Chapter 15
On the Acceleration of Forward-Backward Splitting via an Inexact Newton Method

Andreas Themelis, Masoud Ahookhosh, and Panagiotis Patrinos

Abstract We propose a Forward-Backward Truncated-Newton method (FBTN) for minimizing the sum of two convex functions, one of which smooth. Unlike other proximal Newton methods, our approach does not involve the employment of variable metrics, but is rather based on a reformulation of the original problem as the unconstrained minimization of a continuously differentiable function, the *forward-backward envelope (FBE)*. We introduce a generalized Hessian for the FBE that *symmetrizes* the generalized Jacobian of the nonlinear system of equations representing the optimality conditions for the problem. This enables the employment of conjugate gradient method (CG) for efficiently solving the resulting (regularized) linear systems, which can be done inexactly. The employment of CG prevents the computation of full (generalized) Jacobians, as it requires only (generalized) directional derivatives. The resulting algorithm is globally (subsequentially) convergent, Q-linearly under an error bound condition, and up to Q-superlinearly and Q-quadratically under regularity assumptions at the possibly non-isolated limit point.

Keywords Forward-backward splitting · Linear Newton approximation · Truncated-Newton method · Backtracking linesearch · Error bound · Superlinear convergence

AMS 2010 Subject Classification 49J52, 49M15, 90C06, 90C25, 90C30

15.1 Introduction

In this work we focus on convex composite optimization problems of the form

$$\text{minimize}_{x\in\mathbb{R}^n}\ \varphi(x) \equiv f(x) + g(x), \tag{15.1}$$

A. Themelis · M. Ahookhosh · P. Patrinos (✉)
Department of Electrical Engineering (ESAT-STADIUS), KU Leuven, Leuven, Belgium
e-mail: andreas.themelis@esat.kuleuven.be; masoud.ahookhosh@esat.kuleuven.be; panos.patrinos@esat.kuleuven.be

H. H. Bauschke et al. (eds.), *Splitting Algorithms, Modern Operator Theory, and Applications*, https://doi.org/10.1007/978-3-030-25939-6_15

where $f : \mathbb{R}^n \to \mathbb{R}$ is convex, twice continuously differentiable and with L_f-Lipschitz-continuous gradient, and $g : \mathbb{R}^n \to \mathbb{R} \cup \{\infty\}$ has a cheaply computable proximal mapping [51]. To ease the notation, throughout the chapter we indicate

$$\varphi_\star := \inf \varphi \quad \text{and} \quad \mathcal{X}_\star := \operatorname{argmin} \varphi. \tag{15.2}$$

Problems of the form (15.1) are abundant in many scientific areas such as control, signal processing, system identification, machine learning, and image analysis, to name a few. For example, when g is the indicator of a convex set then (15.1) becomes a constrained optimization problem, while for $f(x) = \frac{1}{2}\|Ax - b\|^2$ and $g(x) = \lambda\|x\|_1$ it becomes the ℓ_1-regularized least-squares problem (lasso) which is the main building block of compressed sensing. When g is equal to the nuclear norm, then (15.1) models low-rank matrix recovery problems. Finally, conic optimization problems such as linear, second-order cone, and semidefinite programs can be brought into the form of (15.1), see [31].

Perhaps the most well-known algorithm for problems in the form (15.1) is the forward-backward splitting (FBS) or proximal gradient method [16, 40], that interleaves gradient descent steps on the smooth function and *proximal* steps on the nonsmooth one, see Section 15.3.1. Accelerated versions of FBS, based on the work of Nesterov [5, 54, 77], have also gained popularity. Although these algorithms share favorable global convergence rate estimates of order $O(\varepsilon^{-1})$ or $O(\varepsilon^{-1/2})$ (where ε is the solution accuracy), they are first-order methods and therefore usually effective at computing solutions of low or medium accuracy only. An evident remedy is to include second-order information by replacing the Euclidean norm in the proximal mapping with that induced by the Hessian of f at x or some approximation of it, mimicking Newton or quasi-Newton methods for unconstrained problems [6, 32, 42]. However, a severe limitation of the approach is that, unless Q has a special structure, the computation of the proximal mapping becomes very hard. For example, if φ models a lasso problem, the corresponding subproblem is as hard as the original problem.

In this work we follow a different approach by reformulating the nonsmooth constrained problem (15.1) into the smooth unconstrained minimization of the *forward-backward envelope (FBE)* [57], a real-valued, continuously differentiable, exact penalty function for φ. Although the FBE might fail to be twice continuously differentiable, by using tools from nonsmooth analysis we show that one can design Newton-like methods to address its minimization, that achieve Q-superlinear asymptotic rates of convergence under nondegeneracy and (generalized) smoothness conditions on the set of solutions. Furthermore, by suitably interleaving FBS and Newton-like iterations the proposed algorithm also enjoys good complexity guarantees provided by a global (non-asymptotic) convergence rate. Unlike the approaches of [6, 32], where the corresponding subproblems are expensive to solve, our algorithm only requires the inexact solution of a linear system to compute the Newton-type direction, which can be done efficiently with a memory-free CG method.

Our approach combines and extends ideas stemming from the literature on merit functions for *variational inequalities* (VIs) and *complementarity problems* (CPs), specifically the reformulation of a VI as a constrained continuously differentiable optimization problem via the regularized gap function [23] and as an unconstrained continuously differentiable optimization problem via the D-gap function [79] (see [19, §10] for a survey and [38, 58] for applications to constrained optimization and model predictive control of dynamical systems).

15.1.1 Contributions

We propose an algorithm that addresses problem (15.1) by means of a Newton-like method on the FBE. Differently from a direct application of the classical Newton method, our approach does not require twice differentiability of the FBE (which would impose additional properties on f and g), but merely twice differentiability of f. This is possible thanks to the introduction of an *approximate generalized Hessian* which only requires access to $\nabla^2 f$ and to the generalized (Clarke) Jacobian of the proximal mapping of g, as opposed to third-order derivatives and classical Jacobian, respectively. Moreover, it allows for inexact solutions of linear systems to compute the update direction, which can be done efficiently with a truncated CG method; in particular, no computation of full (generalized) Hessian matrices is necessary, as only (generalized) directional derivatives are needed. The method is thus particularly appealing when the Clarke Jacobians are sparse and/or well structured, so that the implementation of CG becomes extremely efficient. Under an error bound condition and a (semi)smoothness assumption at the limit point, which is not required to be isolated, the algorithm exhibits asymptotic Q-superlinear rates. For the reader's convenience we collect explicit formulas of the needed Jacobians of the proximal mapping for a wide range of frequently encountered functions, and discuss when they satisfy the needed semismoothness requirements that enable superlinear rates.

15.1.2 Related Work

This work is a revised version of the unpublished manuscript [59] and extends ideas proposed in [57], where the FBE is first introduced. Other FBE-based algorithms are proposed in [69, 71, 75]; differently from the truncated-CG type of approximation proposed here, they all employ quasi-Newton directions to mimic second-order information. The underlying ideas can also be extended to enhance other popular proximal splitting algorithms: the Douglas Rachford splitting (DRS) and the alternating direction method of multipliers (ADMM) [74], and for strongly convex problems also the alternating minimization algorithm (AMA) [70].

The algorithm proposed in this chapter adopts the recent techniques investigated in [71, 75] to enhance and greatly simplify the scheme in [59]. In particular, Q-linear and Q-superlinear rates of convergence are established under an error bound condition, as opposed to uniqueness of the solution. The proofs of superlinear convergence with an error bound pattern the arguments in [82, 83], although with less conservative requirements.

15.1.3 Organization

The work is structured as follows. In Section 15.2 we introduce the adopted notation and list some known facts on generalized differentiability needed in the sequel. Section 15.3 offers an overview on the connections between FBS and the proximal point algorithm, and serves as a prelude to Section 15.4 where the forward-backward envelope function is introduced and analyzed. Section 15.5 deals with the proposed truncated-Newton algorithm and its convergence analysis. In Section 15.6 we collect explicit formulas for the generalized Jacobian of the proximal mapping of a rich list of nonsmooth functions, needed for computing the update directions in the proposed algorithm. Finally, Section 15.7 draws some conclusions.

15.2 Preliminaries

15.2.1 Notation and Known Facts

Our notation is standard and follows that of convex analysis textbooks [2, 8, 28, 63]. For the sake of clarity we now properly specify the adopted conventions, and briefly recap known definitions and facts in convex analysis. The interested reader is referred to the above-mentioned textbooks for the details.

Matrices and Vectors The $n \times n$ identity matrix is denoted as I_n, and the $\mathbb{R}^n$ vector with all elements equal to 1 is as $\mathbf{1}_n$; whenever n is clear from context we simply write I or $\mathbf{1}$, respectively. We use the Kronecker symbol $\delta_{i,j}$ for the (i, j)-th entry of I. Given $v \in \mathbb{R}^n$, with $\operatorname{diag} v$ we indicate the $n \times n$ diagonal matrix whose i-th diagonal entry is v_i. With $\mathrm{S}(\mathbb{R}^n)$, $\mathrm{S}_+(\mathbb{R}^n)$, and $\mathrm{S}_{++}(\mathbb{R}^n)$ we denote respectively the set of symmetric, symmetric positive semidefinite, and symmetric positive definite matrices in $\mathbb{R}^{n\times n}$.

The minimum and maximum eigenvalues of $H \in \mathrm{S}(\mathbb{R}^n)$ are denoted as $\lambda_{\min}(H)$ and $\lambda_{\max}(H)$, respectively. For $Q, R \in \mathrm{S}(R^n)$ we write $Q \succeq R$ to indicate that $Q - R \in \mathrm{S}_+(\mathbb{R}^n)$, and similarly $Q \succ R$ indicates that $Q - R \in \mathrm{S}_{++}(\mathbb{R}^n)$. Any matrix $Q \in \mathrm{S}_+(\mathbb{R}^n)$ induces the semi-norm $\| \cdot \|_Q$ on $\mathbb{R}^n$, where $\|x\|_Q^2 \coloneqq \langle x, Qx\rangle$; in case $Q = \mathrm{I}$, that is, for the Euclidean norm, we omit the subscript and simply write $\|\cdot\|$. No ambiguity occurs in adopting the same notation for the induced matrix norm, namely $\|M\| := \max\{\|Mx\| | x \in \mathbb{R}^n,\ \|x\| = 1|\}$ for $M \in \mathbb{R}^{n\times n}$.

Topology The *convex hull* of a set $E \subseteq \mathbb{R}^n$, denoted as $\operatorname{conv} E$, is the smallest convex set that contains E (the intersection of convex sets is still convex). The *affine hull* $\operatorname{aff} E$ and the *conic hull* $\operatorname{cone} E$ are defined accordingly. Specifically,

$$\operatorname{conv} E := \left\{\textstyle\sum_{i=1}^k \alpha_i x_i \mid k \in \mathbb{N},\ x_i \in E,\ \alpha_i \geq 0,\ \sum_{i=1}^k \alpha_i = 1\right\},$$

$$\operatorname{cone} E := \left\{\textstyle\sum_{i=1}^k \alpha_i x_i \mid k \in \mathbb{N},\ x_i \in E,\ \alpha_i \geq 0\right\},$$

$$\operatorname{aff} E := \left\{\textstyle\sum_{i=1}^k \alpha_i x_i \mid k \in \mathbb{N},\ x_i \in E,\ \alpha_i \in \mathbb{R},\ \sum_{i=1}^k \alpha_i = 1\right\}.$$

The *closure* and *interior* of E are denoted as $\operatorname{cl} E$ and $\operatorname{int} E$, respectively, whereas its *relative interior*, namely the interior of E as a subspace of $\operatorname{aff} E$, is denoted as $\operatorname{relint} E$. With $\mathrm{B}(x;r)$ and $\overline{\mathrm{B}}(x;r)$ we indicate, respectively, the open and closed balls centered at x with radius r.

Sequences The notation $(a^k)_{k\in K}$ represents a sequence indexed by elements of the set K, and given a set E we write $(a^k)_{k\in K} \subset E$ to indicate that $a^k \in E$ for all indices $k \in K$. We say that $(a^k)_{k\in K} \subset \mathbb{R}^n$ is *summable* if $\sum_{k\in K} \|a^k\|$ is finite, and *square-summable* if $(\|a^k\|^2)_{k\in K}$ is summable. We say that the sequence converges to a point $a \in \mathbb{R}^n$ *superlinearly* if either $a^k = a$ for some $k \in \mathbb{N}$, or $\|a^{k+1}-a\|/\|a^k-a\| \to 0$; if $\|a^{k+1}-a\|/\|a^k-a\|^q$ is bounded for some $q > 1$, then we say that the sequence converges *superlinearly with order* q, and in case $q = 2$ we say that the convergence is *quadratic*.

Extended-Real Valued Functions The extended-real line is $\overline{\mathbb{R}} = \mathbb{R} \cup \{\infty\}$. Given a function $h : \mathbb{R}^n \to [-\infty, \infty]$, its *epigraph* is the set

$$\operatorname{epi} h := \left\{(x,\alpha) \in \mathbb{R}^n \times \mathbb{R} \mid h(x) \leq \alpha\right\},$$

while its *domain* is

$$\operatorname{dom} h := \left\{x \in \mathbb{R}^n \mid h(x) < \infty\right\},$$

and for $\alpha \in \mathbb{R}$ its α*-level set* is

$$\operatorname{lev}_{\leq\alpha} h := \left\{x \in \mathbb{R}^n \mid h(x) \leq \alpha\right\}.$$

Function h is said to be *lower semicontinuous (lsc)* if $\operatorname{epi} h$ is a closed set in $\mathbb{R}^{n+1}$ (equivalently, h is said to be *closed*); in particular, all level sets of an lsc function are closed. We say that h is *proper* if $h > -\infty$ and $\operatorname{dom} h \neq \emptyset$, and that it is *level bounded* if for all $\alpha \in \mathbb{R}$ the level set $\operatorname{lev}_{\leq\alpha} h$ is a bounded subset of $\mathbb{R}^n$.

Continuity and Smoothness A function $G : \mathbb{R}^n \to \mathbb{R}^m$ is *ϑ-Hölder continuous* for some $\vartheta > 0$ if there exists $L \geq 0$ such that

$$\|G(x) - G(x')\| \leq L\|x - x'\|^{\vartheta} \tag{15.3}$$

for all x, x'. In case $\vartheta = 1$ we say that G is (L-)Lipschitz continuous. G is *strictly differentiable* at $\bar{x} \in \mathbb{R}^n$ if the Jacobian matrix $JG(\bar{x}) \coloneqq \left(\frac{\partial G_i}{\partial x_j}(\bar{x})\right)_{i,j}$ exists and

$$\lim_{\substack{x,x' \to \bar{x} \\ x \neq x'}} \frac{\|G(x') - JG(\bar{x})(x' - x) - G(x)\|}{\|x' - x\|} = 0. \tag{15.4}$$

The class of functions $h : \mathbb{R}^n \to \mathbb{R}$ that are k times continuously differentiable is denoted as $C^k(\mathbb{R}^n)$. We write $h \in C^{1,1}(\mathbb{R}^n)$ to indicate that $h \in C^1(\mathbb{R}^n)$ and that ∇h is Lipschitz continuous with modulus L_h. To simplify the terminology, we will say that such an h is *L_h-smooth*. If h is L_h-smooth and convex, then for any $u, v \in \mathbb{R}^n$

$$0 \leq h(v) - \left[h(u) + \langle \nabla h(u), v - u\rangle\right] \leq \tfrac{L_h}{2}\|v - u\|^2. \tag{15.5}$$

Moreover, having $h \in C^{1,1}(\mathbb{R}^n)$ and μ_h-strongly convex is equivalent to having

$$\mu_h\|v - u\|^2 \leq \langle \nabla h(v) - \nabla h(u), v - u\rangle \leq L_h\|v - u\|^2 \tag{15.6}$$

for all $u, v \in \mathbb{R}^n$.

Set-Valued Mappings We use the notation $H : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$ to indicate a point-to-set function $H : \mathbb{R}^n \to \mathcal{P}(\mathbb{R}^m)$, where $\mathcal{P}(\mathbb{R}^m)$ is the power set of $\mathbb{R}^m$ (the set of all subsets of $\mathbb{R}^m$). The *graph* of H is the set

$$\operatorname{gph} H \coloneqq \left\{(x, y) \in \mathbb{R}^n \times \mathbb{R}^m \mid y \in H(x)\right\},$$

while its *domain* is

$$\operatorname{dom} H \coloneqq \left\{x \in \mathbb{R}^n \mid H(x) \neq \emptyset\right\}.$$

We say that H is *outer semicontinuous (osc)* at $\bar{x} \in \operatorname{dom} H$ if for any $\varepsilon > 0$ there exists $\delta > 0$ such that $H(x) \subseteq H(\bar{x}) + \mathrm{B}(0; \varepsilon)$ for all $x \in \mathrm{B}(\bar{x}; \delta)$. In particular, this implies that whenever $(x^k)_{k \in \mathbb{N}} \subseteq \operatorname{dom} H$ converges to x and $(y^k)_{k \in \mathbb{N}}$ converges to y with $y^k \in H(x^k)$ for all k, it holds that $y \in H(x)$. We say that H is osc (without mention of a point) if H is osc at every point of its domain or, equivalently, if gph H is a closed subset of $\mathbb{R}^n \times \mathbb{R}^m$ (notice that this notion does *not* reduce to lower semicontinuity for a single-valued function H).

Convex Analysis The *indicator function* of a set $S \subseteq \mathbb{R}^n$ is denoted as $\delta_S : \mathbb{R}^n \to \overline{\mathbb{R}}$, namely

$$\delta_S(x) = \begin{cases} 0 & \text{if } x \in S, \\ \infty & \text{otherwise.} \end{cases} \tag{15.7}$$

If S is nonempty closed and convex, then δ_S is proper convex and lsc, and both the projection $\mathscr{P}_S : \mathbb{R}^n \to \mathbb{R}^n$ and the distance $\operatorname{dist}(\,\cdot\,, S) : \mathbb{R}^n \to [0, \infty)$ are well-defined functions, given by $\mathscr{P}_S(x) = \operatorname{argmin}_{z \in S} \|z - x\|$ and $\operatorname{dist}(x, S) = \min_{z \in S} \|z - x\|$, respectively.

The *subdifferential* of h is the set-valued mapping $\partial h : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ defined as

$$\partial h(x) := \left\{ v \in \mathbb{R}^n \mid h(z) \geq h(x) + \langle v, z - x \rangle \ \ \forall z \in \mathbb{R}^n \right\}. \tag{15.8}$$

A vector $v \in \partial h(x)$ is called a *subgradient* of h at x. It holds that $\operatorname{dom} \partial h \subseteq \operatorname{dom} h$, and if h is proper and convex, then $\operatorname{dom} \partial h$ is a nonempty convex set containing $\operatorname{relint} \operatorname{dom} h$, and $\partial h(x)$ is convex and closed for all $x \in \mathbb{R}^n$.

A function h is said to be μ-strongly convex for some $\mu \geq 0$ if $h - \frac{\mu}{2}\| \cdot \|^2$ is convex. Unless differently specified, we allow for $\mu = 0$ which corresponds to h being convex but not strongly so. If $\mu > 0$, then h has a unique (global) minimizer.

15.2.2 *Generalized Differentiability*

Due to its inherent nonsmooth nature, classical notions of differentiability may not be directly applicable in problem (15.1). This subsection contains some definitions and known facts on generalized differentiability that will be needed later on in the chapter. The interested reader is referred to the textbooks [15, 19, 65] for the details.

Definition 15.2.1 (Bouligand and Clarke Subdifferentials) Let $G : \mathbb{R}^n \to \mathbb{R}^m$ be locally Lipschitz continuous, and let $C_G \subseteq \mathbb{R}^n$ be the set of points at which G is differentiable (in particular $\mathbb{R}^n \setminus C_G$ has measure zero). The B-subdifferential (also known as *Bouligand* or *limiting Jacobian*) of G at $\bar{x}$ is the set-valued mapping $\partial_B G : \mathbb{R}^n \rightrightarrows \mathbb{R}^{m \times n}$ defined as

$$\partial_B G(\bar{x}) := \left\{ H \in \mathbb{R}^{m \times n} \mid \exists (x^k)_{k \in \mathbb{N}} \subset C_G \text{ with } x^k \to \bar{x},\ JG(x^k) \to H \right\}, \tag{15.9}$$

whereas the (*Clarke*) *generalized Jacobian* of G at $\bar{x}$ is $\partial_C G : \mathbb{R}^n \rightrightarrows \mathbb{R}^{m \times n}$ given by

$$\partial_C G(\bar{x}) := \operatorname{conv}(\partial_B G(\bar{x})). \tag{15.10}$$

If $G : \mathbb{R}^n \to \mathbb{R}^m$ is locally Lipschitz on $\mathbb{R}^n$, then $\partial_C G(x)$ is a nonempty, convex, and compact subset of $\mathbb{R}^{m \times n}$ matrices, and as a set-valued mapping it is osc at every

$x \in \mathbb{R}^n$. *Semismooth* functions [60] are precisely Lipschitz-continuous mappings for which the generalized Jacobian (and consequently the B-subdifferential) furnishes a first-order approximation.

