SIAM J. IMAGING SCIENCES
Vol. 8, No. 3, pp. 1752–1797

Generic Half-Quadratic Optimization for Image Reconstruction*

Marc C. Robini† and Yuemin Zhu†

Abstract. We study the global and local convergence of a generic half-quadratic optimization algorithm inspired from the dual energy formulation of Geman and Reynolds [*IEEE Trans. Pattern Anal. Mach. Intell.*, 14 (1992), pp. 367–383]. The target application is the minimization of C^1 convex and nonconvex objective functionals arising in regularized image reconstruction. Our global convergence proofs are based on a monotone convergence theorem of Meyer [*J. Comput. System Sci.*, 12 (1976), pp. 108–121]. Compared to existing results, ours extend to a larger class of objectives and apply under weaker conditions; in particular, we cover the case where the set of stationary points is not discrete. Our local convergence results use a majorization-minimization interpretation to derive an insightful characterization of the basins of attraction; this new perspective grounds a formal description of the intuitive water-flooding analogy. We conclude with image restoration experiments to illustrate the efficiency of the algorithm under various nonconvex scenarios.

Key words. half-quadratic optimization, nonconvex optimization, inverse problems, ill-posedness, regularization, image reconstruction, image restoration, edge preservation

AMS subject classifications. 49N45, 65F22, 65K05, 65K10, 68U10, 68W40, 90C26, 94A08

DOI. 10.1137/140987845

1. Introduction.

We focus on a general problem which includes regularized linear least squares and regularized linear robust estimation as special cases. The issue it to minimize a functional $\Theta : \mathbb{R}^N \longrightarrow \mathbb{R}$ of the form

$$\Theta(\boldsymbol{x}) := \sum_{k=1}^{K} \theta_k\big(\|\boldsymbol{A}_k\boldsymbol{x} - \boldsymbol{a}_k\|\big), \tag{1.1}$$

where $\|\cdot\|$ denotes the ℓ_2-norm and, for each k, $\boldsymbol{A}_k$ is a matrix in $\mathbb{R}^{N_k \times N}$, $\boldsymbol{a}_k$ is a vector in $\mathbb{R}^{N_k}$, and θ_k belongs to the set of functions $\theta : \mathbb{R}_+ \longrightarrow \mathbb{R}$ satisfying the following conditions:

- (C1) θ is increasing and nonconstant;
- (C2) θ is C^1 on $(0, +\infty)$ and continuous at zero;
- (C3) the function $\theta^\dagger : t \in (0, +\infty) \longmapsto t^{-1}\theta'(t)$ is decreasing and bounded.

In particular, the functions θ_k can be nonconvex and even eventually constant. Note that the boundedness assumption in (C3) implies that the derivatives θ_k' vanish at zero; this excludes for example the identity function and the strictly concave C^1 functions considered in [1]. However, in practice, such functions can be approximated arbitrarily closely by functions satisfying (C1)–(C3).

*Received by the editors September 19, 2014; accepted for publication (in revised form) June 24, 2015; published electronically August 27, 2015. This work was partly supported by the French ANR (ANR-13-MONU-0009).
http://www.siam.org/journals/siims/8-3/98784.html
†CREATIS (CNRS research unit UMR5220 and INSERM research unit U1044), INSA-Lyon, 69621 Villeurbanne Cedex, France (marc.robini@creatis.insa-lyon.fr, zhu@creatis.insa-lyon.fr).

To perform this optimization task, we consider a generalization of the half-quadratic approach that emerged from the dual energy formulation of Geman and Reynolds [2] and whose convergence was subsequently studied in [3, 4, 5, 6, 7, 8]. The process consists in generating a sequence $(\boldsymbol{x}^{(p)})_{p\in\mathbb{N}}$ by using the recurrence relation

$$\boldsymbol{x}^{(p+1)} := \underset{\boldsymbol{y}\in\mathbb{R}^N}{\arg\min} \sum_{k=1}^{K} \theta_k^{\ddagger}\big(\|\boldsymbol{A}_k\boldsymbol{x}^{(p)} - \boldsymbol{a}_k\|\big)\,\|\boldsymbol{A}_k\boldsymbol{y} - \boldsymbol{a}_k\|^2, \tag{1.2}$$

where $\theta_k^{\ddagger}$ is the continuous extension of $\theta^{\dagger}$ to $\mathbb{R}_+$ (later, we give a necessary and sufficient condition for the argument of the minimum to be a singleton, and we prove that the iteration map $\boldsymbol{x}^{(p)} \longmapsto \boldsymbol{x}^{(p+1)}$ is continuous). Depending on the assumptions made, this algorithm has different interpretations: alternating minimization [3, 5], majorization-minimization [7], quasi-Newton optimization [7, 9], and fixed-point iteration [8]. Its simplicity and ease of implementation make it attractive for rapid and efficient testing of solutions to minimizing objective functionals of the form of (1.1). Our motivation—signal reconstruction from noisy linear measurements—is discussed in section 1.1. We then review existing convergence results in section 1.2, and we outline our contributions in section 1.3.

1.1. Motivation. We are interested in the general problem of reconstructing a signal $\boldsymbol{x}^{\sharp} \in \mathbb{R}^N$ given some data

$$\boldsymbol{d} := \chi(\boldsymbol{D}\boldsymbol{x}^{\sharp} + \boldsymbol{\nu}), \tag{1.3}$$

where $\boldsymbol{D} \in \mathbb{R}^{M\times N}$ models the deterministic part of the observation process, $\boldsymbol{\nu} \in \mathbb{R}^M$ is a noise term consisting of a realization of M independent zero-mean random variables, and $\chi : \mathbb{R}^M \longrightarrow \mathbb{R}^M$ is a componentwise function that stands for possible contamination by impulse noise. The original signal $\boldsymbol{x}^{\sharp}$ can represent a temporal sequence, an image, a volume, or more generally a set of samples on a point lattice. When $\boldsymbol{D}$ is known, estimates of $\boldsymbol{x}^{\sharp}$ are usually defined as global minimizers of functionals of the form

$$\Theta(\boldsymbol{x}) = \underbrace{\sum_{m=1}^{M} \theta_{\mathrm{fid}}\big(|(\boldsymbol{D}\boldsymbol{x} - \boldsymbol{d})_m|\big)}_{=:\,\Theta_{\mathrm{fid}}(\boldsymbol{x})} + \underbrace{\sum_{l=1}^{L} \theta_{\mathrm{reg}}\big(\|\boldsymbol{R}_l\boldsymbol{x}\|\big)}_{=:\,\Theta_{\mathrm{reg}}(\boldsymbol{x})}, \tag{1.4}$$

where the data-fidelity term Θ_{fid} favors solutions consistent with the observations (the notation $(\boldsymbol{D}\boldsymbol{x} - \boldsymbol{d})_m$ stands for the mth component of the residual vector), and the regularization term Θ_{reg} imposes prior constraints modeled by the matrices $\boldsymbol{R}_l$. The functional (1.4) is a special case of (1.1) obtained by setting $K = M + L$ and

$$(\theta_k, \boldsymbol{A}_k, \boldsymbol{a}_k) = \begin{cases} (\theta_{\mathrm{fid}}, \boldsymbol{D}(k,:), d_k) & \text{if } k \in [1\,..\,M], \\ (\theta_{\mathrm{reg}}, \boldsymbol{R}_{k-M}, \boldsymbol{0}) & \text{if } k \in [M+1\,..\,K], \end{cases} \tag{1.5}$$

where $\boldsymbol{D}(k,:)$ denotes the kth row of the observation matrix $\boldsymbol{D}$.

Usually, $\{\boldsymbol{R}_l\}_{l\in[1..L]}$ is a discrete approximation to the gradient operator, which favors piecewise-smooth solutions [10, 11]. In image restoration, a recent trend is toward second-

and higher-order operators that reduce staircase artifacts and improve contour regularity; examples include tight framelet filters [12], isotropic second-order total variation [13], Hessian Frobenius norm regularization [14], and wavelet-domain edge-continuation [8]. The set $\{\boldsymbol{R}_l\}_l$ can also consist of projections on a set of vectors to promote sparsity either in a linear transform domain [15] or with respect to a redundant dictionary [16]. We can of course combine different types of operators: our results cover functionals of the form of (1.4) with additional regularization terms similar to Θ_{reg}.

The function θ_{reg} is called a *potential function* in reference to the Bayesian interpretation of regularization [17]—we use the same designation for the functions θ_k in (1.1). In practice, a potential function θ is derived from a *mother potential function* ϑ by scaling the value and the argument of ϑ, that is,

$$\theta(t) = \gamma\vartheta(t/\delta),$$

where γ and δ are free positive parameters. Obviously, θ satisfies conditions (C1)–(C3) if and only if ϑ does. The mother potential functions used for regularization can be divided into three categories:

- convex functions with affine behavior at infinity, such as the function of minimal surfaces

$$\vartheta_{\text{MS}}(t) := (1+t^2)^{1/2} - 1 \tag{1.6}$$

  and the Huber function

$$\vartheta_{\text{Hu}}(t) := \begin{cases} t^2/2 & \text{if } t \leqslant 1, \\ t - 1/2 & \text{if } t > 1; \end{cases} \tag{1.7}$$

- nonconvex unbounded functions, such as the Lorentzian error function

$$\vartheta_{\text{LE}}(t) := \ln(1+t^2); \tag{1.8}$$

- bounded functions, such as the Geman and McClure function

$$\vartheta_{\text{GM}}(t) := \frac{t^2}{1+t^2} \tag{1.9}$$

  and the Tukey biweight function

$$\vartheta_{\text{TB}}(t) := \begin{cases} 1 - \left(1 - t^2/6\right)^3 & \text{if } t \leqslant \sqrt{6}, \\ 1 & \text{if } t > \sqrt{6}. \end{cases} \tag{1.10}$$

The frequent dichotomy between convex and nonconvex potential functions is best illustrated when the original signal is an image and $\{\boldsymbol{R}_l\}_l$ is a discrete gradient operator. In this case, convex potential functions with affine behavior at infinity reduce—but do not eliminate—smoothing at edge locations [3], while nonconvex potential functions yield solutions whose gradient magnitudes are all outside a nonempty interval [18]. Within the nonconvex class, unbounded potential functions are distinguished from bounded ones in that the latter produce sharper discontinuities at the expense of increased optimization difficulty. Additional

examples of edge-preserving potential functions satisfying conditions (C1)–(C3) can be found in [19, 20, 21, 22]. It is not a coincidence that all the examples of mother potential functions given above have a quadratic behavior near zero, as this property is a consequence of (C1)–(C3) (see Lemma 2.1 in section 2.1). If a potential function satisfies (C1)–(C3) except that $\lim_{t\to 0^+} \theta^\dagger(t) = +\infty$, we can instead use the function $t \longmapsto \theta\big((\epsilon^2 + t^2)^{1/2}\big)$, where ϵ is a "small enough" positive constant. This approximation is useful, for example, for concave functions derived from $t/(1+t)$ [2] and for functions proportional to t^α with $\alpha \in [1, 2)$ [23].

In the data-fidelity term, the most commonly used potential functions are the square function $\theta_{\text{fid}}(t) = t^2$ and the identity function $\theta_{\text{fid}}(t) = t$. The square function yields the squared ℓ_2-norm of the residual; it is fully justified when there is no impulse noise and the components of the noise $\boldsymbol{\nu}$ are realizations of independent and identically distributed Gaussian random variables. The identity function yields the ℓ_1-norm of the residual and is appropriate for impulse noise [24, 25]. To fit our framework, the ℓ_1-norm must be replaced by a smooth approximation, such as that obtained with the potential function

$$\theta_{\text{fid}}(t) \;=\; \epsilon\vartheta_{\text{MS}}(t/\epsilon) \;=:\; \text{id}_\epsilon(t), \tag{1.11}$$

where $\epsilon > 0$ is small compared to the range of values covered by the components of the original signal. This technique traces back to the pioneering work of Acar and Vogel [26] on total variation regularization and was popularized by Vogel and Oman [27, 28]. If the original signal is corrupted by both Gaussian and impulse noise, we can use a potential function which behaves quadratically below a certain threshold and linearly above it. For example, Li et al. [12] use a function derived from $t^2/(1+t)$. Another possible choice is the Huber function (1.7) with its argument scaled proportionally to the standard deviation of the Gaussian noise. Note, however, that even if the current trend in image reconstruction is to use convex data-fidelity terms, our conditions on θ_{fid} are the same as those on θ_{reg}, and so θ_{fid} can be nonconvex.

1.2. Previous work. The convergence of the recurrence (1.2) is studied in [3, 4, 5, 6, 7, 8] under different requirements in addition to conditions (C1)–(C3) on the potential functions.

The results established in [3, 5, 6, 7] are restricted to the case where the objective functional Θ is strictly convex and its potential functions $\theta_1, \ldots, \theta_K$ are convex. (It is not necessary that all the potential functions be convex to ensure that Θ is strictly convex: see Proposition B.2 in Appendix B.) Among these contributions, the most general result is due to Allain, Idier, and Goussard [7], who showed that a sequence generated by (1.2) converges to the global minimizer of Θ if either condition (C′1a) or (C′1b) below is satisfied:

$$\text{(C′1a)} \quad \begin{cases} \boldsymbol{A}_1 \text{ has full column rank,} \\ \theta_1 \text{ is the square function, and } \theta_2, \ldots, \theta_K \text{ are convex;} \end{cases}$$

$$\text{(C′1b)} \quad \begin{cases} \bigcap_{k=1}^K \text{null}(\boldsymbol{A}_k) = \{\boldsymbol{0}\}, \\ \theta_1, \ldots, \theta_K \text{ are strictly convex.} \end{cases}$$

Charbonnier et al. [3] impose (C′1a) and two extra conditions:

(C′2) $\theta_2^\dagger, \ldots, \theta_K^\dagger$ are strictly decreasing and vanish at infinity;

(C′3) $\theta_2, \ldots, \theta_K$ are C^2 on $(0, +\infty)$ and C^3 near zero.

Idier [5] imposes (C′1a) and (C′2) (indeed, (C′2) is equivalent to the condition that for every $k \in [2\,..\,K]$ the composite function $\theta_k \circ \sqrt{\ }$ be strictly concave and $\theta_k(t) = o(t^2)$ as $t \longrightarrow +\infty$). Nikolova and Ng [6] assume that either (C′1a) or (C′1b) is satisfied and that (C′3) holds for *all* potential functions (meaning including θ_1).

Delaney and Bresler [4] and Robini, Zhu, and Luo [8] established global and local convergence results without convexity assumptions. They impose that $\bigcap_k \operatorname{null}(\boldsymbol{A}_k) = \{\mathbf{0}\}$ and that (C′2) and (C′3) hold for all potential functions with the class C^2 condition relaxed to twice differentiability. Their conclusions can be summarized as follows. Let $\mathcal{S}$ and $\mathcal{L}_0$ respectively denote the set of stationary points of Θ and the sublevel set of Θ at the initial height $\Theta(\boldsymbol{x}^{(0)})$, that is,

$$\mathcal{S} := \big\{\boldsymbol{x} \in \mathbb{R}^N : \nabla\Theta(\boldsymbol{x}) = \mathbf{0}\big\}, \tag{1.12}$$

$$\mathcal{L}_0 := \big\{\boldsymbol{x} \in \mathbb{R}^N : \Theta(\boldsymbol{x}) \leqslant \Theta(\boldsymbol{x}^{(0)})\big\}. \tag{1.13}$$

Let $\mathcal{X} := (\boldsymbol{x}^{(p)})_p$ be a sequence generated by the recurrence (1.2). If $\mathcal{L}_0$ is bounded, then $\mathcal{X}$ converges to $\mathcal{S}$ in terms of point-to-set distance, and so $\mathcal{X}$ converges to a stationary point if $\mathcal{S}$ is discrete. Furthermore, every isolated minimizer is an attractor: if $\boldsymbol{x}^*$ is an isolated stationary point and a local minimizer, then $\mathcal{X}$ converges to $\boldsymbol{x}^*$ provided $\boldsymbol{x}^{(0)}$ is close enough to $\boldsymbol{x}^*$.

To sum up, the minimal conditions for convergence include that the potential functions satisfy conditions (C1)–(C3) and that $\bigcap_k \operatorname{null}(\boldsymbol{A}_k) = \{\mathbf{0}\}$. If Θ is strictly convex, convergence to the global minimizer is not guaranteed if one (or more) of the potential functions is nonconvex. If Θ is nonconvex or nonstrictly convex, the potential functions must be strictly increasing and at least C^2 (which excludes, for example, the potential functions derived from the Huber function and the Tukey biweight function).

