

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

Error bounds, quadratic growth, and linear convergence of proximal methods

Dmitriy Drusvyatskiy

University of Washington, Department of Mathematics, Seattle, WA 98195;
ddrusv@uw.edu, www.math.washington.edu/~ddrusv

Adrian S. Lewis

School of Operations Research and Information Engineering, Cornell University, Ithaca, New York, USA;
aslewis@orie.cornell.edu, people.orie.cornell.edu/~aslewis

The proximal gradient algorithm for minimizing the sum of a smooth and a nonsmooth convex function often converges linearly even without strong convexity. One common reason is that a multiple of the step length at each iteration may linearly bound the “error” – the distance to the solution set. We explain the observed linear convergence intuitively by proving the equivalence of such an error bound to a natural quadratic growth condition. Our approach generalizes to linear and quadratic convergence analysis for proximal methods (of Gauss-Newton type) for minimizing compositions of nonsmooth functions with smooth mappings. We observe incidentally that short step-lengths in the algorithm indicate near-stationarity, suggesting a reliable termination criterion.

Key words: Proximal algorithm, error bound, quadratic growth, linear convergence, subregularity, subdifferential, tilt-stability

MSC2000 subject classification: 90C25, 90C31, 90C55, 65K10

1. Introduction Under favorable conditions, many fundamental optimization algorithms converge linearly: the distance of the iterates to the optimal solution set (the “error”) is bounded by a decreasing geometric sequence. Classical optimization literature highlights how quadratic growth properties of the objective function, typically guaranteed through second-order optimality conditions, ensure such linear convergence. Central examples traditionally include the method of steepest descent for smooth minimization [5, Theorem 3.4] and, more abstractly, the proximal point method for nonsmooth convex problems [48, Theorem 2, Proposition 7].

More recent techniques, originally highlighted in the work of Luo and Tseng [31], postulate that the step length at each iteration of the algorithm linearly bounds the error. Such “error bounds” are commonly used in the analysis of first-order methods for strongly convex functions, popular in modern applications such as machine learning and high-dimensional statistics, including in particular the proximal gradient method and its variants; see for example Nesterov [40] and Beck-Teboulle [7]. Convergence analysis based only on the error bound property is appealingly simple even without strong convexity, but the underlying assumption on the optimization problem is opaque at least at first sight.

Some recent developments have focused on linear convergence guarantees based on more intuitive, geometric properties, akin to the classical quadratic growth condition. Two interesting examples

are [8] and [39]. Our aim here is to take a systematic and transparent approach, that is generalizable to problems with more complex structure. Our aim is to show, in several interesting contemporary optimization frameworks, the equivalence between, on the one hand, the intuitive notion of quadratic growth of the objective function away from the set of minimizers, and on the other hand, the powerful analytic tool furnished by an error bound. Rockafellar already foreshadowed this possibility with his original analysis of the proximal point method [48]. We extend that relationship here to the *proximal gradient method* for problems

$$\min_x f(x) + g(x),$$

with g convex and f convex and smooth, and more generally to the *prox-linear algorithm* (a variant of Gauss-Newton) for convex-composite problems

$$\min_x g(x) + h(c(x)), \tag{1.1}$$

where g is an extended-real-valued closed convex function, h is a finite-valued convex function, and c is a smooth mapping. Acceleration strategies for the prox-linear algorithm have recently appeared in [24]. In parallel, we show how the error bound property quickly yields linear convergence guarantees. In essence, our analysis depends on viewing these two methods as approximations of the original proximal point algorithm – a perspective of an independent interest. Our assumptions are mild: we rely primarily on a natural strict complementarity condition. In particular, we simplify and extend some of the novel convergence guarantees established in the recent preprint [55] for the prox-gradient method.¹

The iterative algorithms we consider assign a “gradient-like” step to each potential iterate, as in the analysis of proximal methods in [40, Section 2.1.5]; the step length is zero at stationary points and otherwise serves as a surrogate measure of optimality. For steepest descent, the step is simply a multiple of the negative gradient, for the proximal point method it is determined by a subdifferential relationship, while the prox-gradient and prox-linear methods combine the two. In the language of variational analysis, the existence of an error bound is exactly “metric subregularity” of the gradient-like mapping; see Dontchev-Rockafellar [49]. We will show that subregularity of the gradient-like mapping is equivalent to subregularity of the subdifferential of the objective function itself, thereby allowing us to call on extensive literature relating the quadratic growth of a function to metric subregularity of its subdifferential [19, 21, 23, 27]. Given the generality of these techniques, we expect that the approach we describe here, rooted in understanding linear convergence through quadratic growth, should extend broadly. We note, in particular, some parallel developments influenced by the first version of this work [22], in the recent manuscript [54].

When analyzing the prox-linear algorithm, we encounter a surprise. The error bound condition yields a linear convergence rate that is an order of magnitude worse than the natural rate for the prox-gradient method in the convex setting. The difficulty is that in the nonconvex case, the “linearizations” used by the method do not *lower-bound* the objective function. Nonetheless, we show that the method does converge with the natural rate if the objective function satisfies the stronger condition of quadratic growth that is uniform with respect to tilt-perturbations – a property equivalent to the well-studied notions of tilt-stability [21] and strong metric regularity of the subdifferential [49]. Concretely, these notions reduce to strong second-order sufficient conditions in nonlinear programming [38], whenever the active gradients are linearly independent. Moreover, we show that under a sharp growth condition [11], the scheme locally converges quadratically.

¹ While finalizing a first version of this work, the authors became aware of a concurrent, independent and nicely complementary approach [30], based on a related calculus of Kurdyka-Łojasiewicz exponents.

An important byproduct of our analysis, worthy of independent interest, relates the step-lengths taken by the prox-linear method to near-stationarity of the objective function at the iterates. Therefore, short step-lengths can be used to terminate the scheme, with explicit guarantees on the quality of the final solution. We end by studying to what extent the tools we have developed generalize to composite optimization where the outer function h may be neither convex nor continuous – an arena of growing recent interest (e.g. [29]). While considerably more technical, key ingredients of our analysis extend to this very general setting.

There have been a number of closely related works that appeared on arXiv shortly before or concurrently with the initial version of this manuscript, notably [8, 30, 39, 54, 55]. These papers pertain to the first four sections on convex additive composite minimization. We postpone a discussion of these articles and their overlap with the current work to Section 4.1, after having introduced relevant notation and our main results.

The outline of the manuscript is as follows. Section 2 briefly records some elementary preliminaries. Section 3 contains a detailed analysis of linear convergence of the prox-gradient method for convex functions through the lens of error bounds and quadratic growth. In section 4, we show that quadratic growth holds in concrete applications under a mild condition of dual strict complementarity; our analysis aims to illuminate and extend some of the results in [55] by dispensing with strong convexity of component functions. Section 5 is dedicated to the local linear convergence of the prox-linear algorithm for minimizing compositions of convex functions with smooth mappings. Section 6 explains how a uniform notion of quadratic growth implies linear convergence of the prox-linear method with the natural rate. In Section 7, we show that under a sharpness condition, the prox-linear method locally converges quadratically. Section 8 explains the consequences of the previous section on the prox-gradient method when the smooth component is not convex. The final section 9 shows the equivalence between the error bound property of the prox-linear map and subdifferential subregularity when both the component functions g and h may be infinite-valued and non-convex.

2. Preliminaries Unless otherwise stated, we follow the terminology and notation of [40, 49, 50]. Throughout \mathbf{R}^n will denote an n -dimensional Euclidean space with inner-product $\langle \cdot, \cdot \rangle$ and corresponding norm $\|\cdot\|$. The closed unit ball will be written as \mathbf{B} , while the open ball of radius r around a point x will be denoted by $B_r(x)$. For any set $Q \subset \mathbf{R}^n$, we define the *distance function*

$$\text{dist}(x; Q) := \inf_{z \in Q} \|z - x\|.$$

The functions we consider will take values in the extended real line $\overline{\mathbf{R}} := \mathbf{R} \cup \{\pm\infty\}$. The *domain* and the *epigraph* of a function $f: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ are defined by

$$\begin{aligned} \text{dom } f &:= \{x \in \mathbf{R}^n : f(x) < +\infty\}, \\ \text{epi } f &:= \{(x, r) \in \mathbf{R}^n \times \mathbf{R} : f(x) \leq r\}, \end{aligned}$$

respectively. We say that f is *closed* if the inequality $\liminf_{x \rightarrow \bar{x}} f(x) \geq f(\bar{x})$ holds for any point $\bar{x} \in \mathbf{R}^n$. The symbol $[f \leq \nu] := \{x : f(x) \leq \nu\}$ will denote the ν -sublevel set of f . For any set $Q \subset \mathbf{R}^n$, the indicator function $\delta_Q: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ evaluates to zero on Q and to $+\infty$ elsewhere.

The Fenchel conjugate of a convex function $f: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ is the closed convex function $f^*: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ defined by

$$f^*(y) = \sup_x \{\langle y, x \rangle - f(x)\}.$$

The *subdifferential* of a convex function f at a point x , denoted by $\partial f(x)$, is the set consisting of all vectors $v \in \mathbf{R}^n$ satisfying $f(z) \geq f(x) + \langle v, z - x \rangle$ for all $z \in \mathbf{R}^n$. For any function $f: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ and a real number $t > 0$, we define the *Moreau envelope*

$$g^t(x) := \min_y \left\{ g(y) + \frac{1}{2t} \|y - x\|^2 \right\},$$

and the *proximal mapping*

$$\text{prox}_{tg}(x) := \arg \min_{y \in \mathbf{R}^n} \left\{ g(y) + \frac{1}{2t} \|y - x\|^2 \right\}.$$

The proximal map $x \mapsto \text{prox}_{tg}(x)$ is always 1-Lipschitz continuous.

A set-valued mapping $F: \mathbf{R}^n \rightrightarrows \mathbf{R}^m$ is a mapping assigning to each point $x \in \mathbf{R}^n$ the subset $F(x)$ of \mathbf{R}^m . The *graph* of such a mapping is the set

$$\text{gph } F := \{(x, y) \in \mathbf{R}^n \times \mathbf{R}^m : y \in F(x)\}.$$

The inverse map $F^{-1}: \mathbf{R}^m \rightrightarrows \mathbf{R}^n$ is defined by setting $F^{-1}(y) = \{x : y \in F(x)\}$. Every mapping $F: \mathbf{R}^n \rightrightarrows \mathbf{R}^m$ obeys the identity [50, Lemma 12.14]:

$$(I + F)^{-1} + (I + F^{-1})^{-1} = I. \quad (2.1)$$

Note that for any convex function g and a real $t > 0$, equality $\text{prox}_{tg} = (I + t\partial g)^{-1}$ holds.

3. Linear convergence of the prox-gradient method To motivate the discussion, consider the optimization problem

$$\min_x \varphi(x) := f(x) + g(x) \quad (3.1)$$

where $g: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ is a closed convex function and $f: \mathbf{R}^n \rightarrow \mathbf{R}$ is a convex C^1 -smooth function with a β -Lipschitz continuous gradient:

$$\|\nabla f(x) - \nabla f(y)\| \leq \beta \|x - y\| \quad \text{for all } x, y \in \mathbf{R}^n.$$

The *proximal gradient method* is the recurrence

$$x_{k+1} = \arg \min_{x \in \mathbf{R}^n} f(x_k) + \langle \nabla f(x_k), x - x_k \rangle + \frac{1}{2t} \|x - x_k\|^2 + g(x),$$

where the constant $t > 0$ is appropriately chosen. More succinctly, the method simply iterates

$$x_{k+1} = \text{prox}_{tg}(x_k - t\nabla f(x_k)).$$

In order to see the parallel between the proximal gradient method and classical gradient descent for smooth minimization, it is convenient to rewrite the recurrence yet again as $x_{k+1} = x_k - t\mathcal{G}_t(x_k)$ where

$$\mathcal{G}_t(x) := t^{-1} \left(x - \text{prox}_{tg}(x - t\nabla f(x)) \right)$$

is the *prox-gradient mapping*. In particular, equality $\mathcal{G}_t(x) = 0$ holds if and only if x minimizes φ .

Let S be the set of minimizers of φ and let φ^* be the minimal value of φ . Supposing now $t \leq \beta^{-1}$, the following two inequalities are standard [40, Theorem 2.2.7, Corollary 2.2.1] and [7, Lemma 2.3]:

$$\varphi(x_k) - \varphi(x_{k+1}) \geq \frac{1}{2\beta} \|\mathcal{G}_t(x_k)\|^2, \quad (3.2)$$

$$\varphi(x_{k+1}) - \varphi^* \leq \langle \mathcal{G}_t(x_k), x_k - x^* \rangle - \frac{1}{2\beta} \|\mathcal{G}_t(x_k)\|^2. \quad (3.3)$$

Here x^* denotes an arbitrary element of S . Hence equation (3.3) immediately implies

$$\varphi(x_{k+1}) - \varphi^* \leq \|\mathcal{G}_t(x_k)\|^2 \left(\frac{\|x^* - x_k\|}{\|\mathcal{G}_t(x_k)\|} - \frac{1}{2\beta} \right).$$

Defining $\gamma_k := \frac{\|x^* - x_k\|}{\|\mathcal{G}_t(x_k)\|}$ and using inequality (3.2), along with some trivial algebraic manipulations, yields the geometric decrease guarantee

$$\varphi(x_{k+1}) - \varphi^* \leq \left(1 - \frac{1}{2\beta\gamma_k}\right) (\varphi(x_k) - \varphi^*). \quad (3.4)$$

Hence if the quantities γ_k are bounded for all large k , asymptotic Q-linear convergence in function values is assured. This observation motivates the following definition, originating in [31].

DEFINITION 3.1 (Error bound condition). Given real numbers $\gamma, \nu > 0$, we say that the *error bound condition holds with parameters* (γ, ν) if the inequality

$$\text{dist}(x, S) \leq \gamma \|\mathcal{G}_t(x)\| \quad \text{is valid for all } x \in [\varphi \leq \varphi^* + \nu].$$

Hence we have established the following.

THEOREM 3.2 (Linear convergence). *Suppose the error bound condition holds with parameters (γ, ν) . Then the proximal gradient method with $t \leq \beta^{-1}$ satisfies $\varphi(x_k) - \varphi^* \leq \epsilon$ after at most*

$$k \leq \frac{\beta}{2\nu} \text{dist}^2(x_0, S) + 2\beta\gamma \ln \left(\frac{\varphi(x_0) - \varphi^*}{\epsilon} \right) \quad \text{iterations.} \quad (3.5)$$

Moreover, the iterate sequence x_k converges to some $x^* \in S$ and there exists an index r such that the inequality

$$\|x_{r+k} - x^*\|^2 \leq \left(1 - \frac{1}{2\beta\gamma}\right)^k C \cdot (\varphi(x_r) - \varphi^*),$$

holds for all $k \geq 1$, where we set $C := \frac{2}{\beta(1 - \sqrt{1 - (2\beta\gamma)^{-1}})^2}$.

Proof. From the the standard sublinear estimate $\varphi(x_k) - \varphi^* \leq \frac{\beta \cdot \text{dist}^2(x_0, S)}{2k}$ (see e.g. [7, Theorem 3.1]), we deduce that after $k \leq \frac{\beta}{2\nu} \text{dist}^2(x_0, S)$ iterations the inequality $\varphi(x_k) - \varphi^* \leq \nu$ holds. The second summand in inequality (3.5) is then immediate from the linear rate (3.4) and the fact that the values $\varphi(x_k)$ decrease monotonically.

It is standard that the sequence x_k for $k \geq 1$ is convergent. This directly follows from the inequality $\|x_{k+1} - \bar{x}\| \leq \|x_k - \bar{x}\|$ holding for all k and $\bar{x} \in S$.² Now, set $x^* := \lim_{k \rightarrow \infty} x_k$. Note that if an iterate x_r lies in the set $[\varphi \leq \varphi^* + \nu]$, then we have

$$\begin{aligned} \|x_{r+k} - x^*\| &\leq \sum_{i=r+k}^{\infty} \|x_i - x_{i+1}\| \leq \sqrt{2/\beta} \sum_{i=r+k}^{\infty} \sqrt{\varphi(x_i) - \varphi(x_{i+1})} \\ &\leq \sqrt{2/\beta} \sqrt{\varphi(x_r) - \varphi^*} \sum_{i=r+k}^{\infty} \left(1 - \frac{1}{2\beta\gamma}\right)^{\frac{i-r}{2}} \leq \left(1 - \frac{1}{2\beta\gamma}\right)^{k/2} D \sqrt{\varphi(x_r) - \varphi^*}, \end{aligned}$$

where we set $D := \frac{\sqrt{2}}{\sqrt{\beta(1 - \sqrt{1 - (2\beta\gamma)^{-1}})}} \cdot$. Squaring both sides, the result follows. \square

Convergence guarantees of Theorem 3.2 are expressed in terms of the error bound parameters (γ, ν) – quantities not stated in terms of the initial data of the problem, f and g . Indeed, the error bound condition is a property of the prox-gradient mapping $\mathcal{G}_t(x)$, a nontrivial object to understand. In contrast, in the current work we will show that the error bound condition is simply equivalent to the objective function φ growing quadratically away from its minimizing set S – a familiar, transparent, and largely classical property in nonsmooth optimization.

² See for example the first displayed equation of the proof of [7, Theorem 3.1]

To gain some intuition, consider the simplest case $g = 0$. Then the prox-gradient method reduces to gradient descent $x_{k+1} = x_k - t\nabla f(x_k)$. Suppose now that f grows quadratically (globally) away from its minimizing set, meaning there is a real number $\alpha > 0$ such that

$$f(x) \geq f^* + \frac{\alpha}{2} \text{dist}^2(x, S) \quad \text{for all } x \in \mathbf{R}^n.$$

Notice this property is weaker than strong convexity even for C^1 -smooth functions; e.g. $f(x) = (\max\{|x| - 1, 0\})^2$. Then convexity implies

$$\frac{\alpha}{2} \text{dist}^2(x, S) \leq f(x) - f^* \leq \langle \nabla f(x), x^* - x \rangle \leq \|\nabla f(x)\| \|x - x^*\|.$$

Thus the error bound condition holds with parameters $(\gamma, \nu) = (\frac{2}{\alpha}, \infty)$, and the complexity bound of Theorem 3.2 becomes $k \leq \frac{4\beta}{\alpha} \ln\left(\frac{\varphi(x_0) - \varphi^*}{\epsilon}\right)$. This is the familiar linear rate of gradient descent (up to a constant).

Our goal is to elucidate the quantitative relationship between quadratic growth and the error bound condition in full generality. The strategy we follow is very natural; we will interpret the proximal gradient method as an approximation to the true proximal point algorithm $y_{k+1} = \text{prox}_{t\varphi}(y_k)$ on the function $\varphi = f + g$, and show a linear relationship between the corresponding step sizes (Theorem 3.5). This will allow us to ignore the linearization appearing in the definition of the proximal gradient method and focus on the relationship between quadratic growth of φ , properties of the mapping $\text{prox}_{t\varphi}$, and of the subdifferential $\partial\varphi$ (Theorems 3.3 and 3.4). We believe this interpretation of the proximal gradient method is of interest in its own right.