Definition 15.2.2 (Semismooth Mappings) Let $G : \mathbb{R}^n \to \mathbb{R}^m$ be locally Lipschitz continuous at $\bar{x}$. We say that G is *semismooth* at $\bar{x}$ if

$$\limsup_{\substack{x \to \bar{x} \\ H \in \partial_C G(x)}} \frac{\|G(x) + H(\bar{x} - x) - G(\bar{x})\|}{\|x - \bar{x}\|} = 0. \tag{15.11a}$$

We say that G is ϑ*-order semismooth* for some $\vartheta > 0$ if the condition can be strengthened to

$$\limsup_{\substack{x \to \bar{x} \\ H \in \partial_C G(x)}} \frac{\|G(x) + H(\bar{x} - x) - G(\bar{x})\|}{\|x - \bar{x}\|^{1+\vartheta}} < \infty, \tag{15.11b}$$

and in case $\vartheta = 1$ we say that G is *strongly semismooth*.

To simplify the notation, we adopt the small-o and big-O convention to write expressions as (15.11a) in the compact form $G(x) + H(\bar{x} - x) - G(\bar{x}) = o(\|x - \bar{x}\|)$, and similarly (15.11b) as $G(x) + H(\bar{x} - x) - G(\bar{x}) = O(\|x - \bar{x}\|^{1+\vartheta})$. We remark that the original definition of semismoothness given by [49] requires G to be directionally differentiable at x. The definition given here is the one employed by [25]. It is also worth remarking that $\partial_C G(x)$ can be replaced with the smaller set $\partial_B G(x)$ in Definition 15.2.2. Fortunately, the class of semismooth mappings is rich enough to include many functions arising in interesting applications. For example, *piecewise smooth (PC^1) mappings* are semismooth everywhere. Recall that a continuous mapping $G : \mathbb{R}^n \to \mathbb{R}^m$ is PC^1 if there exists a finite collection of smooth mappings $G_i : \mathbb{R}^n \to \mathbb{R}^m$, $i = 1, \ldots, N$, such that

$$G(x) \in \{G_1(x), \ldots, G_N(x)\} \quad \forall x \in \mathbb{R}^n. \tag{15.12}$$

The definition of PC^1 mapping given here is less general than the one of, e.g., [66, §4] but it suffices for our purposes. For every $x \in \mathbb{R}^n$ we introduce the set of essentially active indices

$$I_G^e(x) := \left\{i \mid x \in \operatorname{cl}\left(\operatorname{int}\{w \mid G(w) = G_i(w)\}\right)\right\}. \tag{15.13}$$

In other words, $I_G^e(x)$ contains only indices of the pieces G_i for which there exists a full-dimensional set on which G agrees with G_i. In accordance with Definition 15.2.1, the generalized Jacobian of G at x is the convex hull of the Jacobians of the essentially active pieces, *i.e.*, [66, Prop. 4.3.1]

$$\partial_C G(x) = \operatorname{conv}\left\{JG_i(x) \mid i \in I_G^e(x)\right\}. \tag{15.14}$$

The following definition is taken from [19, Def. 7.5.13].

Definition 15.2.3 (Linear Newton Approximation) Let $G : \mathbb{R}^n \to \mathbb{R}^m$ be continuous on $\mathbb{R}^n$. We say that G admits a *linear Newton approximation (LNA)* at $\bar{x} \in \mathbb{R}^n$ if there exists a set-valued mapping $\mathcal{H} : \mathbb{R}^n \rightrightarrows \mathbb{R}^{m\times n}$ that has nonempty compact images, is outer semicontinuous at $\bar{x}$, and

$$\limsup_{\substack{x\to\bar{x}\\ H\in\mathcal{H}(x)}} \frac{\|G(x)+H(\bar{x}-x)-G(\bar{x})\|}{\|x-\bar{x}\|} = 0.$$

If for some $\vartheta > 0$ the condition can be strengthened to

$$\limsup_{\substack{x\to\bar{x}\\ H\in\mathcal{H}(x)}} \frac{\|G(x)+H(\bar{x}-x)-G(\bar{x})\|}{\|x-\bar{x}\|^{1+\vartheta}} < \infty,$$

then we say that $\mathcal{H}$ is a ϑ*-order LNA*, and if $\vartheta = 1$ we say that $\mathcal{H}$ is a *strong LNA*.

Functions G as in Definition 15.2.3 are also referred to as $\mathcal{H}$-semismooth in the literature, see, *e.g.*, [78], however we prefer to stick to the terminology of [19] and rather say that $\mathcal{H}$ is a LNA for G. Arguably the most notable example of a LNA for semismooth mappings is the generalized Jacobian, cf. Definition 15.2.1. However, semismooth mappings can admit LNAs different from the generalized Jacobian. More importantly, mappings that are not semismooth may also admit a LNA.

Lemma 15.2.4 ([19, Prop. 7.4.10]) *Let $h \in C^1(\mathbb{R}^n)$ and suppose that $\mathcal{H} : \mathbb{R}^n \rightrightarrows \mathbb{R}^{n\times n}$ is a LNA for ∇h at $\bar{x}$. Then,*

$$\lim_{\substack{x\to\bar{x}\\ H\in\mathcal{H}(x)}} \frac{h(x)-h(\bar{x})-\langle\nabla h(\bar{x}), x-\bar{x}\rangle-\frac{1}{2}\langle H(x-\bar{x}), x-\bar{x}\rangle}{\|x-\bar{x}\|^2} = 0. \tag{15.15}$$

We remark that although [19, Prop. 7.4.10] assumes semismoothness of ∇h at $\bar{x}$ and uses $\partial_C(\nabla h)$ in place of $\mathcal{H}$; however, exactly the same arguments apply for any LNA of ∇h at $\bar{x}$ even without the semismoothness assumption.

15.3 Proximal Algorithms

15.3.1 Proximal Point and Moreau Envelope

The *proximal mapping* of a proper closed and convex function $h : \mathbb{R}^n \to \overline{\mathbb{R}}$ with parameter $\gamma > 0$ is $\operatorname{prox}_{\gamma h} : \mathbb{R}^n \to \mathbb{R}^n$, given by

$$\operatorname{prox}_{\gamma h}(x) \coloneqq \operatorname{argmin}_{w\in\mathbb{R}^n}\Big\{\overbrace{h(w)+\tfrac{1}{2\gamma}\|w-x\|^2}^{\mathcal{M}_\gamma^h(w;x)}\Big\}. \tag{15.16}$$

The *majorization model* $\mathcal{M}_\gamma^h(x;\ \cdot\ )$ is a proper and strongly convex function, and therefore has a unique minimizer. The value function, as opposed to the minimizer, defines the *Moreau envelope* $h^\gamma : \mathbb{R}^n \to \mathbb{R}$, namely

$$h^\gamma(x) := \min_{w\in\mathbb{R}^n} \left\{ h(w) + \tfrac{1}{2\gamma}\|w - x\|^2 \right\}, \tag{15.17}$$

which is real valued and Lipschitz differentiable, despite the fact that h might be extended-real valued. Properties of the Moreau envelope and the proximal mapping are well documented in the literature, see, *e.g.*, [2, §24]. For example, $\mathrm{prox}_{\gamma h}$ is nonexpansive (Lipschitz continuous with modulus 1) and is characterized by the implicit inclusion

$$\hat{x} = \mathrm{prox}_{\gamma h}(x) \quad \Leftrightarrow \quad \tfrac{1}{\gamma}(x - \hat{x}) \in \partial h(\hat{x}). \tag{15.18}$$

For the sake of a brief recap, we now list some other important known properties. Theorem 15.3.1 provides some relations between h and its Moreau envelope h^γ, which we informally refer to as *sandwich property* for apparent reasons, cf. Figure 15.1. Theorem 15.3.2 highlights that the minimization of a (proper, lsc and) convex function can be expressed as the convex smooth minimization of its Moreau envelope.

Theorem 15.3.1 (Moreau Envelope: Sandwich Property [2, 12]) *For all $\gamma > 0$ the following hold for the cost function φ:*

(i) $\varphi^\gamma(x) \leq \varphi(x) - \frac{1}{2\gamma}\|x - \hat{x}\|^2$ *for all* $x \in \mathbb{R}^n$ *where* $\hat{x} := \mathrm{prox}_{\gamma\varphi}(x)$*;*
(ii) $\varphi(\hat{x}) = \varphi^\gamma(x) - \frac{1}{2\gamma}\|x - \hat{x}\|^2$ *for all* $x \in \mathbb{R}^n$ *where* $\hat{x} := \mathrm{prox}_{\gamma\varphi}(x)$*;*
(iii) $\varphi^\gamma(x) = \varphi(x)$ *iff* $x \in \operatorname{argmin}\varphi$.

Proof

- 15.3.1*(i)*. This fact is shown in [12, Lem. 3.2] for a more general notion of proximal point operator; namely, the square Euclidean norm appearing in (15.16) and (15.17) can be replaced by arbitrary Bregman divergences. In this simpler case, since $\frac{1}{\gamma}(x - \hat{x})$ is a subgradient of φ at $\hat{x}$, cf. (15.18), we have

$$\varphi(x) \geq \varphi(\hat{x}) + \langle \tfrac{1}{\gamma}(x - \hat{x}), x - \hat{x}\rangle = \varphi(\hat{x}) + \tfrac{1}{\gamma}\|x - \hat{x}\|^2. \tag{15.19}$$

  The claim now follows by subtracting $\frac{1}{2\gamma}\|x - \hat{x}\|^2$ from both sides.
- 15.3.1*(ii)*. Follows by definition, cf. (15.16) and (15.17).
- 15.3.1*(iii)*. See [2, Prop. 17.5]. □

□

Theorem 15.3.2 (Moreau Envelope: Convex Smooth Minimization Equivalence [2]) *For all $\gamma > 0$ the following hold for the cost function φ:*

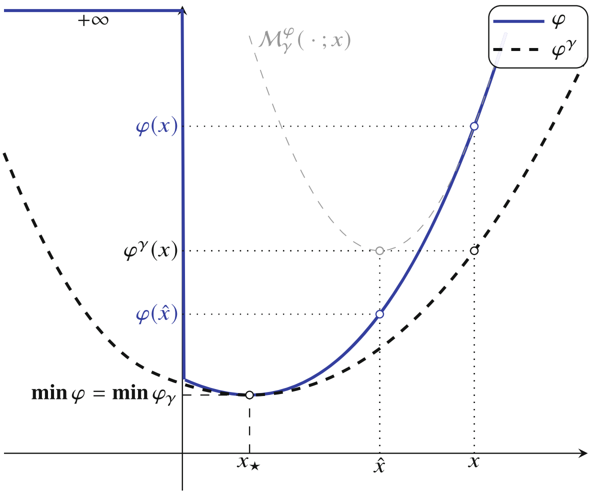

Fig. 15.1 *Moreau envelope of the function* $\varphi(x) = \frac{1}{3}x^3 + x^2 - x + 1 + \delta_{[0,\infty)}(x)$ *with parameter* $\gamma = 0.2$. At each point x, the Moreau envelope φ^γ is the minimum of the quadratic majorization model $\mathcal{M}_\gamma^\varphi = \varphi + \frac{1}{2\gamma}(\,\cdot\, - x)^2$, the unique minimizer being, by definition, the proximal point $\hat{x} := \operatorname{prox}_{\gamma\varphi}(x)$. It is a convex smooth lower bound to φ, despite the fact that φ might be extended-real valued. Function φ and its Moreau envelope φ^γ have same inf and argmin; in fact, the two functions agree (only) on the set of minimizers. In general, φ^γ is *sandwiched* as $\varphi \circ \operatorname{prox}_{\gamma\varphi} \le \varphi^\gamma \le \varphi$

(i) φ^γ *is convex and smooth with* $L_{\varphi^\gamma} = \gamma^{-1}$ *and* $\nabla\varphi^\gamma(x) = \gamma^{-1}\big(x - \operatorname{prox}_{\gamma\varphi}(x)\big)$*;*
(ii) $\inf\varphi = \inf\varphi^\gamma$*;*
(iii) $x_\star \in \operatorname{argmin}\varphi$ *iff* $x_\star \in \operatorname{argmin}\varphi^\gamma$ *iff* $\nabla\varphi^\gamma(x_\star) = 0$.

Proof

- 15.3.2*(i)*. See [2, Prop.s 12.15 and 12.30].
- 15.3.2*(ii)*. See [2, Prop. 12.9(iii)].
- 15.3.2*(iii)*. See [2, Prop. 17.5]. □

□

As a consequence of Theorem 15.3.2, one can address the minimization of the convex but possibly nonsmooth and extended-real-valued function φ by means of gradient descent on the smooth envelope function φ^γ with stepsize $0 < \tau < 2/L_{\varphi_\gamma} = 2\gamma$. As first noticed by Rockafellar [64], this simply amounts to (relaxed) fixed-point iterations of the proximal point operator, namely

$$x^+ = (1-\lambda)x + \lambda \operatorname{prox}_{\gamma\varphi}(x), \tag{15.20}$$

where $\lambda = \tau/\gamma \in (0, 2)$ is a possible relaxation parameter. The scheme, known as *proximal point algorithm* (PPA) and first introduced by Martinet [45], is well covered by the broad theory of monotone operators, where convergence properties can be easily derived with simple tools of Fejérian monotonicity, see, *e.g.*, [2, Thm.s 23.41 and 27.1]. Nevertheless, not only does the interpretation as gradient method provide a beautiful theoretical link, but it also enables the employment of acceleration techniques exclusively stemming from smooth unconstrained optimization, such as Nesterov's extrapolation [26] or quasi-Newton schemes [13], see also [7] for extensions to the dual formulation.

15.3.2 Forward-Backward Splitting

While it is true that every convex minimization problem can be smoothened by means of the Moreau envelope, unfortunately it is often the case that the computation of the proximal operator (which is needed to evaluate the envelope) is as hard as solving the original problem. For instance, evaluating the Moreau envelope of the cost of modeling a convex QP at one point amounts to solving another QP with same constraints and augmented cost. To overcome this limitation there comes the idea of *splitting schemes*, which decompose a complex problem in small components which are easier to operate onto. A popular such scheme is the *forward-backward splitting* (FBS), which addresses minimization problems of the form (15.1).

Given a point $x \in \mathbb{R}^n$, one iteration of *forward-backward splitting* (FBS) for problem (15.1) with stepsize $\gamma > 0$ and relaxation $\lambda > 0$ consists in

$$x^+ = (1 - \lambda)x + \lambda T_\gamma(x), \tag{15.21}$$

where

$$T_\gamma(x) \coloneqq \operatorname{prox}_{\gamma g}(x - \gamma \nabla f(x)) \tag{15.22}$$

is the *forward-backward operator*, characterized as

$$\bar{x} = T_\gamma(x) \quad \Leftrightarrow \quad \tfrac{1}{\gamma}(x - \bar{x}) - (\nabla f(x) - \nabla f(\bar{x})) \in \partial\varphi(\bar{x}), \tag{15.23}$$

as it follows from (15.18). FBS interleaves a gradient descent step on f and a proximal point step on g, and as such it is also known as *proximal gradient method*. If both f and g are (lsc, proper and) convex, then the solutions to (15.1) are exactly the fixed points of the forward-backward operator T_γ. In other words,

$$x_\star \in \operatorname{argmin} \varphi \quad \text{iff} \quad R_\gamma(x_\star) = 0, \tag{15.24}$$

where

$$R_\gamma(x) := \tfrac{1}{\gamma}\big(x - \operatorname{prox}_{\gamma g}(x - \gamma\nabla f(x))\big) \tag{15.25}$$

is the *forward-backward residual*.[1] FBS iterations (15.21) are well known to converge to a solution to (15.1) provided that f is smooth and that the parameters are chosen as $\gamma \in (0, 2/L_f)$ and $\lambda \in (0, 2 - \gamma L_f/2)$ [2, Cor. 28.9] ($\lambda = 1$, which is always feasible, is the typical choice).

15.3.3 Error Bounds and Quadratic Growth

We conclude the section with some inequalities that will be useful in the sequel.

Lemma 15.3.3 *Suppose that $\mathcal{X}_\star$ is nonempty. Then,*

$$\varphi(x) - \varphi_\star \le \operatorname{dist}(0, \partial\varphi(x)) \operatorname{dist}(x, \mathcal{X}_\star) \quad \forall x \in \mathbb{R}^n. \tag{15.26}$$

Proof From the subgradient inequality it follows that for all $x_\star \in \mathcal{X}_\star$ and $v \in \partial\varphi(x)$ we have

$$\varphi(x) - \varphi_\star = \varphi(x) - \varphi(x_\star) \le \langle v, x - x_\star\rangle \le \|v\|\|x - x_\star\| \tag{15.27}$$

and the claimed inequality follows from the arbitrariness of $x_\star$ and v. □

Lemma 15.3.4 *Suppose that $\mathcal{X}_\star$ is nonempty. For all $x \in \mathbb{R}^n$ and $\gamma > 0$ the following holds:*

$$\|R_\gamma(x)\| \ge \tfrac{1}{1+\gamma L_f} \operatorname{dist}\big(0, \partial\varphi(T_\gamma(x))\big) \tag{15.28}$$

Proof Let $\bar{x} := T_\gamma(x)$. The characterization (15.23) of T_γ implies that

$$\|R_\gamma(x)\| \ge \operatorname{dist}\big(0, \partial\varphi(\bar{x})\big) - \|\nabla f(x) - \nabla f(\bar{x})\| \ge \operatorname{dist}\big(0, \partial\varphi(\bar{x})\big) - \gamma L_f \|R_\gamma(x)\|. \tag{15.29}$$

After trivial rearrangements the sought inequality follows. □

Further interesting inequalities can be derived if the cost function φ satisfies an *error bound*, which can be regarded as a generalization of strong convexity that does not require uniqueness of the minimizer. The interested reader is referred to [3, 17, 43, 55] and the references therein for extensive discussions.

[1] Due to apparent similarities with gradient descent iterations, having $x^+ = x - \gamma R_\gamma(x)$ in FBS, R_γ is also referred to as (generalized) *gradient mapping*, see, *e.g.*, [17]. In particular, if $g = 0$, then $R_\gamma = \nabla f$ whereas if $f = 0$ then $R_\gamma = \nabla g^\gamma$. The analogy will be supported by further evidence in the next section where we will see that, up to a change of metric, indeed R_γ is the gradient of the *forward-backward envelope function*.

Definition 15.3.5 (Quadratic Growth and Error Bound) Suppose that $\mathcal{X}_\star \neq \emptyset$. Given $\mu, \nu > 0$, we say that

(a) φ satisfies the *quadratic growth* with constants (μ, ν) if

$$\varphi(x) - \varphi_\star \geq \tfrac{\mu}{2} \operatorname{dist}(x, \mathcal{X}_\star)^2 \quad \forall x \in \operatorname{lev}_{\leq \varphi_\star + \nu} \varphi; \tag{15.30}$$

(b) φ satisfies the *error bound* with constants (μ, ν) if

$$\operatorname{dist}(0, \partial\varphi(x)) \geq \tfrac{\mu}{2} \operatorname{dist}(x, \mathcal{X}_\star) \quad \forall x \in \operatorname{lev}_{\leq \varphi_\star + \nu} \varphi. \tag{15.31}$$

In case $\nu = \infty$ we say that the properties are satisfied *globally*.