1.3. Contributions. We establish global and local convergence properties that generalize the works of Allain, Idier, and Goussard [7] and Robini, Zhu, and Luo [8] (and thus also the earlier results in [3, 4, 5, 6]). The starting point of our global convergence analysis is the application of the monotone convergence theorem of Meyer [29, Theorem 3.1] using a fixed-point interpretation of the recurrence (1.2). The local convergence analysis is built on a majorization-minimization interpretation and is initiated using the asymptotic stationarity result and the capture property established by Jacobson and Fessler [30, Theorem 4.1 and Proposition 6.3]. In addition to conditions (C1)–(C3) on the potential functions, we suppose that

$$\text{(C4)} \quad \bigcap_{k \in \mathcal{J}_+} \operatorname{null}(\boldsymbol{A}_k) = \{\mathbf{0}\}, \qquad \mathcal{J}_+ := \big\{k \in [1\,..\,K] : \theta_k \text{ is strictly increasing}\big\}.$$

This condition holds in the convex and nonconvex settings considered in [7, 8]. Contrary to the conditions in [7], some or all of the potential functions can be nonconvex, and contrary to the conditions in [8], the nonconvex potential functions can eventually be constant and of class C^1 only.

Our global convergence results require that the sublevel set $\mathcal{L}_0$ be bounded, as in [8]. This condition holds in particular when Θ is coercive (that is, $\Theta(\boldsymbol{x}) \longrightarrow +\infty$ as $\|\boldsymbol{x}\| \longrightarrow +\infty$), which is for example the case in the convex setting. We first establish that a sequence $\mathcal{X} := (\boldsymbol{x}^{(p)})_p$ generated by (1.2) converges to a stationary point if the intersections of $\mathcal{S}$ with the level sets of Θ are discrete. This result extends those in [7, 8] to a wider class of objectives which includes convex and nonconvex functionals with a single minimizer, nonconvex functionals with isolated stationary points, and even some nonconvex functionals with nonisolated stationary points. Our second result provides weaker convergence guarantees for the remaining case where a level set contains a bounded infinite set of stationary points: $\mathcal{X}$ converges to the boundary of $\mathcal{S}$ in terms of point-to-set distance, and the gradient sequence $(\nabla\Theta(\boldsymbol{x}^{(p)}))_p$ converges to zero.

Regarding local convergence, we show that the basin of attraction of an isolated minimizer $\boldsymbol{x}^*$ includes the bounded open sets that contain no stationary point other than $\boldsymbol{x}^*$ and whose boundaries are flat—we call such sets *cups*. We also show that local convergence is still guaranteed under a condition weaker than (C4), and we conclude with a characterization of unions of cups which leads to a formal water-flooding interpretation of the basins of attraction.

1.4. Organization. In section 2, we describe the fixed-point construction of the half-quadratic iterative scheme (1.2) and we discuss other fruitful interpretations of this algorithm. Section 3 is devoted to global convergence: we first establish monotone convergence, and then we derive sufficient conditions for convergence to stationary points with a distinction between convergence in norm and point-to-set convergence. Local convergence is studied in section 4, where we characterize the basins of attraction of isolated minimizers and discuss the water-flooding analogy. Section 5 concludes the paper with detailed image restoration experiments which illustrate the behavior of the algorithm for various nonconvex objectives. (We introduce our notation as the paper progresses. The main symbols we use are listed in Appendix A for convenience.)

2. Half-quadratic optimization.

2.1. Basic properties of potential functions. The following two lemmas state basic properties of the potential functions satisfying our assumptions. These properties are used in the construction of the half-quadratic algorithm discussed next.

Lemma 2.1. *Let $\theta : \mathbb{R}_+ \longrightarrow \mathbb{R}$ be a function satisfying conditions* (C1)–(C3)*:*

(i) *θ is right-differentiable at zero, and $\theta'_+(0) = \lim_{t\to 0^+} \theta'(t) = 0$.*

(ii) *If θ is not strictly increasing, then there is real number $\tau > 0$ such that θ is strictly increasing on $[0, \tau)$ and constant on $[\tau, +\infty)$.*

(iii) *θ' is right-differentiable at zero, and $\theta''_+(0) = \lim_{t\to 0^+} \theta^\dagger(t) > 0$.*

Proof. (i) Suppose θ' does not tend to zero as $t \longrightarrow 0^+$. Then there are a constant $c > 0$ and a sequence $(t_n)_n$ in $(0, +\infty)$ such that $t_n \longrightarrow 0^+$ and $\theta'(t_n) \geqslant c$ for all n, which contradicts the condition that $\theta^\dagger$ be bounded. So $\lim_{t\to 0^+} \theta'(t) = 0$. Since θ is differentiable on $(0, +\infty)$ and continuous at zero, it follows that θ is right-differentiable at zero and $\theta'_+(0) = 0$.

(ii) Assume θ is not strictly increasing. Then its derivative is eventually zero (because $\theta^\dagger$ is nonnegative and decreasing), and thus the set

$$\mathcal{T} := \big\{t > 0 : \theta' \text{ is zero on } [t, +\infty)\big\}$$

is nonempty. Let $\tau := \inf \mathcal{T}$. We have $\tau > 0$ since θ is nonconstant. By definition of the infimum, the set $(0,\tau) \cap \mathcal{T}$ is empty and θ' is zero on $(\tau, +\infty)$. Furthermore, θ' must be positive on $(0,\tau)$, for otherwise $(0,\tau) \cap \mathcal{T}$ is nonempty.

(iii) $\theta^\dagger$ is decreasing and bounded, and it is positive near zero by property (ii). So $\theta^\dagger$ has a positive right-limit at zero, and since $\theta'_+(0) = 0$ by property (i),

$$\lim_{t\to 0^+} \theta^\dagger(t) \;=\; \lim_{t\to 0^+} \frac{\theta'(t) - \theta'_+(0)}{t} \;=\; \theta''_+(0). \qquad \blacksquare$$

Notation. By Lemma 2.1(iii), the function $\theta^\dagger$ has a continuous extension which we denote by $\theta^\ddagger$, that is,

$$\theta^\ddagger(t) \;:=\; \begin{cases} t^{-1}\theta'(t) & \text{if } t > 0, \\ \theta''_+(0) & \text{if } t = 0. \end{cases} \tag{2.1}$$

Lemma 2.2. *Let $\theta : \mathbb{R}_+ \longrightarrow \mathbb{R}$ be a function satisfying conditions* (C1)–(C3). *The composite function $\eta := \theta \circ \sqrt{\ }$ is C^1 with derivative $\eta' = \frac{1}{2}\,\theta^\ddagger \circ \sqrt{\ }$.*

Proof. Clearly, η is C^1 on $(0,+\infty)$ and $\eta' = \frac{1}{2}\,\theta^\dagger \circ \sqrt{\ }$ on this interval. Hence, since η is continuous at zero, it remains to check that $\lim_{t\to 0^+} \eta'(t) = \frac{1}{2}\theta''_+(0)$, which follows directly from Lemma 2.1(iii). $\blacksquare$

2.2. Construction of the algorithm. The half-quadratic iterative scheme (1.2) can be interpreted as a sequence of fixed-point iterations toward stationary points of the objective functional. We give here the details of this construction, and we show that the resulting algorithm is well-defined under our conditions.

Notation. We let $\boldsymbol{A}$ and $\boldsymbol{a}$ be the vertical concatenations of the matrices $\boldsymbol{A}_k \in \mathbb{R}^{N_k \times N}$ and of the vectors $\boldsymbol{a}_k \in \mathbb{R}^{N_k}$, that is,

$$\boldsymbol{A} \;:=\; \begin{pmatrix} \boldsymbol{A}_1 \\ \vdots \\ \boldsymbol{A}_K \end{pmatrix} \in \mathbb{R}^{N' \times N} \quad \text{and} \quad \boldsymbol{a} \;:=\; \begin{pmatrix} \boldsymbol{a}_1 \\ \vdots \\ \boldsymbol{a}_K \end{pmatrix} \in \mathbb{R}^{N'},$$

where $N' := \sum_{k=1}^K N_k$. For every $\boldsymbol{x} \in \mathbb{R}^N$, we put

$$\varepsilon_k(\boldsymbol{x}) \;:=\; \theta_k^\ddagger\big(\|\boldsymbol{A}_k\boldsymbol{x} - \boldsymbol{a}_k\|\big), \qquad k \in [1\,..\,K], \tag{2.2}$$

and we define the $N' \times N'$ nonnegative diagonal matrix

$$\boldsymbol{E}(\boldsymbol{x}) \;:=\; \operatorname{diag}\big(\varepsilon_1(\boldsymbol{x})\boldsymbol{I}_{N_1}, \ldots, \varepsilon_K(\boldsymbol{x})\boldsymbol{I}_{N_K}\big), \tag{2.3}$$

where $\boldsymbol{I}_n$ denotes the identity matrix of order n.

Lemma 2.3. *The objective functional Θ is C^1, and its gradient is given by*

$$\nabla\Theta(\boldsymbol{x}) \;=\; \boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})(\boldsymbol{A}\boldsymbol{x} - \boldsymbol{a}). \tag{2.4}$$

Proof. Set $\eta_k := \theta_k \circ \sqrt{\ }$ for all k. Then $\Theta(\boldsymbol{x}) = \sum_{k=1}^{K} \eta_k\big(\|\boldsymbol{A}_k\boldsymbol{x} - \boldsymbol{a}_k\|^2\big)$, and so

$$\nabla\Theta(\boldsymbol{x}) \;=\; 2\sum_{k=1}^{K} \eta_k'\big(\|\boldsymbol{A}_k\boldsymbol{x} - \boldsymbol{a}_k\|^2\big)\boldsymbol{A}_k^T(\boldsymbol{A}_k\boldsymbol{x} - \boldsymbol{a}_k).$$

Using Lemma 2.2 and notation (2.2), we have $\eta_k'\big(\|\boldsymbol{A}_k\boldsymbol{x}-\boldsymbol{a}_k\|^2\big) = \frac{1}{2}\varepsilon_k(\boldsymbol{x})$, and hence (2.4) follows. Furthermore, since the functions $\theta_k^{\ddagger}$ are continuous, the map $\boldsymbol{x} \longmapsto \boldsymbol{E}(\boldsymbol{x})$ is continuous, and thus so is $\nabla\Theta$. ∎

Assume for now that the matrix $\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})\boldsymbol{A}$ is positive definite for all $\boldsymbol{x} \in \mathbb{R}^N$. We deduce from (2.4) that the stationary points of Θ can be approached by the following fixed-point algorithm.

Algorithm 1. *Given a starting point $\boldsymbol{x}^{(0)} \in \mathbb{R}^N$, generate the sequence $(\boldsymbol{x}^{(p)})_{p\in\mathbb{N}}$ via the recurrence relation*

$$\boldsymbol{x}^{(p+1)} \;:=\; (\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x}^{(p)})\boldsymbol{A})^{-1}\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x}^{(p)})\boldsymbol{a}, \tag{2.5}$$

or equivalently,

$$\boldsymbol{x}^{(p+1)} \;:=\; \operatorname*{arg\,min}_{\boldsymbol{y}\in\mathbb{R}^N} \Theta^{\ddagger}(\boldsymbol{x}^{(p)}, \boldsymbol{y}), \tag{2.6a}$$

$$\Theta^{\ddagger}(\boldsymbol{x}, \boldsymbol{y}) \;:=\; \sum_{k=1}^{K} \varepsilon_k(\boldsymbol{x})\,\|\boldsymbol{A}_k\boldsymbol{y} - \boldsymbol{a}_k\|^2. \tag{2.6b}$$

Each iteration consists of two steps: first compute the weighting coefficients $\varepsilon_1(\boldsymbol{x}^{(p)}), \ldots, \varepsilon_K(\boldsymbol{x}^{(p)})$ according to (2.2), and then minimize the positive definite quadratic functional $\boldsymbol{y} \longmapsto \Theta^{\ddagger}(\boldsymbol{x}^{(p)}, \boldsymbol{y})$. The second step can be performed efficiently by either a direct method (such as Cholesky decomposition) or an iterative method (such as conjugate gradient), the choice depending on the dimension of the solution space and on the sparsity ratio of the matrix $\boldsymbol{A}$.

Proposition 2.4 makes it clear that Algorithm 1 is well-defined under our conditions, and Proposition 2.5 shows that condition (C4) is superfluous if Θ is coercive.

Proposition 2.4. *Suppose the potential functions θ_k involved in the definition of the matrices $\boldsymbol{E}(\boldsymbol{x})$ satisfy conditions* (C1)–(C3). *Then $\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})\boldsymbol{A}$ is positive definite for all $\boldsymbol{x} \in \mathbb{R}^N$ if and only if condition* (C4) *holds.*

Proof. Let $\boldsymbol{x} \in \mathbb{R}^N$. For every $\boldsymbol{y} \in \mathbb{R}^N$,

$$\boldsymbol{y}^T\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})\boldsymbol{A}\boldsymbol{y} \;=\; \sum_{k=1}^{K} \varepsilon_k(\boldsymbol{x})\|\boldsymbol{A}_k\boldsymbol{y}\|^2 \;\geqslant\; 0$$

with equality if and only if

$$\boldsymbol{y} \;\in\; \bigcap_{k\in\mathcal{J}_{\boldsymbol{\varepsilon}}(\boldsymbol{x})} \operatorname{null}(\boldsymbol{A}_k), \qquad \mathcal{J}_{\boldsymbol{\varepsilon}}(\boldsymbol{x}) := \big\{k \in [1\,..\,K] : \varepsilon_k(\boldsymbol{x}) > 0\big\}.$$

Therefore, since $\mathcal{J}_+ \subseteq \mathcal{J}_{\boldsymbol{\varepsilon}}(\boldsymbol{x})$, the matrix $\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})\boldsymbol{A}$ is positive definite if (C4) is satisfied.

Now assume $\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})\boldsymbol{A}$ is always positive definite, that is,

$$\bigcap_{k\in\mathcal{J}_{\boldsymbol{\varepsilon}}(\boldsymbol{x})} \operatorname{null}(\boldsymbol{A}_k) = \{\boldsymbol{0}\} \quad \text{for all } \boldsymbol{x}\in\mathbb{R}^N. \tag{2.7}$$

Seeking a contradiction, suppose there is a nonzero vector $\boldsymbol{y}$ such that

$$\mathcal{J}_+ \subseteq \big\{k\in[1\,..\,K] : \boldsymbol{y}\in\operatorname{null}(\boldsymbol{A}_k)\big\} =: \mathcal{J}_0(\boldsymbol{y}).$$

Then the set $[1\,..\,K]\setminus\mathcal{J}_0(\boldsymbol{y})$ is nonempty (for otherwise $\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})\boldsymbol{A}$ is never positive definite), and for every index k in this set, θ_k is eventually constant (by Lemma 2.1(ii)) and $\lim_{\alpha\to+\infty}\|\alpha\boldsymbol{A}_k\boldsymbol{y}-\boldsymbol{a}_k\| = +\infty$. Hence, for each k such that $\boldsymbol{y}\notin\operatorname{null}(\boldsymbol{A}_k)$, there is an $\alpha_k>0$ such that $\varepsilon_k(\alpha\boldsymbol{y})=0$ for all $\alpha\geqslant\alpha_k$. In other words, $\mathcal{J}_{\boldsymbol{\varepsilon}}(\alpha\boldsymbol{y})\subseteq\mathcal{J}_0(\boldsymbol{y})$ for α sufficiently large. Therefore, there exists some $\alpha>0$ such that

$$\boldsymbol{y}\in\bigcap_{k\in\mathcal{J}_{\boldsymbol{\varepsilon}}(\alpha\boldsymbol{y})}\operatorname{null}(\boldsymbol{A}_k),$$

which contradicts (2.7). ∎

Proposition 2.5. *The objective functional Θ is coercive if and only if*

$$\bigcap_{k\in\mathcal{J}_\infty}\operatorname{null}(\boldsymbol{A}_k)=\{\boldsymbol{0}\},\qquad \mathcal{J}_\infty := \{k\in[1\,..\,K] : \lim_{t\to+\infty}\theta_k(t)=+\infty\}. \tag{2.8}$$

In particular, condition (C4) *holds if Θ is coercive.*

Proof. The equivalence is trivial if $\mathcal{J}_\infty$ is empty, for in this case Θ is bounded and $\bigcap_{k\in\mathcal{J}_\infty}\operatorname{null}(\boldsymbol{A}_k)=\mathbb{R}^N$. So we assume $\mathcal{J}_\infty$ is nonempty.