The following is a central result we will need. It establishes a relationship between quadratic growth properties and a “global error bound property” of the function $x \mapsto d(0, \partial\varphi(x))$. Variants of this result have appeared in [1, Theorem 3.3], [4, Theorem 6.1], [19, Theorem 4.3], [23, Theorem 3.1], and [27].

THEOREM 3.3 (Subdifferential error bound and quadratic growth). *Consider a closed convex function $h: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ with minimal value h^* and let S be its set of minimizers. Consider the conditions*

$$h(x) \geq h(\bar{x}) + \frac{\alpha}{2} \cdot \text{dist}^2(x; S) \quad \text{for all } x \in [h \leq h^* + \nu] \quad (3.6)$$

and

$$\text{dist}(x; S) \leq L \cdot \text{dist}(0; \partial h(x)) \quad \text{for all } x \in [h \leq h^* + \nu]. \quad (3.7)$$

If condition (3.6) holds, then so does condition (3.7) with $L = 2\alpha^{-1}$. Conversely, condition (3.7) implies condition (3.6) with any $\alpha \in (0, \frac{1}{L}]$.

The proof of the implication (3.6) \Rightarrow (3.7) is identical to the proof of the analogous implication in [1, Theorem 3.3]; the proof of the implication (3.7) \Rightarrow (3.6) is the same as that of [19, Theorem 4.3], [23, Theorem 3.1]. Hence we omit the arguments.

Given the equality $\text{prox}_{th} = (I + t\partial h)^{-1}$, it is clear that the subdifferential error bound condition (3.7) is related to an analogous property of the proximal mapping. This is the content of the following elementary result.

THEOREM 3.4 (Proximal and subdifferential error bounds).

Consider a closed convex function $h: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ with minimal value h^ and let S be its set of minimizers. Consider the conditions*

$$\text{dist}(x; S) \leq L \cdot \text{dist}(0; \partial h(x)) \quad \text{for all } x \in [h \leq h^* + \nu]. \quad (3.8)$$

and

$$\text{dist}(x; S) \leq \widehat{L} \cdot t^{-1} \|x - \text{prox}_{th}(x)\| \quad \text{for all } x \in [h \leq h^* + \nu]. \quad (3.9)$$

If condition (3.8) holds, then so does condition (3.9) with $\widehat{L} = L + t$. Conversely, condition (3.9) implies condition (3.8) with $L = \widehat{L}$.

Proof. Suppose condition (3.8) holds and consider a point $x \in [h \leq h^* + \nu]$. Then clearly the inequality $h(\text{prox}_{th}(x)) \leq h(x) \leq h^* + \nu$ holds. Taking into account the inclusion $t^{-1}(x - \text{prox}_{th}(x)) \in \partial h(\text{prox}_{th}(x))$, we obtain

$$\begin{aligned} \text{dist}(x, S) &\leq \|x - \text{prox}_{th}(x)\| + \text{dist}(\text{prox}_{th}(x), S) \\ &\leq \|x - \text{prox}_{th}(x)\| + L \cdot \text{dist}(0; \partial h(\text{prox}_{th}(x))) \\ &\leq (1 + t^{-1}L)\|x - \text{prox}_{th}(x)\|, \end{aligned}$$

as claimed. Conversely suppose condition (3.9) holds and fix a point $x \in [h \leq h^* + \nu]$. Then for any subgradient $v \in \partial h(x)$, equality $\text{prox}_{th}(x + tv) = x$ holds. Hence, we obtain

$$\text{dist}(x, S) \leq \widehat{L}t^{-1}\|x - \text{prox}_{th}(x)\| \leq \widehat{L}t^{-1}\|\text{prox}_{th}(x + tv) - \text{prox}_{th}(x)\| \leq \widehat{L}\|v\|,$$

where we have used the fact that the proximal mapping is 1-Lipschitz continuous. Since the subgradient $v \in \partial h(x)$ is arbitrary, the result follows. \square

The final step is to relate the step sizes taken by the proximal gradient and the proximal point methods. The ensuing arguments are best stated in terms of monotone operators. To this end, observe that our running problem (3.1) is equivalent to solving the inclusion

$$0 \in \nabla f(x) + \partial g(x).$$

More generally, consider monotone operators $F: \mathbf{R}^n \rightarrow \mathbf{R}^n$ and $G: \mathbf{R}^n \rightrightarrows \mathbf{R}^n$, meaning that F and G satisfy the inequalities $\langle v_1 - v_2, x_1 - x_2 \rangle \geq 0$ and $\langle F(x_1) - F(x_2), x_1 - x_2 \rangle \geq 0$ for all $x_i \in \mathbf{R}^n$ and $v_i \in G(x_i)$ with $i = 1, 2$. We now further assume that G is *maximal monotone*, meaning that the graph $\text{gph} G$ is not a proper subset of the graph of any other monotone operator. Along with the operator G and a real $t > 0$, we associate the *resolvent*

$$\text{prox}_{tG} := (I + tG)^{-1}.$$

The mapping $\text{prox}_{tG}: \mathbf{R}^n \rightarrow \mathbf{R}^n$ is then single-valued and nonexpansive (1-Lipschitz continuous [50, Theorem 12.12]). We aim to solve the inclusion

$$0 \in \Phi(x) := F(x) + G(x)$$

by the *Forward-Backward algorithm*

$$x_{k+1} = \text{prox}_{tG}(x_k - tF(x_k)),$$

Equivalently we may write $x_{k+1} = x_k - t\mathcal{G}_t(x_k)$ where $\mathcal{G}_t(x)$ is the *prox-gradient* mapping

$$\mathcal{G}_t(x) := \frac{1}{t}(x - \text{prox}_{tG}(x - tF(x))).$$

Setting $F = \nabla f$, $G := \partial g$, $\Phi := \partial \varphi$ recovers the proximal gradient method for the problem (3.1).

The following key result shows that the step lengths of the Forward-Backward algorithm and those taken by the proximal point algorithm $z_{k+1} = \text{prox}_{t\Phi}(z_k)$ are proportional. During the final stages of the review process, we became aware of the earlier paper [17], where essentially the same result is proved.

THEOREM 3.5 (Step-lengths comparison). *Consider two maximal monotone operators $G: \mathbf{R}^n \rightrightarrows \mathbf{R}^n$ and $\Phi: \mathbf{R}^n \rightrightarrows \mathbf{R}^n$, with the difference $F := \Phi - G$ that is single-valued. Then the inequality*

$$\|\mathcal{G}_t(x)\| \leq \text{dist}(0; \Phi(x)) \quad \text{holds.} \quad (3.10)$$

Supposing that F is in addition β -Lipschitz continuous, the inequalities hold:

$$(1 - \beta t) \cdot \|\mathcal{G}_t(x)\| \leq \|t^{-1}(x - \text{prox}_{t\Phi}(x))\| \leq (1 + \beta t) \cdot \|\mathcal{G}_t(x)\| \quad (3.11)$$

Proof. Fix a point $x \in \mathbf{R}^n$ and a vector $v \in \Phi(x)$. Then clearly the inclusion

$$(x - tF(x)) + tv \in x + tG(x) \quad \text{holds,}$$

or equivalently $x = \text{prox}_{tG}((x - tF(x)) + tv)$. Since the proximal mapping is nonexpansive, we deduce

$$t\|\mathcal{G}_t(x)\| = \|x - \text{prox}_{tG}(x - tF(x))\| \leq t\|v\|.$$

Letting v be the minimal norm element of $\Phi(x)$, we deduce the claimed inequality $\|\mathcal{G}_t(x)\| \leq \text{dist}(0; \Phi(x))$.

Now suppose that F is β -Lipschitz continuous. Consider a point $x \in \mathbf{R}^n$ and define $z := \mathcal{G}_t(x)$. Observe the chain of equivalences:

$$\begin{aligned} z = \mathcal{G}_t(x) &\Leftrightarrow tz = x - \text{prox}_{tG}(x - tF(x)) \\ &\Leftrightarrow x - tF(x) \in (x - tz) + tG(x - tz) \\ &\Leftrightarrow x + t(F(x - tz) - F(x)) \in (I + t\Phi)(x - tz) \end{aligned}$$

Define now the vector $w = F(x - tz) - F(x)$ and note $\|w\| \leq \beta t\|z\|$. Hence taking into account that resolvents are nonexpansive, we obtain

$$x - tz = \text{prox}_{t\Phi}(x + tw) \subset \text{prox}_{t\Phi}(x) + t\|w\|\mathbf{B},$$

and so deduce

$$z \in \frac{1}{t}(x - \text{prox}_{t\Phi}(x)) + \beta t\|z\|\mathbf{B}.$$

Hence

$$\left| \|z\| - t^{-1}\|x - \text{prox}_{t\Phi}(x)\| \right| \leq \|z - t^{-1}(x - \text{prox}_{t\Phi}(x))\| \leq \beta t\|z\|.$$

The two inequalities in (3.11) follow immediately. \square

We now arrive at the main result of this section.

COROLLARY 3.6 (Error bound and quadratic growth). *Consider a closed, convex function $g: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ and a C^1 -smooth convex function $f: \mathbf{R}^n \rightarrow \mathbf{R}$ with β -Lipschitz continuous gradient. Suppose that the function $\varphi := f + g$ has a nonempty set S of minimizers and consider the following conditions:*

- (Quadratic growth)

$$\varphi(x) \geq \varphi^* + \frac{\alpha}{2} \cdot \text{dist}^2(x; S) \quad \text{for all } x \in [\varphi \leq \varphi^* + \nu] \quad (3.12)$$

- (Error bound condition)

$$\text{dist}(x, S) \leq \gamma \|\mathcal{G}_t(x)\| \quad \text{is valid for all } x \in [\varphi \leq \varphi^* + \nu], \quad (3.13)$$

Then property (3.12) implies property (3.13) with $\gamma = (2\alpha^{-1} + t)(1 + \beta t)$. Conversely, condition (3.13) implies condition (3.12) with any $\alpha \in (0, \gamma^{-1}]$.

Proof. Suppose condition (3.12) holds. Then for any $x \in [\varphi \leq \varphi^* + \nu]$, we deduce

$$\begin{aligned} \text{dist}(x, S) &\leq 2\alpha^{-1} \cdot \text{dist}(0; \partial\varphi(x)) && \text{(Theorem 3.3)} \\ &\leq (2\alpha^{-1} + t)\|t^{-1}(x - \text{prox}_{t\varphi}(x))\| && \text{(Theorem 3.4)} \\ &\leq (2\alpha^{-1} + t)(1 + \beta t)\|\mathcal{G}_t(x)\| && \text{(Inequality 3.11)} \end{aligned}$$

This establishes (3.13) with $\gamma = (2\alpha^{-1} + t)(1 + \beta t)$. Conversely suppose (3.13) holds. Then for any $x \in [\varphi \leq \varphi^* + \nu]$ we deduce using Theorem 3.5 the inequality $\text{dist}(x, S) \leq \gamma \|\mathcal{G}_t(x)\| \leq \gamma \cdot \text{dist}(0, \partial\varphi(x))$. An application of Theorem 3.3 completes the proof. \square

The following convergence result is now immediate from Theorem 3.2 and Corollary 3.6. Notice that the complexity bound matches (up to a constant) the linear rate of convergence of the proximal gradient method when applied to strongly convex functions.

COROLLARY 3.7 (Quadratic growth and linear convergence). *Consider a closed, convex function $g: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ and a C^1 -smooth function $f: \mathbf{R}^n \rightarrow \mathbf{R}$ with β -Lipschitz continuous gradient. Suppose that the function $\varphi := f + g$ has a nonempty set S of minimizers and that the quadratic growth condition holds:*

$$\varphi(x) \geq \varphi^* + \frac{\alpha}{2} \cdot \text{dist}^2(x; S) \quad \text{for all } x \in [\varphi \leq \varphi^* + \nu].$$

Then the proximal gradient method with $t \leq \beta^{-1}$ satisfies $\varphi(x_k) - \varphi^* \leq \epsilon$ after at most

$$k \leq \frac{\beta}{2\nu} \text{dist}^2(x_0, S) + 12 \cdot \frac{\beta}{\alpha} \ln \left(\frac{\varphi(x_0) - \varphi^*}{\epsilon} \right) \quad \text{iterations.}$$

Moreover, the iterate sequence x_k converges to some $x^* \in S$ and there exists an index r such that the inequality

$$\|x_{r+k} - x^*\|^2 \leq \left(1 - \frac{1}{2\beta\gamma}\right)^k C \cdot (\varphi(x_r) - \varphi^*),$$

holds for all $k \geq 1$, where we set $C := \frac{2}{\beta(1 - \sqrt{1 - (2\beta\gamma)^{-1}})^2}$ and $\gamma := (2\alpha^{-1} + t)(1 + \beta t)$.

4. Quadratic growth in structured optimization Recently, the authors of [55] proved that the error bound condition holds under very mild assumptions, thereby explaining asymptotic linear convergence of the proximal gradient method often observed in practice. In this section, we aim to use the equivalence between the error bound condition and quadratic growth, established in Theorem 3.6, to streamline and illuminate the arguments in [55], while also extending their results to a wider setting. To this end, consider the problem

$$\inf_x \varphi(x) := f(Ax) + g(x) \tag{4.1}$$

where $f: \mathbf{R}^m \rightarrow \mathbf{R}$ is convex and C^1 -smooth, $g: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ is closed and convex, and $A: \mathbf{R}^n \rightarrow \mathbf{R}^m$ is a linear mapping. We assume that g is proper, meaning that its domain is nonempty and it does not take the value $-\infty$. Consider now the Fenchel dual problem

$$\sup_y -\Psi(y) := -f^*(y) - g^*(-A^T y). \tag{4.2}$$

By [47, Corollary 31.2.1(a)], the optimal values of the primal (4.1) and of the dual (4.2) are equal, and the dual optimal value is attained. To make progress, we assume that the dual problem (4.2) admits a strictly feasible point:

1. (dual nondegeneracy) $0 \in A^T(\text{ri dom } f^*) + \text{ri dom } g^*$,

Then by [47, Corollary 31.2.1(b)], the optimal value of the primal (4.1) is also attained. From the Kuhn-Tucker conditions [47, p. 333], any optimal solution of the dual coincides with $\nabla f(Ax)$ for any primal optimal solution x . In particular, the dual has a unique optimal solution and we will denote it by \bar{y} . Clearly, the inclusion $0 \in \partial\Psi(\bar{y})$ holds. We now assume the mildly stronger property:

2. (dual strict complementarity) $0 \in \text{ri } \partial\Psi(\bar{y})$.

Taken together, these two standard conditions (dual nondegeneracy and dual strict complementarity) immediately imply

$$0 \in \text{ri } \partial\Psi(\bar{y}) = \text{ri} \left(\partial f^*(\bar{y}) - A \partial g^*(-A^T \bar{y}) \right) = \text{ri } \partial f^*(\bar{y}) - A \left(\text{ri } \partial g^*(-A^T \bar{y}) \right), \tag{4.3}$$

where the last equality follows for example from [47, Theorem 6.6].

Let S be the solution set of the primal problem (4.1). To highlight the impact of the inclusion (4.3) on error bounds, recall that we must estimate the distance $\text{dist}(x, S)$ for an arbitrary point x . To this end, the Kuhn-Tucker conditions again directly imply that S admits the description

$$S = \partial g^*(-A^T \bar{y}) \cap A^{-1} \partial f^*(\bar{y}).$$

The inclusion (4.3), combined with [47, Theorem 6.7], guarantees that the relative interiors of the two sets $\partial g^*(-A^T \bar{y})$ and $A^{-1} \partial f^*(\bar{y})$ meet. Then appealing to [6, Corollary 4.5], we deduce that for any compact set $\mathcal{X} \subset \mathbf{R}^n$ there exists a constant $\kappa \geq 0$ satisfying³

$$\text{dist}(x, S) \leq \kappa \left(\text{dist}(x, \partial g^*(-A^T \bar{y})) + \text{dist}(Ax, \partial f^*(\bar{y})) \right) \quad \text{for all } x \in \mathcal{X}. \quad (4.4)$$

This type of an inequality is often called linear regularity; see for example [6, Corollary 4.5]. The final assumption we need to deduce quadratic growth of φ , not surprisingly, is a quadratic growth condition on the individual functions f and g after tilt perturbations.

DEFINITION 4.1 (Firm convexity). A closed convex function $h: \mathbf{R}^n \rightarrow \bar{\mathbf{R}}$ is *firmly convex relative to a vector* $v \in \mathbf{R}^n$ if the tilted function $h_v(x) := h(x) - \langle v, x \rangle$ satisfies the quadratic growth condition: for any compact set $\mathcal{X} \subset \mathbf{R}^n$ there is a constant α satisfying

$$h_v(x) \geq (\inf h_v) + \frac{\alpha}{2} \text{dist}^2(x, (\partial h_v)^{-1}(0)) \quad \text{for all } x \in \mathcal{X}.$$

We say that h is *firmly convex* if h is firmly convex relative to any $v \in \mathbf{R}^n$, for which $(\partial h_v)^{-1}(0) \neq \emptyset$.

Note that not all convex functions are firmly convex; for example $f(x) = x^4$ is not firmly convex at $x = 0$ relative to $v = 0$. We are now ready to prove the main theorem of this section; note that unlike in [55], we do not require strong convexity of the function f . This generalization is convenient since it allows to capture “robust” formulations where f is a translate of the Huber penalty or its asymmetric extensions.

THEOREM 4.2 (Quadratic growth in composite optimization). *Consider a closed, convex function $g: \mathbf{R}^n \rightarrow \bar{\mathbf{R}}$ and a C^1 -smooth convex function $f: \mathbf{R}^m \rightarrow \mathbf{R}$. Suppose that the sum $\varphi(x) := f(Ax) + g(x)$ has a nonempty set S of minimizers and let \bar{y} be the optimal solution of the dual problem (4.2). Suppose the conditions hold:*

1. (Compactness) *The solution set S is bounded.*
2. (Dual nondegeneracy and strict complementarity) *Assumptions 1 and 2.*
3. (Quadratic growth of components) *The functions f and g are firmly convex relative to \bar{y} and $-A^T \bar{y}$, respectively.*

Then the error bound condition holds with some parameters (γ, ν) .