Theorem 15.3.6 ([17, Thm. 3.3]) *For a proper convex and lsc function, the quadratic growth with constants* (μ, ν) *is equivalent to the error bound with same constants.*

Lemma 15.3.7 (Globality of Quadratic Growth) *Suppose that* φ *satisfies the quadratic growth with constants* (μ, ν)*. Then, for every* $\nu' > \nu$ *it satisfies the quadratic growth with constants* (μ', ν')*, where*

$$\mu' := \tfrac{\mu}{2} \min\left\{1, \tfrac{\nu}{\nu' - \nu}\right\}. \tag{15.32}$$

Proof Let $\nu' > \nu$ be fixed, and let $x \in \operatorname{lev}_{\leq \varphi_\star + \nu'}$ be arbitrary. Since $\mu' \leq \mu$, the claim is trivial if $\varphi(x) \leq \varphi_\star + \nu$; we may thus suppose that $\varphi(x) > \varphi_\star + \nu$. Let z be the projection of x onto the (nonempty closed and convex) level set $\operatorname{lev}_{\leq \varphi_\star + \nu}$, and observe that $\varphi(z) = \varphi_\star + \nu$. With Lemma 15.3.3 and Theorem 15.3.6 we can upper bound ν as

$$\nu = \varphi(z) - \varphi_\star \leq \operatorname{dist}(0, \partial\varphi(z)) \operatorname{dist}(z, \mathcal{X}_\star) \leq \tfrac{2}{\mu} \operatorname{dist}(0, \partial\varphi(z))^2. \tag{15.33}$$

Moreover, it follows from [28, Thm. 1.3.5] that there exists a subgradient $v \in \partial\varphi(z)$ such that $\langle v, x - z \rangle = \|v\| \|x - z\|$. Then,

$$\begin{aligned} \varphi(x) &\geq \varphi(z) + \langle v, x - z \rangle = \varphi(z) + \|v\| \|x - z\| \geq \varphi(z) + \operatorname{dist}(0, \partial\varphi(z)) \|x - z\| \\ &\overset{(15.33)}{\geq} \varphi(z) + \sqrt{\tfrac{\mu\nu}{2}} \|x - z\|. \end{aligned} \tag{15.34}$$

By subtracting $\varphi(z)$ from the first and last terms we obtain

$$\|x - z\| \leq \sqrt{\tfrac{2}{\mu\nu}} \big(\varphi(x) - \varphi(z)\big) \leq \sqrt{\tfrac{2}{\mu\nu}} (\nu' - \nu), \tag{15.35}$$

which implies

$$\|x - z\| \geq \sqrt{\tfrac{\mu\nu}{2}} \tfrac{1}{\nu' - \nu} \|x - z\|^2. \tag{15.36}$$

Thus,

$$\varphi(x) - \varphi_\star \overset{(15.34)}{\geq} \varphi(z) - \varphi_\star + \sqrt{\tfrac{\mu\nu}{2}}\|x - z\|$$

using the quadratic growth at z and the inequality (15.36)

$$\begin{aligned} &\geq \tfrac{\mu}{2}\operatorname{dist}(z, \mathcal{X}_\star)^2 + \tfrac{\mu\nu}{2(\nu'-\nu)}\|x - z\|^2 \\ &\geq \tfrac{\mu}{2}\min\left\{1, \tfrac{\nu}{\nu'-\nu}\right\}\left[\operatorname{dist}(z, \mathcal{X}_\star)^2 + \|x - z\|^2\right]. \end{aligned}$$

By using the fact that $a^2 + b^2 \geq \frac{1}{2}(a + b)^2$ for any $a, b \in \mathbb{R}$ together with the triangular inequality $\operatorname{dist}(x, \mathcal{X}_\star) \leq \|x - z\| + \operatorname{dist}(z, \mathcal{X}_\star)$, we conclude that $\varphi(x) - \varphi_\star \geq \frac{\mu'}{2}\operatorname{dist}(x, \mathcal{X}_\star)^2$, with μ' as in the statement. Since μ' depends only on μ, ν, and ν', from the arbitrariness of $x \in \operatorname{lev}_{\leq \varphi_\star + \nu'}$ the claim follows. □

Theorem 15.3.8 ([17, Cor. 3.6]) *Suppose that φ satisfies the quadratic growth with constants (μ, ν). Then, for all $\gamma \in (0, 1/L_f)$ and $x \in \operatorname{lev}_{\leq \varphi_\star + \nu}\varphi$ we have*

$$\operatorname{dist}(x, \mathcal{X}_\star) \leq (\gamma + 2/\mu)(1 + \gamma L_f)\|R_\gamma(x)\|. \tag{15.37}$$

15.4 Forward-Backward Envelope

There are clearly infinite ways of representing the (proper, lsc and) convex function φ in (15.1) as the sum of two convex functions f and g with f smooth, and each of these choices leads to a different FBS operator T_γ. If $f = 0$, for instance, then T_γ reduces to $\operatorname{prox}_{\gamma\varphi}$, and consequently FBS (15.21) to the PPA (15.20). A natural question then arises, whether a function exists that serves as "envelope" for FBS in the same way that φ_γ does for $\operatorname{prox}_{\gamma\varphi}$. We will now provide a positive answer to this question by reformulating the nonsmooth problem (15.1) as the minimization of a differentiable function. To this end, the following requirements on f and g will be assumed throughout the chapter without further mention.

Assumption I (Basic Requirements) *In problem (15.1),*

(i) $f : \mathbb{R}^n \to \mathbb{R}$ is convex, twice continuously differentiable and L_f-smooth;
(ii) $g : \mathbb{R}^n \to \overline{\mathbb{R}}$ is lsc, proper, and convex.

Compared to the classical FBS assumptions, the only additional requirement is twice differentiability of f. This ensures that the *forward operator* $x \mapsto x - \gamma\nabla f(x)$ is differentiable; we denote its Jacobian as $Q_\gamma : \mathbb{R}^n \to \mathbb{R}^{n\times n}$, namely

$$Q_\gamma(x) := \mathrm{I} - \gamma\nabla^2 f(x). \tag{15.38}$$

Notice that, due to the bound $\nabla^2 f(x) \preceq L_f\mathrm{I}$ (which follows from L_f-smoothness of f, see [53, Lem. 1.2.2]) $Q_\gamma(x)$ is invertible (in fact, positive definite) whenever $\gamma < {}^1\!/_{L_f}$. Moreover, due to the chain rule and Theorem 15.3.2*(i)* we have that

$$\begin{aligned}\nabla\big[g^\gamma \circ (\mathrm{id} - \gamma\nabla f)\big](x) &= \gamma^{-1} Q_\gamma(x)\big[x - \gamma\nabla f(x) - \operatorname{prox}_{\gamma g}(x - \gamma\nabla f(x))\big]\\ &= Q_\gamma(x)\big[R_\gamma(x) - \nabla f(x)\big].\end{aligned}$$

Rearranging,

$$\begin{aligned}Q_\gamma(x)R_\gamma(x) &= \nabla f(x) - \gamma\nabla^2 f(x)\nabla f(x) + \nabla\big[g^\gamma \circ (\mathrm{id} - \gamma\nabla f)\big](x)\\ &= \nabla f(x) - \nabla\left[\tfrac{\gamma}{2}\|\nabla f\|^2\right](x) + \nabla\big[g^\gamma \circ (\mathrm{id} - \gamma\nabla f)\big](x)\\ &= \nabla\left[f - \tfrac{\gamma}{2}\|\nabla f\|^2 + g^\gamma \circ (\mathrm{id} - \gamma\nabla f)\right](x)\end{aligned}$$

we obtain the gradient of a real-valued function, which we define as follows.

Definition 15.4.1 (Forward-Backward Envelope) The *forward-backward envelope* (FBE) for the composite minimization problem (15.1) is the function $\varphi_\gamma : \mathbb{R}^n \to \mathbb{R}$ defined as

$$\varphi_\gamma(x) := f(x) - \tfrac{\gamma}{2}\|\nabla f(x)\|^2 + g^\gamma(x - \gamma\nabla f(x)). \tag{15.39}$$

In the next section we discuss some of the favorable properties enjoyed by the FBE.

15.4.1 Basic Properties

We already verified that the FBE is differentiable with gradient

$$\nabla\varphi_\gamma(x) = Q_\gamma(x)R_\gamma(x). \tag{15.40}$$

In particular, for $\gamma < 1/L_f$ one obtains that a FBS step is a (scaled) gradient descent step on the FBE, similarly as the relation between Moreau envelope and PPA; namely,

$$T_\gamma(x) = x - \gamma Q_\gamma(x)^{-1}\nabla\varphi_\gamma(x). \tag{15.41}$$

To take the analysis of the FBE one step further, let us consider the equivalent expression of the operator T_γ as

$$T_\gamma(x) = \operatorname{argmin}_{w\in\mathbb{R}^n}\left\{\overbrace{f(x) + \langle\nabla f(x), w - x\rangle + \tfrac{1}{2\gamma}\|w - x\|^2 + g(w)}^{\mathcal{M}_\gamma^{f,g}(w;x)}\right\}. \tag{15.42}$$

Differently from the quadratic model $\mathcal{M}_\gamma^\varphi$ in (15.16), $\mathcal{M}_\gamma^{f,g}$ replaces the differentiable component f with a linear approximation. Building upon the idea of the Moreau envelope, instead of the minimizer $T_\gamma(x)$ we consider the value attained in the subproblem (15.42), and with simple algebra one can easily verify that this gives rise once again to the FBE:

$$\varphi_\gamma(x) = \min_{w\in\mathbb{R}^n}\left\{f(x) + \langle\nabla f(x), w - x\rangle + \tfrac{1}{2\gamma}\|w - x\|^2 + g(w)\right\}. \tag{15.43}$$

Starting from this expression we can easily mirror the properties of the Moreau envelope stated in Theorems 15.3.1 and 15.3.2. These results appeared in the independent works [54] and [57], although the former makes no mention of an "envelope" function and simply analyzes the *majorization-minimization* model $\mathcal{M}_\gamma^{f,g}$.

Theorem 15.4.2 (FBE: Sandwich Property) *Let $\gamma > 0$ and $x \in \mathbb{R}^n$ be fixed, and denote $\bar{x} = T_\gamma(x)$. The following hold:*

(i) $\varphi_\gamma(x) \le \varphi(x) - \frac{1}{2\gamma}\|x - \bar{x}\|^2$;

(ii) $\varphi_\gamma(x) - \frac{1}{2\gamma}\|x - \bar{x}\|^2 \le \varphi(\bar{x}) \le \varphi_\gamma(x) - \frac{1-\gamma L_f}{2\gamma}\|x - \bar{x}\|^2$.

In particular,

(iii) $\varphi_\gamma(x_\star) = \varphi(x_\star)$ *iff* $x_\star \in \operatorname{argmin}\varphi$.

In fact, the assumption of twice continuous differentiability of f can be dropped.

Proof

- 15.4.2*(i)* Since the minimum in (15.43) is attained at $w = \bar{x}$, cf. (15.42), we have

$$\varphi_\gamma(x) = f(x) + \langle\nabla f(x), \bar{x} - x\rangle + \tfrac{1}{2\gamma}\|\bar{x} - x\|^2 + g(\bar{x}) \tag{15.44}$$

$$
\begin{aligned}
&\leq f(x) + \langle \nabla f(x), \bar{x} - x \rangle + \tfrac{1}{2\gamma}\|\bar{x} - x\|^2 + g(x) \\
&\quad + \langle \tfrac{1}{\gamma}(x - \bar{x}) - \nabla f(x), \bar{x} - x \rangle \\
&= f(x) + g(x) - \tfrac{1}{2\gamma}\|x - \bar{x}\|^2
\end{aligned}
$$

where in the inequality we used the fact that $\frac{1}{\gamma}(x - \bar{x}) - \nabla f(x) \in \partial g(\bar{x})$, cf. (15.23).

- 15.4.2*(ii)* Follows by using (15.5) (with $h = f$, $u = x$ and $v = \bar{x}$) in (15.44).
- 15.4.2*(iii)* Follows by 15.4.2*(i)* and the optimality condition (15.24). □

□

Notice that by combining Theorems 15.4.2*(i)* and 15.4.2*(ii)* we recover the "sufficient decrease" condition of (convex) FBS [54, Thm. 1], that is

$$
\varphi(\bar{x}) \leq \varphi(x) - \tfrac{2-\gamma L_f}{2\gamma}\|x - \bar{x}\|^2 \tag{15.45}
$$

holding for all $x \in \mathbb{R}^n$ with $\bar{x} = T_\gamma(x)$.

Theorem 15.4.3 (FBE: Smooth Minimization Equivalence) *For all* $\gamma > 0$

(i) $\varphi_\gamma \in C^1(\mathbb{R}^n)$ *with* $\nabla\varphi_\gamma = Q_\gamma R_\gamma$.

Moreover, if $\gamma \in (0, 1/L_f)$ *then the following also hold:*

(ii) $\inf \varphi = \inf \varphi_\gamma$*;*

(iii) $x_\star \in \operatorname{argmin} \varphi$ *iff* $x_\star \in \operatorname{argmin} \varphi_\gamma$ *iff* $\nabla\varphi_\gamma(x_\star) = 0$.

Proof

- 15.4.3*(i)*. Since $f \in C^2(\mathbb{R}^n)$ and $g^\gamma \in C^1(\mathbb{R}^n)$ (cf. Theorem 15.3.2*(i)*), from the definition (15.39) it is apparent that φ_γ is continuously differentiable for all $\gamma > 0$. The formula for the gradient was already shown in (15.40).

Suppose now that $\gamma < 1/L_f$.

- 15.4.3*(ii)*. $\inf \varphi \leq \inf_{x\in\mathbb{R}^n} \varphi(T_\gamma(x)) \overset{15.4.2(ii)}{\leq} \inf_{x\in\mathbb{R}^n} \varphi_\gamma(x) = \inf \varphi_\gamma \overset{15.4.2(i)}{\leq} \inf \varphi$.
- 15.4.3*(iii)*. We have

$$
x_\star \in \operatorname{argmin} \varphi \overset{(15.24)}{\Leftrightarrow} R_\gamma(x_\star) = 0 \quad \Leftrightarrow \quad Q_\gamma(x_\star)R_\gamma(x_\star) = 0 \overset{15.4.3(i)}{\Leftrightarrow} \nabla\varphi_\gamma(x_\star) = 0, \tag{15.46}
$$

where the second equivalence follows from the invertibility of Q_γ.
Suppose now that $x_\star \in \operatorname{argmin} \varphi_\gamma$. Since $\varphi_\gamma \in C^1(\mathbb{R}^n)$ the first-order necessary condition reads $\nabla\varphi_\gamma = 0$, and from the equivalence proven above we conclude that $\operatorname{argmin} \varphi_\gamma \subseteq \operatorname{argmin} \varphi$. Conversely, if $x_\star \in \operatorname{argmin} \varphi$, then

$$
\varphi_\gamma(x_\star) \overset{15.4.2(iii)}{=} \varphi(x_\star) = \inf \varphi \overset{15.4.3(ii)}{=} \inf \varphi_\gamma, \tag{15.47}
$$

proving $x_\star \in \operatorname{argmin} \varphi_\gamma$, hence the inclusion $\operatorname{argmin} \varphi_\gamma \supseteq \operatorname{argmin} \varphi$. □

□

Proposition 15.4.4 (FBE and Moreau Envelope [54, Thm. 2]) *For any $\gamma \in (0, 1/L_f)$, it holds that*

$$\varphi^{\frac{\gamma}{1-\gamma L_f}} \leq \varphi_\gamma \leq \varphi^\gamma. \tag{15.48}$$

Proof We have

$$\begin{aligned}
\varphi_\gamma(x) &= \min_{w \in \mathbb{R}^n} \left\{ \underbrace{f(x) + \langle \nabla f(x), w - x \rangle} + \tfrac{1}{2\gamma}\|w - x\|^2 + g(w) \right\} \\
&\overset{(15.5)}{\leq} \min_{w \in \mathbb{R}^n} \left\{ \overbrace{f(w) - \tfrac{L_f}{2}\|w - x\|^2} + \tfrac{1}{2\gamma}\|w - x\|^2 + g(w) \right\} \\
&= \min_{w \in \mathbb{R}^n} \left\{ f(w) + g(w) + \tfrac{1-\gamma L_f}{2\gamma}\|w - x\|^2 \right\} = \varphi^{\frac{\gamma}{1-\gamma L_f}}(x).
\end{aligned}$$

Using the upper bound in (15.5) instead yields the other inequality. □

Since φ_γ is upper bounded by the γ^{-1}-smooth function φ^γ with which it shares the set of minimizers $\mathcal{X}_\star$, from (15.5) we easily infer the following quadratic upper bound.

Corollary 15.4.5 (Global Quadratic Upper Bound) *If $\mathcal{X}_\star \neq \emptyset$, then*

$$\varphi_\gamma(x) - \varphi_\star \leq \tfrac{1}{2\gamma}\operatorname{dist}(x, \mathcal{X}_\star)^2 \quad \forall x \in \mathbb{R}^n. \tag{15.49}$$

Although the FBE may fail to be convex, for $\gamma < 1/L_f$ its stationary points and minimizers coincide and are the same as those of the original function φ (Figure 15.2). That is, the minimization of φ is equivalent to the minimization of the differentiable function φ_γ. This is a clear analogy with the Moreau envelope, which in fact is the special case of the FBE corresponding to $f \equiv 0$ in the decomposition of φ. In the next result we tighten the claims of Theorem 15.4.3*(i)* when f is a convex quadratic function, showing that in this case the FBE is convex and smooth and thus recover all the properties of the Moreau envelope.

Theorem 15.4.6 (FBE: Convexity & Smoothness for Quadratic f [24, Prop. 4.4]) *Suppose that f is convex quadratic, namely $f(x) = \frac{1}{2}\langle x, Hx \rangle + \langle h, x \rangle$ for some $H \in \mathrm{S}_+(\mathbb{R}^n)$ and $h \in \mathbb{R}^n$. Then, for all $\gamma \in (0, 1/L_f]$ the FBE φ_γ is convex and smooth, with*

$$L_{\varphi_\gamma} = \tfrac{1-\gamma\mu_f}{\gamma} \quad \textit{and} \quad \mu_{\varphi_\gamma} = \min\left\{\mu_f(1-\gamma\mu_f), L_f(1-\gamma L_f)\right\}, \tag{15.50}$$

where $L_f = \lambda_{\max}(H)$ and $\mu_f = \lambda_{\min}(H)$. In particular, when f is μ_f-strongly convex the strong convexity of φ_γ is maximized for $\gamma = \frac{1}{\mu_f + L_f}$, in which case

$$L_{\varphi_\gamma} = L_f \quad \textit{and} \quad \mu_{\varphi_\gamma} = \tfrac{L_f \mu_f}{\mu_f + L_f}. \tag{15.51}$$

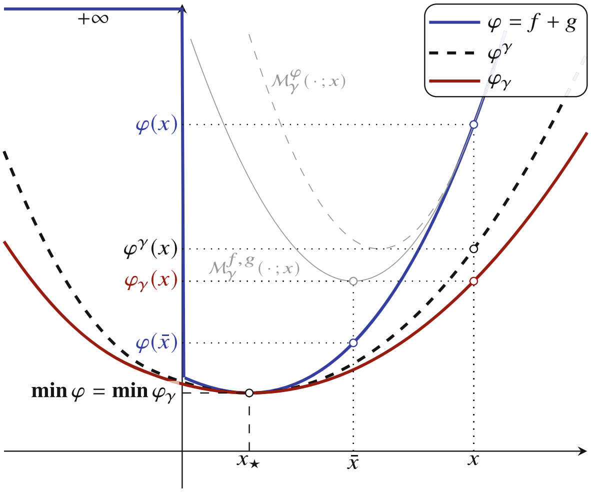

Fig. 15.2 *FBE of the function φ as in Figure 15.1 with same parameter $\gamma = 0.2$, relative to the decomposition as the sum of $f(x) = x^2 + x - 1$ and $g(x) = \frac{1}{3}x^3 + \delta_{[0,\infty)}(x)$.*
For $\gamma < 1/L_f$ ($L_f = 2$ in this example) at each point x the FBE φ_γ is the minimum of the quadratic majorization model $\mathcal{M}_\gamma^{f,g}(\,\cdot\,, x)$ for φ, the unique minimizer being the proximal gradient point $\bar{x} = T_\gamma(x)$. The FBE is a differentiable lower bound to φ and since f is quadratic in this example, it is also smooth and convex (cf. Theorem 15.4.6). In any case, its stationary points and minimizers coincide, and are equivalent to the minimizers of φ

Proof Letting $Q \coloneqq \mathrm{I} - \gamma H$, we have that $Q_\gamma \equiv Q$ and $x - \gamma\nabla f(x) = Qx - \gamma h$. Therefore,

$$\begin{aligned}\gamma\langle\nabla\varphi_\gamma(x) - \nabla\varphi_\gamma(y), x - y\rangle &\overset{(15.40)}{=} \langle Q(R_\gamma(x) - R_\gamma(y)), x - y\rangle \\ &= \langle Q(x-y), x-y\rangle - \langle Q(T_\gamma(x) - T_\gamma(y)), x - y\rangle \\ &= \|x - y\|_Q^2 \\ &\quad - \langle \operatorname{prox}_{\gamma g}(Qx - \gamma h) \\ &\quad - \operatorname{prox}_{\gamma g}(Qy - \gamma h), Q(x - y)\rangle.\end{aligned}$$

From the *firm* nonexpansiveness of $\operatorname{prox}_{\gamma g}$ (see [2, Prop.s 4.35(iii) and 12.28]) it follows that

$$0 \le \langle \operatorname{prox}_{\gamma g}(Qx-\gamma h) - \operatorname{prox}_{\gamma g}(Qy-\gamma h), Q(x-y)\rangle \le \|Q(x-y)\|^2. \tag{15.52}$$

By combining with the previous inequality, we obtain

$$\tfrac{1}{\gamma}\|x - y\|_{Q-Q^2}^2 \le \langle\nabla\varphi_\gamma(x) - \nabla\varphi_\gamma(y), x - y\rangle \le \tfrac{1}{\gamma}\|x - y\|_Q^2. \tag{15.53}$$

Since $\lambda_{\min}(Q) = 1 - \gamma L_f$ and $\lambda_{\max}(Q) = 1 - \gamma\mu_f$, from Lemma 2 we conclude that

$$\mu_{\varphi_\gamma}\|x-y\|^2 \le \langle\nabla\varphi_\gamma(x) - \nabla\varphi_\gamma(y), x-y\rangle \le L_{\varphi_\gamma}\|x-y\|^2 \tag{15.54}$$

with μ_{φ_γ} and L_{φ_γ} as in the statement, hence the claim, cf. (15.6). □

Lemma 15.4.7 *Suppose that φ has the quadratic growth with constants (μ, ν), and let $\varphi_\star := \min\varphi$. Then, for all $\gamma \in (0, 1/L_f]$ and $x \in \operatorname{lev}_{\le\varphi_\star+\nu}\varphi_\gamma$ it holds that*

$$\varphi_\gamma(x) - \varphi_\star \le \gamma\big[\tfrac{1}{2} + (1 + 2/\gamma\mu)(1+\gamma L_f)^2\big]\|R_\gamma(x)\|^2. \tag{15.55}$$