Suppose there is a nonzero vector $\boldsymbol{x}\in\bigcap_{k\in\mathcal{J}_\infty}\operatorname{null}(\boldsymbol{A}_k)$. Then, for every $\alpha\in\mathbb{R}$,

$$\Theta(\alpha\boldsymbol{x}) = \sum_{k\in\mathcal{J}_\infty}\theta_k\big(\|\boldsymbol{a}_k\|\big) + \sum_{k\notin\mathcal{J}_\infty}\theta_k\big(\|\alpha\boldsymbol{A}_k\boldsymbol{x}-\boldsymbol{a}_k\|\big),$$

and thus

$$\lim_{\alpha\to+\infty}\Theta(\alpha\boldsymbol{x}) \leqslant \sum_{k\in\mathcal{J}_\infty}\theta_k\big(\|\boldsymbol{a}_k\|\big) + \sum_{k\notin\mathcal{J}_\infty}\sup_{\mathbb{R}_+}\theta_k < +\infty. \tag{2.9}$$

Therefore, Θ is not coercive. Now, for the converse implication, suppose Θ is not coercive; that is, there exists a sequence $(\boldsymbol{y}^{(p)})_p$ such that $\|\boldsymbol{y}^{(p)}\|\longrightarrow+\infty$ and $\lim_{p\to\infty}\Theta(\boldsymbol{y}^{(p)})<+\infty$. Let f be the quadratic form on $\mathbb{R}^N$ defined by $f(\boldsymbol{y}) := \sum_{k\in\mathcal{J}_\infty}\|\boldsymbol{A}_k\boldsymbol{y}\|^2$. We have

$$\begin{aligned}
\lim_{p\to\infty}\Theta(\boldsymbol{y}^{(p)})<+\infty &\iff \lim_{p\to\infty}\theta_k\big(\|\boldsymbol{A}_k\boldsymbol{y}^{(p)}-\boldsymbol{a}_k\|\big)<+\infty \quad \text{for all } k\in\mathcal{J}_\infty\\
&\iff \sup\big\{\|\boldsymbol{A}_k\boldsymbol{y}^{(p)}\| : k\in\mathcal{J}_\infty,\ p\in\mathbb{N}\big\}<+\infty\\
&\iff \sup_{p\in\mathbb{N}} f(\boldsymbol{y}^{(p)})<+\infty.
\end{aligned}$$

It follows that $\bigcap_{k\in\mathcal{J}_\infty} \operatorname{null}(\boldsymbol{A}_k) \neq \{\boldsymbol{0}\}$, for otherwise f is positive definite and so $f(\boldsymbol{y}^{(p)}) \longrightarrow +\infty$.

Finally, condition (C4) holds if Θ is coercive because $\mathcal{J}_\infty \subseteq \mathcal{J}_+$ by Lemma 2.1(ii). ■

Remark 1. A corollary to Proposition 2.5 is that condition (C4) is not needed if we add a suitably chosen Tikhonov penalty to Θ. Indeed, given an invertible matrix $\boldsymbol{A}_0 \in \mathbb{R}^{N\times N}$, the quadratic form $\boldsymbol{x} \longmapsto \|\boldsymbol{A}_0\boldsymbol{x}\|^2$ is coercive, and thus so is the functional $\boldsymbol{x} \longmapsto \Theta(\boldsymbol{x}) + \|\boldsymbol{A}_0\boldsymbol{x}\|^2$. This technique is used for example in [31], where $\boldsymbol{A}_0$ is proportional to the identity matrix.

2.3. Interpretations. Throughout this section, we assume that condition (C4) holds. For every $\boldsymbol{x} \in \mathbb{R}^N$, the gradient of the quadratic functional $\Theta^\ddagger(\boldsymbol{x},\cdot) : \boldsymbol{y} \longmapsto \Theta^\ddagger(\boldsymbol{x},\boldsymbol{y})$ defined by (2.6a) is

$$\nabla\Theta^\ddagger(\boldsymbol{x},\cdot)(\boldsymbol{y}) \;=\; 2\big(\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})\boldsymbol{A}(\boldsymbol{y}-\boldsymbol{x}) + \nabla\Theta(\boldsymbol{x})\big), \tag{2.10}$$

and thus Algorithm 1 is equivalently defined by the iterative scheme

$$\boldsymbol{x}^{(p+1)} := \boldsymbol{\Phi}(\boldsymbol{x}^{(p)}), \tag{2.11a}$$

$$\boldsymbol{\Phi}(\boldsymbol{x}) := \boldsymbol{x} - (\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})\boldsymbol{A})^{-1}\nabla\Theta(\boldsymbol{x}). \tag{2.11b}$$

Hence Algorithm 1 is a quasi-Newton optimization method where the Hessian matrix of Θ at $\boldsymbol{x}$ is approximated by $\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})\boldsymbol{A}$.

Algorithm 1 can also be viewed as an alternating minimization process and a majorization-minimization scheme. These two interpretations are discussed below; the former shows that Algorithm 1 generalizes the half-quadratic regularization approach inspired by the construction of Geman and Reynolds [2], and the latter is used in section 4 to study local convergence.

2.3.1. Alternating minimization. Let $\theta : \mathbb{R}_+ \longrightarrow \mathbb{R}$ be a function satisfying conditions (C1)–(C3), and suppose that

$$\theta(t) \propto t^2 \quad \text{or} \quad \theta(t) = o(t^2). \tag{2.12}$$

As shown in [9], there is a decreasing convex function $\zeta : \mathbb{R}_+ \longrightarrow (-\infty,+\infty]$ such that for every $t \in \mathbb{R}_+$,

$$\theta(t) \;=\; \min_{u\in\mathbb{R}_+}\left(\frac{t^2u}{2} + \zeta(u)\right) \;=\; \frac{t^2\theta^\ddagger(t)}{2} + \zeta(\theta^\ddagger(t)).$$

Therefore, if the potential functions in the objective functional satisfy (2.12), then

$$\Theta(\boldsymbol{x}) \;=\; \min_{\boldsymbol{u}=(u_1,\ldots,u_K)\in\mathbb{R}_+^K} \Theta^\star(\boldsymbol{x},\boldsymbol{u}),$$

$$\Theta^\star(\boldsymbol{x},\boldsymbol{u}) := \sum_{k=1}^{K}\left(\frac{u_k}{2}\|\boldsymbol{A}_k\boldsymbol{x} - \boldsymbol{a}_k\|^2 + \zeta_k(u_k)\right),$$

and the minimum is attained when $\boldsymbol{u} = (\varepsilon_1(\boldsymbol{x}), \dots, \varepsilon_K(\boldsymbol{x}))$. It follows that the iterative scheme (2.6) of Algorithm 1 is equivalent to the alternating minimization process

$$\begin{aligned} \boldsymbol{u}^{(p+1)} &:= \operatorname*{arg\,min}_{\boldsymbol{u}\in\mathbb{R}_+^K} \Theta^\star(\boldsymbol{x}^{(p)}, \boldsymbol{u}), \\ \boldsymbol{x}^{(p+1)} &:= \operatorname*{arg\,min}_{\boldsymbol{x}\in\mathbb{R}^N} \Theta^\star(\boldsymbol{x}, \boldsymbol{u}^{(p+1)}). \end{aligned}$$

This process is half-quadratic, as the augmented objective functional $\Theta^\star$ is quadratic with respect to the primal variable $\boldsymbol{x}$. Following the terminology of Nikolova and Ng [6], we call it *multiplicative* because the quadratic terms $\boldsymbol{x} \longmapsto \|\boldsymbol{A}_k\boldsymbol{x} - \boldsymbol{a}_k\|^2$, $k \in [1\,..\,K]$, are multiplied by the components of the dual variable $\boldsymbol{u}$. Hence, Algorithm 1 includes the multiplicative half-quadratic algorithms studied in [3, 4, 5, 6, 7, 8] as special cases. Note that some authors call half-quadratic any method based on an augmented formulation involving a quadratic term to simplify the original optimization problem (this is for example the case of the robust sparse representation approach proposed in [32], where the augmented objective is a nondifferentiable convex function of the primal variable when the dual variable is fixed). It is important to be clear that our focus here is only on *fully* half-quadratic algorithms, that is, those in which the augmented objective is quadratic with respect to the primal variable.

2.3.2. Majorization-minimization. By the definition in [33], a majorization-minimization algorithm for minimizing Θ consists of the iterations

$$\boldsymbol{x}^{(p+1)} := \operatorname*{arg\,min}_{\boldsymbol{y}\in\mathbb{R}^N} \overline{\Theta}(\boldsymbol{x}^{(p)}, \boldsymbol{y}), \tag{2.13}$$

where $\overline{\Theta}$ is such that for all $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^N$,

$$\overline{\Theta}(\boldsymbol{x}, \boldsymbol{y}) \geqslant \Theta(\boldsymbol{y}) \text{ with equality if } \boldsymbol{y} = \boldsymbol{x} \tag{2.14}$$

(it is implicitly assumed that for every $\boldsymbol{x}$ the functional $\overline{\Theta}(\boldsymbol{x}, \cdot)$ has a unique global minimizer). We call $\overline{\Theta}$ a *surrogate generator*.

Proposition 2.6. *Suppose condition* (C4) *holds. Then Algorithm 1 is a majorization-minimization algorithm with surrogate generator* $\overline{\Theta}$ *given by*

$$\overline{\Theta}(\boldsymbol{x}, \boldsymbol{y}) = \Theta(\boldsymbol{x}) + (\boldsymbol{y} - \boldsymbol{x})^T \nabla\Theta(\boldsymbol{x}) + \frac{1}{2}(\boldsymbol{y} - \boldsymbol{x})^T \boldsymbol{A}^T \boldsymbol{E}(\boldsymbol{x}) \boldsymbol{A}(\boldsymbol{y} - \boldsymbol{x}). \tag{2.15}$$

Proof. Let $\overline{\Theta}$ be defined by (2.15). Then $\nabla\overline{\Theta}(\boldsymbol{x}, \cdot) = \frac{1}{2}\nabla\Theta^\ddagger(\boldsymbol{x}, \cdot)$, where $\nabla\Theta^\ddagger(\boldsymbol{x}, \cdot)$ is given in (2.10), and thus $\overline{\Theta}(\boldsymbol{x}, \cdot)$ and $\Theta^\ddagger(\boldsymbol{x}, \cdot)$ have the same global minimizer if $\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})\boldsymbol{A}$ is positive definite. So Algorithm 1 is equivalently defined by the iterative scheme (2.13) if condition (C4) holds. It remains to show that $\overline{\Theta}(\boldsymbol{x}, \cdot)$ majorizes Θ (we readily have $\overline{\Theta}(\boldsymbol{x}, \boldsymbol{x}) = \Theta(\boldsymbol{x})$). We can write $\overline{\Theta}$ in the form

$$\overline{\Theta}(\boldsymbol{x}, \boldsymbol{y}) = \Theta(\boldsymbol{x}) + \langle \boldsymbol{A}(\boldsymbol{y} - \boldsymbol{x}), \boldsymbol{A}\boldsymbol{x} - \boldsymbol{a} \rangle_{\boldsymbol{E}(\boldsymbol{x})} + \frac{1}{2}\|\boldsymbol{A}(\boldsymbol{y} - \boldsymbol{x})\|^2_{\boldsymbol{E}(\boldsymbol{x})}, \tag{2.16}$$

where $\langle \cdot, \cdot \rangle_{\boldsymbol{E}(\boldsymbol{x})}$ is the symmetric bilinear form on $\mathbb{R}^{N'}$ given by

$$\langle \boldsymbol{v}, \boldsymbol{w} \rangle_{\boldsymbol{E}(\boldsymbol{x})} := \boldsymbol{v}^T \boldsymbol{E}(\boldsymbol{x}) \boldsymbol{w}$$

and $\|\cdot\|_{\boldsymbol{E}(\boldsymbol{x})}$ is the associated seminorm, that is, $\|\boldsymbol{v}\|_{\boldsymbol{E}(\boldsymbol{x})} := \langle \boldsymbol{v}, \boldsymbol{v}\rangle_{\boldsymbol{E}(\boldsymbol{x})}^{1/2}$. It is easy to check that

$$\|\boldsymbol{w}-\boldsymbol{v}\|_{\boldsymbol{E}(\boldsymbol{x})}^2 = \|\boldsymbol{w}\|_{\boldsymbol{E}(\boldsymbol{x})}^2 - \|\boldsymbol{v}\|_{\boldsymbol{E}(\boldsymbol{x})}^2 - 2\langle \boldsymbol{w}-\boldsymbol{v}, \boldsymbol{v}\rangle_{\boldsymbol{E}(\boldsymbol{x})}.$$

Setting $\boldsymbol{w} = \boldsymbol{A}\boldsymbol{y} - \boldsymbol{a}$ and $\boldsymbol{v} = \boldsymbol{A}\boldsymbol{x} - \boldsymbol{a}$, and substituting into (2.16), we obtain

$$\overline{\Theta}(\boldsymbol{x},\boldsymbol{y}) = \Theta(\boldsymbol{x}) + \frac{1}{2}\big(\|\boldsymbol{A}\boldsymbol{y}-\boldsymbol{a}\|_{\boldsymbol{E}(\boldsymbol{x})}^2 - \|\boldsymbol{A}\boldsymbol{x}-\boldsymbol{a}\|_{\boldsymbol{E}(\boldsymbol{x})}^2\big). \tag{2.17}$$

Consequently,

$$\overline{\Theta}(\boldsymbol{x},\boldsymbol{y}) - \Theta(\boldsymbol{y}) = \sum_{k=1}^{K}\Big[\theta_k\big(\|\boldsymbol{A}_k\boldsymbol{x}-\boldsymbol{a}_k\|\big) - \theta_k\big(\|\boldsymbol{A}_k\boldsymbol{y}-\boldsymbol{a}_k\|\big) + \frac{1}{2}\,\theta_k^{\ddagger}\big(\|\boldsymbol{A}_k\boldsymbol{x}-\boldsymbol{a}_k\|\big)\big(\|\boldsymbol{A}_k\boldsymbol{y}-\boldsymbol{a}_k\|^2 - \|\boldsymbol{A}_k\boldsymbol{x}-\boldsymbol{a}_k\|^2\big)\Big].$$

By Lemma 2.2, each term in the above sum is of the form

$$\theta(t) - \theta(u) + \frac{1}{2}\,\theta^{\ddagger}(t)(u^2-t^2) = \eta(t^2) - \eta(u^2) + \eta'(t^2)(u^2-t^2),$$

where $\eta := \theta\circ\sqrt{\ }$ is concave since $\theta^{\ddagger}$ is decreasing. Therefore, $\overline{\Theta}(\boldsymbol{x},\boldsymbol{y}) - \Theta(\boldsymbol{y})$ is a sum of nonnegative numbers, and hence $\overline{\Theta}(\boldsymbol{x},\boldsymbol{y}) \geqslant \Theta(\boldsymbol{y})$. ∎

3. Global convergence to stationary points. It follows from Proposition 2.6 that Algorithm 1 inherits the properties of majorization-minimization algorithms. In particular, we deduce Theorem 3.2 below from Theorem 4.4 of Jacobson and Fessler in [30]. We first recall the definitions of an isolated point and a discrete set.

Definition 3.1. *Let $\mathcal{G}$ be a nonempty subset of $\mathbb{R}^N$. A point $\boldsymbol{x}\in\mathcal{G}$ is said to be* isolated *in $\mathcal{G}$ (or an isolated point of $\mathcal{G}$) if there is an $\alpha>0$ such that $\|\boldsymbol{y}-\boldsymbol{x}\|\geqslant\alpha$ for all $\boldsymbol{y}\in\mathcal{G}\setminus\{\boldsymbol{x}\}$. The set $\mathcal{G}$ is called* discrete *if all its points are isolated in $\mathcal{G}$.*

Remark 2. A point $\boldsymbol{x}$ is isolated in the singleton $\{\boldsymbol{x}\}$ (indeed, the proposition "$\forall\boldsymbol{y}\in\emptyset$, $\mathfrak{P}(\boldsymbol{y})$" is true whatever the propositional function $\mathfrak{P}$). Thus singletons are discrete.