Proof. Since S is compact, all sublevel sets of φ are compact. Choose a number $\nu > 0$ and set $\mathcal{X} := [\varphi \leq \varphi^* + \nu]$ and $\mathcal{Y} = A(\mathcal{X})$. Let $\bar{x} \in S$ be arbitrary and note the equality $\bar{y} = \nabla f(A\bar{x})$. Then observing that $A\bar{x}$ minimizes $f(\cdot) - \langle \bar{y}, \cdot \rangle$ and \bar{x} minimizes $g(\cdot) + \langle A^T \bar{y}, \cdot \rangle$, property (3) (Quadratic growth of components) guarantees that there exist constants $c, \alpha \geq 0$ such that

$$f(y) \geq f(A\bar{x}) + \langle \bar{y}, y - A\bar{x} \rangle + \frac{c}{2} \text{dist}^2(y, (\partial f)^{-1}(\bar{y})) \quad \text{for all } y \in \mathcal{Y}, \quad (4.5)$$

and

$$g(x) \geq g(\bar{x}) + \langle -A^T \bar{y}, x - \bar{x} \rangle + \frac{\alpha}{2} \text{dist}^2(x, (\partial g)^{-1}(-A^T \bar{y})) \quad \text{for all } x \in \mathcal{X}.$$

³This follows by applying the characterization of linear regularity in [6, Corollary 4.5] first to the pair of sets, $\partial g^*(-A^T \bar{y})$ and $A^{-1} \partial f^*(\bar{y})$, and then to the pair, $\text{gph } A$ and $\mathbf{R}^n \times \partial f^*(\bar{y})$.

Letting κ be the constant from (4.4) and setting $y := Ax$ in (4.5), we deduce

$$\begin{aligned} \varphi(x) = f(Ax) + g(x) &\geq \left(f(A\bar{x}) + \langle \bar{y}, Ax - A\bar{x} \rangle + \frac{c}{2} \text{dist}^2(Ax, \partial f^*(\bar{y})) \right) + \\ &\quad + \left(g(\bar{x}) + \langle -A^T \bar{y}, x - \bar{x} \rangle + \frac{\alpha}{2} \text{dist}^2(x, \partial g^*(-A^T \bar{y})) \right) \\ &\geq \varphi(\bar{x}) + \frac{\min\{\alpha, c\}}{4\kappa^2} \text{dist}^2(x, S). \end{aligned}$$

This completes the proof. \square

Notice that firm convexity requires a certain inequality to hold on compact sets \mathcal{X} , rather than on sublevel sets. In any case, firm convexity is intimately tied to error bounds. For example, analogously to Theorem 3.3, one can show that h is firmly convex relative to v if and only if for any compact set \mathcal{X} there exists a constant $L \geq 0$ satisfying

$$\text{dist}(x, (\partial h)^{-1}(v)) \leq L \cdot \text{dist}(v, \partial h(x)) \quad \text{for all } x \in \mathcal{X}.$$

Indeed this is implicitly shown in the proof of Theorem [1, Theorem 3.3], for example. Moreover, the same argument as in Theorem 3.4 shows that h is firmly convex relative to v if and only if for any compact set \mathcal{X} there exists a constant $\hat{L} \geq 0$ satisfying

$$\text{dist}(x; S) \leq \hat{L} \cdot \|t^{-1}(x - \text{prox}_{th}(x))\| \quad \text{for all } x \in \mathcal{X}.$$

The class of firmly convex functions is large, including for example all strongly convex functions and polyhedral functions h . More generally, all convex Piecewise Linear Quadratic (PLQ) functions [50, Section 10.20] are firmly convex, since their subdifferential graphs are finite unions of polyhedra. Indeed, the subclass of affinely composed PLQ penalties [50, Example 11.18] is ubiquitous in optimization. These are functions of the form

$$h(x) := \sup_{z \in Z} \langle Bx - b, z \rangle - \langle Az, z \rangle$$

where Z is a polyhedron, B is a linear map, and A is a positive-semidefinite matrix. For more details on the PLQ family, see [2, 3]. For example, the elastic net penalty [56], used for group detection, and the soft-insensitive loss [15], used for training Support Vector Machines, fall within this class. We note that the class of PLQ penalties is subsumed by the dual functions considered in [31], which were shown to satisfy the error bound property in [32].

Note that the assumptions of dual nondegeneracy and strict complementarity (Assumptions 1 and 2) were only used in the proof Theorem 4.2 to guarantee inequality (4.4). On the other hand, this inequality holds automatically if the subdifferentials $\partial g^*(-A^T \bar{y})$ and $\partial f^*(\bar{y})$ are polyhedral—a common situation.

COROLLARY 4.3 (Quadratic growth without strict complementarity).

Consider a closed convex function $g: \mathbf{R}^n \rightarrow \bar{\mathbf{R}}$ and a C^1 -smooth convex function $f: \mathbf{R}^m \rightarrow \mathbf{R}$. Suppose that the function $\varphi(x) := f(Ax) + g(x)$ has a nonempty compact set S of minimizers and let \bar{y} be the optimal solution of the dual problem (4.2). Suppose the conditions hold:

1. (Compactness) *The solution set S is bounded.*
2. (Polyhedrality) *Assume that the subdifferential sets, $\partial g^*(-A^T \bar{y})$ and $\partial f^*(\bar{y})$, are polyhedral.*
3. (Quadratic growth of components) *The functions f and g are firmly convex relative to \bar{y} and $-A^T \bar{y}$, respectively.*

Then the error bound condition holds with some parameters (γ, ν) .

Firm convexity is preserved under separable sums.

LEMMA 4.4 (Separable sum). *Consider a family of functions $f_i: \mathbf{R}^{n_i} \rightarrow \overline{\mathbf{R}}$ for $i = 1, \dots, m$ with each f_i firmly convex relative to some $v_i \in \mathbf{R}^{n_i}$. Then the separable function $f: \mathbf{R}^{n_1 + \dots + n_m} \rightarrow \overline{\mathbf{R}}$ defined by $f(x) = \sum_{i=1}^m f_i(x_i)$ is firmly convex relative to the vector (v_1, \dots, v_n) .*

Proof. The proof is immediate from definitions. \square

Combining Corollary 4.3 and Lemma 4.4, one recovers immediately the result of [51, Theorem 2] showing that the group LASSO always satisfies the error bound condition.

Next, we note that firmly convex functions are preserved by the Moreau envelope.

THEOREM 4.5 (Moreau envelope). *Consider a function $h: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ that is firmly convex relative to a vector v . Then the Moreau envelope h^t is itself firmly convex relative to v .*

Proof. Define the tilted functions $h_v(x) := h(x) - \langle v, x \rangle$ and $(h^t)_v(x) := h^t(x) - \langle v, x \rangle$. Observe

$$\begin{aligned} (h^t)_v(x) &= \min_y \left\{ h(y) + \frac{1}{2t} \|y - x\|^2 - \frac{1}{t} \langle tv, x \rangle \right\} \\ &= \min_y \left\{ h(y) - \langle v, y \rangle + \frac{1}{2t} \|y - (x - tv)\|^2 - \frac{t}{2} \|v\|^2 \right\} \\ &= (h_v)^t(x - tv) - \frac{t}{2} \|v\|^2. \end{aligned}$$

Since firm convexity is invariant under translation of the domain, it is now sufficient to show that $(h_v)^t$ is firmly convex relative to the zero vector. To this end, let S be the set of minimizers of h_v , or equivalently the set of minimizers of $(h_v)^t$. Since h is firmly convex relative to v , for any compact set $\mathcal{Z} \subset \mathbf{R}^n$, there exists a constant $\widehat{L} \geq 0$ so that

$$\text{dist}(z, S) \leq \widehat{L} \|t^{-1}(z - \text{prox}_{th_v}(z))\| \quad \text{for all } z \in \mathcal{Z}.$$

Taking into account the equation

$$\nabla(h_v)^t(z) = t^{-1}[z - \text{prox}_{th_v}(z)],$$

we deduce $\text{dist}(z, S) \leq \widehat{L} \cdot \|\nabla(h_v)^t(z)\|$ for all $z \in \mathcal{Z}$, thereby completing the proof. \square

In summary, all typical smooth penalties (e.g. square l_2 -norm, logistic loss), polyhedral functions (e.g. l_1 and l_∞ -penalties, vavnik, hinge loss, check function, anisotropic total variation penalty), Moreau envelopes of polyhedral functions (e.g. Huber and quantile huber [3]), group LASSO regularizer, and general affinely composed PLQ penalties (e.g. soft-insensitive loss [15], elastic net [56]) are firmly convex. Another important example is the nuclear norm [55].

REMARK 4.1 (ROLE OF COMPACTNESS). Notice that the assumptions of Theorem 4.2 require the solution set S to be compact. This assumption is used in two places: (1) the definition of firm convexity has to do with quadratic growth on compact sets and (2) the dual strict complementarity only implies equation (4.4) on compact sets. If instead, the functions f and g are firmly convex globally (i.e. \mathcal{X} in the definition is the whole space) and the subdifferentials $\partial g^*(-A^T \bar{y})$ and $\partial f^*(\bar{y})$ are polyhedral, then the compactness assumption on S can be dropped.

4.1. Discussion of recent related work There have been a number of closely related manuscripts circulating shortly before or concurrently with the initial version of our manuscript [22], namely [8, 30, 39, 54, 55]. These papers pertain to the first four sections on additive composite minimization. Concerning these papers, we here discuss only those results that overlap with ours.

The authors of [39] consider the problem of minimizing a smooth convex function over a simple closed convex set. They define various quadratic growth conditions, and show their equivalence to the error bound of Luo-Tseng, up to a rescaling by problem-independent constants. They then

prove linear convergence guarantees for a number of popular first-order methods, and compute condition numbers when the constraint set is polyhedral. An important distinguishing feature of our work is that we prove an equivalence of the error bound and subdifferential subregularity (3.7), the latter being known to be equivalent to quadratic growth for convex functions. It is this equivalence that extends to Gauss-Newton type methods for problems of the form (1.1) – the setting where quadratic growth alone does not necessarily imply the error bound. The proofs in the two manuscripts are also very different, and our results are stated for additive composite problems.

The paper [8] studies the impact of the KL-inequality on complexity of first-order methods for convex minimization. In particular, the authors show that the KL-inequality with exponent 1/2 is equivalent to quadratic growth. Combining our work with [8], one concludes that for any closed convex function, the KL-inequality with exponent 1/2, quadratic growth, and subdifferential subregularity are equivalent (up to problem-independent constants); if the function is a sum of a smooth and a convex function, these notions are also equivalent to the error bound property. The equivalence proof of the KL-inequality with exponent 1/2 and quadratic growth in [8] relies on a subgradient flow argument. As an aside, it is worth noting that a different proof based instead on a perturbation argument appears in [20, Theorems 3.7, Proposition 3.8]. Indeed, in nonconvex settings, the KL-inequality with a desingularizing power function always implies Hölder subdifferential subregularity, while the converse holds for prox-regular (e.g. weakly convex) functions.

The paper [55] develops sufficient conditions under which the error bound holds for problems of Fenchel form (4.1). Most importantly, the authors recognized the key ingredient: strict dual complementarity. Their strategy follows a careful analysis of subregularity of certain maps associated with the problem. Our paper complements and builds on their work; we show in Section 4 how the equivalence of the error bound with quadratic growth significantly simplifies the analysis.

Manuscript [30] focuses on a calculus of the KL-inequality in nonconvex optimization. Aside from proving a number of positive results, the authors make a compelling case that for nonconvex problems, no reasonable condition implies validity of a sum rule. Taking into account the equivalence of the KL-inequality with growth conditions, it is therefore quiet striking that the sum rule does hold for convex problems (see [55] or Section 4 of the current manuscript).

The paper [54] also seeks to understand linear rates of converge under conditions weaker than strong convexity. Their development, in part using the results developed here, instead focuses on equating restricted strong convexity with the error bound.

5. Prox-linear algorithm We next step away from convex formulations (3.1), and consider the broad class of nonsmooth and nonconvex optimization problems

$$\min_x \varphi(x) := g(x) + h(c(x)), \tag{5.1}$$

where $g: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ is a proper closed convex function, $h: \mathbf{R}^m \rightarrow \mathbf{R}$ is a finite-valued convex function, and $c: \mathbf{R}^n \rightarrow \mathbf{R}^m$ is a C^1 -smooth mapping. Since such problems are typically nonconvex, we seek a point x that is only *first-order stationary*, meaning that the directional derivate of φ at x is nonnegative in all directions. The directional derivate of φ is exactly the support function of the *subdifferential set*, defined by

$$\partial\varphi(x) := \partial g(x) + \nabla c(x)^T \partial h(c(x)),$$

and hence stationary of φ at x simply amounts to the inclusion $0 \in \partial\varphi(x)$.⁴

⁴For those well-versed in variational analysis, we note that $\partial\varphi(x)$ coincides with the Fréchet, limiting, and Clarke subdifferentials of φ at x . The more general definition outside the composite setting will appear in Definition 9.1.

To specify the algorithm we study, define for any points $x, y \in \mathbf{R}^n$ the linearized function

$$\varphi(x; y) := g(y) + h(c(x) + \nabla c(x)(y - x)),$$

and for any real $t > 0$ consider the quadratic perturbation

$$\varphi_t(x; y) := \varphi(x; y) + \frac{1}{2t} \|x - y\|^2.$$

Note that the function $\varphi(x; \cdot)$ is always convex, even though φ typically is not convex. Let x^t be the minimizer of the proximal subproblem

$$x^t := \arg \min_y \varphi_t(x; y).$$

Suppose now that h is L -Lipschitz continuous and the Jacobian $\nabla c(x)$ is β -Lipschitz continuous. It is then immediate that the linearized function $\varphi(x; \cdot)$ is quadratically close to φ itself:

$$-\frac{L\beta}{2} \|x - y\|^2 \leq \varphi(y) - \varphi(x; y) \leq \frac{L\beta}{2} \|x - y\|^2. \quad (5.2)$$

In particular, $\varphi_t(x; \cdot)$ is a quadratic upper estimator of φ for any $t \leq (L\beta)^{-1}$. We now define the prox-gradient mapping in the natural way

$$\mathcal{G}_t(x) := t^{-1}(x - x^t).$$

It is easily verified that equality $\mathcal{G}_t(x) = 0$ holds if and only if x is stationary for φ . In this section, we consider the well-known prox-linear method (Algorithm 1), recently studied for example in [29]; see also [14] for interesting variants. The ideas behind the method (and its trust-region versions) go back a long time, e.g. [9, 10, 12, 13, 25, 45, 46, 52, 53]; see [9] for a historical discussion. We note that Algorithm 1 differs slightly from the one in [29] in the step acceptance criterion.

Algorithm 1 Prox-linear method

Input: A point $x_1 \in \text{dom } g$, and constants $q \in (0, 1)$, $t > 0$, and $\sigma > 0$.

$k \leftarrow 1$

while $\|\mathcal{G}_t(x_k)\| > \epsilon$ **do**

while $\varphi(x_k^t) > \varphi(x_k) - \frac{\sigma}{2} \|\mathcal{G}_t(x_k)\|^2$ **do**

$t \leftarrow qt$

 ▷ Backtracking

end while

$x_{k+1} \leftarrow x_k^t$

 ▷ Iterate update

$k \leftarrow k + 1$

end while

return x_k

Note that the prox-gradient method in Section 3 for the problem $\min_x f(x) + g(x)$ is an example of Algorithm 1 with the decomposition $c(x) = f(x)$ and $h(r) = r$. In this case, we have $L = 1$. Observe also that we do not require f to be convex anymore. Motivated by this observation, we now perform an analysis following the same strategy as for the proximal gradient method; there are important and surprising differences, however, both in the conclusions we make and in the proof techniques. We begin with the following lemma; the proof follows that of [40, Lemma 2.3.2].

LEMMA 5.1 (**Gradient inequality**). *For all points $x, y \in \mathbf{R}^n$, the inequality*

$$\varphi(x; y) \geq \varphi_t(x; x^t) + \langle \mathcal{G}_t(x), y - x \rangle + \frac{t}{2} \|\mathcal{G}_t(x)\|^2, \quad (5.3)$$

holds, and consequently we have

$$\varphi(y) \geq \varphi(x^t) + \langle \mathcal{G}_t(x), y - x \rangle + \frac{t}{2} (2 - L\beta t) \|\mathcal{G}_t(x)\|^2 - \frac{L\beta}{2} \|x - y\|^2. \quad (5.4)$$

and

$$\varphi(x) \geq \varphi(x^t) + \frac{t}{2} (2 - L\beta t) \|\mathcal{G}_t(x)\|^2. \quad (5.5)$$

Proof. Noting that the function $\varphi_t(x; y) := \varphi(x; y) + \frac{1}{2t} \|y - x\|^2$ is strongly convex in the variable y , we deduce

$$\begin{aligned} \varphi(x; y) &= \varphi_t(x; y) - \frac{1}{2t} \|y - x\|^2 \\ &\geq \varphi_t(x; x^t) + \frac{1}{2t} (\|y - x^t\|^2 - \|y - x\|^2) \\ &= \varphi_t(x; x^t) + \langle \mathcal{G}_t(x), y - x \rangle + \frac{t}{2} \|\mathcal{G}_t(x)\|^2, \end{aligned}$$

establishing (5.3). Inequality (5.4) follows by combining (5.2) and (5.3). Finally, we obtain inequality (5.5) from (5.4) by setting $y = x$. \square

For simplicity, we assume that the constants L and β are known and we set $t \leq \frac{1}{L\beta}$ and $\sigma = \frac{1}{L\beta}$ in Algorithm 1, so that the line search always accepts the initial step. The more general setting with the backtracking line-search is entirely analogous. Observe now that the inequality (5.5) yields the functional decrease guarantee

$$\varphi(x_{k+1}) \leq \varphi(x_k) - \frac{1}{2L\beta} \|\mathcal{G}_t(x_k)\|^2, \quad (5.6)$$

and hence we obtain the global convergence rate

$$\min_{i=1, \dots, k} \|\mathcal{G}_t(x_i)\|^2 \leq \frac{2L\beta}{k} \sum_{i=1}^k \varphi(x_i) - \varphi(x_{i+1}) = \frac{2L\beta(\varphi(x_1) - \varphi^*)}{k},$$

where $\varphi^* := \lim_{k \rightarrow \infty} \varphi(x_k)$. Note moreover the prox-gradients $\mathcal{G}_t(x_k)$ tend to zero, since their norms are square-summable. Thus after $2L\beta(\varphi(x_1) - \varphi^*)/\epsilon$ iterations, we can be sure that Algorithm 1 finds a point x_k satisfying $\|\mathcal{G}_t(x_k)\|^2 \leq \epsilon$.

Does a small stepsize $\|\mathcal{G}_t(x)\|$ imply that x is “nearly stationary”? This question is fundamental, and speaks directly to reliability of the termination criterion of Algorithm 1. We will see shortly (Theorem 5.3) that the answer is affirmative: if the quantity $\|\mathcal{G}_t(x)\|$ is small, then x is close to a point that is nearly stationary for φ . Our key tool for establishing this result will be Ekeland’s variational principle.

THEOREM 5.2 (Ekeland’s variational principle).

Consider a closed function $f: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ that is bounded from below. Suppose that for some $\epsilon > 0$ and $\bar{x} \in \mathbf{R}^n$, we have $f(\bar{x}) \leq \inf f + \epsilon$. Then for any $\rho > 0$, there exists a point \bar{u} satisfying

- $f(\bar{u}) \leq f(\bar{x})$,
- $\|\bar{x} - \bar{u}\| \leq \epsilon/\rho$,
- $\{\bar{u}\} = \arg \min_u \{f(u) + \rho\|u - \bar{u}\|\}$.