Proof Fix $x \in \operatorname{lev}_{\le\varphi_\star+\nu}\varphi_\gamma$ and let $\bar{x} := T_\gamma(x)$. We have

$$\begin{aligned}
\varphi_\gamma(x) - \varphi_\star &\overset{15.4.2(ii)}{\le} \tfrac{\gamma}{2}\|R_\gamma(x)\|^2 + \varphi(\bar{x}) - \varphi_\star \\
&\overset{15.3.3}{\le} \tfrac{\gamma}{2}\|R_\gamma(x)\|^2 + \operatorname{dist}(\bar{x}, \mathcal{X}_\star)\operatorname{dist}(0, \partial\varphi(\bar{x})) \\
&\overset{15.3.4}{\le} \big[\tfrac{\gamma}{2}\|R_\gamma(x)\| + (1+\gamma L_f)\operatorname{dist}(\bar{x}, \mathcal{X}_\star)\big]\|R_\gamma(x)\|
\end{aligned}$$

and since $\bar{x} \in \operatorname{lev}_{\le\varphi_\star+\nu}\varphi$ (cf. Theorem 15.4.2*(ii)*), from Theorem 15.3.8 we can bound the quantity $\operatorname{dist}(\bar{x}, \mathcal{X}_\star)$ in terms of the residual as

$$\le \big[\tfrac{\gamma}{2}\|R_\gamma(x)\| + (\gamma + 2/\mu)(1+\gamma L_f)^2\|R_\gamma(\bar{x})\|\big]\|R_\gamma(x)\|.$$

The proof now follows from the inequality $\|R_\gamma(\bar{x})\| \le \|R_\gamma(x)\|$, see [4, Thm. 10.12], after easy algebraic manipulations. □

15.4.2 Further Equivalence Properties

Proposition 15.4.8 (Equivalence of Level Boundedness) *For any $\gamma \in (0, 1/L_f)$, φ has bounded level sets iff φ_γ does.*

Proof Theorem 15.4.2 implies that $\operatorname{lev}_{\le\alpha}\varphi \subseteq \operatorname{lev}_{\le\alpha}\varphi_\gamma$ for all $\alpha \in \mathbb{R}$, therefore level boundedness of φ_γ implies that of φ. Conversely, suppose that φ_γ is not level bounded, and consider $(x_k)_{k\in\mathbb{N}} \subseteq \operatorname{lev}_{\le\alpha}\varphi_\gamma$ with $\|x_k\| \to \infty$. Then from Theorem 15.4.2 it follows that $\varphi(\bar{x}_k) \le \varphi_\gamma(x_k) - \frac{1}{2\gamma}\|x_k - \bar{x}_k\|^2 \le \alpha - \frac{1}{2\gamma}\|x_k - \bar{x}_k\|^2$, where $\bar{x}_k = T_\gamma(x_k)$. In particular, $(\bar{x}_k)_{k\in\mathbb{N}} \subseteq \operatorname{lev}_{\le\alpha}\varphi$. If $(\bar{x}_k)_{k\in\mathbb{N}}$ is bounded, then $\inf\varphi = -\infty$; otherwise, $\operatorname{lev}_{\le\alpha}\varphi$ contains the unbounded sequence $(\bar{x}_k)_{k\in\mathbb{N}}$. Either way, φ cannot be level bounded. □

Proposition 15.4.9 (Equivalence of Quadratic Growth) *Let $\gamma \in (0, 1/L_f)$ be fixed. Then,*

(i) if φ satisfies the quadratic growth condition with constants (μ, ν), then so does φ_γ with constants (μ', ν), where $\mu' := \frac{1-\gamma L_f}{(1+\gamma L_f)^2} \frac{\mu\gamma}{(2+\gamma\mu)^2} \mu$;
(ii) conversely, if φ_γ satisfies the quadratic growth condition, then so does φ with same constants.

Proof Since φ and φ_γ have same infimum and minimizers (cf. Theorem 15.4.3), 15.4.9*(ii)* is a straightforward consequence of the fact that $\varphi_\gamma \leq \varphi$ (cf. Theorem 15.4.2*(i)*).

Conversely, suppose that φ satisfies the quadratic growth with constants (μ, ν). Then, for all $x \in \operatorname{lev}_{\leq \varphi_\star + \nu} \varphi_\gamma$ we have that $\bar{x} := T_\gamma(x) \in \operatorname{lev}_{\leq \varphi_\star + \nu} \varphi$, therefore

$$\varphi_\gamma(x) - \varphi_\star \overset{15.4.2(ii)}{\geq} \varphi(\bar{x}) - \varphi_\star + \gamma \tfrac{1-\gamma L_f}{2} \|R_\gamma(x)\|^2 \geq \tfrac{\mu'}{2} \operatorname{dist}(x, \mathcal{X}_\star), \tag{15.56}$$

where in the last inequality we discarded the term $\varphi(\bar{x}) - \varphi_\star \geq 0$ and used Theorem 15.3.8 to lower bound $\|R_\gamma(x)\|^2$. □

Corollary 15.4.10 (Equivalence of Strong Minimality) *For all $\gamma \in (0, 1/L_f)$, a point $x_\star$ is a (locally) strong minimizer for φ iff it is a (locally) strong minimizer for φ_γ.*

Lastly, having showed that for convex functions the quadratic growth can be extended to arbitrary level sets (cf. Lemma 15.3.7), an interesting consequence of Proposition 15.4.9 is that, although φ_γ may fail to be convex, it enjoys the same property.

Corollary 15.4.11 (FBE: Globality of Quadratic Growth) *Let $\gamma \in (0, 1/L_f)$ and suppose that φ_γ satisfies the quadratic growth with constants (μ, ν). Then, for every $\nu' > \nu$ there exists $\mu' > 0$ such that φ_γ satisfies the quadratic growth with constants (μ', ν').*

15.4.3 Second-Order Properties

Although φ_γ is continuously differentiable over $\mathbb{R}^n$, it fails to be C^2 in most cases; since g is nonsmooth, its Moreau envelope g^γ is hardly ever C^2. For example, if g is real valued, then g^γ is C^2 (and $\operatorname{prox}_{\gamma g}$ is C^1) if and only if g is C^2 [33]. Therefore, we hardly ever have the luxury of assuming continuous differentiability of $\nabla\varphi_\gamma$ and we must resort to generalized notions of differentiability stemming from nonsmooth analysis. Specifically, our analysis is largely based on generalized differentiability properties of $\operatorname{prox}_{\gamma g}$ which we study next.

Theorem 15.4.12 *For all $x \in \mathbb{R}^n$, $\partial_C(\operatorname{prox}_{\gamma g})(x) \neq \emptyset$ and any $P \in \partial_C(\operatorname{prox}_{\gamma g})(x)$ is a symmetric positive semidefinite matrix that satisfies $\|P\| \leq 1$.*

Proof Nonempty-valuedness of $\partial_C(\text{prox}_{\gamma g})$ is due to Lipschitz continuity of $\text{prox}_{\gamma g}$. Moreover, since g is convex, its Moreau envelope is a convex function as well, therefore every element of $\partial_C(\nabla g^\gamma)(x)$ is a symmetric positive semidefinite matrix (see, *e.g.*, [19, §8.3.3]). Due to Theorem 15.3.2*(i)*, we have that $\text{prox}_{\gamma g}(x) = x - \gamma \nabla g^\gamma(x)$, therefore

$$\partial_C(\text{prox}_{\gamma g})(x) = \mathrm{I} - \gamma \partial_C(\nabla g^\gamma)(x). \tag{15.57}$$

The last relation holds with equality (as opposed to inclusion in the general case) due to the fact that one of the summands is continuously differentiable. Now, from (15.57) we easily infer that every element of $\partial_C(\text{prox}_{\gamma g})(x)$ is a symmetric matrix. Since $\nabla g^\gamma(x)$ is Lipschitz continuous with Lipschitz constant γ^{-1}, using [15, Prop. 2.6.2(d)], we infer that every $H \in \partial_C(\nabla g^\gamma)(x)$ satisfies $\|H\| \leq \gamma^{-1}$. Now, according to (15.57) it holds that

$$P \in \partial_C(\text{prox}_{\gamma g})(x) \quad \Leftrightarrow \quad P = I - \gamma H, \quad H \in \partial_C(\nabla g^\gamma)(x). \tag{15.58}$$

Therefore, for every $d \in \mathbb{R}^n$ and $P \in \partial_C(\text{prox}_{\gamma g})(x)$,

$$\langle d, Pd \rangle = \|d\|^2 - \gamma \langle d, Hd \rangle \geq \|d\|^2 - \gamma\gamma^{-1}\|d\|^2 = 0. \tag{15.59}$$

On the other hand, since $\text{prox}_{\gamma g}$ is Lipschitz continuous with Lipschitz constant 1, using [15, Prop. 2.6.2(d)] we obtain that $\|P\| \leq 1$ for all $P \in \partial_C(\text{prox}_{\gamma g})(x)$. □

We are now in a position to construct a generalized Hessian for φ_γ that will allow the development of Newton-like methods with fast asymptotic convergence rates. An obvious route to follow would be to assume that $\nabla\varphi_\gamma$ is semismooth and employ $\partial_C(\nabla\varphi_\gamma)$ as a generalized Hessian for φ_γ. However, this approach would require extra assumptions on f and involve complicated operations to evaluate elements of $\partial_C(\nabla\varphi_\gamma)$. On the other hand, what is really needed to devise Newton-like algorithms with fast local convergence rates is a *linear Newton approximation (LNA)*, cf. Definition 15.2.3, at some stationary point of φ_γ, which by Theorem 15.4.3*(iii)* is also a minimizer of φ, provided that $\gamma \in (0, 1/L_f)$.

The approach we follow is largely based on [72], [19, Prop. 10.4.4]. Without any additional assumptions we can define a set-valued mapping $\hat{\partial}^2\varphi_\gamma : \mathbb{R}^n \rightrightarrows \mathbb{R}^{n\times n}$ with full domain and whose elements have a simpler form than those of $\partial_C(\nabla\varphi_\gamma)$, which serves as a LNA for $\nabla\varphi_\gamma$ at any stationary point $x_\star$ provided $\text{prox}_{\gamma g}$ is semismooth at $x_\star - \gamma\nabla f(x_\star)$. We call it *approximate generalized Hessian* of φ_γ and it is given by

$$\hat{\partial}^2\varphi_\gamma(x) := \left\{\gamma^{-1}Q_\gamma(x)(\mathrm{I} - PQ_\gamma(x)) \mid P \in \partial_C(\text{prox}_{\gamma g})(x - \gamma\nabla f(x))\right\}. \tag{15.60}$$

Notice that if f is quadratic, then $\hat{\partial}^2\varphi_\gamma \equiv \partial_C \nabla\varphi_\gamma$; more generally, the key idea in the definition of $\hat{\partial}^2\varphi_\gamma$, reminiscent of the Gauss-Newton method for nonlinear least-squares problems, is to omit terms vanishing at $x_\star$ that contain third-order derivatives of f.

Proposition 15.4.13 *Let $\bar{x} \in \mathbb{R}^n$ and $\gamma > 0$ be fixed. If $\operatorname{prox}_{\gamma g}$ is (ϑ-order) semismooth at $\bar{x} - \gamma\nabla f(\bar{x})$ (and $\nabla^2 f$ is ϑ-Hölder continuous around $\bar{x}$), then*

$$\mathcal{R}_\gamma(x) := \left\{\gamma^{-1}(\mathrm{I} - PQ_\gamma(x)) \mid P \in \partial_C \operatorname{prox}_{\gamma g}(x - \gamma\nabla f(x))\right\} \tag{15.61}$$

is a (ϑ-order) LNA for R_γ at $\bar{x}$.

Proof We shall prove only the ϑ-order semismooth case, as the other one is shown by simply replacing all occurrences of $O(\|\cdot\|^{1+\vartheta})$ with $o(\|\cdot\|)$ in the proof. Let $q_\gamma = \mathrm{id} - \gamma\nabla f$ be the forward operator, so that the forward-backward operator T_γ can be expressed as $T_\gamma = \operatorname{prox}_{\gamma g} \circ q_\gamma$. With a straightforward adaptation of the proof of [19, Prop. 7.2.9] to include the ϑ-Hölderian case, it can be shown that

$$q_\gamma(x) - q_\gamma(\bar{x}) - Q_\gamma(x)(x - \bar{x}) = O(\|x - \bar{x}\|^{1+\vartheta}). \tag{15.62}$$

Moreover, since ∇f is Lipschitz continuous and thus so is q_γ, we also have

$$q_\gamma(x) - q_\gamma(\bar{x}) = O(\|x - \bar{x}\|). \tag{15.63}$$

Let $U_x \in \mathcal{R}_\gamma(x)$ be arbitrary; then, there exists $P_x \in \partial_C \operatorname{prox}_{\gamma g}(x - \gamma\nabla f(x))$ such that $U_x = \gamma^{-1}(\mathrm{I} - P_x Q_\gamma(x))(\bar{x} - x)$. We have

$$\begin{aligned}
&R_\gamma(x) + U_x(\bar{x} - x) - R_\gamma(\bar{x}) \\
&= R_\gamma(x) + \gamma^{-1}(\mathrm{I} - P_x Q_\gamma(x))(\bar{x} - x) - R_\gamma(\bar{x}) \\
&= \gamma^{-1}\big[\operatorname{prox}_{\gamma g}(q_\gamma(\bar{x})) - \operatorname{prox}_{\gamma g}(q_\gamma(x)) - P_x Q_\gamma(x)(\bar{x} - x)\big]
\end{aligned}$$

due to ϑ-order semismoothness of $\operatorname{prox}_{\gamma g}$ at $q_\gamma(\bar{x})$,

$$\begin{aligned}
&= \gamma^{-1} P_x\big[q_\gamma(\bar{x}) - q_\gamma(x) + O(\|q_\gamma(\bar{x}) - q_\gamma(x)\|^{1+\vartheta}) - Q_\gamma(x)(\bar{x} - x)\big] \\
&\overset{(15.63)}{=} \gamma^{-1} P_x\big[q_\gamma(\bar{x}) - q_\gamma(x) - Q_\gamma(\bar{x})(\bar{x} - x) + O(\|\bar{x} - x\|^{1+\vartheta})\big] \\
&\overset{(15.62)}{=} \gamma^{-1} P_x O(\|\bar{x} - x\|^{1+\vartheta}) = O(\|\bar{x} - x\|^{1+\vartheta}),
\end{aligned}$$

where in the last equality we used the fact that $\|P_x\| \leq 1$, cf. Theorem 15.4.12. □

Corollary 15.4.14 *Let $\gamma \in (0, 1/L_f)$ and $x_\star \in \mathcal{X}_\star$. If $\operatorname{prox}_{\gamma g}$ is (ϑ-order) semismooth at $x_\star - \gamma\nabla f(x_\star)$ (and $\nabla^2 f$ is locally ϑ-Hölder continuous around $x_\star$), then $\hat\partial^2\varphi_\gamma$ is a (ϑ-order) LNA for $\nabla\varphi_\gamma$ at $x_\star$.*

Proof Let $H_x \in \hat\partial^2\varphi_\gamma(x) = \left\{Q_\gamma(x)U \mid U \in \mathcal{R}_\gamma(x)\right\}$, so that $H_x = Q_\gamma(x)U_x$ for some $U_x \in \mathcal{R}_\gamma(x)$. Then,

$$\begin{aligned}\|\nabla\varphi_\gamma(x) + H_x(x_\star - x) - \nabla\varphi_\gamma(x_\star)\| &= \|Q_\gamma(x)R_\gamma(x) + Q_\gamma(x)U_x(x - x_\star))\| \\ &= \|Q_\gamma(x)[R_\gamma(x) + U_x(x - x_\star) - R_\gamma(x_\star)]\| \\ &\le \|R_\gamma(x) + U_x(x - x_\star) - R_\gamma(x_\star)\|,\end{aligned}$$

where in the equalities we used the fact that $\nabla\varphi_\gamma(x_\star) = R_\gamma(x_\star) = 0$, and in the inequality the fact that $\|Q_\gamma\| \le 1$. Since $\mathcal{R}_\gamma$ is a (ϑ-order) LNA of R_γ at $x_\star$, the last term is $o(\|x - x_\star\|)$ (resp. $O(\|x - x_\star\|^{1+\vartheta})$). □

As shown in the next result, although the FBE is in general not convex, for γ small enough every element of $\hat\partial^2\varphi_\gamma(x)$ is a (symmetric and) positive semidefinite matrix. Moreover, the eigenvalues are lower and upper bounded uniformly over all $x \in \mathbb{R}^n$.

Proposition 15.4.15 *Let $\gamma \le 1/L_f$ and $H \in \hat\partial^2\varphi_\gamma(x)$ be fixed. Then, $H \in \mathrm{S}_+(\mathbb{R}^n)$ with*

$$\lambda_{\min}(H) = \min\left\{(1 - \gamma\mu_f)\mu_f, (1 - \gamma L_f)L_f\right\} \;\; \textit{and} \;\; \lambda_{\max}(H) = \gamma^{-1}(1 - \gamma\mu_f), \tag{15.64}$$

where $\mu_f \ge 0$ is the modulus of strong convexity of f.

Proof Fix $x \in \mathbb{R}^n$ and let $Q := Q_\gamma(x)$. Any $H \in \hat\partial^2\varphi_\gamma(x)$ can be expressed as $H = \gamma^{-1}Q(\mathrm{I} - PQ)$ for some $P \in \partial_C(\operatorname{prox}_{\gamma g})(x - \gamma\nabla f(x))$. Since both Q and P are symmetric (cf. Theorem 15.4.12), it follows that so is H. Moreover, for all $d \in \mathbb{R}^n$

$$\begin{aligned}\langle Hd, d\rangle &= \gamma^{-1}\langle Qd, d\rangle - \gamma^{-1}\langle PQd, Qd\rangle \\ &\overset{15.4.12}{\ge} \gamma^{-1}\langle Qd, d\rangle - \gamma^{-1}\|Qd\|^2 \\ &= \langle(\mathrm{I} - \gamma\nabla^2 f(x))\nabla^2 f(x)d, d\rangle \\ &\overset{2}{\ge} \min\left\{(1 - \gamma\mu_f)\mu_f,\ (1 - \gamma L_f)L_f\right\}\|d\|^2.\end{aligned} \tag{15.65}$$

On the other hand, since $P \succeq 0$ (cf. Theorem 15.4.12) and thus $QPQ \succeq 0$, we can upper bound (15.65) as

$$\langle Hd, d\rangle \le \gamma^{-1}\langle Qd, d\rangle \le \|Q\|\|d\|^2 \le \gamma^{-1}(1 - \gamma\mu_f)\|d\|^2.$$

□

The next lemma links the behavior of the FBE close to a solution of (15.1) and a nonsingularity assumption on the elements of $\hat{\partial}^2\varphi_\gamma(x_\star)$. Part of the statement is similar to [19, Lem. 7.2.10]; however, here $\nabla\varphi_\gamma$ is not required to be locally Lipschitz around $x_\star$.

Lemma 15.4.16 *Let* $x_\star \in \operatorname{argmin}\varphi$ *and* $\gamma \in (0, 1/L_f)$. *If* $\operatorname{prox}_{\gamma g}$ *is semismooth at* $x_\star - \gamma\nabla f(x_\star)$, *then the following conditions are equivalent:*

(a) $x_\star$ *is a locally strong minimum for* φ *(or, equivalently, for* φ_γ*);*
(b) every element of $\hat{\partial}^2\varphi_\gamma(x_\star)$ *is nonsingular.*

In any such case, there exist $\delta, \kappa > 0$ *such that*

$$\|x - x_\star\| \leq \kappa\|R_\gamma(x)\| \quad \text{and} \quad \max\left\{\|H\|, \|H^{-1}\|\right\} \leq \kappa, \tag{15.66}$$

for any $x \in \mathrm{B}(x_\star; \delta)$ *and* $H \in \hat{\partial}^2\varphi_\gamma(x)$.

Proof Observe first that Corollary 15.4.14 ensures that $\hat{\partial}^2\varphi_\gamma$ is a LNA of $\nabla\varphi_\gamma$ at $x_\star$, thus semicontinuous and compact valued (by definition). In particular, the last claim follows from [19, Lem. 7.5.2].

