|\nabla f(x_k + t(x_k^+ - x_k)) - \nabla f_k(x_k + t(x_k^+ - x_k))\|}{\|x_k^+ - x_k\|} \,\mathrm{d}t = 0$$

if $\|x_k^+ - x_k\| \to 0$ or equivalently,

$$\int_0^1 \|\nabla f(x_k + t(x_k^+ - x_k)) - \nabla f_k(x_k + t(x_k^+ - x_k))\| \,\mathrm{d}t = o(\|x_k^+ - x_k\|).$$

Hence, if $\|x_k^+ - x_k\| \to 0$, then

$$|\mathrm{ared}_k - \mathrm{pred}_k| \le o(\|x_k^+ - x_k\|)\|x_k^+ - x_k\|.$$

To bound the denominator, we consider the sequence of subproblem iterates $\{x_{k,1}, \dots, x_{k,n_k}\}$. We can bound the model decrease achieved at each iteration, using again the fundamental theorem of calculus, as

$$\begin{aligned}&m_k(x_{k,j+1}) - m_k(x_{k,j}) \\ &\quad = Q_{k,j} + \int_0^1 \langle \nabla f_k(x_{k,j} + t(x_{k,j+1} - x_{k,j})) - \nabla f_k(x_{k,j}), x_{k,j+1} - x_{k,j}\rangle \,\mathrm{d}t.\end{aligned}$$

Employing S1, we have that

$$\int_0^1 \langle \nabla f_k(x_{k,j} + t(x_{k,j+1} - x_{k,j})) - \nabla f_k(x_{k,j}), x_{k,j+1} - x_{k,j}\rangle \,\mathrm{d}t \le -(1-\mu)Q_{k,j}.$$

In addition, M4 ensures that

$$\int_0^1 \langle \nabla f_k(x_{k,j}+t(x_{k,j+1}-x_{k,j}))-\nabla f_k(x_{k,j}), x_{k,j+1}-x_{k,j}\rangle \,\mathrm{d}t \geq \frac{m}{2}\|x_{k,j+1}-x_{k,j}\|^2.$$

Combining these results with S1, then yields

$$\begin{aligned} m_k(x_{k,j}) - m_k(x_{k,j+1}) &\geq -\mu Q_{k,j} \geq -\frac{\mu}{1-\mu}(1-\mu)Q_{k,j} \\ &\geq \frac{\mu}{1-\mu}\frac{m}{2}\|x_{k,j+1}-x_{k,j}\|^2 \end{aligned}$$

Define $\kappa_0 := \frac{\mu}{1-\mu}\frac{m}{2}$. Now, we can reuse some arguments from the proof of [30, Th. 5.3]. Specifically, since $n_k \leq \texttt{maxit}$, we have

$$m_k(x_k) - m_k(x_k^+) \geq \kappa_0 \max_{0\leq j\leq n_k}\{\|x_{k,j+1}-x_{k,j}\|^2\}$$

and

$$\|x_k^+ - x_k\| \leq (\texttt{maxit}+1)\max_{0\leq j\leq n_k}\{\|x_{k,j+1}-x_{k,j}\|\},$$

which yield the inequality $\mathrm{pred}_k \geq \kappa_1\|x_k^+ - x_k\|^2$ for $\kappa_1 := \frac{\kappa_0}{(\texttt{maxit}+1)^2}$. Combining the bounds on the numerator and denominator, we arrive at

$$|\rho_k^* - 1| \leq o(\|x_k^+ - x_k\|)\frac{\|x_k^+ - x_k\|}{\kappa_1\|x_k^+ - x_k\|^2} = o(1)$$

and hence, $|\rho_k^* - 1| \to 0$ as long as $\|x_k^+ - x_k\| \to 0$. To conclude, suppose that $\Delta_k \to 0$, then $\|x_k^+ - x_k\| \to 0$ and consequently, $|\rho_k^* - 1| \to 0$. Therefore, after sufficiently many iterations Δ_k is increased, contradicting the assumption that $\Delta_k \to 0$. Hence, there exists $\epsilon > 0$ such that $\Delta_k \geq \epsilon$ for all k. □

Building upon Theorem 2, the next results demonstrates that if S2 and S3 hold, then the trial iterate x_k^+ will satisfy (S2.2) for k sufficiently large. This result is a critical step in obtaining a convergence rate.