Theorem 3.2. *Let $\mathcal{S}$ be the set of stationary points of Θ, let $\mathcal{X} := (\boldsymbol{x}^{(p)})_{p\in\mathbb{N}}$ be a sequence generated by Algorithm 1 under condition* (C4)*, and let $\mathcal{L}_0$ be the sublevel set of Θ at the initial height $\Theta(\boldsymbol{x}^{(0)})$. If $\mathcal{S}$ is discrete, $\mathcal{L}_0$ is bounded, and $\|\boldsymbol{x}^{(p+1)}-\boldsymbol{x}^{(p)}\|\longrightarrow 0$, then $\mathcal{X}$ converges to a point in $\mathcal{S}$.*

Proof. By Proposition 2.6, Algorithm 1 is a majorization-minimization algorithm with surrogate generator given by (2.15). This generator satisfies the conditions of Theorem 4.4 in [30], from which we deduce that $\mathcal{X}$ converges to a stationary point under the stated conditions on $\mathcal{S}$, $\mathcal{X}$, and $\mathcal{L}_0$. ∎

In the following, we use the monotone convergence theorem of Meyer in [29] to show that $\|\boldsymbol{x}^{(p+1)}-\boldsymbol{x}^{(p)}\|\longrightarrow 0$ under the other conditions of Theorem 3.2. We do not actually need Meyer's result to do so, but this approach allows us to relax the condition that $\mathcal{S}$ be discrete, and even to remove it if convergence to a single stationary point is weakened to convergence to $\mathcal{S}$. (The remaining condition that $\mathcal{L}_0$ be bounded holds in particular if Θ is coercive, in which case we can drop condition (C4) according to Proposition 2.5.)

3.1. Monotone convergence. This section paves the way to study global convergence to stationary points. Our results are built on Theorem 3.4, which is a special instance of Theorem 3.1 in [29].

Definition 3.3. *Let $\boldsymbol{\Psi}$ be a map from $\mathcal{G} \subseteq \mathbb{R}^N$ to itself.*

(i) *A sequence $(\boldsymbol{x}^{(p)})_{p\in\mathbb{N}}$ in $\mathbb{R}^N$ is said to be* generated by $\boldsymbol{\Psi}$ *if $\boldsymbol{x}^{(0)} \in \mathcal{G}$ and $\boldsymbol{x}^{(p+1)} = \boldsymbol{\Psi}(\boldsymbol{x}^{(p)})$ for all $p \in \mathbb{N}$.*

(ii) *Let $\mathcal{F}_{\boldsymbol{\Psi}}$ denote the set of fixed points of $\boldsymbol{\Psi}$. The map $\boldsymbol{\Psi}$ is said to be* strictly monotonic *with respect to a functional $\Xi : \mathcal{G} \longrightarrow \mathbb{R}$ if*

$$(\Xi \circ \boldsymbol{\Psi})(\boldsymbol{x}) < \Xi(\boldsymbol{x}) \quad \text{for all } \boldsymbol{x} \in \mathcal{G} \setminus \mathcal{F}_{\boldsymbol{\Psi}}.$$

Theorem 3.4 (Meyer [29]). *Let $\mathcal{G}$ be a closed subset of $\mathbb{R}^N$, and let $\boldsymbol{\Psi} : \mathcal{G} \longrightarrow \mathcal{G}$ be strictly monotonic with respect to Ξ. Let $\mathcal{X} := (\boldsymbol{x}^{(p)})_p$ be a sequence generated by $\boldsymbol{\Psi}$, and denote by $\mathcal{C}_{\mathcal{X}}$ the set of cluster points of $\mathcal{X}$. If $\boldsymbol{\Psi}$ and Ξ are continuous and $\mathcal{X}$ is bounded, then*

(i) $\emptyset \neq \mathcal{C}_{\mathcal{X}} \subseteq \mathcal{F}_{\boldsymbol{\Psi}}$,

(ii) $\|\boldsymbol{x}^{(p+1)} - \boldsymbol{x}^{(p)}\| \longrightarrow 0$, *and*

(iii) $(\Xi(\boldsymbol{x}^{(p)}))_p$ *decreases to* $\Xi(\boldsymbol{x}^*)$ *for some* $\boldsymbol{x}^* \in \mathcal{C}_{\mathcal{X}}$.

Recall that Algorithm 1 consists in iteratively applying the map $\boldsymbol{\Phi} : \mathbb{R}^N \longrightarrow \mathbb{R}^N$ defined in (2.11). The following lemma is used in the proof of Theorem 3.7, where we show monotone convergence to $\mathcal{S}$ by applying Theorem 3.4 to restrictions of $\boldsymbol{\Phi}$ and Θ.

Lemma 3.5. *Suppose condition* (C4) *holds.*

(i) $\boldsymbol{\Phi}$ *is continuous.*

(ii) $\boldsymbol{\Phi}$ *is strictly monotonic with respect to* Θ.

(iii) $\mathcal{F}_{\boldsymbol{\Phi}} = \mathcal{S}$ *(that is, the fixed points of* $\boldsymbol{\Phi}$ *are the stationary points of* Θ*).*

Proof. (i) The matrix inverse function is continuous at every point that represents an invertible matrix (indeed, matrix groups are Lie groups), and so the map $\boldsymbol{x} \longmapsto (\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})\boldsymbol{A})^{-1}$ is continuous. Thus $\boldsymbol{\Phi}$ is continuous since Θ is C^1.

(ii) Let $\boldsymbol{x} \in \mathbb{R}^N \setminus \mathcal{F}_{\boldsymbol{\Phi}}$. Using (2.14) and (2.17), we have

$$\begin{aligned}(\Theta \circ \boldsymbol{\Phi})(x) - \Theta(x) &< \overline{\Theta}(\boldsymbol{x}, \boldsymbol{\Phi}(\boldsymbol{x})) - \Theta(x) \\ &= \frac{1}{2}\big(\|\boldsymbol{A}\boldsymbol{\Phi}(\boldsymbol{x}) - \boldsymbol{a}\|^2_{\boldsymbol{E}(\boldsymbol{x})} - \|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{a}\|^2_{\boldsymbol{E}(\boldsymbol{x})}\big).\end{aligned}$$

Therefore, by the definition of $\Theta^\ddagger$ in (2.6a),

$$(\Theta \circ \boldsymbol{\Phi})(x) - \Theta(x) < \frac{1}{2}\big(\Theta^\ddagger(\boldsymbol{x}, \boldsymbol{\Phi}(\boldsymbol{x})) - \Theta^\ddagger(\boldsymbol{x}, \boldsymbol{x})\big).$$

Since $\boldsymbol{\Phi}(\boldsymbol{x}) = \arg\min\{\Theta^\ddagger(\boldsymbol{x}, \boldsymbol{y}) : \boldsymbol{y} \in \mathbb{R}^N\}$, we deduce that $(\Theta \circ \boldsymbol{\Phi})(x) < \Theta(x)$. So $\boldsymbol{\Phi}$ is strictly monotonic with respect to Θ.

(iii) From the definition of $\boldsymbol{\Phi}$, and since $\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})\boldsymbol{A}$ is positive definite for all $\boldsymbol{x}$ (by Proposition 2.4), we have

$$\begin{aligned}\boldsymbol{\Phi}(\boldsymbol{x}) = \boldsymbol{x} &\iff (\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})\boldsymbol{A})^{-1}\nabla\Theta(\boldsymbol{x}) = \boldsymbol{0}\\ &\iff \nabla\Theta(\boldsymbol{x}) = \boldsymbol{0},\end{aligned}$$

which shows that $\mathcal{F}_{\boldsymbol{\Phi}} = \mathcal{S}$. ∎

Definition 3.6. *A nonempty subset of $\mathbb{R}^N$ is called a* continuum *if it is compact and connected. A continuum $\mathcal{C} \subseteq \mathbb{R}^N$ is said to be* flat *if Θ is constant on $\mathcal{C}$. (In particular, a singleton is a flat continuum.)*

Remark 3. Some authors define a continuum to be a closed connected set. The above definition of a continuum is from [34].

Theorem 3.7. *Let $\mathcal{S}$, $\mathcal{X}$, and $\mathcal{L}_0$ be as stated in Theorem 3.2, and suppose $\mathcal{L}_0$ is bounded. Then*

(i) *$\emptyset \neq \mathcal{C}_{\mathcal{X}} \subseteq \mathcal{S}$ (that is, $\mathcal{X}$ has at least one cluster point and every cluster point is a stationary point of Θ),*

(ii) *$\|\boldsymbol{x}^{(p+1)} - \boldsymbol{x}^{(p)}\| \longrightarrow 0$,*

(iii) *$\mathcal{C}_{\mathcal{X}}$ is a flat continuum, and*

(iv) *$(\Theta(\boldsymbol{x}^{(p)}))_p$ decreases to the value of Θ on $\mathcal{C}_{\mathcal{X}}$.*

Proof. Suppose $\mathcal{L}_0$ is bounded. Then $\mathcal{X}$ is bounded since $\boldsymbol{\Phi}(\mathcal{L}_0) \subseteq \mathcal{L}_0$ by Lemma 3.5(ii). Let $\boldsymbol{\Phi}|_{\mathcal{L}_0}$ and $\Theta|_{\mathcal{L}_0}$ denote the restrictions of $\boldsymbol{\Phi}$ and Θ to $\mathcal{L}_0$. Using Lemma 3.5 again, we have $\mathcal{F}_{\boldsymbol{\Phi}|_{\mathcal{L}_0}} = \mathcal{S} \cap \mathcal{L}_0$, and $\boldsymbol{\Phi}|_{\mathcal{L}_0}$ is continuous and strictly monotonic with respect to $\Theta|_{\mathcal{L}_0}$. So Theorem 3.4 applies with $\mathcal{G} = \mathcal{L}_0$, $\boldsymbol{\Psi} = \boldsymbol{\Phi}|_{\mathcal{L}_0}$, and $\Xi = \Theta|_{\mathcal{L}_0}$. Consequently, $\emptyset \neq \mathcal{C}_{\mathcal{X}} \subseteq \mathcal{F}_{\boldsymbol{\Phi}|_{\mathcal{L}_0}} \subseteq \mathcal{S}$, the difference $\boldsymbol{x}^{(p+1)} - \boldsymbol{x}^{(p)}$ goes to zero, and $\Theta(\boldsymbol{x}^{(p)})$ decreases to $\Theta(\boldsymbol{x}^*)$ for some $\boldsymbol{x}^* \in \mathcal{C}_{\mathcal{X}}$.

It remains to show that $\mathcal{C}_{\mathcal{X}}$ is a flat continuum. This set is compact because $\mathcal{X}$ is bounded and the set of cluster points of a sequence in a topological space is closed. Furthermore, since $\|\boldsymbol{x}^{(p+1)} - \boldsymbol{x}^{(p)}\| \longrightarrow 0$, we deduce from Theorem 26.1 in [35] that $\mathcal{C}_{\mathcal{X}}$ is connected. To prove that $\mathcal{C}_{\mathcal{X}}$ is flat, suppose for contradiction that $\mathcal{X}$ has two cluster points $\boldsymbol{x}^*$ and $\boldsymbol{y}^*$ such that $\Theta(\boldsymbol{x}^*) < \Theta(\boldsymbol{y}^*)$. Let $\beta := (\Theta(\boldsymbol{y}^*) - \Theta(\boldsymbol{x}^*))/3$, and denote by $\mathcal{B}(\boldsymbol{z}, \alpha)$ the open ball with center $\boldsymbol{z}$ and radius α. Since Θ is continuous, there is an $\alpha > 0$ such that $|\Theta(\boldsymbol{x}) - \Theta(\boldsymbol{x}^*)| < \beta$ for all $\boldsymbol{x} \in \mathcal{B}(\boldsymbol{x}^*, \alpha)$ and $|\Theta(\boldsymbol{x}) - \Theta(\boldsymbol{y}^*)| < \beta$ for all $\boldsymbol{x} \in \mathcal{B}(\boldsymbol{y}^*, \alpha)$. Furthermore, since $\boldsymbol{x}^*$ and $\boldsymbol{y}^*$ are cluster points of $\mathcal{X}$, there are two integers p and $q > p$ such that $\boldsymbol{x}^{(p)} \in \mathcal{B}(\boldsymbol{x}^*, \alpha)$ and $\boldsymbol{x}^{(q)} \in \mathcal{B}(\boldsymbol{y}^*, \alpha)$. Therefore,

$$\Theta(\boldsymbol{x}^{(q)}) - \Theta(\boldsymbol{x}^{(p)}) = 3\beta + \Theta(\boldsymbol{x}^{(q)}) - \Theta(\boldsymbol{y}^*) + \Theta(\boldsymbol{x}^*) - \Theta(\boldsymbol{x}^{(p)}) > \beta,$$

which contradicts the fact that $(\Theta(\boldsymbol{x}^{(p)}))_p$ is a decreasing sequence. ∎

3.2. Convergence in norm. Theorem 3.9 provides sufficient conditions for convergence to a stationary point regardless of the convexity of Θ. Instead of the usual requirement that $\mathcal{S}$ be discrete, we impose the weaker condition that every stationary point be isolated from the others in the same level set—we give in Appendix C an example of an objective function satisfying this latter condition but whose set of stationary points is not discrete.

Definition 3.8. *We call* $\mathcal{S}$ level-discrete *if for every* $h \in \mathbb{R}$ *the set* $\{\boldsymbol{x} \in \mathcal{S} : \Theta(\boldsymbol{x}) = h\}$ *(that is, the intersection of* $\mathcal{S}$ *with the* h*-level set of* Θ*) is either discrete or empty.*

Theorem 3.9. *Let* $\mathcal{S}$*,* $\mathcal{X}$*, and* $\mathcal{L}_0$ *be as stated in Theorem 3.2, and suppose* $\mathcal{L}_0$ *is bounded. If* $\mathcal{S}$ *is level-discrete, then* $\mathcal{X}$ *converges to a point in* $\mathcal{S}$*. In particular, if* Θ *has a unique stationary point* $\boldsymbol{x}^*$*, then* $\boldsymbol{x}^*$ *is the unique global minimizer of* Θ *and* $\mathcal{X}$ *converges to* $\boldsymbol{x}^*$*.*

Proof. Suppose $\mathcal{L}_0$ is bounded—so Theorem 3.7 applies—and $\mathcal{S}$ is level-discrete. Since $\mathcal{C}_{\mathcal{X}}$ is flat, there is an $h \in \mathbb{R}$ such that

$$\emptyset \neq \mathcal{C}_{\mathcal{X}} \subseteq \{\boldsymbol{x} \in \mathcal{S} : \Theta(\boldsymbol{x}) = h\},$$

and thus $\mathcal{C}_{\mathcal{X}}$ is discrete. But $\mathcal{C}_{\mathcal{X}}$ is also a continuum; so $\mathcal{C}_{\mathcal{X}}$ is a singleton whose unique element, say $\boldsymbol{x}^*$, is in $\mathcal{S}$. Now suppose for contradiction that $\mathcal{X}$ diverges. Then $\mathcal{X}$ has a subsequence lying outside a neighborhood of $\boldsymbol{x}^*$, and since $\mathcal{X}$ is in the compact set $\mathcal{L}_0$, this subsequence has a further subsequence which converges to a point different from $\boldsymbol{x}^*$. This contradicts the fact that $\boldsymbol{x}^*$ is the only cluster point of $\mathcal{X}$. Therefore, $\mathcal{X}$ converges to $\boldsymbol{x}^*$.

Assume that $\mathcal{S} = \{\boldsymbol{x}^*\}$. Seeking a contradiction, suppose there is an $\boldsymbol{x}$ such that $\Theta(\boldsymbol{x}) < \Theta(\boldsymbol{x}^*)$. Because $\Theta(\boldsymbol{x}^{(p)})$ decreases to $\Theta(\boldsymbol{x}^*)$, we have $\Theta(\boldsymbol{x}) < \Theta(\boldsymbol{x}^{(0)})$, and thus $\inf_{\mathcal{L}_0} \Theta < \Theta(\boldsymbol{x}^{(0)})$. Consequently, since $\mathcal{L}_0$ is compact, there is a $\boldsymbol{y}^*$ in the interior of $\mathcal{L}_0$ such that $\Theta(\boldsymbol{y}^*) = \inf_{\mathcal{L}_0} \Theta$. It follows that $\boldsymbol{y}^* \in \mathcal{S}$ and $\boldsymbol{y}^* \neq \boldsymbol{x}^*$, a contradiction. So $\boldsymbol{x}^*$ is a global minimizer of Θ. Furthermore, this global minimizer is unique since $\mathcal{S} = \{\boldsymbol{x}^*\}$. ∎

Remark 4. In order for $\mathcal{S}$ to be level-discrete but not discrete, Θ must involve a potential function whose derivative oscillates infinitely often around some positive value (as is the case of θ_2 in the example discussed in Appendix C). To our knowledge, at least in the field of image reconstruction, there is currently no practical situation motivating this property. Still, we do not want to exclude possible future applications, and relaxing the requirement that $\mathcal{S}$ be discrete is a further theoretical step toward full understanding of half-quadratic optimization.