- 15.4.16*(a)* $\Rightarrow$ 15.4.16*(b)* It follows from Corollary 15.4.10 that there exists $\mu, \delta > 0$ such that $\varphi_\gamma(x) - \varphi_\star \geq \frac{\mu}{2}\|x - x_\star\|^2$ for all $x \in \mathrm{B}(x_\star; \delta)$. In particular, for all $H \in \hat{\partial}^2\varphi_\gamma(x_\star)$ and $x \in \mathrm{B}(x_\star; \delta)$ we have

$$\tfrac{\mu}{2}\|x - x_\star\|^2 \leq \varphi_\gamma(x) - \varphi_\star = \tfrac{1}{2}\langle H(x - x_\star), x - x_\star\rangle + o(\|x - x_\star\|^2). \tag{15.67}$$

  Let $v_{\min}$ be a unitary eigenvector of H corresponding to the minimum eigenvalue $\lambda_{\min}(H)$. Then, for all $\varepsilon \in (-\delta, \delta)$ the point $x_\varepsilon = x_\star + \varepsilon v_{\min}$ is δ-close to $x_\star$ and thus

$$\tfrac{1}{2}\lambda_{\min}(H)\varepsilon^2 \geq \tfrac{\mu}{2}\varepsilon^2 + o(\varepsilon^2) \geq \tfrac{\mu}{4}\varepsilon^2, \tag{15.68}$$

  where the last inequality holds up to possibly restricting δ (and thus ε). The claim now follows from the arbitrarity of $H \in \hat{\partial}^2\varphi_\gamma(x_\star)$.
- 15.4.16*(a)* $\Leftarrow$ 15.4.16*(b)* Easily follows by reversing the arguments of the other implication. □

□

15.5 Forward-Backward Truncated-Newton Algorithm (FBTN)

Having established the equivalence between minimizing φ and φ_γ, we may recast problem (15.1) into the smooth unconstrained minimization of the FBE. Under some assumptions the elements of $\hat{\partial}^2\varphi_\gamma$ mimic second-order derivatives of φ_γ, suggesting

Algorithm 15.1 (FBTN) Forward-Backward Truncated-Newton method

REQUIRE $\gamma \in (0, 1/L_f)$; $\sigma \in (0, \frac{\gamma(1-\gamma L_f)}{2})$; $\bar{\eta}, \zeta \in (0, 1)$; $\rho, \nu \in (0, 1]$
initial point $x_0 \in \mathbb{R}^n$; accuracy $\varepsilon > 0$

PROVIDE ε-suboptimal solution x^k (*i.e.*, such that $\|R_\gamma(x^k)\| \leq \varepsilon$)

INITIALIZE $k \leftarrow 0$

1: **while** $\|R_\gamma(x^k)\| > \varepsilon$ **do**

2: $\delta_k \leftarrow \zeta \|\nabla\varphi_\gamma(x^k)\|^\nu$, $\eta_k \leftarrow \min\{\bar{\eta}, \|\nabla\varphi_\gamma(x^k)\|^\rho\}$, $\varepsilon_k \leftarrow \eta_k \|\nabla\varphi_\gamma(x^k)\|$

3: Apply CG(Alg. 15.2) to find an ε_k-approximate solution d^k to

$$\big[H_k + \delta_k \mathrm{I}\big] d^k \approx -\nabla\varphi_\gamma(x^k) \tag{15.69}$$

for some $H_k \in \hat{\partial}^2\varphi_\gamma(x^k)$

4: Let τ_k be the maximum in $\{2^{-i} \mid i \in \mathbb{N}\}$ such that

$$\varphi_\gamma(x^{k+1}) \leq \varphi_\gamma(x^k) - \sigma\|R_\gamma(x^k)\|^2 \tag{15.70}$$

where $x^{k+1} \leftarrow (1-\tau_k)T_\gamma(x^k) + \tau_k\big[x^k + d^k\big]$

5: $k \leftarrow k+1$ and go to step 1

6: **end while**

the employment of Newton-like update directions $d = -(H + \delta\mathrm{I})^{-1}\nabla\varphi_\gamma(x)$ with $H \in \hat{\partial}^2\varphi_\gamma(x)$ and $\delta > 0$ (the regularization term $\delta\mathrm{I}$ ensures the well definedness of d, as H is positive semidefinite, see Proposition 15.4.15). If δ and ε are suitably selected, under some nondegeneracy assumptions updates $x^+ = x + d$ are locally superlinearly convergent. Since such d's are directions of descent for φ_γ, a possible globalization strategy is an Armijo-type linesearch. Here, however, we follow the simpler approach proposed in [71, 75] that exploits the basic properties of the FBE investigated in Section 15.4.1. As we will discuss shortly after, this is also advantageous from a computational point of view, as it allows an arbitrary warm starting for solving the underlying linear system.

Let us elaborate on the linesearch. To this end, let x be the current iterate; then, Theorem 15.4.2 ensures that $\varphi_\gamma(T_\gamma(x)) \leq \varphi_\gamma(x) - \gamma\frac{1-\gamma L_f}{2}\|R_\gamma(x)\|^2$. Therefore, unless $R_\gamma(x) = 0$, in which case x would be a solution, for any $\sigma \in (0, \gamma\frac{1-\gamma L_f}{2})$ the strict inequality $\varphi_\gamma(T_\gamma(x)) < \varphi_\gamma(x) - \sigma\|R_\gamma(x)\|^2$ is satisfied. Due to the continuity of φ_γ, all points sufficiently close to $T_\gamma(x)$ will also satisfy the inequality, thus so will the point $x^+ = (1-\tau)T_\gamma(x) + \tau(x + d)$ for small enough stepsizes τ. This fact can be used to enforce the iterates to *sufficiently* decrease the value of the FBE, cf. (15.70), which straightforwardly implies optimality of all accumulation points of the generated sequence. We defer the details to the proof of Theorem 15.5.1. In Theorems 15.5.4 and 15.5.5 we will provide conditions ensuring acceptance of unit stepsizes so that the scheme reduces to a regularized version of the (undamped) linear Newton method [19, Alg. 7.5.14] for solving $\nabla\varphi_\gamma(x) = 0$, which, under due assumptions, converges superlinearly.

In order to ease the computation of d^k, we allow for inexact solutions of the linear system by introducing a tolerance $\varepsilon_k > 0$ and requiring $\|(H_k+\delta_k \mathrm{I})d^k+\nabla\varphi_\gamma(x^k)\| \leq \varepsilon_k$. Since $H_k + \delta_k\mathrm{I}$ is positive definite, inexact solutions of the linear system can be efficiently retrieved by means of CG(Alg. 15.2) , which only requires matrix-vector products and thus only (generalized) directional derivatives, namely, (generalized) derivatives (denoted as $\frac{\partial}{\partial\lambda}$) of the single-variable functions $t \mapsto \operatorname{prox}_{\gamma g}(x + t\lambda)$ and $t \mapsto \nabla f(x + t\lambda)$, as opposed to computing the full (generalized) Hessian matrix. To further enhance computational efficiency, we may warm start the CG method with the previously computed direction, as eventually subsequent update directions are expected to have a small difference. Notice that this warm starting does not ensure that the provided (inexact) solution d^k is a direction of descent for φ_γ; either way, this property is not required by the adopted linesearch, showing a considerable advantage over classical Armijo-type rules. Putting all these facts together we obtain the proposed FBE-based truncated-Newton algorithm FBTN(Alg. 15.1) for convex composite minimization.

Remark 15.1 (Adaptive Variant When L_f Is Unknown) In practice, no prior knowledge of the global Lipschitz constant L_f is required for FBTN. In fact, replacing L_f with an initial estimate $L > 0$, the following instruction can be added at the beginning of each iteration, before step 1:

0: $\bar{x}^k \leftarrow T_\gamma(x^k)$
while $f(\bar{x}^k) > f(x^k) + \langle \nabla f(x^k), \bar{x}^k - x^k\rangle + \frac{L}{2}\|\bar{x}^k - x^k\|^2$ **do**
$\quad\gamma \leftarrow \gamma/2,\ L \leftarrow 2L,\ \bar{x}^k \leftarrow T_\gamma(x^k)$

Algorithm 15.2 (CG) Conjugate Gradient for computing the update direction

---

REQUIRE $\nabla\varphi_\gamma(x^k)$; δ_k; ε_k; d^{k-1} (set to 0 if $k = 0$)
(generalized) directional derivatives $\lambda \mapsto \frac{\partial^{\cdot 0}\operatorname{prox}_{\gamma g}}{\partial\lambda}(x^k - \gamma\nabla f(x^k))$ and
$\lambda \mapsto \frac{\partial^{\cdot 0}\nabla f}{\partial\lambda}(x^k)$
PROVIDE update direction d^k
INITIALIZE $e, p \leftarrow -\nabla\varphi_\gamma(x^k)$; warm start $d^k \leftarrow d^{k-1}$
while $\|e\| > \varepsilon_k$ **do**
$\quad u \leftarrow \frac{\partial\nabla f}{\partial p}(x^k)$
$\quad v \leftarrow p - \gamma u$ ▷ $v = Q_\gamma(x^k)p$
$\quad w \leftarrow p - \frac{\partial\operatorname{prox}_{\gamma g}}{\partial v}(x^k - \gamma\nabla f(x^k))$
$\quad z \leftarrow \delta_k p + w - \gamma\frac{\partial\nabla f}{\partial w}(x^k)$ ▷ $z = H_k p$
$\quad \alpha \leftarrow \|e\|^2/\langle p, z\rangle$
$\quad d^k \leftarrow d^k + \alpha p,\ e^+ \leftarrow e - \alpha z$
$\quad p \leftarrow e^+ + \left(\|e^+\|/\|e\|\right)^2 p$
$\quad e \leftarrow e^+$
end while

Moreover, since positive definiteness of $H_k + \delta_k \mathrm{I}$ is ensured only for $\gamma \leq {}^1\!/L_f$ where L_f is the true Lipschitz constant of $\nabla\varphi_\gamma$ (cf. Proposition 15.4.15), special care should be taken when applying CG in order to find the update direction d^k. Specifically, CG should be stopped prematurely whenever $\langle p, z\rangle \leq 0$ in step 6, in which case $\gamma \leftarrow {}^\gamma\!/2$, $L \leftarrow 2L$ and the iteration should start again from step 1.

Whenever the quadratic bound (15.5) is violated with L in place of L_f, the estimated Lipschitz constant L is increased, γ is decreased accordingly, and the proximal gradient point $\bar{x}^k$ with the new stepsize γ is evaluated. Since replacing L_f with any $L \geq L_f$ still satisfies (15.5), it follows that L is incremented only a finite number of times. Therefore, there exists an iteration k_0 starting from which γ and L are constant; in particular, all the convergence results here presented remain valid starting from iteration k_0, at latest. Moreover, notice that this step does not increase the complexity of the algorithm, since both $\bar{x}^k$ and $\nabla f(x^k)$ are needed for the evaluation of $\varphi_\gamma(x^k)$.

15.5.1 Subsequential and Linear Convergence

Before going through the convergence proofs let us spend a few lines to emphasize that FBTN is a well-defined scheme. First, that a matrix H_k as in line 1 exists is due to the nonemptyness of $\hat{\partial}^2\varphi_\gamma(x^k)$ (cf. Section 15.4.3). Second, since $\delta_k > 0$ and $H_k \succeq 0$ (cf. Proposition 15.4.15) it follows that $H_k + \delta_k\mathrm{I}$ is (symmetric and) positive definite, and thus CG is indeed applicable at line 3.

Having clarified this, the proof of the next result falls as a simplified version of [75, Lem. 5.1 and Thm. 5.6]; we elaborate on the details for the sake of self-inclusiveness. To rule out trivialities, in the rest of the chapter we consider the limiting case of infinite accuracy, that is $\varepsilon = 0$, and assume that the termination criterion $\|R_\gamma(x^k)\| = 0$ is never met. We shall also work under the assumption that a solution to the investigated problem (15.1) exists, thus in particular that the cost function φ is lower bounded.

Theorem 15.5.1 (Subsequential Convergence) *Every accumulation point of the sequence $(x^k)_{k\in\mathbb{N}}$ generated by FBTN(Alg. 15.1) is optimal.*

Proof Observe that

$$\varphi_\gamma\bigl(x^k - \gamma R_\gamma(x^k)\bigr) \overset{15.4.2}{\leq} \varphi_\gamma(x^k) - \gamma\tfrac{1-\gamma L_f}{2}\|R_\gamma(x^k)\|^2 < \varphi_\gamma(x^k) - \sigma\|R_\gamma(x^k)\|^2 \tag{15.71}$$

and that $x^{k+1} \to T_\gamma(x^k)$ as $\tau_k \to 0$. Continuity of φ_γ ensures that for small enough τ_k the linesearch condition (15.70) is satisfied, in fact, regardless of what d^k is. Therefore, for each k the stepsize τ_k is decreased only a finite number of times. By telescoping the linesearch inequality (15.70) we obtain

$$\sigma\sum_{k\in\mathbb{N}}\|R_\gamma(x^k)\|^2 \leq \sum_{k\in\mathbb{N}}\bigl[\varphi_\gamma(x^k) - \varphi_\gamma(x^{k+1})\bigr] \leq \varphi_\gamma(x^0) - \varphi_\star < \infty \tag{15.72}$$

and in particular $R_\gamma(x^k) \to 0$. Since R_γ is continuous we infer that every accumulation point $x_\star$ of $(x^k)_{k\in\mathbb{N}}$ satisfies $R_\gamma(x_\star) = 0$, hence $x_\star \in \operatorname{argmin}\varphi$, cf. (15.24). □

Remark 15.2 Since FBTN is a descent method on φ_γ, as ensured by the linesearch condition (15.70), from Proposition 15.4.8 it follows that a sufficient condition for the existence of cluster points is having φ with bounded level sets or, equivalently, having $\operatorname{argmin}\varphi$ bounded (cf. Lemma 1).

As a straightforward consequence of Lemma 15.4.7, from the linesearch condition (15.70) we infer Q-linear decrease of the FBE along the iterates of FBTN provided that the original function φ has the quadratic growth property. In particular, although the quadratic growth is a local property, Q-linear convergence holds globally, as described in the following result.

Theorem 15.5.2 (Q-Linear Convergence of FBTN Under Quadratic Growth) *Suppose that φ satisfies the quadratic growth with constants (μ, ν). Then, the iterates of FBTN(Alg. 15.1) decrease Q-linearly the value of φ_γ as*

$$\varphi_\gamma(x^{k+1}) - \varphi_\star \le \left(1 - \tfrac{2\sigma\mu'}{\gamma\mu + 2(2+\gamma\mu')(1+\gamma L_f)^2}\right)(\varphi_\gamma(x^k) - \varphi_\star) \quad \forall k \in \mathbb{N}, \tag{15.73}$$

where

$$\mu' := \begin{cases} \mu & \text{if } \varphi_\gamma(x_0) \le \varphi_\star + \nu, \\ \frac{\mu}{2}\min\left\{1, \frac{\nu}{\varphi_\gamma(x_0) - \varphi_\star - \nu}\right\} & \text{otherwise.} \end{cases} \tag{15.74}$$

Proof Since FBTN is a descent method on φ_γ, it holds that $(x^k)_{k\in\mathbb{N}} \subseteq \operatorname{lev}_{\le\alpha}\varphi_\gamma$ with $\alpha = \varphi_\gamma(x^0)$. It follows from Lemma 15.3.7 that φ satisfies the quadratic growth condition with constants $(\mu', \varphi_\gamma(x^0))$, with μ' is as in the statement. The claim now follows from the inequality ensured by linesearch condition (15.70) combined with Lemma 15.4.7. □

15.5.2 Superlinear Convergence

In this section we provide sufficient conditions that enable superlinear convergence of FBTN. In the sequel, we will make use of the notion of *superlinear directions* that we define next.

Definition 15.5.3 (Superlinear Directions) Suppose that $\mathcal{X}_\star \ne \emptyset$ and consider the iterates generated by FBTN(Alg. 15.1). We say that $(d^k)_{k\in\mathbb{N}} \subset \mathbb{R}^n$ are *superlinearly convergent directions* if

$$\lim_{k\to\infty} \frac{\operatorname{dist}(x^k + d^k, \mathcal{X}_\star)}{\operatorname{dist}(x^k, \mathcal{X}_\star)} = 0.$$

If for some $q > 1$ the condition can be strengthened to

$$\limsup_{k\to\infty} \frac{\operatorname{dist}(x^k + d^k, \mathcal{X}_\star)}{\operatorname{dist}(x^k, \mathcal{X}_\star)^q} < \infty$$

then we say that $(d^k)_{k\in\mathbb{N}}$ are *superlinearly convergent directions with order* q.

We remark that our definition of superlinear directions extends the one given in [19, §7.5] to cases in which $\mathcal{X}_\star$ is not a singleton. The next result constitutes a key component of the proposed methodology, as it shows that the proposed algorithm does not suffer from the *Maratos' effect* [44], a well-known obstacle for fast local methods that inhibits the acceptance of the unit stepsize. On the contrary, we will show that whenever the directions $(d^k)_{k\in\mathbb{N}}$ computed in FBTN are superlinear, then indeed the unit stepsize is eventually always accepted, and the algorithm reduces to a regularized version of the (undamped) linear Newton method [19, Alg. 7.5.14] for solving $\nabla\varphi_\gamma(x) = 0$ or, equivalently, $R_\gamma(x) = 0$, and $\operatorname{dist}(x^k, \mathcal{X}_\star)$ converges superlinearly.

Theorem 15.5.4 (Acceptance of the Unit Stepsize and Superlinear Convergence) *Consider the iterates generated by FBTN(Alg. 15.1). Suppose that φ satisfies the quadratic growth (locally) and that $(d^k)_{k\in\mathbb{N}}$ are superlinearly convergent directions (with order q). Then, there exists $\bar{k} \in \mathbb{N}$ such that*

$$\varphi_\gamma(x^k + d^k) \le \varphi_\gamma(x^k) - \sigma\|R_\gamma(x^k)\|^2 \qquad \forall k \ge \bar{k}. \tag{15.75}$$

In particular, eventually the iterates reduce to $x^{k+1} = x^k + d^k$, and $\operatorname{dist}(x^k, \mathcal{X}_\star)$ *converges superlinearly (with order q).*

Proof Without loss of generality we may assume that $(x^k)_{k\in\mathbb{N}}$ and $(x^k + d^k)_{k\in\mathbb{N}}$ belong to a region in which quadratic growth holds. Denoting $\varphi_\star := \min\varphi$, since φ_γ also satisfies the quadratic growth (cf. Proposition 15.4.9*(i)*) it follows that

$$\varphi_\gamma(x^k) - \varphi_\star \ge \tfrac{\mu'}{2}\operatorname{dist}(x^k, \mathcal{X}_\star)^2 \tag{15.76}$$

for some constant $\mu' > 0$. Moreover, we know from Lemma 15.4.7 that

$$\varphi_\gamma(x^k + d^k) - \varphi_\star \le c\|R_\gamma(x^k + d^k)\|^2 \le c'\operatorname{dist}(x^k + d^k, \mathcal{X}_\star)^2 \tag{15.77}$$

for some constants $c, c' > 0$, where in the second inequality we used Lipschitz continuity of R_γ (Lemma 3) together with the fact that $R_\gamma(x_\star) = 0$ for all points $x_\star \in \mathcal{X}_\star$. By combining the last two inequalities, we obtain

$$t_k := \frac{\varphi_\gamma(x^k + d^k) - \varphi_\star}{\varphi_\gamma(x^k) - \varphi_\star} \le \frac{2c'\operatorname{dist}(x^k + d^k, \mathcal{X}_\star)^2}{\mu'\operatorname{dist}(x^k, \mathcal{X}_\star)^2} \to 0 \quad \text{as } k \to \infty. \tag{15.78}$$

Moreover,

$$\varphi_\gamma(x^k) - \varphi_\star \geq \varphi_\gamma(x^k) - \varphi(T_\gamma(x^k)) \overset{15.4.2(ii)}{\geq} \gamma\frac{1-\gamma L_f}{2}\|R_\gamma(x^k)\|^2. \tag{15.79}$$

Thus,

$$\begin{aligned}\varphi_\gamma(x^k + d^k) - \varphi_\gamma(x^k) &= \big[\varphi_\gamma(x^k + d^k) - \varphi_\star\big] - \big[\varphi_\gamma(x^k) - \varphi_\star\big] \\ &= (t_k - 1)\big[\varphi_\gamma(x^k) - \varphi_\star\big]\end{aligned}$$

and since $t_k \to 0$, eventually it holds that $t_k \leq 1 - \frac{2\sigma}{\gamma(1-\gamma L_f)} \in (0, 1)$, resulting in

$$\leq -\sigma\|R_\gamma(x^k)\|^2.$$

□

Theorem 15.5.5 *Consider the iterates generated by FBTN(Alg. 15.1). Suppose that φ satisfies the quadratic growth (locally), and let $x_\star$ be the limit point of $(x^k)_{k\in\mathbb{N}}$.*[2] *Then, $(d^k)_{k\in\mathbb{N}}$ are superlinearly convergent directions provided that*

(i) either R_γ is strictly differentiable at $x_\star$[3] *and there exists $D > 0$ such that $\|d^k\| \leq D\|\nabla\varphi_\gamma(x^k)\|$ for all k's,*

(ii) or $\mathcal{X}_\star = \{x_\star\}$ and $\operatorname{prox}_{\gamma g}$ is semismooth at $x_\star - \gamma\nabla f(x_\star)$. In this case, if $\operatorname{prox}_{\gamma g}$ is ϑ-order semismooth at $x_\star - \gamma\nabla f(x_\star)$ and $\nabla^2 f$ is ϑ-Hölder continuous close to $x_\star$, then the order of superlinear convergence is at least $1 + \min\{\rho, \vartheta, \nu\}$.

Proof Due to Proposition 15.4.9 and Theorem 15.5.2, if $\mathcal{X}_\star = \{x_\star\}$ then the sequence $(x^k)_{k\in\mathbb{N}}$ converges to $x_\star$. Otherwise, the hypothesis ensure that

$$\|x^{k+1} - x^k\| = \tau_k\|d^k\| \leq D\|\nabla\varphi_\gamma(x^k)\| \leq D\|R_\gamma(x^k)\|, \tag{15.80}$$

from which we infer that $(\|x^{k+1} - x^k\|)_{k\in\mathbb{N}}$ is R-linearly convergent, hence that $(x^k)_{k\in\mathbb{N}}$ is a Cauchy sequence, and again we conclude that the limit point $x_\star$ indeed exists. Moreover, in light of Proposition 15.4.9 we have that $(x^k)_{k\in\mathbb{N}}$ is contained in a level set of φ_γ where φ_γ has quadratic growth. To establish a notation, let $e^k := [H_k + \delta_k \mathrm{I}]d^k + \nabla\varphi_\gamma(x^k)$ be the error in solving the linear system at line 3, so that

$$\|e^k\| \leq \varepsilon_k \leq \|\nabla\varphi_\gamma(x^k)\|^{1+\rho}, \tag{15.81}$$

(cf. line 1), and let $H_k = Q_\gamma(x^k)U_k$ for some $U_k \in \mathcal{R}_\gamma(x^k)$, see (15.61). Let us now analyze the two cases separately.