Corollary 2 *Let P1–P3, M1–M5, A1–A3 and S1–S3 hold. Suppose that* $x_k^{\mathrm{n}} \in U_k$ *and* $x_{k,i} \in U_k$ *for* $i = 0, 1, \ldots, n_k$ *and* $k \geq K_0$*, where* U_k *and* K_0 *are defined in M4 and* x_k^{n} *is defined in S3. Then, there exists a positive integer* $K_1 \geq K_0$ *such that* x_k^+ *satisfies* (S2.2) *and* $\Delta_{k+1} \geq \Delta_k$ *for all* $k \geq K_1$*.*

Proof By Theorem 2, there exists $\epsilon > 0$ such that $\Delta_k \geq \epsilon$ for all k. Since $x_k^{\mathrm{n}} \in U_k$ for $k \geq K_0$, Lemma 2 ensures that

$$m\|x_k^{\mathrm{n}} - x_k\| \leq \|G_k(x_k^{\mathrm{n}}, t) - G_k(x_k, t)\| \leq \|G_k(x_k^{\mathrm{n}}, t)\| + \|G_k(x_k, t)\|$$

for fixed $t \in (0, 2m\kappa_{\mathrm{curv}}^{-2})$ and all $k \geq K_0$. If $r_0 < 2m\kappa_{\mathrm{curv}}^{-2}$, then we set $t = r_0$. Otherwise, we have that

$$t\|G_k(x_k, t)\| \leq r_0\|G_k(x_k, r_0)\| \quad\iff\quad \|G_k(x_k, t)\| \leq \frac{r_0}{t}\|G_k(x_k, r_0)\|$$

by [1, Lem. 2]. Thus, (S2.2) ensures that

$$\|x_k^{\mathrm{n}} - x_k\| \leq \max\left\{\frac{1}{m}, \frac{r_0}{tm}\right\}(\tau_k + 1)h_k \tag{15}$$

and consequently, by Theorem 1 there exists a positive integer $K_\epsilon \geq K_0$ such that

$$\|x_k^{\mathrm{n}} - x_k\| \leq \epsilon \quad \forall\, k \geq K_\epsilon.$$

By S3, x_k^+ satisfies (S2.2) for all $k \geq K_\epsilon$. The bound (15) further shows that $\|x_k^+ - x_k\| \to 0$ and hence $|\rho_k - 1| \to 0$ as demonstrated in the proof of Theorem 2. The trust-region update mechanism in Algorithm 1 then ensures that there exists $K_1 \geq K_\epsilon$ for which $\Delta_{k+1} \geq \Delta_k$ for all $k \geq K_1$. □

Our final result provides convergence rates for $\{x_k\}$ generated by Algorithm 1, when the trust-region subproblem (6) is solved using the stopping conditions in S2. We note that the former and latter results require M5, which in turn requires that $g_k = \nabla f(x_k)$. This requirement is extremely difficult to overcome when determining convergence rates. In particular, it would seem that any inexact gradient condition that ensures q-sublinear (or faster) convergence may be difficult to enforce in general [33, Ch. 2.3.1].

Theorem 3 *Let P1–P3, M1–M4, A1–A3, and S1–S3 hold. Suppose that x_k^{n}, $x_{k,i} \in U_k$ for $i = 0, 1, \ldots, n_k$ for all $k \geq K_0$ and $x_k \to \bar{x}$, where $\bar{x}$ is a critical point of (1).*

1. *If M5 holds and the relative tolerance τ_k defined in (S2.2) satisfies $\tau_k \to \bar{\tau}$ with*
$$0 < \bar{\tau} < \frac{m}{r_0 L_1 + 1} \min\left\{r_0, \frac{2m}{\kappa_{\mathrm{curv}}^2}\right\},$$
then x_k converges q-linearly to $\bar{x}$.
2. *If M5 holds and $\tau_k \to 0$, then x_k converges q-superlinearly to $\bar{x}$.*
3. *If P6 holds, f_k is the quadratic model (7), $g_k = \nabla f(x_k)$, $B_k = \nabla^2 f(x_k)$, and $\tau_k \leq \tau h_k^{1+\alpha}$ for fixed $\tau > 0$ and $\alpha \geq 0$, then x_k converges q-quadratically to $\bar{x}$.*