Remark 5. We can prove that Theorem 3.9 holds if we replace the condition that $\mathcal{S}$ be level-discrete by the condition that the boundary of $\mathcal{S}$ be discrete. However, this result is irrelevant when $N \geqslant 2$, for in this case $\mathcal{S}$ is discrete if and only if its boundary is discrete.

The following corollary generalizes the convergence results in [3, 5, 6, 7]. The important point is that some potential functions can be nonconvex as long as Θ is strictly convex—sufficient conditions for strict convexity are given in Appendix B.

Corollary 3.10. *If* Θ *is strictly convex, then any sequence generated by Algorithm 1 converges to the unique global minimizer of* Θ*.*

Proof. Assume Θ is strictly convex. Seeking a contradiction, suppose Θ is not coercive. By Proposition 2.5 and inequality (2.9), there is a nonzero vector $\boldsymbol{x}$ such that the function $\alpha \in \mathbb{R} \longmapsto \Theta(\alpha \boldsymbol{x})$ is bounded above and hence is either constant or nonconvex. This contradicts the strict convexity assumption, and thus Θ is coercive. Consequently, condition (C4) holds, and every sublevel set of Θ is bounded. So Theorem 3.7 applies to every sequence generated by Algorithm 1. The corollary then follows from the fact that Θ has a unique stationary point because $\mathcal{S}$ is nonempty by Theorem 3.7(i) and Θ is strictly convex. ∎

3.3. Convergence in terms of point-to-set distance. We now show that, regardless of the distribution of the stationary points, the sequences generated by Algorithm 1 converge to the boundary of $\mathcal{S}$ with an objective gradient tending to zero. In other words, Algorithm 1

behaves well even when $\mathcal{S}$ is not level-discrete, including in the presence of flat continua (as happens, for example, when Θ is convex and has several global minimizers).

Definition 3.11. *Let $\mathcal{G}$ be a nonempty subset of $\mathbb{R}^N$. The* point-to-set distance *to $\mathcal{G}$ is the functional $d_{\mathcal{G}} : \mathbb{R}^N \longrightarrow \mathbb{R}$ defined by*

$$d_{\mathcal{G}}(\boldsymbol{x}) := \inf_{\boldsymbol{y}\in\mathcal{G}} \|\boldsymbol{y} - \boldsymbol{x}\|.$$

A sequence $(\boldsymbol{x}^{(p)})_p$ is said to converge to $\mathcal{G}$ *if $d_{\mathcal{G}}(\boldsymbol{x}^{(p)}) \longrightarrow 0$.*

Theorem 3.12. *Let $\mathcal{S}$, $\mathcal{X}$, and $\mathcal{L}_0$ be as stated in Theorem 3.2, and suppose $\mathcal{L}_0$ is bounded.*

(i) *If $\mathcal{X}$ does not converge to a single point in $\mathcal{S}$, then $\mathcal{X}$ converges to $\mathcal{S} \setminus \mathcal{S}^\circ$, where $\mathcal{S}^\circ$ denotes the interior of $\mathcal{S}$.*

(ii) $\|\nabla\Theta(\boldsymbol{x}^{(p)})\| \longrightarrow 0$.

Proof. Suppose $\mathcal{L}_0$ is bounded so that the conclusions of Theorem 3.7 hold.

(i) Seeking a contradiction, suppose $\mathcal{X}$ does not converge to $\mathcal{C}_{\mathcal{X}}$. Then there are an $\alpha > 0$ and a subsequence $(\boldsymbol{y}^{(q)})_q =: \mathcal{Y}$ of $\mathcal{X}$ such that $d_{\mathcal{C}_{\mathcal{X}}}(\boldsymbol{y}^{(q)}) \geqslant \alpha$ for all q. It follows that the set $\mathcal{C}_{\mathcal{Y}}$ of cluster points of $\mathcal{Y}$ is empty (since $\mathcal{C}_{\mathcal{Y}} \subseteq \mathcal{C}_{\mathcal{X}}$), which contradicts the fact that $\mathcal{Y}$ lies in the compact set $\mathcal{L}_0$. So $\mathcal{X}$ converges to $\mathcal{C}_{\mathcal{X}}$ and hence to $\mathcal{S}$. Now suppose $\mathcal{X}$ does not converge to a single stationary point. Since $\mathcal{X}$ gets stuck at stationary points by Lemma 3.5(iii), the set $\mathcal{C}_{\mathcal{X}} \cap \mathcal{S}^\circ$ must be empty. Therefore, $\mathcal{C}_{\mathcal{X}} \subseteq \mathcal{S} \setminus \mathcal{S}^\circ$, and so convergence to $\mathcal{C}_{\mathcal{X}}$ implies convergence to $\mathcal{S} \setminus \mathcal{S}^\circ$.

(ii) Since $\mathcal{C}_{\mathcal{X}}$ is compact, there is a sequence $(\boldsymbol{z}^{(p)})_p$ in $\mathcal{C}_{\mathcal{X}}$ such that $\|\boldsymbol{z}^{(p)} - \boldsymbol{x}^{(p)}\| = d_{\mathcal{C}_{\mathcal{X}}}(\boldsymbol{x}^{(p)})$ for all p. So $\nabla\Theta(\boldsymbol{z}^{(p)})$ is always zero and $\|\boldsymbol{z}^{(p)} - \boldsymbol{x}^{(p)}\| \longrightarrow 0$. Since Θ is C^1 and $\mathcal{L}_0$ is compact, $\nabla\Theta$ is uniformly continuous on $\mathcal{L}_0$ by the Heine-Cantor theorem, and hence $\|\nabla\Theta(\boldsymbol{x}^{(p)})\| = \|\nabla\Theta(\boldsymbol{z}^{(p)}) - \nabla\Theta(\boldsymbol{x}^{(p)})\| \longrightarrow 0$. ∎

Corollary 3.13. *If Θ is convex and coercive, then any sequence generated by Algorithm 1 converges to the set of global minimizers of Θ.*

Proof. If Θ is coercive, then Theorem 3.12 applies to every sequence generated by Algorithm 1. If, in addition, Θ is convex, then $\mathcal{S}$ is the set of global minimizers of Θ (see, for example, Theorem 7.4.4 in [36]). ∎

Remark 6. In the convex case, the coercivity of the objective functional does not strengthen the conditions of Theorem 3.12. Indeed, if Θ is convex, the condition that Θ be coercive is equivalent to requiring that the set of global minimizers of Θ be nonempty and bounded, which is the case when $\mathcal{L}_0$ is bounded.

4. Local convergence to isolated minimizers.

The majorization-minimization interpretation given in Proposition 2.6 leads to the capture property stated in Theorem 4.3—the proof of this theorem uses Theorem 4.1 and Proposition 6.3 of Jacobson and Fessler in [30]. First, we recall the definition of a generalized basin given in [30], and we introduce the notion of a feasible initial point which is used later in section 4.2 to characterize the basins of attraction for Algorithm 1.

Notation. Let $\mathcal{G} \subseteq \mathbb{R}^N$. We denote the interior, the closure, and the boundary of $\mathcal{G}$ by $\mathcal{G}^\circ$, $\overline{\mathcal{G}}$, and $\partial\mathcal{G}$, respectively. (The only subsets of $\mathbb{R}^N$ with empty boundary are $\emptyset$ and $\mathbb{R}^N$.)

Definition 4.1. *A set $\mathcal{G} \subseteq \mathbb{R}^N$ is called a* generalized basin *of Θ if there is an $\boldsymbol{x} \in \mathcal{G}$ such that*

$$\Theta(\boldsymbol{x}) < \Theta(\boldsymbol{y}) \quad \text{for all } \boldsymbol{y} \in \partial\mathcal{G}.$$

(In particular, $\mathbb{R}^N$ is a generalized basin of any functional on $\mathbb{R}^N$.) Such an $\boldsymbol{x}$ is said to be well-contained *in $\mathcal{G}$.*

Definition 4.2. *We call $\boldsymbol{x}^{(0)} \in \mathbb{R}^N$ a* feasible initial point *if the matrix $\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x}^{(p)})\boldsymbol{A}$ is positive definite at each step of the recurrence (2.5) starting from $\boldsymbol{x}^{(0)}$; in this case, we say that the sequence $(\boldsymbol{x}^{(p)})_p$ so generated is* feasible.

Theorem 4.3. *Let $\mathcal{X} := (\boldsymbol{x}^{(p)})_p$ be a feasible sequence, and let $\mathcal{G}$ be a bounded generalized basin of Θ such that $\overline{\mathcal{G}}$ contains a unique stationary point $\boldsymbol{x}^*$. If there is an integer $q \geqslant 0$ such that $\boldsymbol{x}^{(q)}$ is well-contained in $\mathcal{G}$, then $\boldsymbol{x}^{(p)} \in \mathcal{G}$ for all $p \geqslant q$ and $\mathcal{X}$ converges to $\boldsymbol{x}^*$.*

Proof. According to the beginning of the proof of Proposition 2.6, if $\boldsymbol{x}^{(0)}$ is a feasible initial point, then the iterative scheme of Algorithm 1 is equivalent to (2.13) with $\overline{\Theta}$ given by (2.15). Consequently, $\mathcal{X}$ is a majorization-minimization sequence whose surrogate generator $\overline{\Theta}$ is continuous on the product space $\mathbb{R}^N \times \mathbb{R}^N$ and such that $\nabla\overline{\Theta}(\boldsymbol{x},\cdot)(\boldsymbol{x}) = \nabla\Theta(\boldsymbol{x})$ for all $\boldsymbol{x}$. Hence, applying Theorem 4.1 in [30], we obtain that $\mathcal{C}_{\mathcal{X}} \subseteq \mathcal{S}$ (that is, any cluster point of $\mathcal{X}$ is a stationary point of Θ).

Let $\mathcal{G}$ be as stated, and suppose there is a $q \in \mathbb{N}$ such that $\boldsymbol{x}^{(q)}$ is well-contained in $\mathcal{G}$. Since $\overline{\Theta}(\boldsymbol{x},\cdot)$ is convex for every $\boldsymbol{x} \in \mathbb{R}^N$, it follows from Proposition 6.3 in [30] that if a point $\boldsymbol{x}$ is well-contained in $\mathcal{G}$, then so is every $\boldsymbol{y}$ such that $\overline{\Theta}(\boldsymbol{x},\boldsymbol{y}) \leqslant \overline{\Theta}(\boldsymbol{x},\boldsymbol{x})$. Therefore, by (2.13), $\boldsymbol{x}^{(p)} \in \mathcal{G}$ implies that $\boldsymbol{x}^{(p+1)} \in \mathcal{G}$, and straightforward induction yields $\boldsymbol{x}^{(p)} \in \mathcal{G}$ for all $p \geqslant q$. Furthermore, since $\overline{\mathcal{G}}$ is compact, it follows that $\emptyset \neq \mathcal{C}_{\mathcal{X}} \subseteq \mathcal{S} \cap \overline{\mathcal{G}}$. But $\mathcal{S} \cap \overline{\mathcal{G}} = \{\boldsymbol{x}^*\}$, and thus $\mathcal{X}$ is eventually in a compact set and has no cluster point other than $\boldsymbol{x}^*$, which completes the proof. ∎

Example 1. Let $\Theta : \mathbb{R} \longrightarrow \mathbb{R}$ be defined by

$$\Theta(x) \;=\; \vartheta_{\mathrm{LE}}\big(|x-1|\big) + \vartheta_{\mathrm{LE}}\big(\sqrt{3}\,|x-3|/\sqrt{2}\big) + 6\,\vartheta_{\mathrm{GM}}\big(|x-6|/\sqrt{2}\big), \tag{4.1}$$

where ϑ_{LE} and ϑ_{GM} are the Lorentzian error function and the Geman and McClure function given in (1.8) and (1.9). The function Θ is plotted in Figure 1(a); it has three stationary points: one isolated maximizer $x_3 \approx 4.274$ and two isolated minimizers $x_2 \approx 2.910$ and $x_5 \approx 5.819$.

The class Υ of intervals $\mathcal{U}$ such that $\mathcal{U}^\circ = (\alpha,\beta)$ with $-\infty \leqslant \alpha < x_2 < \beta < x_4$ or $-\infty \leqslant \alpha < x_5 < \beta \leqslant +\infty$ are generalized basins of Θ, but the intervals $\mathcal{U}$ such that $\sup\mathcal{U} \in [x_4, x_5]$ are not. (Note that the definition of a generalized basin does not impose that it be connected; however, disconnected basins are of no practical or theoretical interest.) The hypotheses of Theorem 4.3 restrict Υ to the subclass Υ^* of intervals that are bounded and whose closure contains a unique stationary point $x^* \in \{x_2, x_5\}$, that is, the intervals with interior (α,β) such that $-\infty < \alpha < x_2 < \beta < x_3$ or $x_3 < \alpha < x_5 < \beta < +\infty$. Furthermore, within an interval $\mathcal{U} \in \Upsilon^*$, the attraction region around x^* is the set of points that are well-contained in $\mathcal{U}$, and thus the capture intervals are

$$\mathcal{O}_h(x_2) \;=\; (x_1, x_3) \cap \{x : \Theta(x) < h\} \quad \text{and} \quad \mathcal{O}_h(x_5) \;=\; (x_3, x_6) \cap \{x : \Theta(x) < h\}$$

with $h \in (\Theta(x_2), \Theta(x_3))$ for $\mathcal{O}_h(x_2)$ and $h \in (\Theta(x_5), \Theta(x_3))$ for $\mathcal{O}_h(x_5)$. We call such sets "cups" by analogy with water-flooding from a bottom source.

Figure 1(b) illustrates the behavior of Algorithm 1 starting from $x^{(0)} = 1$. Each iterate $x^{(p+1)}$ is the minimizer of the majorizing quadratic function $\overline{\Theta}(x^{(p)},\cdot)$ defined by (2.15), and

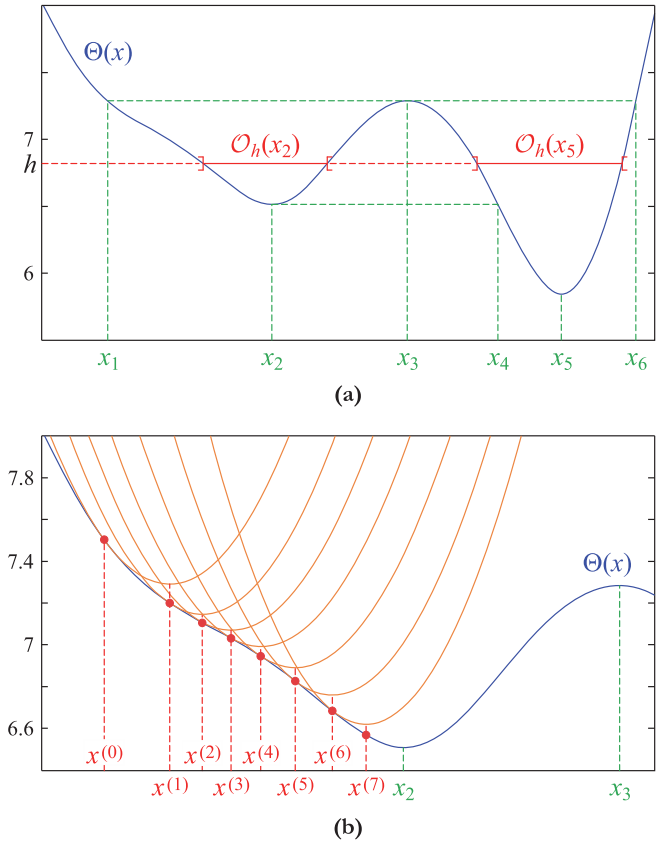

Figure 1. *One-dimensional illustration of local convergence:* (a) *the open intervals* $\mathcal{O}_h(x_2)$ *and* $\mathcal{O}_h(x_5)$ *are basins of attraction of the function* Θ *(defined by* (4.1)*) provided* $h < \Theta(x_3)$*;* (b) *first iterates of Algorithm* 1 *starting from* $x^{(0)} = 1$ *and associated majorizing quadratic functions plotted in orange.*

the sequence so generated converges to x_2. This example illustrates why it is not possible to jump between cups with different bottoms: moving from a point $x^{(p)}$ in a cup $\mathcal{O}_h(x_2)$ to a point $x^{(p+1)}$ in a cup $\mathcal{O}_{h'}(x_5)$ would imply that $\overline{\Theta}(x^{(p)}, \cdot)$ is strictly smaller than Θ on some interval containing x_3, contradicting the fact that $\overline{\Theta}(x^{(p)}, \cdot)$ majorizes Θ.