[2]As detailed in the proof, under the assumptions the limit point indeed exists.

[3]From the chain rule of differentiation it follows that R_γ is strictly differentiable at $x_\star$ if $\operatorname{prox}_{\gamma g}$ is strictly differentiable at $x_\star - \gamma\nabla f(x_\star)$ (strict differentiability is closed under composition).

- 15.5.5*(i)* Let $x_\star^k := \mathscr{P}_{\mathcal{X}_\star} x^k$, so that $\operatorname{dist}(x^k, \mathcal{X}_\star) = \|x^k - x_\star^k\|$. Recall that $\nabla\varphi_\gamma = Q_\gamma R_\gamma$ and that $(1-\gamma L_f)\mathrm{I} \preceq Q_\gamma \preceq \mathrm{I}$. Since $R_\gamma(x_\star^k) = 0$, from Lemma 3 and Theorem 15.3.8 we infer that there exist $r_1, r_2 > 0$ such that

$$\|R_\gamma(x^k)\| \geq r_1 \operatorname{dist}(x^k, \mathcal{X}_\star) \quad \text{and} \quad \|\nabla\varphi_\gamma(x^k)\| \leq r_2 \operatorname{dist}(x^k, \mathcal{X}_\star). \tag{15.82}$$

In particular, the assumption on d^k ensures that $\|d^k\| = O\big(\operatorname{dist}(x^k, \mathcal{X}_\star)\big)$. We have

$$\begin{aligned} r_1 \operatorname{dist}(x^k + d^k, \mathcal{X}_\star) &\overset{(15.82)}{\leq} \|R_\gamma(x^k + d^k)\| \\ &\leq \underbrace{\|R_\gamma(x^k + d^k) - R_\gamma(x^k) - U_k d^k\|}_{\text{(a)}} + \underbrace{\|R_\gamma(x^k) + U_k d^k\|}_{\text{(b)}}. \end{aligned}$$

As to quantity (a), we have

$$\begin{aligned} \text{(a)} &\leq \|R_\gamma(x^k + d^k) - R_\gamma(x^k) - JR_\gamma(x_\star)d^k\| + \|U_k - JR(x_\star)\|\|d^k\| \\ &= o\big(\operatorname{dist}(x^k, \mathcal{X}_\star)\big), \end{aligned}$$

where we used strict differentiability and the fact that $\partial_C R_\gamma(x_\star) = \big\{JR_\gamma(x_\star)\big\}$ [15, Prop. 2.2.4] which implies $U_k \to JR(x_\star)$. In order to bound (b), recall that $\delta_k = \zeta\|\nabla\varphi_\gamma(x^k)\|^\nu$ (cf. line 1). Then,

$$\begin{aligned} \text{(b)} &= \|Q_\gamma(x^k)^{-1}(e^k - \delta_k d^k)\| \\ &\overset{(15.81)}{\leq} \tfrac{1}{1-\gamma L_f}\|\nabla\varphi_\gamma(x^k)\|\left(\|\nabla\varphi_\gamma(x^k)\|^\rho + \zeta\|\nabla\varphi_\gamma(x^k)\|^{\nu-1}\|d^k\|\right). \\ &\overset{(15.82)}{\leq} \tfrac{r_2}{1-\gamma L_f}\operatorname{dist}(x^k, \mathcal{X}_\star)\left(r_2^\rho \operatorname{dist}(x^k, \mathcal{X}_\star)^\rho + \zeta\|\nabla\varphi_\gamma(x^k)\|^{\nu-1}\|d^k\|\right) \\ &= O\big(\operatorname{dist}(x^k, \mathcal{X}_\star)^{1+\min\{\rho,\nu\}}\big), \end{aligned}$$

and we conclude that $\operatorname{dist}(x^k + d^k, \mathcal{X}_\star) \leq \text{(a)} + \text{(b)} \leq o\big(\operatorname{dist}(x^k, \mathcal{X}_\star)\big)$.

- 15.5.5*(ii)* In this case $\operatorname{dist}(x^k, \mathcal{X}_\star) = \|x^k - x_\star\|$ and the assumption of (ϑ-order) semismoothness ensures through Proposition 15.4.13 that $\mathcal{R}_\gamma$ is a (ϑ-order) LNA for R_γ at $x_\star$. Moreover, due to Lemma 15.4.16 there exists $c > 0$ such that $\|[H_k + \delta_k \mathrm{I}]^{-1}\| \leq c$ for all k's. We have

$$\begin{aligned} \|x^k + d^k - x_\star\| &= \big\|x^k + [H_k + \delta_k\mathrm{I}]^{-1}(e^k - \nabla\varphi_\gamma(x^k)) - x_\star\big\| \\ &\leq \big\|[H_k + \delta_k\mathrm{I}]^{-1}\big\|\,\big\|[H_k + \delta_k\mathrm{I}](x^k - x_\star) + e^k - \nabla\varphi_\gamma(x^k)\big\| \\ &\leq c\big\|H_k(x^k - x_\star) - \nabla\varphi_\gamma(x^k)\big\| + c\delta_k\|x^k - x_\star\| + c\|e^k\| \\ &= c\big\|Q_\gamma(x^k)\underbrace{\big(U_k(x^k - x_\star) - R_\gamma(x^k)\big)}\big\| + c\delta_k\|x^k - x_\star\| + c\|e^k\|. \end{aligned}$$

Since $\mathcal{R}_\gamma$ is a LNA at $x_\star$, it follows that the quantity emphasized in the bracket is a $o(\|x^k - x_\star\|)$, whereas in case of a (ϑ-order) LNA the tighter estimate $O(\|x^k - x_\star\|^{1+\vartheta})$ holds. Combined with the fact that $\delta_k = O(\|x^k - x_\star\|^\nu)$ and $\|e^k\| = O(\|x^k - x_\star\|^{1+\rho})$, we conclude that $(d^k)_{k\in\mathbb{N}}$ are superlinearly convergent directions, and with order at least $1 + \min\{\rho, \vartheta, \nu\}$ in case of ϑ-order semismoothness. □

□

Problems where the residual is (ϑ-order) semismooth are quite common. For instance, piecewise affine functions are everywhere strongly semismooth, as it is the case for the residual in lasso problems [67]. On the contrary, when the solution is not unique the condition $\|d^k\| \leq D\|\nabla\varphi_\gamma(x^k)\|$ (or, equivalently, $\|d^k\| \leq D'\|R_\gamma(x^k)\|$) is trickier. As detailed in [82, 83], this bound on the directions is ensured if $\rho = 1$ and for all iterates x^k and points x close enough to the limit point the following smoothness condition holds:

$$\|R_\gamma(x^k) + U_k(x - x^k)\| \leq c\|x - x^k\|^2 \tag{15.83}$$

for some constant $c > 0$. This condition is implied by and closely related to local Lipschitz differentiability of R_γ and thus conservative. We remark that, however, this can be weakened by requiring $\rho \geq \nu$, and a notion of ϑ-order semismoothness at the limit point with some degree of uniformity on the set of solutions $\mathcal{X}_\star$, namely

$$\limsup_{\substack{x,x'\to x_\star \\ x'\in\mathcal{X}_\star,\, x\neq x' \\ U\in\mathcal{R}_\gamma(x)}} \frac{\|R_\gamma(x) + U(x' - x)\|}{\|x' - x\|^{1+\vartheta}} < \infty \tag{15.84}$$

for some $\vartheta \in [\nu, 1]$. This weakened requirement comes from the observation that point x in (15.83) is in fact $x_\star^k$, the projection of x^k onto $\mathcal{X}_\star$, set onto which R_γ is constant (equal to 0). To see this, notice that (15.84) implies that $\|R_\gamma(x^k) + U_k(x_\star^k - x^k)\| \leq c\|x_\star^k - x^k\|^{1+\vartheta}$ for some $c > 0$. In particular, mimicking the arguments in the cited references, since $H_k \succeq 0$ and $\|Q_\gamma\| \leq 1$, observe that

$$\|[H_k + \delta_k \mathrm{I}]^{-1} Q_\gamma(x^k)\| \leq \|[H_k + \delta_k \mathrm{I}]^{-1}\| \leq \delta_k^{-1} \tag{15.85a}$$

and

$$\|[H_k + \delta_k \mathrm{I}]^{-1} H_k\| = \|\mathrm{I} - \delta_k [H_k + \delta_k \mathrm{I}]^{-1}\| \leq 2. \tag{15.85b}$$

Therefore,

$$\begin{aligned}
\|d^k\| &= \left\|[H_k + \delta_k \mathrm{I}]^{-1}\big(e^k - \nabla\varphi_\gamma(x^k)\big)\right\| \\
&\leq \left\|[H_k + \delta_k \mathrm{I}]^{-1}\right\|\|e^k\| + \left\|[H_k + \delta_k \mathrm{I}]^{-1} Q_\gamma(x^k)\big(R_\gamma(x^k) + U_k(x_\star^k - x^k)\big)\right\| \\
&\quad + \left\|[H_k + \delta_k \mathrm{I}]^{-1} H_k(x_\star^k - x^k)\right\|
\end{aligned}$$

$$
\begin{aligned}
&\overset{(15.85)}{\leq} \delta_k^{-1}\|\nabla\varphi_\gamma(x^k)\|^{1+\rho} + c\delta_k^{-1}\|x_\star^k - x^k\|^{1+\vartheta} + 2\|x_\star^k - x^k\| \\
&= \zeta^{-1}\|\nabla\varphi_\gamma(x^k)\|^{1+\rho-\nu} + c\zeta^{-1}\|x_\star^k - x^k\|^{1+\vartheta-\nu} + 2\|x_\star^k - x^k\| \\
&= O\big(\operatorname{dist}(x^k, \mathcal{X}_\star)^{1+\min\{0,\rho-\nu,\vartheta-\nu\}}\big),
\end{aligned}
$$

which is indeed $O(\|\nabla\varphi_\gamma(x^k)\|)$ whenever $\nu \leq \min\{\vartheta, \rho\}$. Some comments are in order to expand on condition (15.84).

(i) If $\mathcal{X}_\star = \{\bar{x}\}$ is a singleton, then x' is fixed to $x_\star$ and the requirement reduces to ϑ-order semismoothness at $x_\star$.
(ii) This notion of uniformity is a *local property*: for any $\varepsilon > 0$ the set $\mathcal{X}_\star$ can be replaced by $\mathcal{X}_\star \cap \mathrm{B}(x_\star; \varepsilon)$.
(iii) The condition $U \in \mathcal{R}_\gamma(x)$ in the limit can be replaced by $U \in \hat{\mathcal{R}}_\gamma(x) := \{\gamma^{-1}(\mathrm{I} - PQ_\gamma(x)) \mid P \in \partial_B \operatorname{prox}_{\gamma g}(x - \gamma\nabla f(x))\}$, since $\mathcal{R}_\gamma(x) = \operatorname{conv}\big(\hat{\mathcal{R}}_\gamma(x)\big)$.

In particular, by exploiting this last condition it can be easily verified that if R_γ is piecewise ϑ-Hölder differentiable around $x_\star$, then (15.84) holds, yet the stronger requirement (15.83) in [82, 83] does not.

15.6 Generalized Jacobians of Proximal Mappings

In many interesting cases $\operatorname{prox}_{\gamma g}$ is PC^1 and thus semismooth. *Piecewise quadratic* (PWQ) functions comprise a special but important class of convex functions whose proximal mapping is PC^1. A convex function g is called PWQ if $\operatorname{dom} g$ can be represented as the union of finitely many polyhedral sets, relative to each of which $g(x)$ is given by an expression of the form $\frac{1}{2}\langle x, Hx\rangle + \langle q, x\rangle + c$ ($H \in \mathbb{R}^{n\times n}$ must necessarily be symmetric positive semidefinite) [65, Def. 10.20]. The class of PWQ functions is quite general since it includes *e.g.*, polyhedral norms, indicators and support functions of polyhedral sets, and it is closed under addition, composition with affine mappings, conjugation, inf-convolution and inf-projection [65, Prop.s 10.22 and 11.32]. It turns out that the proximal mapping of a PWQ function is *piecewise affine* (PWA) [65, 12.30] ($\mathbb{R}^n$ is partitioned in polyhedral sets relative to each of which $\operatorname{prox}_{\gamma g}$ is an affine mapping), hence strongly semismooth [19, Prop. 7.4.7]. Another example of a proximal mapping that is strongly semismooth is the projection operator over symmetric cones [73].

A big class with semismooth proximal mapping is formed by the semi-algebraic functions. We remind that a set $A \subseteq \mathbb{R}^n$ is *semi-algebraic* if it can be expressed as

$$
A = \bigcup_{i=1}^{p}\bigcap_{j=1}^{q}\big\{x \in \mathbb{R}^m \mid P_{ij}(x) = 0,\ Q_{ij}(x) < 0\big\} \tag{15.86}
$$

for some polynomial functions $P_{ij}, Q_{ij} : \mathbb{R}^n \to \mathbb{R}$, and that a function $h : \mathbb{R}^n \to \overline{\mathbb{R}}^m$ is *semi-algebraic* if $\operatorname{gph} h$ is a semi-algebraic subset of $\mathbb{R}^{n+m}$.

Proposition 15.6.1 *If $g : \mathbb{R}^n \to \overline{\mathbb{R}}$ is semi-algebraic, then so are g^γ and $\operatorname{prox}_{\gamma g}$. In particular, g^γ and $\operatorname{prox}_{\gamma g}$ are semismooth.*

Proof Since g^γ and $\operatorname{prox}_{\gamma g}$ are both Lipschitz continuous, semismoothness will follow once we show that they are semi-algebraic [11, Rem. 4]. Every polynomial is clearly semi-algebraic, and since the property is preserved under addition [10, Prop. 2.2.6(ii)], the function $(x, w) \mapsto g(w) + \frac{1}{2\gamma}\|w - x\|^2$ is semi-algebraic. Moreover, since parametric minimization of a semi-algebraic function is still semi-algebraic (see, *e.g.*, [1, §2]), it follows that the Moreau envelope g^γ is semi-algebraic and therefore so is $h(x, w) := g(w) + \frac{1}{2\gamma}\|w - x\|^2 - g^\gamma(x)$. Notice that $\operatorname{prox}_{\gamma g}(x) = \{w \in \mathbb{R}^n \mid h(x, w) \le 0\}$, therefore

$$\begin{aligned}\operatorname{gph}\operatorname{prox}_{\gamma g} &= \left\{(x, \bar{x}) \in \mathbb{R}^n \times \mathbb{R}^n \mid \operatorname{prox}_{\gamma g}(x) = \bar{x}\right\}\\ &= \left\{(x, \bar{x}) \in \mathbb{R}^n \times \mathbb{R}^n \mid h(x, \bar{x}) \le 0\right\}\\ &= h^{-1}((-\infty, 0])\end{aligned}$$

is a semi-algebraic set, since the interval $(-\infty, 0]$ is clearly semi-algebraic and thus so is $h^{-1}((-\infty, 0])$ [10, Prop. 2.2.7]. □

In fact, with the same arguments it can be shown that the result still holds if "semi-algebraic" is replaced with the broader notion of "*tame*", see [11]. Other conditions that guarantee semismoothness of the proximal mapping can be found in [46–48, 50]. The rest of the section is devoted to collecting explicit formulas of $\partial_C \operatorname{prox}_{\gamma g}$ for many known useful instances of convex functions g.

15.6.1 Properties

a. Separable functions.
Whenever g is (block) separable, *i.e.*, $g(x) = \sum_{i=1}^N g_i(x_i)$, $x_i \in \mathbb{R}^{n_i}$, $\sum_{i=1}^N n_i = n$, then every $P \in \partial_C(\operatorname{prox}_{\gamma g})(x)$ is a (block) diagonal matrix. This has favorable computational implications especially for large-scale problems. For example, if g is the ℓ_1 norm or the indicator function of a box, then the elements of $\partial_C \operatorname{prox}_{\gamma g}(x)$ (or $\partial_B \operatorname{prox}_{\gamma g}(x)$) are diagonal matrices with diagonal elements in $[0, 1]$ (or in $\{0, 1\}$).

b. Convex conjugate.
With a simple application of the Moreau's decomposition [2, Thm. 14.3(ii)], all elements of $\partial_C \operatorname{prox}_{\gamma g^*}$ are readily available as long as one can compute $\partial_C \operatorname{prox}_{g/\gamma}$. Specifically,

$$\partial_C(\operatorname{prox}_{\gamma g^*})(x) = \mathrm{I} - \partial_C(\operatorname{prox}_{g/\gamma})(x/\gamma). \tag{15.87}$$

c. Support function.
The *support function* of a nonempty closed and convex set D is the proper convex and lsc function $\sigma_D(x) := \sup_{y \in D}\langle x, y\rangle$. Alternatively, σ_D can be expressed as the convex conjugate of the indicator function δ_D, and one can use the results of Section §15.6.1b to find that

$$\partial_C(\mathrm{prox}_{\gamma g})(x) = \mathrm{I} - \partial_C(\mathscr{P}_D)(x/\gamma). \tag{15.88}$$

Section 15.6.2 offers a rich list of sets D for which a closed form expression exists.

d. Spectral functions.
The eigenvalue function $\lambda : \mathrm{S}(\mathbb{R}^{n\times n}) \to \mathbb{R}^n$ returns the vector of eigenvalues of a symmetric matrix in nonincreasing order. *Spectral functions* are of the form

$$G := h \circ \lambda : \mathrm{S}(\mathbb{R}^{n\times n}) \to \overline{\mathbb{R}}. \tag{15.89}$$

where $h : \mathbb{R}^n \to \overline{\mathbb{R}}$ is proper, lsc, convex, and *symmetric*, *i.e.*, invariant under coordinate permutations [35]. Such G inherits most of the properties of h [36, 37]; in particular, its proximal mapping is [56, §6.7]

$$\mathrm{prox}_{\gamma G}(X) = Q\,\mathrm{diag}(\mathrm{prox}_{\gamma h}(\lambda(X)))\,Q^\top, \tag{15.90}$$

where $X = Q\,\mathrm{diag}(\lambda(X))Q^\top$ is the spectral decomposition of X (Q is an orthogonal matrix). If, additionally,

$$h(x) = g(x_1) + \cdots + g(x_N) \tag{15.91}$$

for some $g : \mathbb{R} \to \overline{\mathbb{R}}$, then

$$\mathrm{prox}_{\gamma h}(x) = (\mathrm{prox}_{\gamma g}(x_1), \ldots, \mathrm{prox}_{\gamma g}(x_N)), \tag{15.92}$$

and therefore the proximal mapping of G can be expressed as

$$\mathrm{prox}_{\gamma G}(X) = Q\,\mathrm{diag}(\mathrm{prox}_{\gamma g}(\lambda_1(X)), \ldots, \mathrm{prox}_{\gamma g}(\lambda_n(X)))\,Q^\top, \tag{15.93}$$

[9, Chap. V], [29, Sec. 6.2]. Now we can use the theory of nonsmooth symmetric matrix-valued functions developed in [14] to analyze differentiability properties of $\mathrm{prox}_{\gamma G}$. In particular, $\mathrm{prox}_{\gamma G}$ is (strongly) semismooth at X iff $\mathrm{prox}_{\gamma g}$ is (strongly) semismooth at the eigenvalues of X [14, Prop. 4.10]. Moreover, for any $X \in \mathrm{S}(\mathbb{R}^{n\times n})$ and $P \in \partial_B(\mathrm{prox}_{\gamma G})(X)$ we have [14, Lem. 4.7]

$$P(X) = Q\left(\Omega^{\gamma g}_{\lambda,\lambda} \odot (Q^\top X Q)\right)Q^\top, \tag{15.94}$$

where $\odot$ denotes the Hadamard product and for vectors $u, v \in \mathbb{R}^n$ we defined $\Omega^{\gamma g}_{u,v}$ as the $n \times n$ matrix

$$(\Omega^{\gamma g}_{u,v})_{ij} := \begin{cases} \partial_B \operatorname{prox}_{\gamma g}(u_i) & \text{if } u_i = v_j, \\ \left\{ \frac{\operatorname{prox}_{\gamma g}(u_i) - \operatorname{prox}_{\gamma g}(v_j)}{u_i - v_j} \right\} & \text{otherwise.} \end{cases} \tag{15.95}$$

e. Orthogonally invariant functions. A function $G : \mathbb{R}^{m\times n} \to \overline{\mathbb{R}}$ is called *orthogonally invariant* if $G(UXV^\top) = G(X)$ for all $X \in \mathbb{R}^{m\times n}$ and orthogonal matrices $U \in \mathbb{R}^{m\times m}$, $V \in \mathbb{R}^{n\times n}$.[4]

A function $h : \mathbb{R}^q \to \overline{\mathbb{R}}$ is *absolutely symmetric* if $h(Qx) = h(x)$ for all $x \in \mathbb{R}^q$ and any generalized permutation matrix Q, *i.e.*, a matrix $Q \in \mathbb{R}^{q\times q}$ that has exactly one nonzero entry in each row and each column, that entry being ± 1 [34]. There is a one-to-one correspondence between orthogonally invariant functions on $\mathbb{R}^{m\times n}$ and absolutely symmetric functions on $\mathbb{R}^q$. Specifically, if G is orthogonally invariant, then

$$G(X) = h(\sigma(X)) \tag{15.96}$$

for the absolutely symmetric function $h(x) = G(\operatorname{diag}(x))$. Here, for $X \in \mathbb{R}^{m\times n}$ and $q := \min\{m, n\}$ the spectral function $\sigma : \mathbb{R}^{m\times n} \to \mathbb{R}^q$ returns the vector of its singular values in nonincreasing order. Conversely, if h is absolutely symmetric, then $G(X) = h(\sigma(X))$ is orthogonally invariant. Therefore, convex analytic and generalized differentiability properties of orthogonally invariant functions can be easily derived from those of the corresponding absolutely symmetric functions [34]. For example, assuming for simplicity that $m \le n$, the proximal mapping of G is given by [56, Sec. 6.7]

$$\operatorname{prox}_{\gamma G}(X) = U \operatorname{diag}(\operatorname{prox}_{\gamma h}(\sigma(X))) V_1^\top, \tag{15.97}$$

where $X = U \left[\operatorname{diag}(\sigma(X))\ 0\right] \left[V_1\ V_2\right]^\top$ is the singular value decomposition of X. If we further assume that h has a separable form as in (15.91), then

$$\operatorname{prox}_{\gamma G}(X) = U \Sigma_g(X) V_1^\top, \tag{15.98}$$

where $\Sigma_g(X) = \operatorname{diag}(\operatorname{prox}_{\gamma g}(\sigma_1(X)), \ldots, \operatorname{prox}_{\gamma g}(\sigma_n(X)))$. Functions of this form are called *nonsymmetric matrix-valued functions*. We also assume that g is a non-negative function such that $g(0) = 0$. This implies that $\operatorname{prox}_{\gamma g}(0) = 0$ and guarantees that the nonsymmetric matrix-valued function (15.98) is well defined [80, Prop. 2.1.1]. Now we can use the results of [80, §2] to draw conclusions about generalized differentiability properties of $\operatorname{prox}_{\gamma G}$.