Proof Recall that M5 holds in case 3 by Lemma 1. By Corollary 2, we have that $x_{k+1} = x_k^+$ satisfy (S2.2) for $k \geq K_1$. By Lemma 2, M3 and the Cauchy-Schwarz inequality, we have that

$$\|x_{k+1} - \bar{x}\| \leq m^{-1} \|G_k(x_{k+1}, t) - G_k(\bar{x}, t)\|$$

for fixed $t \in (0, 2m\kappa_{\mathrm{curv}}^{-1})$ and $k \geq K_1$. Using (S2.2), we have that

$$\|x_{k+1} - \bar{x}\| \leq m^{-1} \left(\|G_k(\bar{x}, t)\| + \max\left\{1, \frac{r_0}{t}\right\} \tau_k h_k\right),$$

where the upper bound follows from similar arguments as in (15). In cases 1 and 2, M5 and the nonexpansivity of the proximity operator ensure that

$$\|G_k(\bar{x}, t)\| = \|G_k(\bar{x}, t) - G(\bar{x}, t)\| = o(\|x_k - \bar{x}\|).$$

In case 3, Lemma 1 ensures that $\|G_k(\bar{x}, t)\| \leq L_2/2\|x_k - \bar{x}\|^2$. Moreover, P1–P2 and the equality $\nabla f_k(x_k) = \nabla f(x_k)$ ensure that

$$h_k = \|G_k(x_k, r_0) - G(\bar{x}, r_0)\| \leq (L_1 + r_0^{-1})\|x_k - \bar{x}\|,$$

where the inequality follows from the arguments in the proof of [23, Lem. 2]. In cases 1 and 2, these results yield

$$\|x_{k+1} - \bar{x}\| \leq o(\|x_k - \bar{x}\|) + \frac{L_1 + r_0^{-1}}{m} \max\left\{1, \frac{r_0}{t}\right\} \tau_k \|x_k - \bar{x}\|,$$

for fixed $t \in (0, 2m\kappa_{\mathrm{curv}}^{-2})$ and similarly for case 3. The desired results follow directly from this bound and the stated assumptions. □

Acknowledgments. This research was sponsored by the U.S. Department of Energy Office of Science and the U.S. Air Force Office of Scientific Research. This article has been authored by an employee of National Technology & Engineering Solutions of Sandia, LLC under Contract No. DE-NA0003525 with the U.S. Department of Energy (DOE). The employee owns all right, title and interest in and to the article and is solely responsible for its contents. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this article or allow others to do so, for United States Government purposes. The DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan.

Declarations. Data sharing not applicable to this article as no datasets were generated or analyzed during the current study.

References

[1] Baraldi, R.J., Kouri, D.P.: A proximal trust-region method for nonsmooth optimization with inexact function and gradient evaluations. Mathematical Programming, 1–40 (2022)

[2] Clever, D., Lang, J., Ulbrich, S., Ziems, C.: In: Leugering, G., Engell, S., Griewank, A., Hinze, M., Rannacher, R., Schulz, V., Ulbrich, M., Ulbrich, S. (eds.) Generalized Multilevel SQP-methods for PDAE-constrained Optimization Based on Space-Time Adaptive PDAE Solvers, pp. 51–74. Springer, Basel (2012). https://doi.org/10.1007/978-3-0348-0133-1_4

[3] Garreis, S., Ulbrich, M.: An inexact trust-region algorithm for constrained problems in Hilbert space and its application to the adaptive solution of optimal control problems with PDEs. Preprint, submitted, Technical University of Munich (2019)

[4] Kouri, D.P., Heinkenschloss, M., Ridzal, D., van Bloemen Waanders, B.G.: A trust-region algorithm with adaptive stochastic collocation for PDE optimization under uncertainty. SIAM Journal on Scientific Computing **35**(4), 1847–1879 (2013)

[5] Kouri, D.P., Heinkenschloss, M., Ridzal, D., van Bloemen Waanders, B.G.: Inexact objective function evaluations in a trust-region algorithm for PDE-constrained optimization under uncertainty. SIAM Journal on Scientific Computing **36**(6), 3011–3029 (2014)

[6] Kouri, D.P., Ridzal, D.: Inexact trust-region methods for PDE-constrained optimization. In: Frontiers in PDE-Constrained Optimization, pp. 83–121. Springer, New York, NY (2018)

[7] Muthukumar, R., Kouri, D.P., Udell, M.: Randomized sketching algorithms for low-memory dynamic optimization. SIAM Journal on Optimization **31**(2), 1242–1275 (2021)

[8] Zahr, M.J., Carlberg, K.T., Kouri, D.P.: An efficient, globally convergent method for optimization under uncertainty using adaptive model reduction and sparse grids. SIAM/ASA Journal on Uncertainty Quantification **7**(3), 877–912 (2019)

[9] Ziems, J.C., Ulbrich, S.: Adaptive multilevel inexact SQP methods for PDE-constrained optimization. SIAM Journal on Optimization **21**(1), 1–40 (2011). https://doi.org/10.1137/080743160