We can now explain the relation of the quantity $\|\mathcal{G}_t(x)\|$ to approximate stationarity. Aside from its immediate appeal, this result will play an important role in the proof of Theorem 5.11 and in sections 6, 7.

THEOREM 5.3 (Prox-gradient and near-stationarity).

Consider the convex-composite problem (5.1), where h is L -Lipschitz continuous and the Jacobian ∇c is β -Lipschitz. Then for any real $t > 0$, there exists a point \hat{x} satisfying the properties

1. (point proximity) $\|x^t - \hat{x}\| \leq \|x^t - x\|$,
2. (value proximity) $\varphi(\hat{x}) - \varphi(x^t) \leq \frac{t}{2}(L\beta t + 1)\|\mathcal{G}_t(x)\|^2$,
3. (near-stationarity) $\text{dist}(0; \partial\varphi(\hat{x})) \leq (3L\beta t + 2)\|\mathcal{G}_t(x)\|$.

Proof. Define the function $\zeta(y) := \varphi(y) + \frac{1}{2}(L\beta + t^{-1})\|x - y\|^2$ and note that the inequality $\zeta(y) \geq \varphi_t(x, y) \geq \varphi_t(x, x^t)$ holds for all $y \in \mathbf{R}^n$. Letting ζ^* be the infimal value of ζ , we have

$$\zeta(x^t) - \zeta^* \leq \varphi(x^t) - \varphi(x; x^t) + \frac{L\beta}{2}\|x^t - x\|^2 \leq L\beta\|x^t - x\|^2,$$

where the last inequality follows from (5.2). Define the constants $\epsilon := L\beta\|x^t - x\|^2$ and $\rho := L\beta\|x^t - x\|$. Applying Ekeland's variational principle, we obtain a point \hat{x} satisfying the inequalities $\zeta(\hat{x}) \leq \zeta(x^t)$ and $\|x^t - \hat{x}\| \leq \epsilon/\rho$, and the inclusion $0 \in \partial\zeta(\hat{x}) + \rho\mathbf{B}$. The proximity conditions (1) and (2) are immediate. To see near-stationarity (3), observe

$$\begin{aligned} \text{dist}(0, \partial\varphi(\hat{x})) &\leq \rho + (L\beta + t^{-1})\|\hat{x} - x\| \\ &\leq L\beta\|x^t - x\| + (L\beta + t^{-1})(\|x^t - \hat{x}\| + \|x^t - x\|) \\ &\leq (L\beta + 2(L\beta + t^{-1}))\|x^t - x\|. \end{aligned}$$

The result follows. \square

Following the general outline of the paper and armed with Theorem 5.3, we now turn to linear convergence. To this end, let $\{x_k\}$ be a sequence generated by Algorithm 1 and suppose that x^* is a limit point of $\{x_k\}$. Then inequality (5.4) (with $y = x^*$) and lower-semicontinuity of φ immediately imply that the values $\varphi(x_k)$ converge to $\varphi(x^*)$. A standard argument also shows that x^* is a stationary point of φ . Appealing to inequality (5.4) with $y = x^*$, we deduce

$$\varphi(x_{k+1}) - \varphi(x^*) \leq \|\mathcal{G}_t(x_k)\|^2 \left(\frac{\|x_k - x^*\|}{\|\mathcal{G}_t(x_k)\|} + \frac{L\beta}{2} \frac{\|x_k - x^*\|^2}{\|\mathcal{G}_t(x_k)\|^2} + \frac{t}{2}(L\beta t - 2) \right).$$

Defining $\gamma_k := \max\{1, \|x_k - x^*\|/\|\mathcal{G}_t(x_k)\|\}$ we deduce

$$\varphi(x_{k+1}) - \varphi(x^*) \leq \|\mathcal{G}_t(x_k)\|^2 \left(\left(1 + \frac{L\beta}{2}\right) \gamma_k^2 + \frac{t}{2}(L\beta t - 2) \right). \quad (5.7)$$

Combining this inequality with (5.6) yields the geometric decay

$$\varphi(x_{k+1}) - \varphi(x^*) \leq \left(1 - \frac{1}{L\beta(2 + L\beta)\gamma_k^2}\right) (\varphi(x_k) - \varphi(x^*)).$$

Hence provided that γ_k are bounded for all large k , the function values asymptotically converge Q-linearly. This motivates the following definition, akin to Definition 3.1.

DEFINITION 5.4 (Error bound condition). We say that the *error bound condition holds around a point \bar{x} with parameter $\gamma > 0$* if there exists a real number $\epsilon > 0$ so that the inequality

$$\text{dist}(x, (\partial\varphi)^{-1}(0)) \leq \gamma\|\mathcal{G}_t(x)\| \quad \text{is valid for all } x \in B_\epsilon(\bar{x}).$$

Hence we arrive at the following convergence guarantee.

THEOREM 5.5 (Linear convergence of the proximal method). *Consider the sequence x_k generated by Algorithm 1 with $t \leq (L\beta)^{-1}$, and suppose that x_k has some limit point x^* around which the error bound condition holds with parameter $\gamma > 0$. Define now the fraction*

$$q := 1 - \frac{1}{L\beta(2 + L\beta)\gamma^2}.$$

Then for all large k , function values converge Q -linearly

$$\varphi(x_{k+1}) - \varphi(x^*) \leq q \cdot (\varphi(x_k) - \varphi(x^*)),$$

while x_k asymptotically converge R -linearly: there exists an index r such that the inequality

$$\|x_{r+k} - x^*\|^2 \leq C \cdot (\varphi(x_r) - \varphi^*) \cdot q^k,$$

holds for all $k \geq 1$, where we set $C := \frac{2}{L\beta(1 - \sqrt{1 - (1-q)^{-1}})^2}$.

Proof. Let $\epsilon, \nu > 0$ be as in Definition 5.4. As observed above, the function values $\varphi(x_k)$ converge to $\varphi(x^*)$. Hence we may assume all the iterates x_k lie in $[\varphi \leq \varphi(x^*) + \nu]$. We aim now to show that if x_r is sufficiently close to x^* , then all following iterates never leave the ball $B_\epsilon(x^*)$. To this end, let r be an index such that x_r lies in $B_\epsilon(x^*)$ and let $k \geq 1$ be the smallest index satisfying $x_{r+k} \notin B_\epsilon(x^*)$. Defining $\zeta := L\beta(2 + L\beta)\gamma^2$ and using inequality (5.6), we deduce

$$\begin{aligned} \|x_k - x_r\| &\leq \sum_{i=r}^{k-1} \|x_i - x_{i+1}\| \leq \sqrt{2/(L\beta)} \sum_{i=r}^{k-1} \sqrt{\varphi(x_i) - \varphi(x_{i+1})} \\ &\leq \sqrt{2/(L\beta)} \sqrt{\varphi(x_r) - \varphi(x^*)} \sum_{i=r}^k \left(1 - \frac{1}{\zeta}\right)^{\frac{i-r}{2}} \leq \sqrt{\frac{2(\varphi(x_r) - \varphi(x^*))}{L\beta(1 - \zeta^{-1})}}. \end{aligned}$$

Hence if x_r lies in the ball $B_{\epsilon/2}(x^*)$ and is sufficiently close to x^* so that the right-hand-side is smaller than $\frac{\epsilon}{2}$, we obtain a contradiction. Thus there exists an index r so that for all $k \geq r$, the iterates x_k lie in $B_\epsilon(x^*)$. The claimed Q -Linear rate follows immediately. To obtain the R -linear rate of the iterates, we argue as in the proof of Theorem 3.2:

$$\begin{aligned} \|x_{r+k} - x^*\| &\leq \sum_{i=r+k}^{\infty} \|x_i - x_{i+1}\| \leq \sqrt{2/(L\beta)} \sum_{i=r+k}^{\infty} \sqrt{\varphi(x_i) - \varphi(x_{i+1})} \\ &\leq \sqrt{2/(L\beta)} \sqrt{\varphi(x_r) - \varphi(x^*)} \sum_{i=r+k}^{\infty} \left(1 - \frac{1}{\zeta}\right)^{\frac{i-r}{2}} \leq \left(1 - \frac{1}{\zeta}\right)^{k/2} D \sqrt{\varphi(x_r) - \varphi^*}, \end{aligned}$$

where $D = \frac{\sqrt{2}}{\sqrt{L\beta(1 - \sqrt{1 - \zeta^{-1}})}}$. Squaring both sides, the result follows. \square

Theorem 5.5 already marks a point of departure from the convex setting. Gradient descent for an α -strongly convex function with β -Lipschitz gradient converges at the linear rate $1 - \alpha/\beta$. Treating this setting as a special case of composite minimization with $g = 0$ and $h(t) = t$, Theorem 5.5 guarantees the linear rate only on the order of $1 - (\alpha/\beta)^2$. The difference is the lack of convexity; the linearizations $\varphi(x; \cdot)$ no longer lower bound the objective function φ , but only do so up to a quadratic deviation, thereby leading to a worse linear rate of convergence. We put this issue aside for the moment and will revisit it in section 6, where we will show that Algorithm 1 accelerates under a natural *uniform* quadratic growth condition. The ensuing discussion relating the error bound condition and quadratic growth will drive that analysis as well.

Following the pattern of the current work, we seek to interpret the error bound condition in terms of a natural property of the subdifferential $\partial\varphi$. It is tempting to proceed by showing that the step-lengths of Algorithm 1 are proportional to the step-lengths of the proximal point method $z_{k+1} \in (I + t\partial\varphi)^{-1}(z_k)$, as in Theorem 3.5. The following Theorem 5.7 attempts to do just that. We will need the following lemma.

LEMMA 5.6. *The inequality, $\varphi(y) \geq \varphi(x) + \langle v, y - x \rangle - \frac{L\beta}{2}\|y - x\|^2$, holds for all $x, y \in \text{dom}\varphi$ and $v \in \partial\varphi(x)$.*

Proof. Fix a point x and a subgradient $v = z + \nabla c(x)^*w$ for some $z \in \partial g(x)$ and $w \in \partial h(c(x))$. Then for all $y \in \mathbf{R}^n$ we successively deduce

$$\begin{aligned} \varphi(y) &= g(y) + h(c(y)) \geq (g(x) + \langle z, y - x \rangle) + (h(c(x)) + \langle w, c(y) - c(x) \rangle) \\ &\geq \varphi(x) + \langle z, y - x \rangle + \langle w, \nabla c(x)(y - x) \rangle - \frac{L\beta}{2}\|y - x\|^2 \\ &= \varphi(x) + \langle v, y - x \rangle - \frac{L\beta}{2}\|y - x\|^2, \end{aligned}$$

as claimed. \square

First, we establish inequality (5.8), showing that the norm $\|\mathcal{G}_t(x)\|$ is always bounded by twice the stationarity measure $\text{dist}(0; \partial\varphi(x))$. This is a direct analogue of (3.10), though we arrive at it through a different argument. Second, seeking to imitate the key inequality (3.11), we arrive at the inequality (5.9) below. The difficulty is that in this more general setting, the proportionality constant we need (left-hand-side of (5.9)) tends to $+\infty$ as t tends to $(L\beta)^{-1}$ – the most interesting regime. We will circumvent this difficulty in the proof of our main result (Theorem 5.11) by a separate argument using Theorem 5.3; nonetheless, we believe the proportionality inequality (5.9) is interesting in its own right.

THEOREM 5.7 (Step-lengths comparison).

Consider the convex-composite problem (5.1). Then the inequality

$$\frac{1}{2}\|\mathcal{G}_t(x)\| \leq \text{dist}(0; \partial\varphi(x)) \quad \text{holds.} \quad (5.8)$$

Suppose in addition that h is L -Lipschitz continuous and ∇c is β -Lipschitz continuous, and set $r := L\beta$. Then for $t < r^{-1}$ and any point $x^+ \in (I + t\partial\varphi)^{-1}(x)$ the inequality holds:

$$1 + (1 - tr)^{-1} \geq \frac{\|x^t - x\|}{\|x^+ - x\|} + \frac{\|x^+ - x\|}{\|x^t - x\|}. \quad (5.9)$$

Proof. Fix a vector $v \in \partial\varphi(x)$. Then there exist vectors $z \in \partial g(x)$ and $w \in \partial h(c(x))$ satisfying $v = z + \nabla c(x)^*w$. Convexity yields

$$\begin{aligned} g(x^t) + h(c(x) + \nabla c(x)(x^t - x)) &\geq \\ (g(x) + \langle z, x^t - x \rangle) + (h(c(x)) + \langle w, \nabla c(x)(x^t - x) \rangle) &= \varphi(x) + \langle v, x^t - x \rangle. \end{aligned}$$

Appealing to the inequality $\varphi(x) \geq \varphi(x; x^t) + \frac{1}{2t}\|x - x^t\|^2$, we deduce $\|v\| \cdot \|x^t - x\| \geq \frac{1}{2t}\|x^t - x\|^2$, completing the proof of (5.8).

Next, again fix a point x and consider a point $x^+ \in (I + t\partial\varphi)^{-1}(x)$. To simplify notation, set $r := L\beta$. Applying Lemma 5.6 with x replaced by x^+ yields

$$\begin{aligned} \varphi(y) &\geq \varphi(x^+) + \langle t^{-1}(x - x^+), y - x^+ \rangle - \frac{L\beta}{2}\|y - x^+\|^2 \\ &= \varphi(x^+) + \frac{1}{2t}\|x^+ - x\|^2 + \frac{t^{-1} - r}{2}\|x^+ - y\|^2 - \frac{1}{2t}\|y - x\|^2. \end{aligned}$$

Plugging in $y = x^t$, we deduce

$$\begin{aligned} \frac{t^{-1} - r}{2} \|x^+ - x^t\|^2 &\leq (\varphi(x^t) + \frac{1}{2t} \|x^t - x\|^2) - (\varphi(x^+) + \frac{1}{2t} \|x^+ - x\|^2) \\ &\leq \varphi_t(x; x^t) - \varphi_t(x; x^+) + \frac{r}{2} (\|x^t - x\|^2 + \|x^+ - x\|^2). \end{aligned}$$

Taking into account the strong convexity inequality $\varphi_t(x; x^+) \geq \varphi_t(x; x^t) + \frac{1}{2t} \|x^+ - x^t\|^2$, we deduce

$$(2t^{-1} - r) \|x^+ - x^t\|^2 \leq r (\|x^t - x\|^2 + \|x^+ - x\|^2).$$

A short computation then shows

$$1 + (1 - tr)^{-1} \geq \frac{\|x^t - x\|}{\|x^+ - x\|} + \frac{\|x^+ - x\|}{\|x^t - x\|},$$

and the result follows. \square

As alluded to prior to the theorem, the inequality (5.9) is meaningless for $t = 1/L\beta$ – the most interesting case – since the left-hand-side becomes infinite. To circumvent this difficulty, instead of relying on (5.9), we will prove our main result (Theorem 5.11) by using Theorem 5.3.

Next, we introduce the “natural property” of the subdifferential with which we will equate the error bound.

DEFINITION 5.8 (Subregularity). A set-valued mapping $F: \mathbf{R}^n \rightrightarrows \mathbf{R}^m$ is *subregular* at $(\bar{x}, \bar{y}) \in \text{gph } F$ with constant $l > 0$ if there exists a neighborhood \mathcal{X} of \bar{x} satisfying

$$\text{dist}(x; F^{-1}(\bar{y})) \leq l \cdot \text{dist}(\bar{y}; F(x)) \quad \text{for all } x \in \mathcal{X}.$$

Clearly, the error bound property around a stationary point \bar{x} of φ with parameter γ amounts to subregularity of the prox-gradient mapping $\mathcal{G}_t(\cdot)$ at $(\bar{x}, 0)$ with constant γ . We aim to show that the error bound property is equivalent to subregularity of the subdifferential $\partial\varphi$ itself – a transparent notion closely tied to quadratic growth [19, 21, 23]. We first record the following elementary lemma; we omit the proof, as it quickly follows from definitions.

LEMMA 5.9 (Perturbation by identity). Consider a set-valued mapping $S: \mathbf{R}^n \rightrightarrows \mathbf{R}^m$ and a pair $(\bar{x}, 0) \in \text{gph } S$. Then if S is subregular at $(\bar{x}, 0)$ with constant l , the mapping $(I + S^{-1})^{-1}$ is subregular at $(\bar{x}, 0)$ with constant $1 + l$. Conversely, if $(I + S^{-1})^{-1}$ is subregular at $(\bar{x}, 0)$ with constant \hat{l} , then S is subregular at $(\bar{x}, 0)$ with constant $1 + \hat{l}$.

The following result analogous to Theorem 3.4 is now immediate.

THEOREM 5.10 (Proximal and subdifferential subregularity).

Consider the convex-composite problem (5.1) and let \bar{x} be a stationary point of φ . Consider the conditions

1. the subdifferential $\partial\varphi$ is subregular at $(\bar{x}, 0)$ with constant l .
2. the mapping $T := t^{-1}(I - (I + t\partial\varphi)^{-1})$ is subregular at $(\bar{x}, 0)$ with constant \hat{l} .

If condition (1) holds, then so does condition (2) with $\hat{l} = l + t$. Conversely, condition (2) implies condition (1) with $l = \hat{l} + t$.

Proof. Note first the equality $tT = (I + (t\partial\varphi)^{-1})^{-1}$ (equation (2.1)). Suppose that $\partial\varphi$ is l -subregular at $(\bar{x}, 0)$ with constant l . Then by Lemma 5.9, the mapping tT is subregular at $(\bar{x}, 0)$ with constant $1 + l/t$, and hence T is subregular at $(\bar{x}, 0)$ with constant $l + t$, as claimed. The converse argument is analogous. \square

We are now ready to prove the main result of this section.

THEOREM 5.11 (Subdifferential subregularity and the error bound property).

Consider the convex-composite problem (5.1), where h is L -Lipschitz continuous and the Jacobian ∇c is β -Lipschitz. Let \bar{x} be a stationary point of φ and consider the conditions:

1. the subdifferential $\partial\varphi$ is subregular at $(\bar{x}, 0)$ with constant l .
2. the prox-gradient mapping $\mathcal{G}_t(\cdot)$ is subregular at $(\bar{x}, 0)$ with constant \hat{l} .

If condition (2) holds, then condition (1) holds with $l := 2\hat{l}$. Conversely, if condition (1) holds, then condition (2) holds with $\hat{l} := (3L\beta t + 2)l + 2t$.