[4] In case of complex-valued matrices, functions of this form are known as *unitarily invariant* [34].

For example, through [80, Thm. 2.27] we have that $\mathrm{prox}_{\gamma G}$ is continuously differentiable at X if and only if $\mathrm{prox}_{\gamma g}$ is continuously differentiable at the singular values of X. Furthermore, $\mathrm{prox}_{\gamma G}$ is (strongly) semismooth at X if $\mathrm{prox}_{\gamma g}$ is (strongly) semismooth at the singular values of X [80, Thm. 2.3.11]. For any $X \in \mathbb{R}^{m\times n}$ the generalized Jacobian $\partial_B(\mathrm{prox}_{\gamma G})(X)$ is well defined and nonempty, and any $P \in \partial_B(\mathrm{prox}_{\gamma G})(X)$ acts on $H \in \mathbb{R}^{m\times n}$ as [80, Prop. 2.3.7]

$$P(H) = U\left[\left(\Omega^{\gamma g}_{\sigma,\sigma} \odot \left(\tfrac{H_1+H_1^\top}{2}\right) + \Omega^{\gamma g}_{\sigma,-\sigma} \odot \left(\tfrac{H_1-H_1^\top}{2}\right)\right), (\Omega^{\gamma g}_{\sigma,0} \odot H_2)\right] V^\top, \tag{15.99}$$

where $V = [V_1\ V_2]$, $H_1 = U^\top H V_1 \in \mathbb{R}^{m\times m}$, $H_2 = U^\top H V_2 \in \mathbb{R}^{m\times(n-m)}$ and matrices Ω are as in (15.95).

15.6.2 Indicator Functions

Smooth constrained convex problems

$$\mathrm{minimize}_{x\in\mathbb{R}^n}\ f(x) \quad \text{subject to } x \in D \tag{15.100}$$

can be cast in the composite form (15.1) by encoding the feasible set D with the indicator function $g = \delta_D$. Whenever $\mathscr{P}_D$ is efficiently computable, then algorithms like the forward-backward splitting (15.21) can be conveniently considered. In the following we analyze the generalized Jacobian of some of such projections.

a. Affine sets. $D = \{x \in \mathbb{R}^n \mid Ax = b\}$ for some $A \in \mathbb{R}^{m\times n}$ and $b \in \mathbb{R}^m$.
In this case, $\mathscr{P}_D(x) = x - A^\dagger(Ax - b)$ where $A^\dagger$ is the Moore-Penrose pseudoinverse of A. For example, if A is surjective (*i.e.*, it has full row rank and thus $m \le n$), then $A^\dagger = A^\top(AA^\top)^{-1}$, whereas if it is injective (*i.e.*, it has full column rank and thus $m \ge n$), then $A^\dagger = (A^\top A)^{-1}A^\top$. Obviously $\mathscr{P}_D$ is an affine mapping, thus everywhere differentiable with

$$\partial_C(\mathscr{P}_D)(x) = \partial_B(\mathscr{P}_D)(x) = \{\nabla \mathscr{P}_D(x)\} = \left\{\mathrm{I} - A^\dagger A\right\}. \tag{15.101}$$

b. Polyhedral sets. $D = \{x \in \mathbb{R}^n \mid Ax = b,\ Cx \le d\}$, for some $A \in \mathbb{R}^{p\times n}$, $b \in \mathbb{R}^p$, $C \in \mathbb{R}^{q\times n}$ and $d \in \mathbb{R}^q$.
It is well known that $\mathscr{P}_D$ is piecewise affine. In particular, let

$$\mathfrak{I}_D := \left\{I \subseteq \{1\ldots q\} \mid \exists x\in\mathbb{R}^n : Ax = b,\ C_{i\cdot}x = d_i\ i\in I,\ C_{j\cdot}x < d_j\ j\notin I\right\}. \tag{15.102}$$

Then, the *faces* of D can be indexed with the elements of $\mathfrak{I}$ [66, Prop. 2.1.3]: for each $I \in \mathfrak{I}_D$ let

$$F_I \coloneqq \{x \in D \mid C_{i\cdot}x = d_i,\ i \in I\}$$

be the I-th face of D,

$$S_I \coloneqq \operatorname{aff} F_I = \left\{x \in \mathbb{R}^n \mid Ax = b,\ C_{i\cdot}x = d_i,\ i \in I\right\}$$

be the hyperplane containing the I-th face of D,

$$N_I \coloneqq \operatorname{ran} A^\top + \operatorname{cone}\left\{C_{I\cdot}^\top\right\}$$

be the normal cone to any point in the relative interior of F_I [66, Eq. (2.44)],[5] and

$$R_I \coloneqq F_I + N_I.$$

We then have $\mathscr{P}_D(x) \in \left\{\mathscr{P}_{S_I}(x) \mid I \in \mathfrak{I}_D\right\}$, *i.e.*, $\mathscr{P}_D$ is a piecewise affine function. The affine pieces of $\mathscr{P}_D$ are the projections on the corresponding affine subspaces S_I (cf. Section §15.6.2a). In fact, for each $x \in R_I$ we have $\mathscr{P}_D(x) = \mathscr{P}_{S_I}(x)$, each R_I is full dimensional and $\mathbb{R}^n = \bigcup_{I\in\mathfrak{I}_D} R_I$ [66, Prop.s 2.4.4 and 2.4.5]. For each $I \in \mathfrak{I}_D$ let

$$P_I \coloneqq \nabla\, \mathscr{P}_{S_I} = \mathrm{I} - \begin{pmatrix} A \\ C_I \end{pmatrix}^\dagger \begin{pmatrix} A \\ C_I \end{pmatrix}, \tag{15.103}$$

and for each $x \in \mathbb{R}^n$ let

$$\mathfrak{I}_D(x) \coloneqq \{I \in \mathfrak{I}_D \mid x \in R_I\}. \tag{15.104}$$

Then,

$$\partial_C(\mathscr{P}_D)(x) = \operatorname{conv} \partial_B(\mathscr{P}_D)(x) = \operatorname{conv}\{P_I \mid I \in \mathfrak{I}_D(x)\}. \tag{15.105}$$

Therefore, an element of $\partial_B \mathscr{P}_D(x)$ is P_I as in (15.103) where $I = \{i \mid C_i\bar{x} = d_i\}$ is the set of active constraints of $\bar{x} = \mathscr{P}_D(x)$. For a more general analysis we refer the reader to [27, 39].

[5]Consistently with the definition in [66], the polyhedron P can equivalently be expressed by means of only inequalities as $P = \{x \in \mathbb{R}^n \mid Ax \le b,\ -Ax \le -b,\ Cx \le b\}$, resulting indeed in $\operatorname{cone}[A^\top,\ -A^\top,\ C^\top] = \operatorname{ran} A^\top + \operatorname{cone} C^\top$.

c. Halfspaces. $H = \{x \in \mathbb{R}^n \mid \langle a, x\rangle \le b\}$ for some $a \in \mathbb{R}^n$ and $b \in \mathbb{R}$.
Then, denoting the *positive part* of $r \in \mathbb{R}$ as $[r]_+ := \max\{0, r\}$,

$$\mathscr{P}_H(x) = x - \tfrac{[\langle a,x\rangle - b]_+}{\|a\|^2} a$$

and

$$\partial_C(\mathscr{P}_H)(x) = \begin{cases} \left\{\mathrm{I} - \|a\|^{-2} a a^\top\right\} & \text{if } x \notin H, \\ \{\mathrm{I}\} & \text{if } \langle a, x\rangle < b, \\ \operatorname{conv}\left\{\mathrm{I}, \mathrm{I} - \|a\|^{-2} a a^\top\right\} & \text{if } \langle a, x\rangle = b. \end{cases}$$

d. Boxes. $D = \{x \in \mathbb{R}^n \mid \ell \le x \le u\}$ for some $\ell, u \in [-\infty, \infty]^n$.
We have

$$\mathscr{P}_D(x) = \min\{\max\{x, \ell\}, u\},$$

and since the corresponding indicator function δ_D is separable, every element of $\partial_C(\mathscr{P}_D)(x)$ is diagonal with (cf. Section §15.6.1a)

$$\partial_C(\mathscr{P}_D)(x)_{ii} = \begin{cases} [0, 1] & \text{if } x_i \in \{\ell_i, u_i\}, \\ \{1\} & \text{if } \ell_i < x_i < u_i, \\ \{0\} & \text{otherwise,} \end{cases}$$

e. Unit simplex. $D = \left\{x \in \mathbb{R}^n \mid x \ge 0, \; \sum_{i=1}^n x_i = 1\right\}$.
By writing down the optimality conditions for the corresponding projection problem, one can easily see that

$$\mathscr{P}_D(x) = [x - \lambda \mathbf{1}]_+, \tag{15.106}$$

where λ solves $\langle \mathbf{1}, [x - \lambda\mathbf{1}]_+\rangle = 1$. Since the unit simplex is a polyhedral set, we are dealing with a special case of Section §15.6.2b, where $A = \mathbf{1}^\top$, $b = 1$, $C = -\mathrm{I}$, and $d = 0$. Therefore, in order to calculate an element of the generalized Jacobian of the projection, we first compute $\mathscr{P}_D(x)$ and then determine the set of active indices $J := \{i \mid \mathscr{P}_D(x)_i = 0\}$. An element $P \in \partial_B(\mathscr{P}_D)(x)$ is given by

$$P_{ij} = \begin{cases} \delta_{i,j} - \frac{1}{n-|J|} & \text{if } i, j \notin J, \\ 0 & \text{otherwise,} \end{cases} \tag{15.107}$$

where $|J|$ denotes the cardinality of the set J. Notice that P is block diagonal after a permutation of rows and columns.

f. Euclidean unit ball. $B = \overline{\mathrm{B}}(0; 1)$.
We have

$$\mathscr{P}_B(x) = \begin{cases} x & \text{if } x \in B, \\ {}^{x}/_{\|x\|} & \text{otherwise,} \end{cases}$$

and

$$\partial_C(\mathscr{P}_B)(x) = \begin{cases} \{\mathrm{I}\} & \text{if } \|x\| < 1, \\ \operatorname{conv}\left\{\|x\|^{-1}(\mathrm{I} - ww^\top),\ \mathrm{I}\right\} & \text{if } \|x\| = 1, \\ \left\{\|x\|^{-1}(\mathrm{I} - ww^\top)\right\} & \text{if } x \notin B, \end{cases}$$

where $w := {}^{x}/_{\|x\|}$.

g. Second-order cone. $\mathcal{K} = \left\{(x_0, \bar{x}) \in \mathbb{R} \times \mathbb{R}^{n-1} \mid x_0 \geq \|\bar{x}\|\right\}$.
Let $x := (x_0, \bar{x})$, and for $w \in \mathbb{R}^n$ and $\alpha \in \mathbb{R}$ define

$$M_{w,\alpha} := \tfrac{1}{2}\begin{bmatrix} 1 & w^\top \\ w & (1-\alpha)\mathrm{I}_{n-1} + \alpha ww^\top \end{bmatrix}. \tag{15.108}$$

Then, $\partial_C(\mathscr{P}_{\mathcal{K}})(x) = \operatorname{conv}(\partial_B(\mathscr{P}_{\mathcal{K}})(x))$ where, for $\bar{w} := {}^{\bar{x}}/_{\|\bar{x}\|}$ and $\bar{\alpha} := -{}^{x_0}/_{\|\bar{x}\|}$, we have [30, Lem. 2.6]

$$\partial_B(\mathscr{P}_{\mathcal{K}})(x) = \begin{cases} \{0\} & \text{if } x_0 < -\|\bar{x}\|, \\ \{\mathrm{I}_n\} & \text{if } x_0 > \|\bar{x}\|, \\ \left\{M_{\bar{w},\bar{\alpha}}\right\} & \text{if } -\|\bar{x}\| < x_0 < \|\bar{x}\|, \\ \left\{\mathrm{I}_n,\ M_{\bar{w},\bar{\alpha}}\right\} & \text{if } x_0 = \|\bar{x}\| \neq 0, \\ \left\{0,\ M_{\bar{w},\bar{\alpha}}\right\} & \text{if } x_0 = -\|\bar{x}\| \neq 0, \\ \{0, \mathrm{I}_n\} \cup \left\{M_{w,\alpha} \mid |\alpha| \leq 1,\ \|w\| \leq 1\right\} & \text{if } x_0 = \bar{x} = 0. \end{cases} \tag{15.109}$$

h. Positive semidefinite cone. $\mathcal{S}_+ = \mathrm{S}_+(\mathbb{R}^{n\times n})$.
For any symmetric matrix M it holds that

$$\mathscr{P}_{\mathcal{S}_+}(M) = Q[\operatorname{diag}(\lambda)]_+ Q^\top, \tag{15.110}$$

where $M = Q\operatorname{diag}(\lambda)Q^\top$ is any spectral decomposition of M. This coincides with (15.93), as $\delta_{\mathcal{S}_+}$ can be expressed as in (15.89), where h has the separable form (15.91) with $g = \delta_{\mathbb{R}_+}$, so that for $r \in \mathbb{R}$ we have

$$\operatorname{prox}_{\gamma g}(r) = [r]_+ \quad \text{and} \quad \partial_B(\operatorname{prox}_{\gamma g})(r) = \begin{cases} \{0\} & \text{if } r < 0, \\ \{0, 1\} & \text{if } r = 0, \\ \{1\} & \text{if } r > 0. \end{cases} \tag{15.111}$$

An element of $\partial_B \mathscr{P}_{\mathrm{S}_+(\mathbb{R}^{n\times n})}(X)$ is thus given by (15.94).

15.6.3 Norms

a. ℓ_1 **norm.** $g(x) = \|x\|_1$.
The proximal mapping is the well-known soft-thresholding operator

$$(\mathrm{prox}_{\gamma g}(x))_i = \mathrm{sign}(x_i)[|x_i| - \gamma]_+, \quad i = 1, \dots, n. \tag{15.112}$$

Function g is separable, and thus every element of $\partial_B(\mathrm{prox}_{\gamma g})$ is a diagonal matrix, cf. Section §15.6.1a. Specifically, the nonzero elements are

$$\partial_C(\mathrm{prox}_{\gamma g})(x)_{ii} = \begin{cases} \{1\} & \text{if } |x_i| > \gamma, \\ [0, 1] & \text{if } |x_i| = \gamma, \\ \{0\} & \text{if } |x_i| < \gamma. \end{cases} \tag{15.113}$$

We could also arrive to the same conclusion by applying the Moreau decomposition of Section §15.6.1b to the function of Section §15.6.2d with $u = -\ell = \mathbf{1}_n$, since the ℓ_1 norm is the conjugate of the indicator of the ℓ_∞-norm ball.

b. ℓ_∞ **norm.** $g(x) = \|x\|_\infty$.
Function g is the convex conjugate of the indicator of the unit simplex D analyzed in Section §15.6.2e. From the Moreau decomposition, see Section §15.6.1b, we obtain

$$\partial_C(\mathrm{prox}_{\gamma g})(x) = \mathrm{I} - \partial_C(\mathscr{P}_D)(x/\gamma). \tag{15.114}$$

Then, $\mathscr{P}_D(x/\gamma) = [x/\gamma - \lambda\mathbf{1}]_+$ where $\lambda \in \mathbb{R}$ solves $\langle \mathbf{1}, [x/\gamma - \lambda\mathbf{1}]_+ \rangle = 1$. Let $J = \{i \mid \mathscr{P}_D(x/\gamma)_i = 0\}$, then an element of $\partial_B(\mathrm{prox}_{\gamma g})(x)$ is given by

$$P_{ij} = \begin{cases} \frac{1}{n-|J|} & \text{if } i, j \notin J, \\ \delta_{i,j} & \text{otherwise.} \end{cases} \tag{15.115}$$

c. Euclidean norm. $g(x) = \|x\|$.
The proximal mapping is given by

$$\mathrm{prox}_{\gamma g}(x) = \begin{cases} (1 - \gamma\|x\|^{-1})x & \text{if } \|x\| \geq \gamma, \\ 0 & \text{otherwise.} \end{cases} \tag{15.116}$$

Since $\mathrm{prox}_{\gamma g}$ is a PC^1 mapping, its B-subdifferential can be computed by simply computing the Jacobians of its smooth pieces. Specifically, denoting $w = x/\|x\|$ we have

$$\partial_C(\mathrm{prox}_{\gamma g})(x) = \begin{cases} \left\{\mathrm{I} - \gamma\|x\|^{-1}(\mathrm{I} - ww^\top)\right\} & \text{if } \|x\| > \gamma, \\ \{0\} & \text{if } \|x\| < \gamma, \\ \mathrm{conv}\left\{\mathrm{I} - \gamma\|x\|^{-1}(\mathrm{I} - ww^\top), 0\right\} & \text{otherwise.} \end{cases} \tag{15.117}$$

d. Sum of Euclidean norms. $g(x) = \sum_{s\in\mathcal{S}} \|x_s\|$, where $\mathcal{S}$ is a partition of $\{1,\dots,n\}$.
Differently from the ℓ_1-norm which induces sparsity on the whole vector, this function serves as regularizer to induce group sparsity [81]. For $s \in \mathcal{S}$, the components of the proximal mapping indexed by s are

$$(\operatorname{prox}_{\gamma g}(x))_s = (1 - \gamma\|x_s\|^{-1})_+ x_s. \tag{15.118}$$

Any $P \in \partial_B(\operatorname{prox}_{\gamma g})(x)$ is block diagonal with the s-block equal to

$$P_s = \begin{cases} \mathrm{I} - \gamma\|x_s\|^{-1}\big(\mathrm{I} - \|x_s\|^{-2} x_s x_s^\top\big) & \text{if } \|x_s\| > \gamma, \\ \mathrm{I} & \text{if } \|x_s\| < \gamma, \\ \text{any of these two matrices} & \text{if } \|x_s\| = \gamma. \end{cases} \tag{15.119}$$

e. Matrix nuclear norm. $G(X) = \|X\|_\star$ for $X \in \mathbb{R}^{m\times n}$.
The *nuclear norm* returns the sum of the singular values of a matrix $X \in \mathbb{R}^{m\times n}$, *i.e.*, $G(X) = \sum_{i=1}^m \sigma_i(X)$ (for simplicity we are assuming that $m \le n$). It serves as a convex surrogate for the rank, and has found many applications in systems and control theory, including system identification and model reduction [20–22, 41, 61]. Other fields of application include *matrix completion problems* arising in machine learning [62, 68] and computer vision [52, 76], and *nonnegative matrix factorization problems* arising in data mining [18].
The nuclear norm can be expressed as $G(X) = h(\sigma(X))$, where $h(x) = \|x\|_1$ is absolutely symmetric and separable. Specifically, it takes the form (15.91) with $g = |\cdot|$, for which $g(0) = 0$ and $0 \in \partial g(0)$, and whose proximal mapping is the soft-thresholding operator. In fact, since the case of interest here is $x \ge 0$ (because $\sigma_i(X) \ge 0$), we have $\operatorname{prox}_{\gamma g}(x) = [x - \gamma]_+$, cf. (15.116). Consequently, the proximal mapping of $\|X\|_\star$ is given by (15.98) with

$$\Sigma_g(X) = \operatorname{diag}([\sigma_1(X) - \gamma]_+, \dots, [\sigma_m(X) - \gamma]_+). \tag{15.120}$$

For $x \in \mathbb{R}_+$ we have that

$$\partial_C(\operatorname{prox}_{\gamma g})(x) = \begin{cases} 0 & \text{if } 0 \le x < \gamma, \\ [0,1] & \text{if } x = \gamma, \\ 1 & \text{if } x > \gamma, \end{cases} \tag{15.121}$$

then $\partial_B(\operatorname{prox}_{\gamma G})(X)$ takes the form as in (15.99).