[10] Zou, Z., Kouri, D.P., Aquino, W.: A locally adapted reduced-basis method for solving risk-averse PDE-constrained optimization problems. SIAM/ASA Journal on Uncertainty Quantification **10**(4), 1629–1651 (2022)

[11] Josephy, N.H.: Newton's method for generalized equations. Technical report, Wisconsin Univ-Madison Mathematics Research Center (1979)

[12] Josephy, N.H.: Quasi-Newton methods for generalized equations. Technical report, Wisconsin Univ-Madison Mathematics Research Center (1979)

[13] Dennis, J.E., Moré, J.J.: A characterization of superlinear convergence and its application to quasi-Newton methods. Mathematics of computation **28**(126), 549–560 (1974)

[14] Dontchev, A.L.: Generalizations of the Dennis–Moré theorem. SIAM Journal on Optimization **22**(3), 821–830 (2012)

[15] Aragón Artacho, F.J., Belyakov, A., Dontchev, A.L., López, M.: Local convergence of quasi-Newton methods under metric regularity. Computational Optimization and Applications **58**(1), 225–247 (2014)

[16] Dontchev, A.L.: Local convergence of the Newton method for generalized equations. Comptes rendus de l'Académie des sciences. Série 1, Mathématique **322**(4), 327–331 (1996)

[17] Dontchev, A.L., Rockafellar, R.T.: Implicit Functions and Solution Mappings vol. 543. Springer, ??? (2009)

[18] Cibulka, R., Dontchev, A., Geoffroy, M.H.: Inexact Newton methods and Dennis–Moré theorems for nonsmooth generalized equations. SIAM Journal on Control and Optimization **53**(2), 1003–1019 (2015)

[19] Dontchev, A.L., Rockafellar, R.T.: Convergence of inexact Newton methods for generalized equations. Mathematical Programming **139**(1), 115–137 (2013)

[20] Izmailov, A.F., Kurennoy, A.S., Solodov, M.V.: The Josephy–Newton method for semismooth generalized equations and semismooth SQP for optimization. Set-Valued and Variational Analysis **21**(1), 17–45 (2013)

[21] Izmailov, A.F., Solodov, M.V.: Inexact Josephy–Newton framework for generalized equations and its applications to local analysis of Newtonian methods for constrained optimization. Computational Optimization and Applications **46**(2), 347–368 (2010)

[22] Kanzow, C., Lechner, T.: Globalized inexact proximal Newton-type methods for nonconvex composite functions. Computational Optimization and Applications **78**(2), 377–410 (2021)

[23] Lee, J.D., Sun, Y., Saunders, M.A.: Proximal Newton-type methods for minimizing composite functions. SIAM J. Optim. **24**(3), 1420–1443 (2014). https://doi.org/10.1137/130921428

[24] Byrd, R.H., Nocedal, J., Oztoprak, F.: An inexact successive quadratic approximation method for L-1 regularized optimization. Mathematical Programming **157**(2), 375–396 (2016)

[25] Bauschke, H.H., Combettes, P.L.: Convex Analysis and Monotone Operator Theory in Hilbert Spaces. CMS Books in Mathematics. Springer, Cham, Switzerland (2018)

[26] Conn, A.R., Gould, N.I.M., Toint, P.L.: Trust Region Methods. SIAM, Philadelphia, PA (2000)

[27] Beck, A.: First Order Methods in Optimization. Society for Industrial and Applied Mathematics, Philadelphia, PA (2017). https://doi.org/10.1137/1.9781611974997

[28] Cartis, C., Gould, N.I.M., Toint, Ph.L.: On the evaluation complexity of composite function minimization with applications to nonconvex nonlinear programming. SIAM J. Optim. **21**(4), 1721–1739 (2011). https://doi.org/10.1137/11082381X

[29] Dembo, R.S., Eisenstat, S.C., Steihaug, T.: Inexact Newton methods. SIAM Journal on Numerical analysis **19**(2), 400–408 (1982)

[30] Lin, C.-J., Moré, J.J.: Newton's method for large bound-constrained optimization problems. SIAM Journal on Optimization **9**(4), 1100–1127 (1999)

[31] Nocedal, J., Wright, S.: Numerical Optimization. Springer Series in Operations Research and Financial Engineering. Springer, ??? (2006)

[32] Dennis Jr., J.E., Mei, H.H.W.: Two new unconstrained optimization algorithms which use function and gradient values. J. Optim. Theory and Applics. **28**, 453–482 (1979). https://doi.org/10.1007/BF00932218

[33] Kelley, C.T.: Iterative Methods for Optimization vol. 18. SIAM, ??? (1999)