Proof. Suppose first that the gradient mapping $\mathcal{G}_t(x)$ is subregular at $(\bar{x}, 0)$ with constant \hat{l} . Then by Theorems 5.7 and 5.10, we deduce for all x near \bar{x} , the inequalities

$$\text{dist}(x, (\partial\varphi)^{-1}(0)) \leq \hat{l} \cdot \|\mathcal{G}_t(x)\| \leq 2\hat{l} \cdot \text{dist}(0; \partial\varphi(x)).$$

Conversely, suppose that $\partial\varphi$ is subregular at $(\bar{x}, 0)$ with constant l . Fix a point x , and let \hat{x} be the point guaranteed to exist by Theorem 5.3. For the purpose of establishing subregularity of $\mathcal{G}_t(\cdot)$, we can suppose the x and x_t are arbitrarily close to \bar{x} . Then \hat{x} is close to \bar{x} , and we deduce

$$\begin{aligned} l \cdot \text{dist}(0, \partial\varphi(\hat{x})) &\geq \text{dist}(\hat{x}; (\partial\varphi)^{-1}(0)) \geq \text{dist}(x; (\partial\varphi)^{-1}(0)) - \|x^t - \hat{x}\| - \|x^t - x\| \\ &\geq \text{dist}(x; (\partial\varphi)^{-1}(0)) - 2\|x^t - x\|. \end{aligned}$$

We conclude $\text{dist}(x; (\partial\varphi)^{-1}(0)) \leq (l(3L\beta + 2t^{-1}) + 2)\|x^t - x\|$, as claimed. \square

Thus subregularity of the subdifferential $\partial\varphi$ and the error bound property are identical notions, with a precise relationship between the constants. Subdifferential subregularity at a unique minimizer, on the other hand, is equivalent to the natural quadratic growth condition when the functions in question are semi-algebraic (or more generally tame) [19] or convex (Theorem 3.3).

REMARK 5.1. As mentioned above, the inequality (5.9) is meaningless for $t = 1/L\beta$, and therefore the proof of our main result Theorem 5.11 followed by an indirect argument based on Ekeland's variational principle. The follow-up paper [24, Theorem 4.3] overcomes this difficulty directly. It shows that the step-lengths of Algorithm 1 indeed linearly bound the step-lengths of the proximal point method: $4 \cdot \|\mathcal{G}_{1/L\beta}\| \geq \|2L\beta(x - \text{prox}_{F/(2L\beta)}(x))\|$. Since this result appeared after the current manuscript was well into the reviewing process, we do not use it in the current work.

6. Natural linear rate of convergence under tilt-stability As we alluded to in section 5, the linear rate at which Algorithm 1 converges under the error bound condition is an order of magnitude slower than the rate that one would expect. In section 5, we highlighted the equivalence between the error bound condition and subregularity of the subdifferential, and their close relationships to quadratic growth. We will now show that when these properties hold *uniformly* relative to tilt-perturbations, the algorithm accelerates to the natural rate.

We begin with the following definition.

DEFINITION 6.1 (Stable strong local minimizer). We say that \bar{x} is a *stable strong local minimizer with constant* $\alpha > 0$ of a function $f: \mathbf{R}^n \rightarrow \bar{\mathbf{R}}$ if there exists a neighborhood \mathcal{X} of \bar{x} so that for each vector v near the origin, there is a point x_v (necessarily unique) in \mathcal{X} , with $x_0 = \bar{x}$, so that in terms of the perturbed functions $f_v := f(\cdot) - \langle v, \cdot \rangle$, the inequality

$$f_v(x) \geq f_v(x_v) + \frac{\alpha}{2}\|x - x_v\|^2 \quad \text{holds for all } x \in \mathcal{X}.$$

This type of uniform quadratic growth is known to be equivalent to a number of influential notions, such as tilt-stability [42] and strong metric regularity of the subdifferential [21]. Here, we specialize the discussion to the convex-composite case, though the relationships hold much more generally. The following theorem appears in [23, Theorem 3.7, Proposition 4.5]; some predecessors were proved in [21, 34].

THEOREM 6.2 (Uniform growth, tilt-stability, & strong regularity).

Consider the convex-composite problem (5.1) and let x^* be a local minimizer of φ . Then the following properties are equivalent.

1. **(uniform quadratic growth)** The point x^* is a stable strong local minimizer of φ with constant α .

2. **(local subdifferential convexity)** There exists a neighborhood \mathcal{X} of x^* so that for any sufficiently small vector v , there is a point $x_v \in \mathcal{X} \cap (\partial\varphi)^{-1}(v)$ so that the inequality

$$\varphi(x) \geq \varphi(x_v) + \langle v, x - x_v \rangle + \frac{\alpha}{2} \|x - x_v\|^2 \quad \text{holds for all } x \in \mathcal{X}.$$

3. **(tilt-stability)** There exists a neighborhood \mathcal{X} of x^* so that the mapping

$$v \mapsto \arg \min_{x \in \mathcal{X}} \{\varphi(x) - \langle v, x \rangle\}$$

is single-valued and $1/\alpha$ -Lipschitz continuous on some neighborhood of the origin.

4. **(strong regularity of the subdifferential)** There exist neighborhoods \mathcal{X} of x^* and \mathcal{V} of $\bar{v} = 0$ so that the restriction $(\partial\varphi)^{-1}: \mathcal{V} \rightrightarrows \mathcal{X}$ is a single-valued $1/\alpha$ -Lipschitz continuous mapping.

There has been a lot of recent work aimed at characterizing the above properties in concrete circumstances; see e.g. [35, 36, 37, 38]. Suppose that the equivalent conditions in Theorem 6.2 hold. We will now investigate the impact of such an assumption on the linear convergence of Algorithm 1. Assume $t \leq (L\beta)^{-1}$. Observe that property 4 directly implies that $\partial\varphi$ is subregular at $(x^*, 0)$ with constant $1/\alpha$. Consequently, local linear convergence with the rate on the order of $1 - (\frac{\alpha}{\beta L})^2$ is already assured by Theorems 5.5 and 5.11; our goal is to derive a faster rate.

Consider a point x near \hat{x} . Let \hat{x} then be the point guaranteed to exist by Theorem 5.3, and set v to be the minimal norm vector in the subdifferential $\partial\varphi(\hat{x})$. Note the inequality $\|v\| \leq (3L\beta t + 2)\|\mathcal{G}_t(x)\| \leq 5\|\mathcal{G}_t(x)\|$. Hence for any point y near x^* , we have

$$\begin{aligned} \varphi(y) &\geq \varphi(\hat{x}) + \langle v, y - \hat{x} \rangle && \text{(strong-regularity)} \\ &\geq \varphi_t(x, \hat{x}) - L\beta \|x - \hat{x}\|^2 + \langle v, y - \hat{x} \rangle && \text{(inequality (5.2))} \\ &\geq \varphi_t(x, x^t) - L\beta \|x - \hat{x}\|^2 + \langle v, y - \hat{x} \rangle && \text{(definition of } x^t) \\ &\geq \varphi(x^t) - L\beta \|x - \hat{x}\|^2 + \langle v, y - \hat{x} \rangle && \text{(inequality (5.2)).} \end{aligned}$$

Plugging in $y = x^*$, we deduce

$$\begin{aligned} \varphi(x^t) - \varphi(x^*) &\leq L\beta \|x - \hat{x}\|^2 + \langle v, \hat{x} - x^* \rangle \\ &\leq 4L\beta \|x^t - x\|^2 + \|v\| \cdot (\|\hat{x} - x^t\| + \|x^t - x\| + \|x - x^*\|) \\ &\leq \frac{4}{L\beta} \|\mathcal{G}_t(x)\|^2 + 5\|\mathcal{G}_t(x)\| \cdot (2\|x^t - x\| + \|x - x^*\|) \\ &\leq \frac{14}{L\beta} \|\mathcal{G}_t(x)\|^2 + 5\|\mathcal{G}_t(x)\| \cdot \|x - x^*\| \\ &= \frac{\|\mathcal{G}_t(x)\|^2}{L\beta} \left(14 + 5L\beta \frac{\|x - x^*\|}{\|\mathcal{G}_t(x)\|} \right). \end{aligned}$$

Hence if while Algorithm 1 is running, the fractions $\frac{\|x_k - x^*\|}{\|\mathcal{G}_t(x_k)\|}$ remain bounded by a constant γ , appealing to the descent inequality (5.6), we obtain the Q-linear convergence guarantee

$$\varphi(x_{k+1}) - \varphi(x^*) \leq \left(1 - \frac{1}{25 + 10L\beta\gamma} \right) (\varphi(x_k) - \varphi(x^*)). \quad (6.1)$$

We have thus established the main result of this section.

THEOREM 6.3 (Natural rate of convergence). *Consider the convex-composite problem (5.1), where h is L -Lipschitz continuous and the Jacobian ∇c is β -Lipschitz. Let x_k be the sequence generated by Algorithm 1 with $t \leq (L\beta)^{-1}$. Suppose that x_k has some limit point x^* around which one of the equivalent properties in Theorem 6.2 hold. Define the fraction*

$$q := 1 - \frac{1}{45 + 50L\beta/\alpha}.$$

Then for all large k , function values converge Q -linearly

$$\varphi(x_{k+1}) - \varphi(x^*) \leq q \cdot (\varphi(x_k) - \varphi(x^*)),$$

while the x_k asymptotically converge R -linearly: there exists an index r such that the inequality

$$\|x_{r+k} - x^*\|^2 \leq q^k \cdot C \cdot (\varphi(x_r) - \varphi^*),$$

holds for all $k \geq 1$, where we set $C := \frac{2}{L\beta(1 - \sqrt{1 - (1-q)^{-1}})^2}$.

Proof. Strong regularity of the subdifferential in Theorem 6.2, in particular, implies that the subdifferential mapping $\partial\varphi$ is subregular at $(x^*, 0)$ with constant $1/\alpha$. Theorem 5.11 then implies that the error bound condition holds near x^* with constant $\frac{5}{\alpha} + \frac{2}{L\beta}$. Theorem 5.5 then shows that the sequence x_k converges to x^* . Consequently for all large indices k , the ratios $\frac{\|x_k - x^*\|}{\|\mathcal{G}_t(x_k)\|}$ are bounded by $\frac{5}{\alpha} + \frac{2}{L\beta}$. The inequality (6.1) immediately yields the claimed Q -linear rate. The R -linear rate follows easily by a standard argument, as in the proof of Theorem 5.5. \square

In summary, Theorem 6.3 shows that if the prox-linear method is initialized sufficiently close to a stable strong local minimizer x^* of φ with constant α , then the function values converge at a linear rate on the order of $1 - \frac{\alpha}{L\beta}$. This is in contrast to the slower rate $1 - (\frac{\alpha}{L\beta})^2$ established in Theorem 5.5 under the weaker condition that $\partial\varphi$ is subregular at $(x^*, 0)$ with constant $1/\alpha$.

7. Quadratic convergence of the prox-linear method In this section, we discuss local quadratic convergence of the prox-linear scheme under a sharpness condition. Local quadratic convergence under such conditions for the trust region variant of the scheme is already known from [12]. The proximal setting, explored here, requires a different analysis. Before we proceed, some commentary is in order.

First, in the previous sections, we have seen that local linear convergence is assured around \bar{x} when the subdifferential $\partial\varphi$ is subregular at $(\bar{x}, 0)$. Local quadratic convergence, however, requires a stronger condition. One such condition is that of \bar{x} being an “isolated sharp minimizer” of φ – a notion originating in [16, 44]. Indeed, under this condition, the proof of quadratic convergence is immediate.

THEOREM 7.1 (Local quadratic convergence: isolated sharpness). *Consider the sequence x_k generated by Algorithm 1 with $t \leq (L\beta)^{-1}$, and suppose that x_k has some limit point x^* . Suppose there is a neighborhood \mathcal{X} of x^* and a real $c > 0$ such that isolated sharpness holds:*

$$\varphi(x) \geq \varphi(x^*) + c \cdot \|x - x^*\| \quad \text{for all } x \in \mathcal{X}. \quad (7.1)$$

Then the iterates x_k asymptotically converge quadratically to x^* :

$$\|x_{k+1} - x^*\| \leq \frac{L\beta + t^{-1}}{2c} \cdot \|x_k - x^*\|^2 \quad \text{holds whenever } x_{k+1} \in \mathcal{X}.$$

Proof. Assuming x_{k+1} lies in \mathcal{X} , observe

$$\begin{aligned}\varphi(x_{k+1}) &\leq \varphi_t(x_k; x_{k+1}) \leq \varphi_t(x_k; x^*) = \varphi(x_k; x^*) + \frac{1}{2t}\|x_k - x^*\|^2 \\ &\leq \varphi(x^*) + \frac{L\beta+t^{-1}}{2}\|x_k - x^*\|^2.\end{aligned}$$

Rearranging and using isolated sharpness yields

$$\frac{L\beta+t^{-1}}{2}\|x_k - x^*\|^2 \geq \varphi(x_{k+1}) - \varphi(x^*) \geq c \cdot \|x_{k+1} - x^*\|.$$

The result follows. \square

Despite the elementary proof, the theorem above has a serious deficiency. The isolated sharpness assumption (7.1) is very restrictive, in particular implying that \bar{x} is an isolated local minimizer of φ . A much looser condition is that there exists a neighborhood \mathcal{X} of \bar{x} such that

$$\varphi(x) \geq \varphi(\text{proj}_S(x)) + c \cdot \text{dist}(x, S) \quad \text{for all } x \in \mathcal{X},$$

where S is the set of critical points of φ . This notion originates in [11], under the name of “weak sharp minima”. This condition in the composite setting is equivalent to requiring the quantity $\text{dist}(0, \partial f(x))$ to be bounded away from zero for all $x \notin S$ near \bar{x} , and thus is a natural strengthening of subdifferential subregularity. Our proof of quadratic convergence is based on two pillars: (1) sharpness implies that the $\|\mathcal{G}_t(x)\|$ linearly upper bounds $\text{dist}(x, S)$, and conversely (2) the distance $\text{dist}(x, S)$ always linearly upper bounds $\|\mathcal{G}_t(x)\|$. We have already established the first pillar (Theorem 5.11); the following lemma establishes the second pillar. The argument is based on Ekeland’s principle and follows closely the proof of Theorem 5.3.

LEMMA 7.1 (Upper bounding the prox-gradient). *Consider the convex-composite problem (5.1), where h is L -Lipschitz continuous and the Jacobian ∇c is β -Lipschitz. Suppose $t \leq (L\beta)^{-1}$ and let S be the set of critical points of φ . Then the inequality*

$$\|\mathcal{G}_t(x)\| \leq (3t^{-1} + 4L\beta) \cdot \text{dist}(x; S) \quad \text{holds for all } x \in \mathbf{R}^n.$$

Proof. Define the function $\psi(y) := \varphi(y) + L\beta\|y - x\|^2$. Notice that ψ is $L\beta$ -strongly convex and therefore has some minimizer $\bar{x} := \text{argmin}_y \psi(y)$. Let $x^* \in S$ be arbitrary. We deduce

$$(\varphi(\bar{x}) + L\beta\|\bar{x} - x\|^2) \leq (\varphi(x^*) + L\beta\|x^* - x\|^2) - \frac{L\beta}{2}\|x^* - \bar{x}\|^2$$

Since x^* is a critical point of φ , we deduce $\varphi(x^*) \leq \varphi(\bar{x}) + \frac{L\beta}{2}\|x^* - \bar{x}\|^2$. Plugging this back into the inequality above and rearranging, we deduce $\|\bar{x} - x\| \leq \|x^* - x\|$. Thus to prove the theorem, we must only upper bound $\|\mathcal{G}_t(x)\|$ by a multiple of $\|\bar{x} - x\|$.

The argument proceeds analogously to Theorem 5.3. The inequalities (5.2) imply

$$\varphi_t(x; y) \geq \psi(\bar{x}) + \frac{t^{-1} - 3L\beta}{2}\|y - x\|^2 \geq \psi(\bar{x}) - L\beta\|y - x\|^2 \quad \text{for all } y \in \mathbf{R}^n.$$

Defining $\zeta(y) := \varphi_t(x; y) + L\beta\|y - x\|^2$, a short computation using (5.2) then shows

$$\zeta(\bar{x}) - \inf \zeta \leq \left(\varphi_t(x; \bar{x}) + L\beta\|\bar{x} - x\|^2 \right) - \psi(\bar{x}) \leq t^{-1}\|\bar{x} - x\|^2.$$

Setting $\epsilon := t^{-1}\|\bar{x} - x\|^2$ and $\rho = t^{-1}\|\bar{x} - x\|$ and applying the variational principle (Theorem 5.2) to ζ , we deduce that there exists a point \hat{x} satisfying $\|\bar{x} - \hat{x}\| \leq \|\bar{x} - x\|$ and $\text{dist}(0; \partial\zeta(\hat{x})) \leq t^{-1}\|\bar{x} - x\|$. The triangle inequality in turn implies

$$\text{dist}(0; \partial\varphi_t(x; \hat{x})) \leq 2L\beta \cdot \|\hat{x} - x\| + \text{dist}(0; \partial\zeta(\hat{x})) \leq 2L\beta\|\hat{x} - x\| + t^{-1}\|\bar{x} - x\|. \quad (7.2)$$

Next since $\varphi_t(x; \cdot)$ is t^{-1} -strongly convex, we deduce

$$\text{dist}(0; \partial\varphi_t(x; \hat{x})) \geq t^{-1}\|\hat{x} - x^t\| \geq t^{-1}\|x^t - x\| - t^{-1}\|x - \hat{x}\|.$$

Combining this estimate with (7.2), we obtain

$$\|\mathcal{G}_t(x)\| \leq (t^{-1} + 2L\beta)\|x - \hat{x}\| + t^{-1}\|\bar{x} - x\| \leq (3t^{-1} + 4L\beta)\|\bar{x} - x\|.$$

The result follows. \square

We can now establish the desired quadratic convergence guarantee.

THEOREM 7.2 (Local quadratic convergence: sharpness). *Consider the sequence x_k generated by Algorithm 1 with $t \leq (L\beta)^{-1}$, and suppose that x_k has some limit point x^* . Let S be the set of critical points of φ and suppose there is a neighborhood \mathcal{X} of x^* and a real $c > 0$ such that the sharpness condition holds:*

$$\varphi(x) \geq \varphi(\text{proj}_S(x)) + c \cdot \text{dist}(x, S) \quad \text{for all } x \in \mathcal{X}.$$

Then the iterates x_k converge quadratically: there is a constant q such that the inequality

$$\|x_{k+1} - x^*\| \leq q \cdot \|x_k - x^*\|^2 \quad \text{holds for all large } k.$$

Proof. We first claim that for all $x \notin S$ near \bar{x} , the quantity $\text{dist}(0; \partial\varphi(x))$ is bounded away from zero. Indeed, fix a point $x \in \mathcal{X} \setminus S$ and $v \in \partial f(x)$ and set $z := \text{proj}_S(x)$. Then we deduce

$$\varphi(x) - c \cdot \|x - z\| \geq \varphi(z) \geq \varphi(x) + \langle v, z - x \rangle - \frac{L\beta}{2}\|z - x\|^2,$$

where the last inequality follows from Lemma 5.6. Canceling $\varphi(x)$ from both sides and dividing through by $\|z - x\|$, we deduce

$$\|v\| - c \geq -\frac{L\beta}{2}\|z - x\| \geq -\frac{L\beta}{2}\|x - x^*\|.$$

Therefore we can be sure that v must be bounded away from zero whenever x is close to x^* . In particular, ∂f is subregular at $(x^*, 0)$. It follows from Theorems 5.5 and 5.11 that the iterates x_k converge to x^* .