In the following three sections, we characterize the basins of attraction of isolated minimizers in terms of the special sets called "cups" exemplified above, which are bounded open sets with a flat boundary and containing a single stationary point. We start with elementary properties of cups in section 4.1. Next, in section 4.2, we show that Theorem 4.3 can be reformulated by saying that the basin of attraction of an isolated minimizer $\boldsymbol{x}^*$ contains every cup around $\boldsymbol{x}^*$. The union of all such cups can be viewed as the N-dimensional region covered by flooding from the source $\boldsymbol{x}^*$. We develop this water-flooding interpretation in section 4.3.

4.1. Basic properties of cups. We consider here a functional $\Xi : \mathbb{R}^N \longrightarrow \mathbb{R}$ which is differentiable and has at least one local minimizer that is also an isolated stationary point—we call such a point an *isolated local minimizer* (an isolated local minimizer is a strict local minimizer, but the converse is not necessarily true).

Definition 4.4. *Let $\mathcal{S}_\Xi$ denote the set of stationary points of Ξ. We call a* cup *of Ξ a bounded open set $\mathcal{O} \subset \mathbb{R}^N$ such that*

(i) $\mathcal{O} \cap \mathcal{S}_\Xi = \{\boldsymbol{x}^*\}$, *where $\boldsymbol{x}^*$ is a local minimizer of Ξ, and*

(ii) Ξ *is constant on $\partial\mathcal{O}$.*

The isolated local minimizer $\boldsymbol{x}^*$ is called the bottom *of $\mathcal{O}$. The value of Ξ on $\partial\mathcal{O}$ is called the* height *of $\mathcal{O}$ and is denoted by $H(\mathcal{O})$.*

We begin with two general lemmas and proceed with more specific ones to characterize cups.

Lemma 4.5. *Let $\mathcal{U}$ be a bounded subset of $\mathbb{R}^N$ containing a single stationary point $\boldsymbol{x}^* \in \mathcal{S}_\Xi$. If there is an $\boldsymbol{x} \in \mathcal{U}$ such that $\Xi(\boldsymbol{x}) < \min_{\partial\mathcal{U}} \Xi$, then $\boldsymbol{x}^*$ is the unique global minimizer of Ξ on $\overline{\mathcal{U}}$.*

Proof. The functional Ξ has a global minimizer on $\overline{\mathcal{U}}$ because this set is compact. Let $\boldsymbol{y}^* \in \arg\min_{\overline{\mathcal{U}}} \Xi$, and suppose there is an $\boldsymbol{x} \in \mathcal{U}$ such that $\Xi(\boldsymbol{x}) < \min_{\partial\mathcal{U}} \Xi$. Then $\Xi(\boldsymbol{y}^*) < \min_{\partial\mathcal{U}} \Xi$, and thus $\boldsymbol{y}^* \in \mathcal{U}^\circ$. So $\boldsymbol{y}^*$ is a local minimizer and hence a stationary point. But $\boldsymbol{x}^*$ is the only stationary point in $\mathcal{U}$, which completes the proof. ∎

Notation. Given a binary relation $\mathcal{R}$ on $\mathbb{R}$ and a real number h, we let

$$\operatorname{lev}_{\mathcal{R}h}\Xi := \big\{\boldsymbol{x} \in \mathbb{R}^N : \Xi(\boldsymbol{x})\mathcal{R}h\big\}. \tag{4.2}$$

For example, $\operatorname{lev}_{=h}\Xi$ and $\operatorname{lev}_{\leqslant h}\Xi$ are the level and sublevel sets of Ξ at height h, respectively.

Lemma 4.6. *Let $\mathcal{U}$ be an open subset of $\mathbb{R}^N$, and let $h \in \mathbb{R}$. If $\partial\mathcal{U} \subseteq \operatorname{lev}_{\geqslant h}\Xi$, then $\partial(\mathcal{U} \cap \operatorname{lev}_{<h}\Xi) \subseteq \operatorname{lev}_{=h}\Xi$.*

Proof. Let $\mathcal{U}_h := \mathcal{U} \cap \operatorname{lev}_{<h}\Xi$. We have

$$\partial\mathcal{U}_h \subseteq \overline{\mathcal{U}_h} \subseteq \overline{\operatorname{lev}_{<h}\Xi} \subseteq \operatorname{lev}_{\leqslant h}\Xi \quad \text{and} \quad \partial\mathcal{U}_h \subseteq \partial\mathcal{U} \cup \partial\operatorname{lev}_{<h}\Xi \subseteq \partial\mathcal{U} \cup \operatorname{lev}_{=h}\Xi$$

(the two rightmost inclusions follow from the continuity of Ξ). Hence $\partial\mathcal{U}_h \subseteq \operatorname{lev}_{=h}\Xi$ if $\partial\mathcal{U} \subseteq \operatorname{lev}_{\geqslant h}\Xi$. ∎

Lemma 4.7. *Let $\boldsymbol{x}^*$ be an isolated local minimizer of Ξ. Every open ball around $\boldsymbol{x}^*$ contains a cup of Ξ with bottom $\boldsymbol{x}^*$.*

Proof. Let $r > 0$, and let $\mathcal{B}(\boldsymbol{x}^*, r)$ denote the open ball with center $\boldsymbol{x}^*$ and radius r. Since $\boldsymbol{x}^*$ is an isolated local minimizer, there is an $\alpha \in (0, 2r]$ such that

$$\mathcal{B}(\boldsymbol{x}^*, \alpha) \cap \mathcal{S}_\Xi = \{\boldsymbol{x}^*\} \quad \textit{and} \quad \Xi(\boldsymbol{x}) > \Xi(\boldsymbol{x}^*) \quad \text{for all } \boldsymbol{x} \in \mathcal{B}(\boldsymbol{x}^*, \alpha) \setminus \{\boldsymbol{x}^*\}.$$

Consider the bounded open set $\mathcal{O} \subseteq \mathcal{B}(\boldsymbol{x}^*, r)$ defined by

$$\mathcal{O} := \mathcal{B}(\boldsymbol{x}^*, \alpha/2) \cap \operatorname{lev}_{<h}\Xi \quad \text{with} \quad h := \min_{\partial\mathcal{B}(\boldsymbol{x}^*, \alpha/2)} \Xi.$$

Since $\partial\mathcal{B}(\boldsymbol{x}^*, \alpha/2)$ is a compact subset of $\mathcal{B}(\boldsymbol{x}^*, \alpha)$, there is an $\boldsymbol{x} \in \mathcal{B}(\boldsymbol{x}^*, \alpha)$ such that $\Xi(\boldsymbol{x}) = h$. So $\Xi(\boldsymbol{x}^*) < h$, and it follows that $\mathcal{O} \cap \mathcal{S}_\Xi = \{\boldsymbol{x}^*\}$. Furthermore, by Lemma 4.6 with $\mathcal{U} = \mathcal{B}(\boldsymbol{x}^*, \alpha/2)$, we have $\partial\mathcal{O} \subseteq \operatorname{lev}_{=h}\Xi$. Therefore, $\mathcal{O}$ is a cup with bottom $\boldsymbol{x}^*$. ∎

Lemma 4.8. *Let $\mathcal{O}$ be a cup of Ξ with bottom $\boldsymbol{x}^*$.*

(i) $\mathcal{O} \subseteq \operatorname{lev}_{<H(\mathcal{O})}\Xi$ *(therefore,* $\max_{\overline{\mathcal{O}}} \Xi = H(\mathcal{O})$ *and* $\arg\max_{\overline{\mathcal{O}}} \Xi = \partial\mathcal{O}$*).*

(ii) $\arg\min_{\overline{\mathcal{O}}} \Xi = \{\boldsymbol{x}^*\}$ *(that is, $\boldsymbol{x}^*$ is the unique global minimizer of Ξ on $\overline{\mathcal{O}}$).*

(iii) $\mathcal{O}$ *is connected.*

Proof. (i) Suppose for contradiction that $\mathcal{O} \cap \mathrm{lev}_{\geqslant H(\mathcal{O})} \Xi \neq \emptyset$. Since $\overline{\mathcal{O}}$ is compact and Ξ is equal to $H(\mathcal{O})$ on $\partial\mathcal{O}$, there is a $\boldsymbol{y}^* \in \mathcal{O}$ such that $\Xi(\boldsymbol{y}^*) = \max_{\overline{\mathcal{O}}} \Xi$. Hence Ξ has a local maximizer in $\mathcal{O}$, a contradiction with the fact that $\mathcal{O}$ contains a single stationary point which is a strict local minimizer.

(ii) By the definition of a cup and the above property (i), Lemma 4.5 applies with $\mathcal{U} = \mathcal{O}$.

(iii) Suppose for contradiction that $\mathcal{O}$ is not connected; that is, there are open sets $\mathcal{O}_1$ and $\mathcal{O}_2$ such that $\{\mathcal{O} \cap \mathcal{O}_1, \mathcal{O} \cap \mathcal{O}_2\}$ is a partition of $\mathcal{O}$. Without loss of generality, we assume that $\boldsymbol{x}^* \in \mathcal{O}_1$. The set $\overline{\mathcal{O} \cap \mathcal{O}_2}$ is compact, its interior $\mathcal{O} \cap \mathcal{O}_2$ is a subset of $\mathrm{lev}_{<H(\mathcal{O})} \Xi$ by property (i), and Ξ is equal to $H(\mathcal{O})$ on $\partial(\mathcal{O} \cap \mathcal{O}_2)$ since $\partial\mathcal{O} = \partial(\mathcal{O} \cap \mathcal{O}_1) \cup \partial(\mathcal{O} \cap \mathcal{O}_2)$. It follows that Ξ has a local minimizer in $\mathcal{O} \cap \mathcal{O}_2$, which contradicts the fact that Ξ has no stationary point in $\mathcal{O} \cap \mathcal{O}_2$. ∎

Lemma 4.9. *Let $\mathcal{O}_1$ and $\mathcal{O}_2$ be cups of Ξ with the same bottom. Then $\mathcal{O}_1 \subseteq \mathcal{O}_2$ if and only if $H(\mathcal{O}_1) \leqslant H(\mathcal{O}_2)$.*

Proof. Suppose $\mathcal{O}_1$ is not a subset of $\mathcal{O}_2$. If $\mathcal{O}_1 \cap \partial\mathcal{O}_2$ is empty, the sets $\mathcal{O}_1 \cap \mathcal{O}_2$ and $\mathcal{O}_1 \cap (\mathbb{R}^N \setminus \overline{\mathcal{O}_2})$ form a partition of $\mathcal{O}_1$. But $\mathcal{O}_1$ is connected by Lemma 4.8(iii); so $\mathcal{O}_1 \cap \partial\mathcal{O}_2$ must be nonempty. Let $\boldsymbol{x} \in \mathcal{O}_1 \cap \partial\mathcal{O}_2$. We have $\Xi(\boldsymbol{x}) < H(\mathcal{O}_1)$ by Lemma 4.8(i), and $\Xi(\boldsymbol{x}) = H(\mathcal{O}_2)$, which proves the converse implication. The forward implication follows immediately: if $H(\mathcal{O}_1) > H(\mathcal{O}_2)$, then $\mathcal{O}_1$ is not a subset of $\mathcal{O}_2$ because $\mathcal{O}_2 \subset \mathcal{O}_1$ by the converse implication ($\mathcal{O}_1 \neq \mathcal{O}_2$ since $H(\mathcal{O}_1) \neq H(\mathcal{O}_2)$). ∎

Lemma 4.10. *Let $\mathcal{O}$ be a cup of Ξ with bottom $\boldsymbol{x}^*$. For every $h \in (\Xi(\boldsymbol{x}^*), H(\mathcal{O}))$, the set $\mathcal{O} \cap \mathrm{lev}_{<h} \Xi$ is the unique cup of Ξ with bottom $\boldsymbol{x}^*$ and height h.*

Proof. Let $h \in (\Xi(\boldsymbol{x}^*), H(\mathcal{O}))$. The set $\mathcal{O}_h := \mathcal{O} \cap \mathrm{lev}_{<h} \Xi$ is clearly open and bounded, and $\mathcal{O}_h \cap \mathcal{S}_\Xi = \{\boldsymbol{x}^*\} \cap \mathrm{lev}_{<h} \Xi = \{\boldsymbol{x}^*\}$. Furthermore, since $\partial\mathcal{O} \subseteq \mathrm{lev}_{=H(\mathcal{O})} \Xi \subset \mathrm{lev}_{>h} \Xi$, it follows from Lemma 4.6 that $\partial\mathcal{O}_h \subseteq \mathrm{lev}_{=h} \Xi$. So $\mathcal{O}_h$ is a cup with bottom $\boldsymbol{x}^*$ and height h. The uniqueness follows from Lemma 4.9. ∎

4.2. Basins of attraction. We prove Theorem 4.11 by showing that its hypotheses imply those of Theorem 4.3, and we establish the converse in the proof of Proposition 4.13. Corollaries of this new formulation provide a clear characterization of the basins of attraction for Algorithm 1.

Theorem 4.11. *Let $\mathcal{X} := (\boldsymbol{x}^{(p)})_p$ be a feasible sequence. If there are an integer $q \geqslant 0$ and a cup $\mathcal{O}$ of Θ such that $\boldsymbol{x}^{(q)} \in \mathcal{O}$, then $\boldsymbol{x}^{(p)} \in \mathcal{O}$ for all $p \geqslant q$ and $\mathcal{X}$ converges to the bottom of $\mathcal{O}$.*

Proof. Let $\boldsymbol{x}^*$ be the bottom of $\mathcal{O}$, and let q be such that $\boldsymbol{x}^{(q)} \in \mathcal{O}$. Then $\Theta(\boldsymbol{x}^*) \leqslant \Theta(\boldsymbol{x}^{(q)}) < H(\mathcal{O})$ by Lemma 4.8. Let $h \in (\Theta(\boldsymbol{x}^{(q)}), H(\mathcal{O}))$ and $\mathcal{O}_h := \mathcal{O} \cap \mathrm{lev}_{<h} \Theta$. We show that Theorem 4.3 applies with $\mathcal{G} = \mathcal{O}_h$ (hence the conclusions that $\boldsymbol{x}^{(p)} \in \mathcal{O}_h$ for all $p \geqslant q$ and $\mathcal{X}$ converges to $\boldsymbol{x}^*$). By Lemma 4.10, the set $\mathcal{O}_h$ is a cup with bottom $\boldsymbol{x}^*$ and height h. So $\mathcal{O}_h$ is bounded, and for every $\boldsymbol{x} \in \partial\mathcal{O}_h$, $\Theta(\boldsymbol{x}) = h > \Theta(\boldsymbol{x}^{(q)})$. It remains to check that $\overline{\mathcal{O}_h} \cap \mathcal{S} = \{\boldsymbol{x}^*\}$. We have

$$\overline{\mathcal{O}_h} \subseteq \overline{\mathcal{O}} \cap \overline{\mathrm{lev}_{<h} \Theta} \subseteq \overline{\mathcal{O}} \cap \mathrm{lev}_{\leqslant h} \Theta,$$

and since $\partial\mathcal{O} \subseteq \mathrm{lev}_{=H(\mathcal{O})} \Theta \subset \mathrm{lev}_{>h} \Theta$, it follows that

$$\overline{\mathcal{O}_h} \subseteq \mathcal{O} \cap \mathrm{lev}_{\leqslant h} \Theta \subset \mathcal{O}.$$

Therefore, $\boldsymbol{x}^* \in \overline{\mathcal{O}_h} \cap \mathcal{S} \subseteq \mathcal{O} \cap \mathcal{S} = \{\boldsymbol{x}^*\}$. ∎

Lemma 4.12. *Let $\mathcal{G}$ be a bounded generalized basin of Θ such that $\mathcal{G}^\circ$ contains a unique stationary point $\boldsymbol{x}^*$. The set*

$$\mathcal{W}_{\mathcal{G}} := \left\{\boldsymbol{x} \in \mathbb{R}^N : \boldsymbol{x} \text{ is well-contained in } \mathcal{G}\right\}$$

is a cup of Θ with bottom $\boldsymbol{x}^*$.