15.7 Conclusions

A forward-backward truncated-Newton method (FBTN) is proposed that minimizes the sum of two convex functions one of which Lipschitz continuous and twice continuously differentiable. Our approach is based on the forward-backward envelope (FBE), a continuously differentiable tight lower bound to the original (nonsmooth and extended-real valued) cost function sharing minima and minimizers. The method requires forward-backward steps, Hessian evaluations of the smooth function and Clarke Jacobians of the proximal map of the nonsmooth term. Explicit formulas of Clarke Jacobians of a wide variety of useful nonsmooth functions are collected from the literature for the reader's convenience. The higher-order operations are needed for the computation of symmetric and positive semidefinite matrices that serve as surrogate for the Hessian of the FBE, allowing for a generalized (regularized, truncated-) Newton method for its minimization. The algorithm exhibits global Q-linear convergence under an error bound condition, and Q-superlinear or even Q-quadratic if an additional semismoothness assumption at the limit point is satisfied.

Acknowledgements This work was supported by the *Research Foundation Flanders (FWO)* research projects G086518N and G086318N; *KU Leuven internal funding* StG/15/043; *Fonds de la Recherche Scientifique—FNRS and the Fonds Wetenschappelijk Onderzoek—Vlaanderen* under EOS Project no 30468160 (SeLMA).

Appendix: Auxiliary Results

Lemma 1 *Any proper lsc convex function with nonempty and bounded set of minimizers is level bounded.*

Proof Let h be such function; to avoid trivialities we assume that $\operatorname{dom} h$ is unbounded. Fix $x_\star \in \operatorname{argmin} h$ and let $R > 0$ be such that $\operatorname{argmin} h \subseteq B := \mathrm{B}(x_\star; R)$. Since $\operatorname{dom} h$ is closed, convex, and unbounded, it holds that h attains a minimum on the compact set $\operatorname{bdry} B$, be it m, which is strictly larger than $h(x_\star)$ (since $\operatorname{dist}(\operatorname{argmin} h, \operatorname{bdry} B) > 0$ due to compactness of $\operatorname{argmin} h$ and openness of B). For $x \notin B$, let $s_x = x_\star + R\frac{x-x_\star}{\|x-x_\star\|}$ denote its projection onto $\operatorname{bdry} B$, and let $t_x := \frac{\|x-x_\star\|}{R} \geq 1$. Then,

$$h(x) = h\big(x_\star + t_x(s_x - x_\star)\big) \geq h(x_\star) + t_x\big(h(s_x) - h(x_\star)\big) \geq h(x_\star) + t_x\big(m - h(x_\star)\big)$$

where in the first inequality we used the fact that $t_x \geq 1$. Since $m - h(x_\star) > 0$ and $t_x \to \infty$ as $\|x\| \to \infty$, we conclude that h is coercive, and thus level bounded. □

Lemma 2 *Let* $H \in \mathrm{S}_+(\mathbb{R}^n)$ *with* $\lambda_{\max}(H) \leq 1$. *Then* $H - H^2 \in \mathrm{S}_+(\mathbb{R}^n)$ *with*

$$\lambda_{\min}(H - H^2) = \min\left\{\lambda_{\min}(H)(1 - \lambda_{\min}(H)),\ \lambda_{\max}(H)(1 - \lambda_{\max}(H))\right\}. \tag{15.122}$$

Proof Consider the spectral decomposition $H = S^\top D S$ for some orthogonal matrix S and diagonal D. Then, $H - H^2 = S^\top \tilde{D} S$ where $\tilde{D} = D - D^2$. Apparently, $\tilde{D}$ is diagonal, hence the eigenvalues of $H - H^2$ are exactly $\left\{\lambda - \lambda^2 \mid \lambda \in \operatorname{eigs}(H)\right\}$.

The function $\lambda \mapsto \lambda - \lambda^2$ is concave, hence the minimum in $\operatorname{eigs}(\tilde{H})$ is attained at one extremum, that is, either at $\lambda = \lambda_{\min}(H)$ or $\lambda = \lambda_{\max}(H)$, which proves the claim. □

Lemma 3 *For any $\gamma \in (0, 2/L_f)$ the forward-backward operator T_γ (15.22) is nonexpansive (in fact, $\frac{2}{4-\gamma L_f}$-averaged), and the residual R_γ is Lipschitz continuous with modulus $\frac{4}{\gamma(4-\gamma L_f)}$.*

Proof By combining [2, Prop. 4.39 and Cor. 18.17] it follows that the *gradient descent operator $x \mapsto x - \gamma \nabla f(x)$* is $\gamma L_f/2$-averaged. Moreover, since the proximal mapping is $1/2$-averaged [2, Prop. 12.28] we conclude from [2, Prop. 4.44] that the forward-backward operator T_γ is α-averaged with $\alpha = \frac{2}{4-\gamma L_f}$, thus nonexpansive [2, Rem. 4.34(i)]. Therefore, by definition of α-averagedness there exists a 1-Lipschitz continuous operator S such that $T_\gamma = (1-\alpha)\operatorname{id} + \alpha S$ and consequently the residual $R_\gamma = \frac{1}{\gamma}\big(\operatorname{id} - T_\gamma\big) = \frac{\alpha}{\gamma}(\operatorname{id} - S)$ is $(2\alpha/\gamma)$-Lipschitz continuous. □

References

1. Attouch, H., Bolte, J., Svaiter, B.F.: Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward-backward splitting, and regularized Gauss-Seidel methods. Mathematical Programming **137**(1), 91–129 (2013). DOI 10.1007/s10107-011-0484-9
2. Bauschke, H.H., Combettes, P.L.: Convex analysis and monotone operator theory in Hilbert spaces. CMS Books in Mathematics. Springer (2017). DOI 10.1007/978-3-319-48311-5
3. Bauschke, H.H., Noll, D., Phan, H.M.: Linear and strong convergence of algorithms involving averaged nonexpansive operators. Journal of Mathematical Analysis and Applications **421**(1), 1–20 (2015)
4. Beck, A.: First-Order Methods in Optimization. Society for Industrial and Applied Mathematics, Philadelphia, PA (2017). DOI 10.1137/1.9781611974997
5. Beck, A., Teboulle, M.: A fast iterative shrinkage-thresholding algorithm for linear inverse problems. SIAM Journal on Imaging Sciences **2**(1), 183–202 (2009). DOI 10.1137/080716542
6. Becker, S., Fadili, J.: A quasi-Newton proximal splitting method. In: Advances in Neural Information Processing Systems, pp. 2618–2626 (2012)
7. Bertsekas, D.P.: Constrained optimization and lagrange multiplier methods. Computer Science and Applied Mathematics, Boston: Academic Press, 1982 (1982)
8. Bertsekas, D.P.: Convex Optimization Algorithms. Athena Scientific (2015)
9. Bhatia, R.: Matrix Analysis. Graduate Texts in Mathematics. Springer New York (1997)
10. Bochnak, J., Coste, M., Roy, M.F.: Real Algebraic Geometry. Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge / A Series of Modern Surveys in Mathematics. Springer Berlin Heidelberg (2013)
11. Bolte, J., Daniilidis, A., Lewis, A.: Tame functions are semismooth. Mathematical Programming **117**(1), 5–19 (2009). DOI 10.1007/s10107-007-0166-9
12. Chen, G., Teboulle, M.: Convergence analysis of a proximal-like minimization algorithm using Bregman functions. SIAM Journal on Optimization **3**(3), 538–543 (1993). DOI 10.1137/0803026
13. Chen, X., Fukushima, M.: Proximal quasi-Newton methods for nondifferentiable convex optimization. Mathematical Programming **85**(2), 313–334 (1999). DOI 10.1007/s101070050059

14. Chen, X., Qi, H., Tseng, P.: Analysis of nonsmooth symmetric-matrix-valued functions with applications to semidefinite complementarity problems. SIAM Journal on Optimization **13**(4), 960–985 (2003). DOI 10.1137/S1052623400380584
15. Clarke, F.H.: Optimization and Nonsmooth Analysis. Society for Industrial and Applied Mathematics (1990). DOI 10.1137/1.9781611971309
16. Combettes, P.L., Pesquet, J.C.: Proximal Splitting Methods in Signal Processing, pp. 185–212. Springer New York, New York, NY (2011). DOI 10.1007/978-1-4419-9569-8_10
17. Drusvyatskiy, D., Lewis, A.S.: Error bounds, quadratic growth, and linear convergence of proximal methods. Mathematics of Operations Research (2018)
18. Eldén, L.: Matrix Methods in Data Mining and Pattern Recognition. Society for Industrial and Applied Mathematics (2007). DOI 10.1137/1.9780898718867
19. Facchinei, F., Pang, J.S.: Finite-dimensional variational inequalities and complementarity problems, vol. II. Springer (2003)
20. Fazel, M.: Matrix rank minimization with applications. Ph.D. thesis, Stanford University (2002)
21. Fazel, M., Hindi, H., Boyd, S.P.: A rank minimization heuristic with application to minimum order system approximation. In: Proceedings of the 2001 American Control Conference, vol. 6, pp. 4734–4739 (2001). DOI 10.1109/ACC.2001.945730
22. Fazel, M., Hindi, H., Boyd, S.P.: Rank minimization and applications in system theory. In: Proceedings of the 2004 American Control Conference, vol. 4, pp. 3273–3278 vol.4 (2004). DOI 10.23919/ACC.2004.1384521
23. Fukushima, M.: Equivalent differentiable optimization problems and descent methods for asymmetric variational inequality problems. Mathematical Programming **53**(1), 99–110 (1992). DOI 10.1007/BF01585696
24. Giselsson, P., Fält, M.: Envelope functions: Unifications and further properties. Journal of Optimization Theory and Applications (2018). DOI 10.1007/s10957-018-1328-z
25. Gowda, M.S.: Inverse and implicit function theorems for H-differentiable and semismooth functions. Optimization Methods and Software **19**(5), 443–461 (2004). DOI 10.1080/10556780410001697668
26. Güler, O.: New proximal point algorithms for convex minimization. SIAM Journal on Optimization **2**(4), 649–664 (1992). DOI 10.1137/0802032
27. Han, J., Sun, D.: Newton and quasi-Newton methods for normal maps with polyhedral sets. Journal of Optimization Theory and Applications **94**(3), 659–676 (1997). DOI 10.1023/A:1022653001160
28. Hiriart-Urruty, J.B., Lemaréchal, C.: Fundamentals of Convex Analysis. Grundlehren Text Editions. Springer Berlin Heidelberg (2004)
29. Horn, R.A., Horn, R.A., Johnson, C.R.: Topics in Matrix Analysis. Cambridge University Press (1994)
30. Kanzow, C., Ferenczi, I., Fukushima, M.: On the local convergence of semismooth Newton methods for linear and nonlinear second-order cone programs without strict complementarity. SIAM Journal on Optimization **20**(1), 297–320 (2009). DOI 10.1137/060657662
31. Lan, G., Lu, Z., Monteiro, R.D.C.: Primal-dual first-order methods with $O(1/\varepsilon)$ iteration-complexity for cone programming. Mathematical Programming **126**(1), 1–29 (2011). DOI 10.1007/s10107-008-0261-6
32. Lee, J.D., Sun, Y., Saunders, M.: Proximal Newton-type methods for minimizing composite functions. SIAM Journal on Optimization **24**(3), 1420–1443 (2014). DOI 10.1137/130921428
33. Lemaréchal, C., Sagastizábal, C.: Practical aspects of the Moreau-Yosida regularization: Theoretical preliminaries. SIAM Journal on Optimization **7**(2), 367–385 (1997). DOI 10.1137/S1052623494267127
34. Lewis, A.S.: The convex analysis of unitarily invariant matrix functions. Journal of Convex Analysis **2**(1), 173–183 (1995)
35. Lewis, A.S.: Convex analysis on the Hermitian matrices. SIAM Journal on Optimization **6**(1), 164–177 (1996). DOI 10.1137/0806009

36. Lewis, A.S.: Derivatives of spectral functions. Mathematics of Operations Research **21**(3), 576–588 (1996)
37. Lewis, A.S., Sendov, H.S.: Twice differentiable spectral functions. SIAM Journal on Matrix Analysis and Applications **23**(2), 368–386 (2001). DOI 10.1137/S089547980036838X
38. Li, W., Peng, J.: Exact penalty functions for constrained minimization problems via regularized gap function for variational inequalities. Journal of Global Optimization **37**(1), 85–94 (2007). DOI 10.1007/s10898-006-9038-8
39. Li, X., Sun, D., Toh, K.C.: On the efficient computation of a generalized Jacobian of the projector over the Birkhoff polytope. ArXiv e-prints (2017)
40. Lions Lions, P.L., Mercier, B.: Splitting algorithms for the sum of two nonlinear operators. SIAM Journal on Numerical Analysis **16**(6), 964–979 (1979). DOI 10.1137/0716071
41. Liu, Z., Vandenberghe, L.: Interior-point method for nuclear norm approximation with application to system identification. SIAM Journal on Matrix Analysis and Applications **31**(3), 1235–1256 (2010). DOI 10.1137/090755436
42. Lu, Z.: Randomized block proximal damped Newton method for composite self-concordant minimization. SIAM Journal on Optimization **27**(3), 1910–1942 (2017). DOI 10.1137/16M1082767
43. Luo, Z.Q., Tseng, P.: Error bounds and convergence analysis of feasible descent methods: a general approach. Annals of Operations Research **46**(1), 157–178 (1993). DOI 10.1007/BF02096261
44. Maratos, N.: Exact penalty function algorithms for finite dimensional and control optimization problems (1978)
45. Martinet, B.: Brève communication. Régularisation d'inéquations variationnelles par approximations successives. Revue française d'informatique et de recherche opérationnelle. Série rouge **4**(R3), 154–158 (1970)
46. Meng, F.: Moreau-Yosida regularization of Lagrangian-dual functions for a class of convex optimization problems. Journal of Global Optimization **44**(3), 375 (2008). DOI 10.1007/s10898-008-9333-7
47. Meng, F., Sun, D., Zhao, G.: Semismoothness of solutions to generalized equations and the Moreau-Yosida regularization. Mathematical Programming **104**(2), 561–581 (2005). DOI 10.1007/s10107-005-0629-9
48. Meng, F., Zhao, G., Goh, M., De Souza, R.: Lagrangian-dual functions and Moreau-Yosida regularization. SIAM Journal on Optimization **19**(1), 39–61 (2008). DOI 10.1137/060673746
49. Mifflin, R.: Semismooth and semiconvex functions in constrained optimization. SIAM Journal on Control and Optimization **15**(6), 959–972 (1977). DOI 10.1137/0315061
50. Mifflin, R., Qi, L., Sun, D.: Properties of the Moreau-Yosida regularization of a piecewise C^2 convex function. Mathematical Programming **84**(2), 269–281 (1999). DOI 10.1007/s10107980029a
51. Moreau, J.J.: Proximité et dualité dans un espace hilbertien. Bulletin de la Société Mathématique de France **93**, 273–299 (1965)
52. Morita, T., Kanade, T.: A sequential factorization method for recovering shape and motion from image streams. IEEE Transactions on Pattern Analysis and Machine Intelligence **19**(8), 858–867 (1997). DOI 10.1109/34.608289
53. Nesterov, Y.: Introductory lectures on convex optimization: A basic course, vol. 87. Springer (2003)
54. Nesterov, Y.: Gradient methods for minimizing composite functions. Mathematical Programming **140**(1), 125–161 (2013). DOI 10.1007/s10107-012-0629-5
55. Pang, J.S.: Error bounds in mathematical programming. Mathematical Programming **79**(1), 299–332 (1997). DOI 10.1007/BF02614322
56. Parikh, N., Boyd, S.: Proximal algorithms. Found. Trends Optim. **1**(3), 127–239 (2014). DOI 10.1561/2400000003
57. Patrinos, P., Bemporad, A.: Proximal Newton methods for convex composite optimization. In: IEEE Conference on Decision and Control, pp. 2358–2363 (2013)

58. Patrinos, P., Sopasakis, P., Sarimveis, H.: A global piecewise smooth Newton method for fast large-scale model predictive control. Automatica **47**(9), 2016–2022 (2011)
59. Patrinos, P., Stella, L., Bemporad, A.: Forward-backward truncated Newton methods for convex composite optimization. ArXiv e-prints (2014)
60. Qi, L., Sun, J.: A nonsmooth version of Newton's method. Mathematical Programming **58**(1), 353–367 (1993). DOI 10.1007/BF01581275
61. Recht, B., Fazel, M., Parrilo, P.A.: Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. SIAM Review **52**(3), 471–501 (2010). DOI 10.1137/070697835
62. Rennie, J.D.M., Srebro, N.: Fast maximum margin matrix factorization for collaborative prediction. In: Proceedings of the 22Nd International Conference on Machine Learning, ICML '05, pp. 713–719. ACM, New York, NY, USA (2005). DOI 10.1145/1102351.1102441
63. Rockafellar, R.: Convex analysis (1970)
64. Rockafellar, R.T.: Monotone operators and the proximal point algorithm. SIAM Journal on Control and Optimization **14**(5), 877–898 (1976). DOI 10.1137/0314056
65. Rockafellar, R.T., Wets, R.J.B.: Variational analysis, vol. 317. Springer Science & Business Media (2011)
66. Scholtes, S.: Piecewise Differentiable Functions, pp. 91–111. Springer New York, New York, NY (2012). DOI 10.1007/978-1-4614-4340-7_4
67. Sopasakis, P., Freris, N., Patrinos, P.: Accelerated reconstruction of a compressively sampled data stream. In: 2016 24th European Signal Processing Conference (EUSIPCO), pp. 1078–1082 (2016). DOI 10.1109/EUSIPCO.2016.7760414
68. Srebro, N.: Learning with matrix factorizations. Ph.D. thesis, Cambridge, MA, USA (2004)
69. Stella, L., Themelis, A., Patrinos, P.: Forward-backward quasi-Newton methods for nonsmooth optimization problems. Computational Optimization and Applications **67**(3), 443–487 (2017). DOI 10.1007/s10589-017-9912-y
70. Stella, L., Themelis, A., Patrinos, P.: Newton-type alternating minimization algorithm for convex optimization. IEEE Transactions on Automatic Control (2018). DOI 10.1109/TAC.2018.2872203
71. Stella, L., Themelis, A., Sopasakis, P., Patrinos, P.: A simple and efficient algorithm for nonlinear model predictive control. In: 2017 IEEE 56th Annual Conference on Decision and Control (CDC), pp. 1939–1944 (2017). DOI 10.1109/CDC.2017.8263933
72. Sun, D., Fukushima, M., Qi, L.: A computable generalized Hessian of the D-gap function and Newton-type methods for variational inequality problems. Complementarity and Variational Problems: State of the Art, MC Ferris and JS Pang (eds.), SIAM, Philadelphia, PA pp. 452–472 (1997)
73. Sun, D., Sun, J.: Semismooth matrix-valued functions. Mathematics of Operations Research **27**(1), 150–169 (2002). DOI 10.1287/moor.27.1.150.342
74. Themelis, A., Patrinos, P.: Douglas-Rachford splitting and ADMM for nonconvex optimization: tight convergence results. ArXiv e-prints (2017)
75. Themelis, A., Stella, L., Patrinos, P.: Forward-backward envelope for the sum of two nonconvex functions: Further properties and nonmonotone linesearch algorithms. SIAM Journal on Optimization **28**(3), 2274–2303 (2018). DOI 10.1137/16M1080240
76. Tomasi, C., Kanade, T.: Shape and motion from image streams under orthography: a factorization method. International Journal of Computer Vision **9**(2), 137–154 (1992). DOI 10.1007/BF00129684
77. Tseng, P.: On accelerated proximal gradient methods for convex-concave optimization. Tech. rep. (2008)
78. Ulbrich, M.: Optimization Methods in Banach Spaces, pp. 97–156. Springer Netherlands, Dordrecht (2009). DOI 10.1007/978-1-4020-8839-1_2
79. Yamashita, N., Taji, K., Fukushima, M.: Unconstrained optimization reformulations of variational inequality problems. Journal of Optimization Theory and Applications **92**(3), 439–456 (1997). DOI 10.1023/A:1022660704427

80. Yang, Z.: A study on nonsymmetric matrix-valued functions. Master's thesis, National University of Singapore (2009)
81. Yuan, M., Lin, Y.: Model selection and estimation in regression with grouped variables. Journal of the Royal Statistical Society. Series B (Statistical Methodology) **68**(1), 49–67 (2006)
82. Zhou, G., Qi, L.: On the convergence of an inexact Newton-type method. Oper. Res. Lett. **34**(6), 647–652 (2006). DOI 10.1016/j.orl.2005.11.001
83. Zhou, G., Toh, K.C.: Superlinear convergence of a Newton-type algorithm for monotone equations. Journal of Optimization Theory and Applications **125**(1), 205–221 (2005). DOI 10.1007/s10957-004-1721-7