Moreover, taking into account (5.7) and Theorem 5.11, we deduce that there exists a constant $a > 0$ such that

$$\varphi(x_{k+1}) - \varphi(x^*) \leq a \cdot \|\mathcal{G}_t(x_k)\|^2,$$

for all large k . Taking into account the sharpness assumption and Lemma 7.1, we deduce

$$\|x_{k+1} - x^*\| \leq a \cdot (3t^{-1} + 4L\beta)^2 \|x_k - x^*\|^2,$$

for all large k , as claimed. \square

8. Proximal gradient method without convexity In this section, we revisit the proximal gradient method, discussed in Section 3 in absence of convexity in f . To this end, consider the optimization problem

$$\min_{x \in \mathbf{R}^n} \varphi(x) := f(x) + g(x),$$

where $f: \mathbf{R}^n \rightarrow \mathbf{R}$ is a C^1 -smooth function with β -Lipschitz gradient and $g: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ is a closed convex function. Observe that the proximal gradient method:

$$x_{k+1} = x_k - t\mathcal{G}_t(x_k),$$

with the gradient mapping

$$\mathcal{G}_t(x) := t^{-1} \left(x - \text{prox}_{tg}(x - t\nabla f(x)) \right),$$

is simply an instance of the prox-linear algorithm (Algorithm 1) applied to the function $\varphi = g + \text{Id} \circ f$. Here Id is the identity map on \mathbf{R} . Hence all the results of sections 5 and 6 apply immediately with $L = 1$. The arguments in this additive case are much simpler, starting with the fact that Ekeland’s variational principle is no longer required to prove that subdifferential subregularity is equivalent to the error bound property.

Indeed, the key perturbation result Theorem 5.3 simplifies drastically: in the notation of the theorem, we can set $\hat{x} := x^t$. Then the optimality conditions for the proximal subproblem

$$\mathcal{G}_t(x) \in \nabla f(x) + \partial g(x^t)$$

immediately imply the slightly improved estimate in Theorem 5.3:

$$\text{dist}(0, \partial\varphi(x^t)) \leq (1 + \beta t) \|\mathcal{G}_t(x)\|. \quad (8.1)$$

As in Theorem 5.11, consider now the two properties:

1. the subdifferential $\partial\varphi$ is subregular at $(\bar{x}, 0)$ with constant l .
2. the prox-gradient mapping $\mathcal{G}_t(\cdot)$ is subregular at $(\bar{x}, 0)$ with constant \hat{l} .

The same argument as in Theorem 5.11 shows that if (1) holds, then (2) holds with $\hat{l} = l(\beta t + 1) + t$. Conversely if (2) is valid, then (1) holds with $l = 2\hat{l}$.

Next, we reevaluate convergence guarantees under tilt-stability. To this end, suppose that one of the equivalent conditions in Theorem 6.2 holds. Set for simplicity $t = \beta^{-1}$ and let v be the minimal norm element of $\partial\varphi(x^t)$. We then deduce for all points x and y near x^* the inequality

$$\varphi(y) \geq \varphi(x^t) + \langle v, y - x^t \rangle.$$

Plugging in $y = x^*$ yields

$$\varphi(x^t) - \varphi(x^*) \leq \|\mathcal{G}_t(x)\|^2 \left(\frac{\|v\| \cdot \|x^* - x^t\|}{\|\mathcal{G}_t(x)\|^2} \right).$$

Appealing to Theorem 6.2 and the equivalence of subdifferential subregularity and the error bound property above, we deduce $\frac{\|x^* - x^t\|}{\|\mathcal{G}_t(x)\|} \leq \frac{\|x^* - x\|}{\|\mathcal{G}_t(x)\|} + t \leq \alpha^{-1}(\beta t + 1) + 2t$. Taking into account the inequality (8.1), we obtain

$$\begin{aligned} \varphi(x^t) - \varphi(x^*) &\leq \|\mathcal{G}_t(x)\|^2 (1 + \beta t)(\alpha^{-1}(\beta t + 1) + 2t) \\ &\leq 8\beta(\varphi(x) - \varphi(x^t)) (\alpha^{-1} + \beta^{-1}). \end{aligned}$$

where the last inequality follows from (5.5). Trivial algebraic manipulations then yield the Q-linear rate of convergence

$$\varphi(x_{k+1}) - \varphi(x^*) \leq \left(1 - \frac{1}{9 + 8\beta/\alpha} \right) (\varphi(x_k) - \varphi(x^*))$$

for all sufficiently large indices k . As an aside, this estimate slightly improves on the constants appearing in Theorem 6.3 for this class of problems.

9. The proximal subproblem in full generality In this final section, we build on the convex composite framework explored in sections 5 and 6 by dropping convexity and finite-valuedness assumptions. Specifically, we begin in great generality with the problem

$$\min_x \varphi(x) := g(x) + h(c(x)), \quad (9.1)$$

where $c: \mathbf{R}^n \rightarrow \mathbf{R}^m$ is a C^1 -smooth mapping and $g, h: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ are merely closed functions. As in section 5, define for any points $x, y \in \mathbf{R}^n$ the linearized function

$$\varphi(x; y) := g(y) + h(c(x) + \nabla c(x)(y - x)),$$

and for any real $t > 0$, the quadratic perturbation

$$\varphi_t(x; y) := \varphi(x; y) + \frac{1}{2t} \|x - y\|^2.$$

It is tempting to simply apply the prox-linear algorithm directly to this setting by iteratively solving the proximal subproblems $\min_y \varphi_t(x; y)$ for appropriately chosen reals $t > 0$. This naive strategy is fundamentally flawed. There are two main difficulties. First, in contrast to the previous sections, the functions $\varphi(x; \cdot)$ and $\varphi_t(x; \cdot)$ are typically nonconvex. Hence when solving the subproblems, one must settle only for “stationary points” of $\varphi_t(x; \cdot)$. Second, and much more importantly, the iterates generated by such a scheme can quickly yield infeasible proximal subproblems, thereby stalling the algorithm. Designing a variant of the prox-linear algorithm that overcomes the latter difficulty is an ongoing research direction and will be pursued in future work. Nonetheless, it is intuitively clear that the proximal subproblems may still enter the picture. In this section, we show that the equivalence between the “error bound property” and subdifferential subregularity in Theorem 5.11 extends to this broader setting. The technical content of this section is based on a careful mix of variational analytic techniques, which we believe is of independent interest.

For simplicity, we will assume $g = 0$ throughout, that is we consider the problem

$$\min_x \varphi(x) = h(c(x)), \quad (9.2)$$

where $c: \mathbf{R}^n \rightarrow \mathbf{R}^m$ is C^1 -smooth and $h: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ is closed. The assumption $g = 0$ is purely for convenience: all results in this section extend verbatim to the more general setting with identical proofs. As alluded to above, we will be interested in “stationary points” of nonsmooth and nonconvex functions φ and $\varphi_t(x; \cdot)$. To make this notion precise, we appeal to a central variational analytic construction, the subdifferential; see e.g. [26, 33, 50].

DEFINITION 9.1 (Subdifferentials and stationary points).

Consider a closed function $f: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ and a point \bar{x} with $f(\bar{x})$ finite.

1. The *proximal subdifferential* of f at \bar{x} , denoted $\partial_p f(\bar{x})$, consists of all $v \in \mathbf{R}^n$ for which there is a neighborhood \mathcal{X} of \bar{x} and a constant $r > 0$ satisfying

$$f(x) \geq f(\bar{x}) + \langle v, x - \bar{x} \rangle - \frac{r}{2} \|x - \bar{x}\|^2 \quad \text{for all } x \in \mathcal{X}.$$

2. The *limiting subdifferential* of f at \bar{x} , denoted $\partial f(\bar{x})$, consists of all vectors $v \in \mathbf{R}^n$ for which there exist sequences $x_i \in \mathbf{R}^n$ and $v_i \in \partial_p f(x_i)$ with $(x_i, f(x_i), v_i)$ converging to $(\bar{x}, f(\bar{x}), v)$.

3. The *horizon subdifferential* of f at \bar{x} , denoted $\partial^\infty f(\bar{x})$, consists of all vectors $v \in \mathbf{R}^n$ for which there exist points $x_i \in \mathbf{R}^n$, vectors $v_i \in \partial f(x_i)$, and real numbers $t_i \searrow 0$ with $(x_i, f(x_i), t_i v_i)$ converging to $(\bar{x}, f(\bar{x}), v)$.

We say that \bar{x} is a *stationary point* of f whenever the inclusion $0 \in \partial f(\bar{x})$ holds.

Proximal and limiting subdifferentials of a convex function coincide with the usual convex subdifferential, as defined in section 2. The horizon subdifferential plays an entirely different role, detecting horizontal normals to the epigraph of the function. For example, a closed function $f: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ is locally Lipschitz continuous around \bar{x} if and only if $\partial^\infty f(\bar{x}) = \{0\}$. Moreover, the horizon subdifferential plays a decisive role in establishing calculus rules and in stability analysis. We introduce the following notation to help the exposition.

DEFINITION 9.2 (Transversality). Consider a C^1 -smooth mapping $c: \mathbf{R}^n \rightarrow \mathbf{R}^m$, a function $h: \mathbf{R}^m \rightarrow \overline{\mathbf{R}}$, and a point \bar{x} with $h(c(\bar{x}))$ finite. We say that c is *transverse to h at \bar{x}* , denoted $c \pitchfork_{\bar{x}} h$, if the condition holds:

$$\partial^\infty h(c(\bar{x})) \cap \text{Null}(\nabla c(\bar{x})^*) = \{0\}, \quad (9.3)$$

If this condition holds with h an indicator function of a set Q , then we say that c is *transverse to Q at \bar{x}* , and denote it by $c \pitchfork_{\bar{x}} Q$.

Transversality unifies both the classical notion of “transverse intersections” in differential manifold theory (e.g. [28, Section 6]) and the Mangasarian-Fromovitz constraint qualification in nonlinear programming (e.g. [50, Example 9.44], [49, Example 4D.3]). Whenever a C^1 -smooth mapping c is transverse to a closed function h at \bar{x} , the key inclusion

$$\partial\varphi(\bar{x}) \subseteq \nabla c(\bar{x})^* \partial h(c(\bar{x})) \quad \text{holds,}$$

resembling the usual chain rule for smooth compositions. Equality is valid under additional assumptions, such as that $\partial_p h(c(\bar{x}))$ and $\partial h(c(\bar{x}))$ coincide for example.

We now define the *stationary point map*

$$\mathcal{S}_t(x) = \{z \in \mathbf{R}^n : z \text{ is a stationary point of } \varphi_t(x, \cdot)\}$$

and the *prox-gradient mapping*

$$\mathcal{G}_t(x) = t^{-1}(x - \mathcal{S}_t(x)).$$

Notice that both \mathcal{S}_t and \mathcal{G}_t are now set-valued operators. Our goal is to relate subregularity of \mathcal{G}_t to subregularity of the subdifferential $\partial\varphi$ itself. The general trend of our arguments follows that of section 5. There are important complications, however, that must be surmounted. An immediate difficulty is that since h is possibly infinite-valued, it is not possible to directly relate the values $\varphi(y)$ and $\varphi(x; y)$, as in inequality (5.2) – the starting point of the analysis in section 5. For example, $\varphi(y)$ can easily be infinite while $\varphi(x; y)$ is finite, or vice versa. The key idea is that such a comparison is possible if we allow a small perturbation of the point y . The following section is dedicated to establishing this result (Theorem 9.6).

9.1. The comparison theorem We will establish Theorem 9.6 by appealing to the fundamental relationship between transversality (9.3), metric regularity of constraint systems, and stability of metric regularity under linear perturbations [18]. We begin by recalling the concept of metric regularity – a uniform version of subregularity (Definition 5.8). For a discussion on the role of metric regularity in nonsmooth optimization, see for example [26, 49].

DEFINITION 9.3 (Metric regularity). A set-valued mapping $G: \mathbf{R}^n \rightrightarrows \mathbf{R}^m$ is *metrically regular around $(\bar{x}, \bar{y}) \in \text{gph } G$* with constant $l > 0$ if there exists a neighborhood \mathcal{X} of \bar{x} and a neighborhood \mathcal{Y} of \bar{y} satisfying

$$\text{dist}(x; G^{-1}(y)) \leq l \cdot \text{dist}(y; G(x)) \quad \text{for all } x \in \mathcal{X}, y \in \mathcal{Y}.$$

Metric regularity, unlike subregularity, is stable under small linear perturbations. The following is a direct consequence of the proof of [18, Theorem 3.3], as described in [29, Theorem 4.2].

THEOREM 9.4 (Metric regularity under perturbation). *Consider a closed set-valued mapping $G: \mathbf{R}^n \rightrightarrows \mathbf{R}^m$ that is metrically regular around $(\bar{x}, \bar{y}) \in \text{gph } G$. Then there exist constants $\epsilon, \gamma > 0$ and neighborhoods \mathcal{X} of \bar{x} and \mathcal{Y} of \bar{y} , so that the estimate*

$$\text{dist}(x, (G + H)^{-1}(y)) \leq \gamma \cdot \text{dist}(y, (G + H)(x)) \quad \text{holds}$$

for any affine mapping H with $\|\nabla H\| < \epsilon$, any $x \in \mathcal{X}$, and any $y \in H(\bar{x}) + \mathcal{Y}$.

The following consequence for stability of constraint systems is now immediate. For any set $Q \subset \mathbf{R}^n$ and a point $x \in Q$, we define the *limiting normal cone* by $N_Q(x) := \partial\delta_Q(x)$.

COROLLARY 9.5 (Uniform metric regularity of constraint systems).

Fix a C^1 -smooth mapping $F: \mathbf{R}^n \rightarrow \mathbf{R}^m$ and a closed set $Q \subset \mathbf{R}^m$, and consider the constraint system

$$F(x) \in Q.$$

Suppose that F is transverse to Q at some point \bar{x} , with $F(\bar{x}) \in Q$. Then there are constants $\epsilon, \gamma > 0$ and a neighborhood \mathcal{X} of \bar{x} so that for any $x \in \mathcal{X}$ and any affine mapping H with $\|\nabla H\| < \epsilon$ and $\|H(\bar{x})\| < \epsilon$, there exists a point z satisfying

$$(F + H)(z) \in Q \quad \text{and} \quad \|x - z\| \leq \gamma \cdot \text{dist}\left((F + H)(x); Q\right).$$

Proof. Consider the set-valued mapping $G(x) := F(x) - Q$. It is well-known (e.g. [50, Example 9.44]) that the condition $F \pitchfork_{\bar{x}} Q$ is equivalent to G being metrically regular around $(\bar{x}, 0)$. Appealing to Theorem 9.4, we deduce that there are constants $\epsilon, \gamma > 0$ and neighborhoods \mathcal{X} of \bar{x} and \mathcal{Y} of 0, so that for any $x \in \mathcal{X}$ and $y \in H(\bar{x}) + \mathcal{Y}$ and any affine mapping H with $\|\nabla H\| < \epsilon$, there exists a point z satisfying

$$F(z) + H(z) \in y + Q \quad \text{and} \quad \|x - z\| \leq \gamma \text{dist}(F(x) + H(x), y + Q).$$

In light of the assumption $\|H(\bar{x})\| < \epsilon$, decreasing ϵ , we can be sure that $-H(\bar{x})$ lies in \mathcal{Y} and hence we can set $y := 0$ above. The result follows immediately. \square

Armed with Corollary 9.5, we can now prove the main result of this section, playing the role of inequalities (5.2).

THEOREM 9.6 (Comparison inequalities). *Consider the composite problem (9.2) satisfying $c \pitchfork_{\bar{x}} h$ for some point $\bar{x} \in \text{dom } \varphi$, around which ∇c is Lipschitz continuous. Then there exist $\epsilon, \gamma > 0$ and a neighborhood \mathcal{X} of \bar{x} such that the following hold.*

1. *For any two points $x, y \in \mathcal{X}$ with $|\varphi(y) - \varphi(\bar{x})| < \epsilon$, there exists a point y^- satisfying*

$$\|y - y^-\| \leq \gamma \|y - x\|^2 \quad \text{and} \quad \varphi(x; y^-) \leq \varphi(y) + \gamma \|y - x\|^2.$$

2. *For any two points $x, y \in \mathcal{X}$ with $|\varphi(x; y) - \varphi(\bar{x})| < \epsilon$, there exists a point y^+ satisfying*

$$\|y - y^+\| \leq \gamma \|y - x\|^2 \quad \text{and} \quad \varphi(y^+) \leq \varphi(x; y) + \gamma \|y - x\|^2.$$

Proof. The second claim appears as [29, Theorem 4.6].⁵ To see the first claim, for any point x define the affine mapping $L_x(z) = c(x) + \nabla c(x)(z - x)$. Consider now the affine mapping $F: \mathbf{R}^{n+1} \rightarrow \mathbf{R}^{m+1}$ given by

$$F(z, t) := (c(\bar{x}) + \nabla c(\bar{x})(z - \bar{x}), t).$$

⁵ In [29, Theorem 4.6], it is assumed that c is C^2 -smooth; however, it is easy to verify from the proof that the same result holds if ∇c is only locally Lipschitz continuous around \bar{x} .

Then the transversality condition $c \pitchfork_{\bar{x}} h$, along with the epigraphical characterization of subgradients [50, Theorem 8.9], implies that F is transverse to $\text{epi } h$ at $(\bar{x}, \varphi(\bar{x}))$. Hence we can apply Corollary 9.5 with $Q := \text{epi } h$. Let $\epsilon, \gamma > 0$ and \mathcal{X} be the resulting constants and a neighborhood of $(\bar{x}, \varphi(\bar{x}))$, respectively. Define now the affine mappings $H_x(z, t) := (L_x(z) - L_{\bar{x}}(z), 0)$. Observe that for all x sufficiently close to \bar{x} , the inequalities $\|\nabla H_x\| < \epsilon$ and $\|H_x(\bar{x})\| < \epsilon$ hold.

We deduce that for all pairs $(y, \varphi(y))$ sufficiently close to $(\bar{x}, \varphi(\bar{x}))$, and for all points x sufficiently close to \bar{x} , there exists a pair (y^-, t) satisfying

$$(L_x(y^-), t) = (F + H_x)(y^-, t) \in \text{epi } h$$

and

$$\begin{aligned} \|(y, \varphi(y)) - (y^-, t)\| &\leq \gamma \text{dist}\left((F + H_x)(y, \varphi(y)); \text{epi } h\right) \\ &= \gamma \text{dist}\left((L_x(y), \varphi(y)); \text{epi } h\right) \\ &\leq \gamma \|L_x(y) - c(y)\| \\ &\leq \frac{\gamma\beta}{2} \|x - y\|^2, \end{aligned}$$

where β is a Lipschitz constant of $\nabla c(\cdot)$ on a neighborhood of \bar{x} . Hence we deduce

$$\varphi(x; y^-) = h(L_x(y^-)) \leq t \leq \varphi(y) + \frac{\gamma\beta}{2} \|x - y\|^2$$

and $\|y - y^-\| \leq \frac{\gamma\beta}{2} \|x - y\|^2$, as claimed. \square

9.2. Prox-regularity of the subproblems The final ingredient we need to study subregularity of the mapping \mathcal{G}_t is the notion of prox-regularity. The idea is that we must focus on functions whose limiting subgradients yields uniformly varying quadratic minorants. These are the prox-regular functions introduced in [41].