Proof. Since $\mathcal{G}$ is bounded, $\partial\mathcal{G}$ is compact, and thus $\mathcal{W}_{\mathcal{G}} = \mathcal{G}^\circ \cap \operatorname{lev}_{<h}\Theta$ with $h := \min_{\partial\mathcal{G}}\Theta$. Hence $\mathcal{W}_{\mathcal{G}}$ is open and bounded. Let $\boldsymbol{x} \in \mathcal{W}_{\mathcal{G}}$ (such a point exists by the definition of a generalized basin). By Lemma 4.5 with $\mathcal{U} = \mathcal{G}^\circ$ and $\Xi = \Theta$, the stationary point $\boldsymbol{x}^*$ is the unique global minimizer of Θ on $\mathcal{G}^\circ$. Therefore, $\Theta(\boldsymbol{x}^*) \leqslant \Theta(\boldsymbol{x}) < h$, and so $\mathcal{W}_{\mathcal{G}} \cap \mathcal{S} = \{\boldsymbol{x}^*\}$. Finally, since $\partial\mathcal{G}^\circ \subseteq \partial\mathcal{G} \subseteq \operatorname{lev}_{\geqslant h}\Theta$, it follows from Lemma 4.6 that $\partial\mathcal{W}_{\mathcal{G}} \subseteq \operatorname{lev}_{=h}\Theta$. ∎

Remark 7. The converse of Lemma 4.12 is also true by the definition of a cup and Lemma 4.8(i). Therefore, a cup with bottom $\boldsymbol{x}^*$ can be equivalently defined to be the set of points that are well-contained in a bounded generalized basin $\mathcal{G}$ such that $\mathcal{G}^\circ \cap \mathcal{S} = \{\boldsymbol{x}^*\}$. Note that we cannot replace $\mathcal{G}^\circ$ by either $\mathcal{G}$ or $\overline{\mathcal{G}}$ in this alternative definition, as cups can have stationary points on their boundaries (this is, for example, the case of (x_1, x_3) and (x_3, x_6) in Figure 1).

Proposition 4.13. *The hypotheses of Theorem 4.11 are equivalent to those of Theorem 4.3.*

Proof. Let $\boldsymbol{x} \in \mathbb{R}^N$. In the proof of Theorem 4.11, we have shown that if $\boldsymbol{x}$ is in a cup $\mathcal{O}$ with bottom $\boldsymbol{x}^*$, then $\boldsymbol{x}$ is well-contained in a bounded generalized basin $\mathcal{G}$ such that $\overline{\mathcal{G}} \cap \mathcal{S} = \{\boldsymbol{x}^*\}$. We must prove the converse. By Lemma 4.5 with $\mathcal{U} = \overline{\mathcal{G}}$ and $\Xi = \Theta$, the stationary point $\boldsymbol{x}^*$ is the unique global minimizer of Θ on $\overline{\mathcal{G}}$. But $\mathcal{G}$ is not a generalized basin if $\boldsymbol{x}^* \in \partial\mathcal{G}$, and thus $\boldsymbol{x}^* \in \mathcal{G}^\circ$. It follows from Lemma 4.12 that $\mathcal{W}_{\mathcal{G}}$ is a cup with bottom $\boldsymbol{x}^*$. ∎

Definition 4.14. *Let $\mathcal{I}^{(0)} \subseteq \mathbb{R}^N$ be the set of feasible initial points, let $\mathcal{X}(\boldsymbol{x})$ denote the sequence $(\boldsymbol{x}^{(p)})_p$ generated by Algorithm 1 starting from $\boldsymbol{x}^{(0)} = \boldsymbol{x} \in \mathcal{I}^{(0)}$, and let $\boldsymbol{x}^*$ be an isolated local minimizer of Θ. The set*

$$\mathcal{A}(\boldsymbol{x}^*) := \left\{\boldsymbol{x} \in \mathcal{I}^{(0)} : \mathcal{X}(\boldsymbol{x}) \text{ converges to } \boldsymbol{x}^*\right\}$$

is called the basin of attraction *of $\boldsymbol{x}^*$.*

Definition 4.15. *Let $\Omega(\boldsymbol{x}^*)$ be the set of all the cups of Θ with bottom $\boldsymbol{x}^*$. We call the union set*

$$\mathcal{O}_{\sup}(\boldsymbol{x}^*) := \bigcup_{\mathcal{O} \in \Omega(\boldsymbol{x}^*)} \mathcal{O}$$

a supremum cup.

Recall that condition (C4) is equivalent to the requirement that the matrix $\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})\boldsymbol{A}$ be positive definite for every $\boldsymbol{x} \in \mathbb{R}^N$ (see Proposition 2.4), which implies that $\mathcal{I}^{(0)} = \mathbb{R}^N$. Hence we have the following corollary to Theorem 4.11.

Corollary 4.16. *If condition* (C4) *holds, then $\mathcal{O}_{\sup}(\boldsymbol{x}^*) \subseteq \mathcal{A}(\boldsymbol{x}^*)$ for every isolated local minimizer $\boldsymbol{x}^*$.*

Example 2. The cups (x_1, x_3) and (x_3, x_6) in Figure 1 are simple examples of supremum cups. Supremum cups can also be unbounded, but we show in Appendix D that this situation does not arise in the one-dimensional case for the class of objective functionals considered. So we focus here on a two-dimensional example.

Let $\Theta : \mathbb{R}^2 \longrightarrow \mathbb{R}$ be defined by

$$\Theta(x_1, x_2) = \vartheta_{\mathrm{GM}}\big(|x_1|\big) + \vartheta_{\mathrm{GM}}\big(|x_2|\big) + \vartheta_{\mathrm{GM}}\big(\|(x_1, x_2) - (3, 3)\|\big), \tag{4.3}$$

where ϑ_{GM} is the Geman and McClure function given in (1.9). This functional has two isolated minimizers $\boldsymbol{x}^*$ and $\boldsymbol{y}^*$ in the vicinity of $\mathbf{0}$ and $(3, 3)$ and one isolated maximizer $\boldsymbol{z}^*$ near $(2, 2)$. Figure 2(a) shows the boundaries of the supremum cups $\mathcal{O}_{\sup}(\boldsymbol{x}^*) =: \mathcal{O}_1$ and $\mathcal{O}_{\sup}(\boldsymbol{y}^*) =: \mathcal{O}_2$ superimposed on Θ. The supremum cups $\mathcal{O}_1$ and $\mathcal{O}_2$ are respectively unbounded and bounded: $\mathcal{O}_1$ is the connected component of $\boldsymbol{x}^*$ in $\mathrm{lev}_{<2}\Theta$, and $\mathcal{O}_2$ is the connected component of $\boldsymbol{y}^*$ in $\mathrm{lev}_{<H(\boldsymbol{y}^*)}\Theta$, where $H(\boldsymbol{y}^*) := \sup\big\{h > \Theta(\boldsymbol{y}^*) : \mathrm{lev}_{<h}\Theta \text{ is disconnected}\big\} \approx 2.26988$.

Figure 2(b) displays the number of iterations for Algorithm 1 to reach either of the open balls $\mathcal{B}(\boldsymbol{x}^*, 10^{-6})$ (dominant blue color) or $\mathcal{B}(\boldsymbol{y}^*, 10^{-6})$ (dominant red color) starting from every location in the square domain $[-2, 5] \times [-2, 5]$. In accordance with Corollary 4.16, we observe convergence to $\boldsymbol{x}^*$ when the starting point $\boldsymbol{x}^{(0)}$ is in $\mathcal{O}_1$ and convergence to $\boldsymbol{y}^*$ when $\boldsymbol{x}^{(0)}$ is in $\mathcal{O}_2$. Note that the absence of convergence guarantee does not imply divergence (the sublevel set $\mathrm{lev}_{<\Theta(\boldsymbol{x})}\Theta$ is unbounded when $\boldsymbol{x} \notin \mathcal{O}_1 \cup \mathcal{O}_2$, and so there is neither global nor local convergence guarantee when starting outside $\mathcal{O}_1 \cup \mathcal{O}_2$).

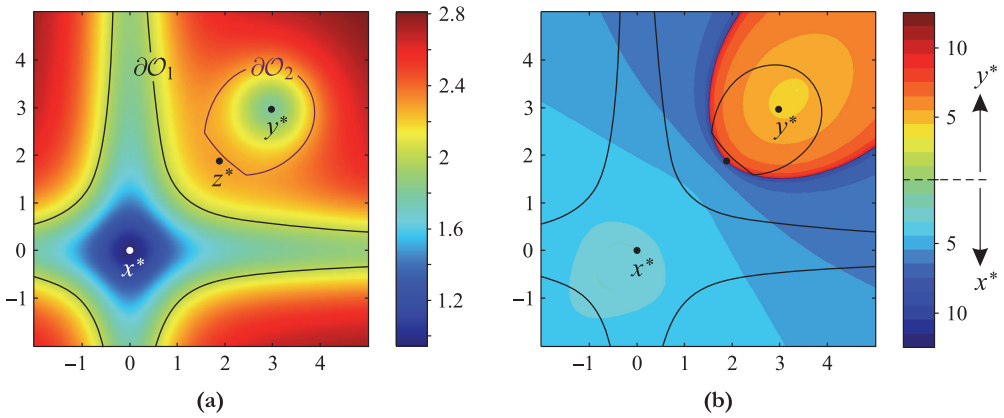

Figure 2. *A two-dimensional objective functional with bounded and unbounded supremum cups:* (a) *functional defined in* (4.3) *and boundaries of the supremum cups* $\mathcal{O}_{\sup}(\boldsymbol{x}^*) =: \mathcal{O}_1$ *and* $\mathcal{O}_{\sup}(\boldsymbol{y}^*) =: \mathcal{O}_2$*;* (b) *number of iterations to reach a* 10^{-6} *distance to either* $\boldsymbol{x}^*$ *or* $\boldsymbol{y}^*$ *as a function of the starting position.*

Corollary 4.17 supplements Corollary 4.16. It states how condition (C4) can be weakened and is illustrated in Example 3 using the motivation example presented in the introduction.

Corollary 4.17. *Let* $\mathcal{O}$ *be a cup of* Θ *with bottom* $\boldsymbol{x}^*$*. For every* $\boldsymbol{x} \in \mathcal{O}$*, define the set* $\mathcal{J}_{\boldsymbol{\varepsilon}}(\boldsymbol{x}) := \big\{k \in [1\,..\,K] : \varepsilon_k(\boldsymbol{x}) > 0\big\}$*, where* ε_k *is given in* (2.2)*. If*

$$\bigcup_{\boldsymbol{x} \in \mathcal{O}} \bigcap_{k \in \mathcal{J}_{\boldsymbol{\varepsilon}}(\boldsymbol{x})} \mathrm{null}(\boldsymbol{A}_k) = \{\mathbf{0}\}, \tag{4.4}$$

then $\mathcal{O} \subseteq \mathcal{A}(\boldsymbol{x}^*)$.

Proof. By the beginning of the proof of Proposition 2.4, condition (4.4) holds if and only if $\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})\boldsymbol{A}$ is positive definite everywhere in $\mathcal{O}$. So we need only show that $\mathcal{O}$ is stable with respect to the recurrence (2.5). Seeking a contradiction, suppose there is an $\boldsymbol{x} \in \mathcal{O}$ such that

$$\boldsymbol{z} := (\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})\boldsymbol{A})^{-1}\boldsymbol{A}^T\boldsymbol{E}(\boldsymbol{x})\boldsymbol{a} \notin \mathcal{O}.$$

Let $\overline{\Theta}$ be defined as in (2.15). By the proof of Proposition 2.6, $\boldsymbol{z}$ is the unique global minimizer of $\overline{\Theta}(\boldsymbol{x}, \cdot)$, and Θ is majorized by $\overline{\Theta}(\boldsymbol{x}, \cdot)$ with equality at $\boldsymbol{x}$. So

$$\overline{\Theta}(\boldsymbol{x}, \boldsymbol{z}) < \overline{\Theta}(\boldsymbol{x}, \boldsymbol{x}) = \Theta(\boldsymbol{x}) \quad \text{and} \quad \Theta(\boldsymbol{y}) \leqslant \overline{\Theta}(\boldsymbol{x}, \boldsymbol{y}) \quad \text{for all } \boldsymbol{y} \in \mathbb{R}^N. \tag{4.5}$$

For every $\alpha \in [0, 1]$, we let $\boldsymbol{y}_\alpha := (1 - \alpha)\boldsymbol{x} + \alpha\boldsymbol{z}$. Since $\boldsymbol{y}_0 = \boldsymbol{x} \in \mathcal{O}$ and $\boldsymbol{y}_1 = \boldsymbol{z} \notin \mathcal{O}$, the supremum $\alpha^* := \sup\{\alpha \in [0, 1] : \boldsymbol{y}_\alpha \in \mathcal{O}\}$ is well-defined. The point $\boldsymbol{y}_{\alpha^*}$ lies on the boundary of $\mathcal{O}$, and thus $\Theta(\boldsymbol{x}) < \Theta(\boldsymbol{y}_{\alpha^*})$ by Lemma 4.8(i). At the same time, using (4.5) and the convexity of $\overline{\Theta}(\boldsymbol{x}, \cdot)$, we have

$$\begin{aligned}
\Theta(\boldsymbol{y}_{\alpha^*}) &\leqslant \overline{\Theta}(\boldsymbol{x}, \boldsymbol{y}_{\alpha^*}) \\
&\leqslant (1 - \alpha^*)\overline{\Theta}(\boldsymbol{x}, \boldsymbol{x}) + \alpha^*\overline{\Theta}(\boldsymbol{x}, \boldsymbol{z}) \\
&\leqslant \max\left\{\overline{\Theta}(\boldsymbol{x}, \boldsymbol{x}), \overline{\Theta}(\boldsymbol{x}, \boldsymbol{z})\right\} \\
&= \Theta(\boldsymbol{x}),
\end{aligned}$$

a contradiction. ∎

Remark 8. By the definition of ε_k and Lemma 2.1, $\mathcal{J}_{\boldsymbol{\varepsilon}}(\boldsymbol{x})$ is the set of indices k such that the potential function θ_k is strictly increasing on $[0, \tau_k)$ for some $\tau_k > \|\boldsymbol{A}_k\boldsymbol{x} - \boldsymbol{a}_k\|$. Therefore, contrary to condition (C4), the intersection in (4.4) is not restricted to the indices of the strictly increasing potential functions: it is also over the indices of the remaining potential functions whose arguments are sufficiently small.