DEFINITION 9.7 (Prox-regularity). A closed function $f: \mathbf{R}^n \rightarrow \bar{\mathbf{R}}$ is *prox-regular* at \bar{x} for $\bar{v} \in \partial f(\bar{x})$ if there exist neighborhoods \mathcal{X} of \bar{x} and \mathcal{V} of \bar{v} , along with constants $\epsilon, r > 0$ so that the inequality

$$f(y) \geq f(x) + \langle v, y - x \rangle - \frac{r}{2} \|y - x\|^2,$$

holds for all $x, y \in \mathcal{X}$ with $|f(x) - f(\bar{x})| < \epsilon$, and for every subgradient $v \in \mathcal{V} \cap \partial f(x)$.

Prox-regular functions are common in nonsmooth optimization, encompassing for example all C^1 -smooth functions with Lipschitz gradients and all closed convex functions. More generally, the authors of [43] showed that the composite function $\varphi = h \circ c$ is prox-regular at \bar{x} for $\bar{v} \in \partial \varphi(\bar{x})$ whenever c is C^1 -smooth with a Lipschitz gradient, c is transverse to h at \bar{x} , and h is prox-regular at $c(\bar{x})$ for every vector $w \in \partial h(c(\bar{x}))$ satisfying $\bar{v} = \nabla c(\bar{x})^* w$. The following proposition shows that under these conditions, the linearized functions $\varphi(x, \cdot)$ are also prox-regular, uniformly in x .

THEOREM 9.8 (Uniform prox-regularity of the subproblems).

Consider the composite problem (9.2) satisfying $c \pitchfork_{\bar{x}} h$ for some point $\bar{x} \in \text{dom } \varphi$. Then there exists a neighborhood \mathcal{X} of \bar{x} and a constant $\epsilon > 0$ so that the affine functions $z \mapsto c(x) + \nabla c(x)(z - x)$ are transverse to h at z , for any $x, z \in \mathcal{X}$ with $|\varphi(x; z) - \varphi(\bar{x})| < \epsilon$.

Consider a vector $\bar{v} \in \partial \varphi(\bar{x})$, and suppose also that ∇c is Lipschitz continuous around \bar{x} and that h is prox-regular at $c(\bar{x})$ for every subgradient $w \in \partial h(c(\bar{x}))$ satisfying $\bar{v} = \nabla c(\bar{x})^* w$. Then the linearized functions $\varphi(x; \cdot)$ are prox-regular at \bar{x} for \bar{v} uniformly in x in the following sense. After

possibly shrinking \mathcal{X} and $\epsilon > 0$, there exists a neighborhood \mathcal{V} of \bar{v} and a constant $\gamma > 0$ so that the inequality

$$\varphi(x; y) \geq \varphi(x; z) + \langle v, y - z \rangle - \frac{\gamma}{2} \|y - z\|^2 \quad (9.4)$$

holds for any $x, y, z \in \mathcal{X}$ with $|\varphi(x; z) - \varphi(\bar{x})| < \epsilon$, and for every subgradient $v \in \mathcal{V} \cap \partial_z \varphi(x; z)$.

Proof. For the sake of contradiction, suppose there exist sequences $x_i \rightarrow \bar{x}$ and $z_i \rightarrow \bar{x}$ along with unit vectors $w_i \in \partial^\infty h(c(x_i) + \nabla c(x_i)(z_i - x_i)) \cap \text{Null}(\nabla c(x_i)^*)$, so that the values $h(c(x_i) + \nabla c(x_i)(z_i - x_i))$ tend to $h(c(\bar{x}))$. Passing to a subsequence, we may suppose that w_i converge to some unit vector $w \in \partial^\infty h(c(\bar{x})) \cap \text{Null}(\nabla c(\bar{x})^*)$, a contradiction. Hence the transversality claim holds.

By the prox-regularity assumption on h , for every subgradient $w \in \partial h(c(\bar{x}))$ satisfying $\bar{v} = \nabla c(\bar{x})^* w$, there exist constants $\delta_w, r_w > 0$ such that the inequality

$$h(\xi_2) \geq h(\xi_1) + \langle \eta, \xi_2 - \xi_1 \rangle - \frac{r_w}{2} \|\xi_2 - \xi_1\|^2$$

holds for all $\xi_1, \xi_2 \in B_{\delta_w}(c(\bar{x}))$ with $|h(\xi_1) - \varphi(\bar{x})| < \delta_w$, and for any subgradient $\eta \in \partial h(\xi_1)$ with $\|\eta - w\| < \delta_w$. We claim that there exist uniform constants $\delta, r > 0$ (independent of w) so that the inequality

$$h(\xi_2) \geq h(\xi_1) + \langle \eta, \xi_2 - \xi_1 \rangle - \frac{r}{2} \|\xi_2 - \xi_1\|^2 \quad (9.5)$$

holds for all $\xi_1, \xi_2 \in B_\delta(c(\bar{x}))$ with $|h(\xi_1) - \varphi(\bar{x})| < \delta$, and for any subgradient $\eta \in \partial h(\xi_1)$ with $\|\nabla c(\bar{x})^* \eta - \bar{v}\| < \delta$. To see that, suppose otherwise and consider sequences $\xi_1, \xi_2 \rightarrow c(\bar{x})$ with $h(\xi_1) \rightarrow \varphi(\bar{x})$, and a sequence $\eta \in \partial h(\xi_1)$ with $\nabla c(\bar{x})^* \eta \rightarrow \bar{v}$ and so that the fractions $\frac{h(\xi_2) - (h(\xi_1) + \langle \eta, \xi_2 - \xi_1 \rangle)}{\|\xi_2 - \xi_1\|^2}$ are not lower-bounded. The transversality condition (9.3) immediately implies that the vectors η are bounded and hence we can assume that η converge to some vector $w \in \partial h(c(\bar{x}))$ with $\nabla c(\bar{x})^* w = \bar{v}$. This immediately yields the contradiction $\frac{h(\xi_2) - (h(\xi_1) + \langle \eta, \xi_2 - \xi_1 \rangle)}{\|\xi_2 - \xi_1\|^2} \geq -r_w$ for the tails of the sequences.

Fix now points $x, y, z \in \mathcal{X}$ with $|\varphi(x; z) - \varphi(\bar{x})| < \epsilon$. Consider a subgradient $v \in \partial_z \varphi(x; z)$. Since we have already proved that the affine function $c(x) + \nabla c(x)(\cdot - x)$ is transverse to h at z , we may write $v = \nabla c(x)^* \eta$ for some subgradient $\eta \in \partial h(c(x) + \nabla c(x)(z - x))$. Define now the points $\xi_1 := c(x) + \nabla c(x)(z - x)$ and $\xi_2 := c(x) + \nabla c(x)(y - x)$. Shrinking \mathcal{X} and $\epsilon > 0$, we can ensure $\xi_1, \xi_2 \in B_\delta(c(\bar{x}))$ with $|h(\xi_1) - \varphi(\bar{x})| < \delta$. Moreover, if v is sufficiently close to \bar{v} we can ensure $\|\nabla c(\bar{x})^* \eta - \bar{v}\| < \delta$. Hence we may apply inequality (9.5), yielding

$$\varphi(x; y) = h(\xi_2) \geq h(\xi_1) + \langle \eta, \nabla c(x)(y - z) \rangle - \frac{r}{2} \|\nabla c(x)(y - z)\|^2,$$

and therefore $\varphi(x; y) \geq \varphi(x; z) + \langle v, y - z \rangle - \frac{rL^2}{2} \|y - z\|^2$, where $L := \max_{x \in \mathcal{X}} \|\nabla c(x)\|$. The result follows. \square

We are ready to prove out main tool generalizing Theorem 5.3.

THEOREM 9.9 (Prox-gradient and near-stationarity). *Consider the composite problem (9.2) satisfying $c \pitchfork_{\bar{x}} h$ for some point \bar{x} with $0 \in \partial \varphi(\bar{x})$. Suppose that ∇c is locally Lipschitz around \bar{x} and that h is prox-regular at $c(\bar{x})$ for every subgradient $w \in \partial h(c(\bar{x}))$ satisfying $0 = \nabla c(\bar{x})^* w$. Then there are constants $\gamma, \epsilon, a, b > 0$ and a neighborhood \mathcal{X} of \bar{x} such that for any $t > 0$, $x \in \mathcal{X}$, and $x^t \in \mathcal{X} \cap S_t(x)$ with $|\varphi(x, x^t) - \varphi(\bar{x})| < \epsilon$, there exists a point \hat{x} satisfying the properties*

1. (point proximity) $\|x^t - \hat{x}\| \leq (1 + \gamma \|x^t - x\|) \cdot \|x^t - x\|$,
2. (functional proximity) $\varphi(\hat{x}) \leq \varphi(x; x^t) + (a + b/t) \cdot \|x^t - x\|^2$,
3. (near-stationarity) $\text{dist}(0; \partial \varphi(\hat{x})) \leq (a + b/t) \cdot \|x^t - x\|$.

Proof. Fix a neighborhood \mathcal{X} of \bar{x} and constants $\epsilon, \gamma > 0$ given by Theorem 9.6. Shrinking \mathcal{X} we can assume that \mathcal{X} is closed and has diameter smaller than one (for simplicity). Suppose x lies in \mathcal{X} and fix a point $x^t \in \mathcal{X} \cap S_t(x)$ with $|\varphi(x, x^t) - \varphi(\bar{x})| < \epsilon$. Then for any $y \in \mathcal{X}$ with $|\varphi(y) - \varphi(\bar{x})| < \epsilon$, there exists a point y^- satisfying

$$\|y - y^-\| \leq \gamma \|y - x\|^2 \quad \text{and} \quad \varphi(y) + \gamma \|y - x\|^2 \geq \varphi(x; y^-). \quad (9.6)$$

Appealing to Theorem 9.8, we can be sure there is a constant $r > 0$ so that for all $x, x^t, y \in \mathcal{X}$ with $|\varphi(y) - \varphi(\bar{x})| < \epsilon$, we have

$$\begin{aligned} \varphi(x; y^-) &\geq \varphi(x, x^t) + \langle t^{-1}(x - x^t), y^- - x^t \rangle - \frac{r}{2} \|y^- - x^t\|^2 \\ &= \varphi_t(x, x^t) - \frac{1}{2t} \left(\|x - x^t\|^2 - 2\langle x - x^t, y^- - x^t \rangle + rt \|y^- - x^t\|^2 \right) \\ &= \varphi_t(x, x^t) - \frac{1}{2t} \left(\|x - y^-\|^2 + (rt - 1) \|y^- - x^t\|^2 \right), \end{aligned} \quad (9.7)$$

where the last inequality follows by completing the square. On the other hand, the triangle inequality, along with the assumption that the diameter of \mathcal{X} is smaller than one, yields

$$\|x - y^-\| \leq (\gamma + 1) \|x - y\| \quad \text{and} \quad \|y^- - x^t\| \leq (\gamma + 1) \|x - y\| + \|x - x^t\|.$$

Defining for notational convenience $q := \gamma + 1$ and $p := \max\{0, rt - 1\}$, we obtain the inequality

$$\varphi(x; y^-) \geq \varphi_t(x, x^t) - \frac{q^2(1+p)}{2t} \|x - y\|^2 - \frac{p}{2t} \left(\|x^t - x\|^2 + 2q \max\{\|y - x\|, \|x^t - x\|\}^2 \right).$$

Define now the function

$$\begin{aligned} \zeta(y) &:= \delta_{\mathcal{X}}(y) + \varphi(y) + \left(\gamma + \frac{q^2(1+p)}{2t} \right) \|x - y\|^2 + \\ &\quad + \frac{p}{2t} \left(\|x^t - x\|^2 + 2q \max\{\|y - x\|, \|x^t - x\|\}^2 \right). \end{aligned}$$

Taking into account (9.6), we deduce that the inequality $\zeta(y) \geq \varphi_t(x, x^t)$ holds for all $y \in \mathcal{X}$ with $|\varphi(y) - \varphi(\bar{x})| < \epsilon$. Moreover, since φ is closed, shrinking \mathcal{X} we can assume $\varphi(y) \geq \varphi(\bar{x}) - \epsilon$ for all $y \in \mathcal{X}$. A trivial argument now shows that again shrinking \mathcal{X} , we can finally ensure $\zeta(y) \geq \varphi_t(x, x^t)$ for all $y \in \mathcal{X}$.

Taking into account the inequality $|\varphi(x, x^t) - \varphi(\bar{x})| < \epsilon$ and applying claim 2 of Theorem 9.6, a quick computation shows

$$\zeta(x^{t+}) - \zeta^* \leq l \cdot \|x^t - x\|^2,$$

where we set

$$l := 2\gamma + \frac{q^2(1+p) + p(1+2q) - 1}{2t},$$

and ζ^* is the minimal value of ζ . Define now the constants $\hat{\epsilon} := l \|x^t - x\|^2$ and $\rho := l \|x^t - x\|$. Applying Ekeland's variational principle, we obtain a point \hat{x} satisfying the inequalities $\zeta(\hat{x}) \leq \zeta(x^{t+})$ and $\|x^{t+} - \hat{x}\| \leq \hat{\epsilon}/\rho$, and the inclusion $0 \in \partial\zeta(\hat{x}) + \rho\mathbf{B}$. The point proximity estimate is now immediate from the inequality

$$\|x^t - \hat{x}\| \leq \|x^t - x^{t+}\| + \|x^{t+} - \hat{x}\| \leq \gamma \|x^t - x\|^2 + \|x^t - x\| = (1 + \gamma \|x^t - x\|) \|x^t - x\|.$$

Using the inequalities $\varphi(\hat{x}) + \frac{p}{2t} \|x^t - x\|^2 \leq \zeta(\hat{x}) \leq \zeta(x^{t+})$, we deduce

$$\varphi(\hat{x}) \leq \varphi(x, x^t) + \left(l - \frac{p-1}{2t} \right) \cdot \|x^t - x\|^2.$$

Finally, we conclude the near-stationarity condition

$$\begin{aligned} \text{dist}(0, \partial\varphi(\hat{x})) &\leq \rho + \left(\gamma + \frac{q^2(1+p) + 2pq}{2t} \right) \|\hat{x} - x\| \\ &\leq \left[l + (2 + \gamma\|x^t - x\|) \left(\gamma + \frac{q^2(1+p) + 2pq}{2t} \right) \right] \|x^t - x\|. \end{aligned}$$

The result follows after noting the inequality $\frac{p}{t} \leq r$. \square

We can now prove the main result of this section comparing subregularity of the subdifferential $\partial\varphi$ and of the prox-gradient mapping \mathcal{G}_t . Naturally, to make such a comparison precise, we must focus on the subgraphs of $\text{gph } \partial\varphi$ and $\text{gph } \mathcal{G}_t$ that arise from points x and $x^t \in \mathcal{S}_t(x)$ at which the function values $\varphi(x)$ and $\varphi(x; x^t)$ are close to $\varphi(\bar{x})$. In most important circumstances, nearness in the graphs, $\text{gph } \partial\varphi$ and $\text{gph } \mathcal{S}_t$, automatically implies nearness in function value. To illustrate, consider the composite problem (9.2) satisfying $c \pitchfork_{\bar{x}} h$ for some point \bar{x} with $0 \in \partial\varphi(\bar{x})$. It follows quickly from [50, Example 13.30] that if h is either convex or continuous on its domain, then h is *subdifferentially continuous* at $(\bar{x}, 0)$: for any sequence $(x_i, v_i) \in \text{gph } \partial\varphi$ converging to $(\bar{x}, 0)$ the values $\varphi(x_i)$ converge to $\varphi(\bar{x})$. Similarly, it is easy to check that if h is either convex or continuous on its domain, then for any sequence $(x_i, y_i) \in \text{gph } \mathcal{S}_t$ converging to (\bar{x}, \bar{x}) the values $\varphi(x_i; y_i)$ converge to $\varphi(\bar{x})$.

When h is not convex, nor is continuous on its domain, we must focus only on the relevant parts of the graphs, $\text{gph } \partial\varphi$ and $\text{gph } \mathcal{G}_t$. This idea of a functionally attentive localization is not new, and goes back at least to [41].

DEFINITION 9.10 (φ and $\varphi(\cdot, \cdot)$ -attentive localizations). Consider the composite problem (9.2) satisfying $c \pitchfork_{\bar{x}} h$ for some point \bar{x} with $0 \in \partial\varphi(\bar{x})$.

1. A set-valued mapping $W: \mathbf{R}^n \rightrightarrows \mathbf{R}^n$ is a φ -attentive localization of the subdifferential $\partial\varphi$ around $(\bar{x}, 0)$ if there exist neighborhoods \mathcal{X} of \bar{x} and \mathcal{V} of 0 and a real number $\epsilon > 0$ so that for any points $x \in \mathcal{X}$ and $v \in \mathcal{V}$ with $|\varphi(x) - \varphi(\bar{x})| < \epsilon$, the equivalence $v \in \partial\varphi(x) \iff v \in W(x)$ holds.

2. A set-valued mapping $W: \mathbf{R}^n \rightrightarrows \mathbf{R}^n$ is a $\varphi(\cdot, \cdot)$ -attentive localization of the stationary point map \mathcal{S}_t around (\bar{x}, \bar{x}) if there exist neighborhoods \mathcal{X} of \bar{x} and \mathcal{Y} of 0 and a real number $\epsilon > 0$ so that for any points $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ with $|\varphi(x, y) - \varphi(\bar{x})| < \epsilon$, the equivalence $y \in \mathcal{S}_t(x) \iff y \in W(x)$ holds.

3. A set-valued mapping $W: \mathbf{R}^n \rightrightarrows \mathbf{R}^n$ is a $\varphi(\cdot, \cdot)$ -attentive localization of the prox-gradient mapping \mathcal{G}_t around $(\bar{x}, 0)$ if it can be written as $W = t^{-1}(I - \widehat{W}(x))$, where \widehat{W} is a $\varphi(\cdot, \cdot)$ -attentive localization of \mathcal{S}_t around (\bar{x}, \bar{x}) .

The following is the main result of this section.

THEOREM 9.11 (Subdifferential subregularity and the error bound property).

Consider the composite problem (9.2) satisfying $c \pitchfork_{\bar{x}} h$ for some point \bar{x} with $0 \in \partial\varphi(\bar{x})$. Suppose that ∇c is Lipschitz around \bar{x} and that h is prox-regular at $c(\bar{x})$ for every subgradient $w \in \partial h(c(\bar{x}))$ satisfying $0 = \nabla c(\bar{x})^* w$. Consider the following two conditions:

1. there exists a φ -attentive localization of $\partial\varphi$ that is metrically subregular at $(\bar{x}, 0)$.
2. there exists a $\varphi(\cdot, \cdot)$ -attentive localization of the mapping \mathcal{G}_t that is metrically subregular at $(\bar{x}, 0)$.

Then the implication (1) \Rightarrow (2) always holds. Conversely, there exists a number $\bar{t} \geq 0$ so that for all $t \in (0, \bar{t})$, the implication (2) \Rightarrow (1) holds. When h is convex, the localizations are not needed: the following two conditions are equivalent for any $t > 0$:

- 1'. The subdifferential $\partial\varphi$ is metrically subregular at $(\bar{x}, 0)$.
- 2'. The prox-gradient \mathcal{G}_t is metrically subregular at $(\bar{x}, 0)$.

Proof. Suppose there exists a φ -attentive localization W of $\partial\varphi$ that is metrically subregular at $(\bar{x}, 0)$ with constant l . Fix a neighborhood \mathcal{X} of \bar{x} and constants $\gamma, \epsilon, a, b > 0$ guaranteed to exist by Theorem 9.9. Then for any points $x \in \mathcal{X}$ and $x^t \in \mathcal{X} \cap \mathcal{S}_t(x)$ with $|\varphi(x, x^t) - \varphi(\bar{x})| < \epsilon$, there exists a point \hat{x} satisfying

$$\|x^t - \hat{x}\| \leq (1 + \gamma\|x^t - x\|) \cdot \|x^t - x\|, \quad (9.8)$$

$$\varphi(\hat{x}) \leq \varphi(x; x^t) + (a + b/t) \cdot \|x^t - x\|^2, \quad (9.9)$$

$$\text{dist}(0; \partial\varphi(\hat{x})) \leq (at + b) \cdot \|t^{-1}(x^t - x)\|. \quad (9.10)$$

Due to inequality (9.9), the subdifferential $\partial\varphi(\hat{x})$ coincides with the localization $W(\hat{x})$ near the origin. Hence we deduce

$$\text{dist}(\hat{x}; W^{-1}(0)) \leq l \cdot \text{dist}(0; W(\hat{x})) = l \cdot \text{dist}(0; \partial\varphi(\hat{x})) \leq l(at + b) \cdot \|t^{-1}(x^t - x)\|.$$

On the other hand, we have

$$\text{dist}(\hat{x}, W^{-1}(0)) \geq \text{dist}(x; W^{-1}(0)) - \|x^t - x\| - \|x^t - \hat{x}\|,$$

and therefore

$$\text{dist}(x; W^{-1}(0)) \leq ((2 + la + \gamma\|x^t - x\|)t + lb) \cdot \|t^{-1}(x^t - x)\|.$$

Hence property (2) holds, as claimed.

Next, we show the converse. Fix the neighborhoods \mathcal{X}, \mathcal{V} and the constants ϵ, γ guaranteed by Theorem 9.8. After possibly shrinking \mathcal{X} , the local existence result [29, Theorem 4.5(a)] guarantees that there is a constant \bar{t} so that provided $t < \bar{t}$, for any $x \in \mathcal{X}$ there exists a point $x^t \in \widehat{W}(x)$ so that the inclusion $x - x^t \in \mathcal{V}$ holds and we have $|\varphi(x; x^t) - \varphi(\bar{x})| < \epsilon$. Henceforth suppose that the map $W = t^{-1}(I - \widehat{W})$ is subregular at $(\bar{x}, 0)$ for some $t < \bar{t}$, where \widehat{W} is a $\varphi(\cdot; \cdot)$ -attentive localization of \mathcal{S}_t around (\bar{x}, \bar{x}) .

Fix a pair $(x, v) \in (\mathcal{X} \times \mathcal{V}) \cap \text{gph } \partial\varphi$ with $|\varphi(x) - \varphi(\bar{x})| < \epsilon$. Then v lies in $\partial_z\varphi(x; z)$ with the choice $z := x$, and hence setting $y := x^t$ in (9.4) we deduce

$$\varphi(x, x^t) \geq \varphi(x) + \langle v, x^t - x \rangle - \frac{\gamma}{2} \|x^t - x\|^2. \quad (9.11)$$

Appealing to (9.4) again with the choices $y = x$, $z = x^t$, and $t^{-1}(x - x^t) \in \partial_z\varphi(x; z)$, we obtain the inequality

$$\begin{aligned} \varphi(x) &\geq \varphi(x; x^t) + t^{-1} \langle x - x^t, x - x^t \rangle - \frac{\gamma}{2} \|x^t - x\|^2 \\ &= \varphi(x, x^t) + \left(t^{-1} - \frac{\gamma}{2}\right) \|x^t - x\|^2. \end{aligned} \quad (9.12)$$

Taking into account (9.11), we deduce $\|v\| \cdot \|x^t - x\| \geq (t^{-1} - \gamma)\|x^t - x\|^2$. Choosing $t < \min\{\bar{t}, \gamma^{-1}\}$ so as to ensure $t^{-1} - \gamma > 0$, we finally conclude

$$(t^{-1} - \gamma)^{-1} \cdot \|v\| \geq \|x^t - x\| \geq t \cdot \text{dist}(0, W(x)) \geq tl \cdot \text{dist}(x, W^{-1}(0)),$$

where l is the constant of subregularity of W at $(\bar{x}, 0)$. On the other hand, observe $z \in W^{-1}(0) \Leftrightarrow z = \widehat{W}(z)$. Hence if $z \in W^{-1}(0)$ is close to \bar{x} , then we can be sure that $\varphi(z)$ is close to $\varphi(\bar{x})$. It follows that the set $W^{-1}(0)$ coincides near \bar{x} with $F^{-1}(0)$, where F is some φ -attentive localization of $\partial\varphi$. Hence Property (1) holds, as claimed.

Finally, when h is convex, the functionally attentive localizations are not needed, as was explained prior to Definition 9.10. Moreover, the inequalities (9.11) and (9.12) hold with $\gamma = 0$ and according to [29, Theorem 4.5(c)], we could have set $\bar{t} = +\infty$ at the onset. Consequently, the implication (2) \Rightarrow (1) holds for arbitrary $t > 0$, as claimed. \square

To summarize, with Theorem 9.11 we have shown an equivalence between subregularity of the subdifferential $\partial\varphi$ and the error bound property, thereby extending Theorem 5.11 to the case where h need not be convex nor finite-valued, but merely prox-regular. We believe that this result can serve the same role as Theorem 5.11 in section 5 for understanding linear convergence of proximal algorithms for composite problems (9.1).

Acknowledgments. We thank the two anonymous referees for many suggestions that have improved the quality of the paper. We also thank Guoyin Li and Ting Kei Pong for a careful reading of an early draft of the manuscript, and for their insightful comments. Research of Drusvyatskiy was partially supported by the AFOSR grant FA9550-15-1-0237 and by the NSF grant DMS-1651851. Research of Lewis was supported in part by National Science Foundation Grants DMS-1208338 and DMS-1613996.

References

- [1] Aragón Artacho F, Geoffroy M (2008) Characterization of metric regularity of subdifferentials. *J. Convex Anal.* 15(2):365–380, ISSN 0944-6532.
- [2] Aravkin A, Burke J, Pillonetto G (2013) Sparse/robust estimation and kalman smoothing with non-smooth log-concave densities: Modeling, computation, and theory. *The Journal of Machine Learning Research* 14(1):2689–2728.
- [3] Aravkin A, Lozano A, Luss R, Kambadur P (2014) Orthogonal matching pursuit for sparse quantile regression. *Data Mining (ICDM), 2014 IEEE International Conference on*, 11–19 (IEEE).
- [4] Azé D, Corvellec JN (2014) Nonlinear local error bounds via a change of metric. *J. Fixed Point Theory Appl.* 16(1-2):351–372, ISSN 1661-7738, URL <http://dx.doi.org/10.1007/s11784-015-0220-9>.
- [5] Basu S, Pollack R, Roy M (2006) *Algorithms in Real Algebraic Geometry (Algorithms and Computation in Mathematics)* (Secaucus, NJ, USA: Springer-Verlag New York, Inc.), ISBN 3540330984.
- [6] Bauschke H, Borwein J (1993) On the convergence of von Neumann’s alternating projection algorithm for two sets. *Set-Valued Anal.* 1(2):185–212, ISSN 0927-6947, URL <http://dx.doi.org/10.1007/BF01027691>.
- [7] Beck A, Teboulle M (2009) A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.* 2(1):183–202, ISSN 1936-4954, URL <http://dx.doi.org/10.1137/080716542>.
- [8] Bolte J, Nguyen T, Peypouquet J, Suter B (2016) From error bounds to the complexity of first-order descent methods for convex functions. *Mathematical Programming* URL <http://dx.doi.org/10.1007/s10107-016-1091-6>.
- [9] Burke J (1985) Descent methods for composite nondifferentiable optimization problems. *Math. Programming* 33(3):260–279, ISSN 0025-5610, URL <http://dx.doi.org/10.1007/BF01584377>.
- [10] Burke J (1987) Second order necessary and sufficient conditions for convex composite NDO. *Math. Programming* 38(3):287–302, ISSN 0025-5610, URL <http://dx.doi.org/10.1007/BF02592016>.
- [11] Burke J, Ferris M (1993) Weak sharp minima in mathematical programming. *SIAM J. Control Optim.* 31(5):1340–1359, ISSN 0363-0129, URL <http://dx.doi.org/10.1137/0331063>.
- [12] Burke J, Ferris M (1995) A Gauss-Newton method for convex composite optimization. *Math. Programming* 71(2, Ser. A):179–194, ISSN 0025-5610, URL <http://dx.doi.org/10.1007/BF01585997>.
- [13] Burke J, Poliquin R (1992) Optimality conditions for non-finite valued convex composite functions. *Math. Programming* 57(1, Ser. B):103–120, ISSN 0025-5610, URL <http://dx.doi.org/10.1007/BF01581075>.
- [14] Cartis C, Gould NIM, Toint PL (2011) On the evaluation complexity of composite function minimization with applications to nonconvex nonlinear programming. *SIAM J. Optim.* 21(4):1721–1739, ISSN 1052-6234, URL <http://dx.doi.org/10.1137/11082381X>.
- [15] Chu W, Keerthi SS, Ong CJ (2004) Bayesian support vector regression using a unified loss function. *Trans. Neur. Netw.* 15(1):29–44, ISSN 1045-9227, URL <http://dx.doi.org/10.1109/TNN.2003.820830>.
- [16] Cromme L (1977/78) Strong uniqueness. A far-reaching criterion for the convergence analysis of iterative procedures. *Numer. Math.* 29(2):179–193, ISSN 0029-599X, URL <http://dx.doi.org/10.1007/BF01390337>.
- [17] Dong Y (2009) An extension of Luque’s growth condition. *Appl. Math. Lett.* 22(9):1390–1393, ISSN 0893-9659, URL <http://dx.doi.org/10.1016/j.aml.2007.07.037>.

- [18] Dontchev A, Lewis A, Rockafellar R (2003) The radius of metric regularity. *Trans. Amer. Math. Soc.* 355(2):493–517 (electronic), ISSN 0002-9947, URL <http://dx.doi.org/10.1090/S0002-9947-02-03088-X>.
- [19] Drusvyatskiy D, Ioffe A (2015) Quadratic growth and critical point stability of semi-algebraic functions. *Math. Program.* 153(2, Ser. A):635–653, ISSN 0025-5610, URL <http://dx.doi.org/10.1007/s10107-014-0820-y>.
- [20] Drusvyatskiy D, Ioffe A, Lewis A (2016) Nonsmooth optimization using Taylor-like models: error bounds, convergence, and termination criteria. *arXiv:1610.03446* .
- [21] Drusvyatskiy D, Lewis A (2013) Tilt stability, uniform quadratic growth, and strong metric regularity of the subdifferential. *SIAM J. Optim.* 23(1):256–267, ISSN 1052-6234, URL <http://dx.doi.org/10.1137/120876551>.
- [22] Drusvyatskiy D, Lewis A (Feb. 2016) Error bounds, quadratic growth, and linear convergence of proximal methods. *Preprint arXiv:1602.06661* .
- [23] Drusvyatskiy D, Mordukhovich B, Nghia T (2014) Second-order growth, tilt stability, and metric regularity of the subdifferential. *J. Convex Anal.* 21(4):1165–1192.
- [24] Drusvyatskiy D, Paquette C (2016) Efficiency of minimizing compositions of convex functions and smooth maps. *Preprint arXiv:1605.00125* .
- [25] Fletcher R (1982) A model algorithm for composite nondifferentiable optimization problems. *Math. Programming Stud.* (17):67–76, ISSN 0303-3929, nondifferential and variational techniques in optimization (Lexington, Ky., 1980).
- [26] Ioffe A (2000) Metric regularity and subdifferential calculus. *Uspekhi Mat. Nauk* 55(3(333)):103–162, ISSN 0042-1316, URL <http://dx.doi.org/10.1070/rm2000v055n03ABEH000292>.
- [27] Klatté D, Kummer B (2002) *Nonsmooth equations in optimization*, volume 60 of *Nonconvex Optimization and its Applications* (Kluwer Academic Publishers, Dordrecht), ISBN 1-4020-0550-4, regularity, calculus, methods and applications.
- [28] Lee J (2013) *Introduction to smooth manifolds*, volume 218 of *Graduate Texts in Mathematics* (Springer, New York), second edition, ISBN 978-1-4419-9981-8.
- [29] Lewis A, Wright S (2015) A proximal method for composite minimization. *Math. Program.* 1–46, URL <http://dx.doi.org/10.1007/s10107-015-0943-9>.
- [30] Li G, Pong T (2016) Calculus of the exponent of Kurdyka-Łojasiewicz inequality and its applications to linear convergence of first-order methods. *arXiv:1602.02915* .
- [31] Luo ZQ, Tseng P (1993) Error bounds and convergence analysis of feasible descent methods: a general approach. *Ann. Oper. Res.* 46/47(1-4):157–178, ISSN 0254-5330, URL <http://dx.doi.org/10.1007/BF02096261>.
- [32] Luo ZQ, Tseng P (1993) On the convergence rate of dual ascent methods for linearly constrained convex minimization. *Math. Oper. Res.* 18(4):846–867, ISSN 0364-765X, URL <http://dx.doi.org/10.1287/moor.18.4.846>.
- [33] Mordukhovich B (2006) *Variational Analysis and Generalized Differentiation I: Basic Theory* (Grundlehren der mathematischen Wissenschaften, Vol 330, Springer, Berlin), ISBN 3540254374.
- [34] Mordukhovich B, Nghia T (2013) Second-order variational analysis and characterizations of tilt-stable optimal solutions in infinite-dimensional spaces. *Nonlinear Anal.* 86:159–180, ISSN 0362-546X, URL <http://dx.doi.org/10.1016/j.na.2013.03.014>.
- [35] Mordukhovich B, Nghia T (2015) Second-order characterizations of tilt stability with applications to nonlinear programming. *Math. Program.* 149(1-2, Ser. A):83–104, ISSN 0025-5610, URL <http://dx.doi.org/10.1007/s10107-013-0739-8>.
- [36] Mordukhovich B, Outrata J (2013) Tilt stability in nonlinear programming under Mangasarian-Fromovitz constraint qualification. *Kybernetika (Prague)* 49(3):446–464, ISSN 0023-5954.
- [37] Mordukhovich B, Outrata J, Ramírez C H (2015) Second-order variational analysis in conic programming with applications to optimality and stability. *SIAM J. Optim.* 25(1):76–101, ISSN 1052-6234, URL <http://dx.doi.org/10.1137/120903221>.

- [38] Mordukhovich B, Rockafellar RT (2012) Second-order subdifferential calculus with applications to tilt stability in optimization. *SIAM J. Optim.* 22(3):953–986, ISSN 1052-6234, URL <http://dx.doi.org/10.1137/110852528>.
- [39] Necoara I, Nesterov Y, Glineur F (2015) Linear convergence of first order methods for non-strongly convex optimization. *Preprint arXiv:1504.06298* .
- [40] Nesterov Y (2004) *Introductory Lectures on Convex Optimization* (Dordrecht, The Netherlands: Kluwer Academic).
- [41] Poliquin R, Rockafellar R (1996) Prox-regular functions in variational analysis. *Trans. Amer. Math. Soc.* 348:1805–1838.
- [42] Poliquin R, Rockafellar R (1998) Tilt stability of a local minimum. *SIAM J. Optim.* 8(2):287–299 (electronic), ISSN 1052-6234, URL <http://dx.doi.org/10.1137/S1052623496309296>.
- [43] Poliquin R, Rockafellar R (2010) A calculus of prox-regularity. *J. Convex Anal.* 17(1):203–210, ISSN 0944-6532.
- [44] Polyak B (1979) Sharp minima. *Institute of Control Sciences Lecture Notes Presented at the IIASA Workshop on Generalized Lagrangians and Their Applications*, IIASA, Laxenburg, Austria.
- [45] Powell M (1983) General algorithms for discrete nonlinear approximation calculations. *Approximation theory, IV (College Station, Tex., 1983)*, 187–218 (Academic Press, New York).
- [46] Powell M (1984) On the global convergence of trust region algorithms for unconstrained minimization. *Math. Programming* 29(3):297–303, ISSN 0025-5610, URL <http://dx.doi.org/10.1007/BF02591998>.
- [47] Rockafellar R (1970) *Convex analysis*. Princeton Mathematical Series, No. 28 (Princeton University Press, Princeton, N.J.).
- [48] Rockafellar R (1976) Monotone operators and the proximal point algorithm. *SIAM J. Control Optimization* 14(5):877–898, ISSN 0363-0129.
- [49] Rockafellar R, Dontchev A (2009) *Implicit functions and solution mappings* (Monographs in Mathematics, Springer-Verlag).
- [50] Rockafellar R, Wets RB (1998) *Variational Analysis* (Grundlehren der mathematischen Wissenschaften, Vol 317, Springer, Berlin).
- [51] Tseng P (2010) Approximation accuracy, gradient methods, and error bound for structured convex optimization. *Math. Program.* 125(2, Ser. B):263–295, ISSN 0025-5610, URL <http://dx.doi.org/10.1007/s10107-010-0394-2>.
- [52] Wright S (1990) Convergence of an inexact algorithm for composite nonsmooth optimization. *IMA J. Numer. Anal.* 10(3):299–321, ISSN 0272-4979, URL <http://dx.doi.org/10.1093/imanum/10.3.299>.
- [53] Yuan Y (1985) On the superlinear convergence of a trust region algorithm for nonsmooth optimization. *Math. Programming* 31(3):269–285, ISSN 0025-5610, URL <http://dx.doi.org/10.1007/BF02591949>.
- [54] Zhang H (2016) The restricted strong convexity: analysis of equivalence to error bound and quadratic growth. *Preprint arXiv:1511.01635* .
- [55] Zhou Z, So AC (2017) A unified approach to error bounds for structured convex optimization problems. *Mathematical Programming* URL <http://dx.doi.org/10.1007/s10107-016-1100-9>.
- [56] Zou H, Hastie T (2005) Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 67(2):301–320, ISSN 1369-7412, URL <http://dx.doi.org/10.1111/j.1467-9868.2005.00503.x>.