Example 3. Let $\boldsymbol{x}^*$ be an isolated local minimizer of a functional of the form of (1.4) with θ_{fid} strictly increasing, θ_{reg} eventually constant, and $\text{null}(\boldsymbol{D}) \neq \{\boldsymbol{0}\}$ (so condition (C4) fails). By Lemma 2.1(ii), there is a $\tau > 0$ such that θ_{reg} is strictly increasing on $[0, \tau)$ and constant on $[\tau, +\infty)$. For every $\boldsymbol{x} \in \mathbb{R}^N$, we let

$$\mathcal{J}_{\text{reg}}(\boldsymbol{x}) := \left\{l \in [1\,..\,L] : \|\boldsymbol{R}_l\boldsymbol{x}\| < \tau\right\} \quad \text{and} \quad \mathcal{N}_{\text{reg}}(\boldsymbol{x}) := \bigcap_{l \in \mathcal{J}_{\text{reg}}(\boldsymbol{x})} \text{null}(\boldsymbol{R}_l).$$

Suppose that

$$\text{null}(\boldsymbol{D}) \cap \mathcal{N}_{\text{reg}}(\boldsymbol{x}^*) = \{\boldsymbol{0}\} \tag{4.6}$$

(in other words, for every $\boldsymbol{y} \neq \boldsymbol{0}$ such that $\boldsymbol{D}\boldsymbol{y} = \boldsymbol{0}$, there is an index l such that $\boldsymbol{R}_l\boldsymbol{y} \neq \boldsymbol{0}$ and $\|\boldsymbol{R}_l\boldsymbol{x}^*\| < \tau$). Since the maps $\boldsymbol{x} \longmapsto \|\boldsymbol{R}_l\boldsymbol{x}\|$ are continuous, there is an $r > 0$ such that $\mathcal{J}_{\text{reg}}(\boldsymbol{x}^*) \subseteq \mathcal{J}_{\text{reg}}(\boldsymbol{x})$ for all $\boldsymbol{x}$ in the open ball $\mathcal{B}(\boldsymbol{x}^*, r)$. Using Lemma 4.7, we deduce that there is a cup $\mathcal{O}$ with bottom $\boldsymbol{x}^*$ such that

$$\bigcup_{\boldsymbol{x} \in \mathcal{O}} \left(\text{null}(\boldsymbol{D}) \cap \mathcal{N}_{\text{reg}}(\boldsymbol{x})\right) = \{\boldsymbol{0}\}.$$

It follows from Corollary 4.17 and Remark 8 that if (4.6) holds, then Algorithm 1 generates feasible sequences converging to $\boldsymbol{x}^*$ when initialized close enough to this point.

Since cups generally do not partition $\mathbb{R}^N$, we are left with the question of whether a cup is actually reached. The most favorable situation is when the three following conditions are satisfied: (C4) holds, the set $\mathcal{S}$ of stationary points is discrete, and Algorithm 1 is initialized in a bounded sublevel set. In this case, Theorem 3.9 guarantees convergence to a point in $\mathcal{S}$, and the majorization-minimization interpretation yields that $\Theta(\boldsymbol{x}^{(p+1)}) < \Theta(\boldsymbol{x}^{(p)})$ whenever $\boldsymbol{x}^{(p)} \notin \mathcal{S}$. Therefore, if a saddle point is encountered, any perturbation will move the iterates away from it, and so any practical implementation will eventually reach a cup because of the round-off errors inherent in floating-point arithmetic (note that maximizers are also unstable but cannot be attained unless chosen for initialization). If the three conditions listed above are not all satisfied, we can still greatly increase the confidence in reaching a cup by using a continuation scheme to guide the first steps of the optimization process (as in our experiments) or by adding a stochastic perturbation sequence to avoid saddle points and shallow minimizers. These two techniques can be combined and are inspired by the graduated nonconvexity framework presented in [37] and the random perturbation method proposed in [38].

4.3. Supremum cups and water-flooding analogy. This section is intended to characterize supremum cups for better understanding the local convergence properties of Algorithm 1. In particular, we establish that supremum cups share the properties of cups stated in Definition 4.4 and Lemma 4.8, except for boundedness. We also show that if the intersection of a supremum cup $\mathcal{C}_{\sup}(\boldsymbol{x}^*)$ with a level set of Θ is nonempty, then this intersection is either $\{\boldsymbol{x}^*\}$ or the boundary of a cup. Together with Corollary 4.16, these results yield a handy water-flooding interpretation of the capture properties of half-quadratic optimization. However, they do not contribute directly to the discussion on practical convergence and hence can be skipped by the reader more interested in numerical experiments.

Proposition 4.18. *Let $\boldsymbol{x}^*$ be an isolated local minimizer of Θ. The supremum cup $\mathcal{C}_{\sup}(\boldsymbol{x}^*)$ is open and connected, and*

$$\mathcal{C}_{\sup}(\boldsymbol{x}^*) \cap \mathcal{S} = \{\boldsymbol{x}^*\}.$$

Proof. By Lemmas 4.7 and 4.8(iii), $\Omega(\boldsymbol{x}^*)$ is nonempty and the cups in $\Omega(\boldsymbol{x}^*)$ are connected open sets containing $\boldsymbol{x}^*$. So $\mathcal{C}_{\sup}(\boldsymbol{x}^*)$ is a nonempty connected open set. Furthermore, $\boldsymbol{x}^*$ is the only stationary point in $\mathcal{C}_{\sup}(\boldsymbol{x}^*)$ because $\mathcal{C} \cap \mathcal{S} = \{\boldsymbol{x}^*\}$ for every $\mathcal{C} \in \Omega(\boldsymbol{x}^*)$. ∎

Notation. We denote by $H(\boldsymbol{x}^*)$ the supremum of the heights of the cups with bottom $\boldsymbol{x}^*$; that is,

$$H(\boldsymbol{x}^*) := \sup\big\{H(\mathcal{C}) : \mathcal{C} \in \Omega(\boldsymbol{x}^*)\big\} \in (0, +\infty]. \tag{4.7}$$

For every $h \in (\Theta(\boldsymbol{x}^*), H(\boldsymbol{x}^*))$, we let $\mathcal{C}_h(\boldsymbol{x}^*)$ be the unique cup with bottom $\boldsymbol{x}^*$ and height h (the existence and uniqueness of such cups follow from Lemma 4.10).

Proposition 4.19 draws an analogy between supremum cups and water-flooded regions. Think of Θ as an altitude function and of $\boldsymbol{x}^*$ as the bottom of a valley to be filled with water. Then $H(\boldsymbol{x}^*)$ is the critical water level above which either water overflows out of the valley

or the volume of water is infinite, and $\mathcal{O}_{\sup}(\boldsymbol{x}^*)$ is the land submerged when the water level reaches $H(\boldsymbol{x}^*)$. Proposition 4.19 is used in the proof of Proposition 4.20 to complete our description of supremum cups.

Proposition 4.19.

$$\mathcal{O}_{\sup}(\boldsymbol{x}^*) = \bigcup_{h\in(\Theta(\boldsymbol{x}^*),H(\boldsymbol{x}^*))} \mathcal{O}_h(\boldsymbol{x}^*), \tag{4.8}$$

and for every $h \in (\Theta(\boldsymbol{x}^*), H(\boldsymbol{x}^*))$,

$$\mathcal{O}_{\sup}(\boldsymbol{x}^*) \cap \operatorname{lev}_{<h}\Theta = \mathcal{O}_h(\boldsymbol{x}^*), \tag{4.9a}$$

$$\mathcal{O}_{\sup}(\boldsymbol{x}^*) \cap \operatorname{lev}_{=h}\Theta = \partial\mathcal{O}_h(\boldsymbol{x}^*). \tag{4.9b}$$

Proof. Let $\omega(\boldsymbol{x}^*) := \{\mathcal{O}_h(\boldsymbol{x}^*) : h \in (\Theta(\boldsymbol{x}^*), H(\boldsymbol{x}^*))\}$. Clearly, $\omega(\boldsymbol{x}^*) \subseteq \Omega(\boldsymbol{x}^*)$ with equality if and only if there is no cup with bottom $\boldsymbol{x}^*$ and height $H(\boldsymbol{x}^*)$, in which case (4.8) holds trivially. Suppose there is cup $\mathcal{O}^* \in \Omega(\boldsymbol{x}^*)$ such that $H(\mathcal{O}^*) = H(\boldsymbol{x}^*)$. We complete the proof of (4.8) by showing that $\mathcal{O}^* \subseteq \bigcup_{\mathcal{O}\in\omega(\boldsymbol{x}^*)} \mathcal{O}$. Let $\boldsymbol{x} \in \mathcal{O}^*$. Since $\mathcal{O}^* \subseteq \operatorname{lev}_{<H(\boldsymbol{x}^*)}\Theta$ by Lemma 4.8(i), we have $\Theta(\boldsymbol{x}) < H(\boldsymbol{x}^*)$. Let $h \in (\Theta(\boldsymbol{x}), H(\boldsymbol{x}^*))$. Then $\mathcal{O}_h(\boldsymbol{x}^*) = \mathcal{O}^* \cap \operatorname{lev}_{<h}\Theta$ by Lemma 4.10, and hence $\boldsymbol{x} \in \mathcal{O}_h(\boldsymbol{x}^*) \in \omega(\boldsymbol{x}^*)$.

Let $h \in (\Theta(\boldsymbol{x}^*), H(\boldsymbol{x}^*))$. Using (4.8), we have

$$\mathcal{O}_{\sup}(\boldsymbol{x}^*) \cap \operatorname{lev}_{<h}\Theta = \bigcup_{l\in(\Theta(\boldsymbol{x}^*),H(\boldsymbol{x}^*))} \big(\underbrace{\mathcal{O}_l(\boldsymbol{x}^*) \cap \operatorname{lev}_{<h}\Theta}_{=:\,\mathcal{O}_l^h(\boldsymbol{x}^*)}\big).$$

If $l > h$, then $\mathcal{O}_l^h(\boldsymbol{x}^*) = \mathcal{O}_h(\boldsymbol{x}^*)$ by Lemma 4.10. If $l \leqslant h$, then $\mathcal{O}_l^h(\boldsymbol{x}^*) = \mathcal{O}_l(\boldsymbol{x}^*) \subseteq \mathcal{O}_h(\boldsymbol{x}^*)$ by Lemmas 4.8(i) and 4.9. Therefore, $\mathcal{O}_{\sup}(\boldsymbol{x}^*) \cap \operatorname{lev}_{<h}\Theta$ is a union of subsets of $\mathcal{O}_h(\boldsymbol{x}^*)$ which includes $\mathcal{O}_h(\boldsymbol{x}^*)$ itself; hence we get (4.9a).

Seeking a contradiction, suppose there is an $\boldsymbol{x} \in \mathcal{O}_{\sup}(\boldsymbol{x}^*) \cap \operatorname{lev}_{=h}\Theta$ such that $\boldsymbol{x} \notin \partial\mathcal{O}_h(\boldsymbol{x}^*)$. Then $\boldsymbol{x} \notin \overline{\mathcal{O}_h(\boldsymbol{x}^*)}$ by (4.9a), and since $\mathcal{O}_{\sup}(\boldsymbol{x}^*)$ and $\mathbb{R}^N \setminus \overline{\mathcal{O}_h(\boldsymbol{x}^*)}$ are open, there is an $\alpha > 0$ such that $\mathcal{B}(\boldsymbol{x}, \alpha) \subset \mathcal{O}_{\sup}(\boldsymbol{x}^*) \setminus \overline{\mathcal{O}_h(\boldsymbol{x}^*)}$. Using (4.9a) again, it follows that $\mathcal{B}(\boldsymbol{x}, \alpha) \cap \operatorname{lev}_{<h}\Theta = \emptyset$. So $\boldsymbol{x}$ is a local minimizer of Θ and $\boldsymbol{x} \neq \boldsymbol{x}^*$, a contradiction with the fact that $\boldsymbol{x}^*$ is the only stationary point in $\mathcal{O}_{\sup}(\boldsymbol{x}^*)$. This proves that $\mathcal{O}_{\sup}(\boldsymbol{x}^*) \cap \operatorname{lev}_{=h}\Theta \subseteq \partial\mathcal{O}_h(\boldsymbol{x}^*)$. For the reverse inclusion, it suffices to show that $\partial\mathcal{O}_h(\boldsymbol{x}^*) \subset \mathcal{O}_{\sup}(\boldsymbol{x}^*)$, as $\partial\mathcal{O}_h(\boldsymbol{x}^*) \subseteq \operatorname{lev}_{=h}\Theta$ by the definition of $\mathcal{O}_h(\boldsymbol{x}^*)$. Let $\boldsymbol{x} \in \partial\mathcal{O}_h(\boldsymbol{x}^*)$ and $l \in (h, H(\boldsymbol{x}^*))$. We have $\mathcal{O}_h(\boldsymbol{x}^*) \subset \mathcal{O}_l(\boldsymbol{x}^*)$ by Lemma 4.9, and thus $\boldsymbol{x} \in \overline{\mathcal{O}_h(\boldsymbol{x}^*)} \subset \overline{\mathcal{O}_l(\boldsymbol{x}^*)}$. But $\boldsymbol{x} \notin \partial\mathcal{O}_l(\boldsymbol{x}^*)$ since $\Theta(\boldsymbol{x}) = h < l$. Hence $\boldsymbol{x} \in \mathcal{O}_l(\boldsymbol{x}^*) \subset \mathcal{O}_{\sup}(\boldsymbol{x}^*)$, which completes the proof of (4.9b). ∎

Proposition 4.20.

(i) *If* $\mathcal{O}_{\sup}(\boldsymbol{x}^*) \neq \mathbb{R}^N$, *then* Θ *is equal to* $H(\boldsymbol{x}^*)$ *on* $\partial\mathcal{O}_{\sup}(\boldsymbol{x}^*)$.

(ii) $\arg\max_{\overline{\mathcal{O}_{\sup}(\boldsymbol{x}^*)}} \Theta = \partial\mathcal{O}_{\sup}(\boldsymbol{x}^*)$.

(iii) $\arg\min_{\overline{\mathcal{O}_{\sup}(\boldsymbol{x}^*)}} \Theta = \{\boldsymbol{x}^*\}$.

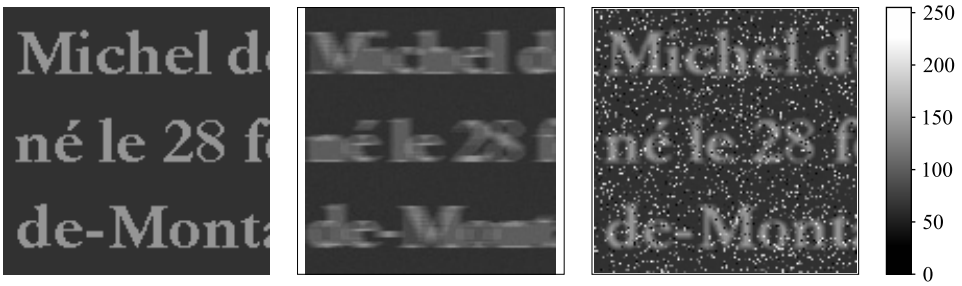

Figure 3. *"Text" image:* (a) *original;* (b) *degraded by* 1×9 *uniform blur and* 25 dB *white Gaussian noise;* (c) *degraded by* 3×3 *uniform blur,* 30 dB *white Gaussian noise, and* 15% *random-valued impulse noise.*

Proof. (i) Suppose $\mathcal{O}_{\sup}(\boldsymbol{x}^*) \neq \mathbb{R}^N$. From (4.8), we have

$$\mathcal{O}_{\sup}(\boldsymbol{x}^*) \subseteq \bigcup_{h \in (\Theta(\boldsymbol{x}^*), H(\boldsymbol{x}^*))} \operatorname{lev}_{<h} \Theta = \operatorname{lev}_{<H(\boldsymbol{x}^*)} \Theta. \tag{4.10}$$

Let $\boldsymbol{y} \in \partial \mathcal{O}_{\sup}(\boldsymbol{x}^*)$. By the definition of the boundary of a set, there is a sequence $(\boldsymbol{y}^{(p)})_p$ in $\mathcal{O}_{\sup}(\boldsymbol{x}^*)$ converging to $\boldsymbol{y}$. The sequence $(\Theta(\boldsymbol{y}^{(p)}))_p$ converges to $\Theta(\boldsymbol{y})$ by the continuity of Θ, and $\Theta(\boldsymbol{y}^{(p)}) < H(\boldsymbol{x}^*)$ for all p by (4.10); thus $\Theta(\boldsymbol{y}) \leqslant H(\boldsymbol{x}^*)$. We now prove by contradiction that $\Theta(\boldsymbol{y}) \geqslant H(\boldsymbol{x}^*)$. Suppose $\Theta(\boldsymbol{y}) < H(\boldsymbol{x}^*)$, and let $h \in (\Theta(\boldsymbol{y}), H(\boldsymbol{x}^*))$. From (4.9), we have

$$\overline{\mathcal{O}_h(\boldsymbol{x}^*)} = \mathcal{O}_{\sup}(\boldsymbol{x}^*) \cap \operatorname{lev}_{\leqslant h} \Theta.$$

Since $\boldsymbol{y} \notin \mathcal{O}_{\sup}(\boldsymbol{x}^*)$, there is an $\alpha > 0$ such that $\mathcal{B}(\boldsymbol{y}, \alpha) \cap \overline{\mathcal{O}_h(\boldsymbol{x}^*)} = \emptyset$. Therefore, if p is sufficiently large,

$$\boldsymbol{y}^{(p)} \in \mathcal{O}_{\sup}(\boldsymbol{x}^*) \setminus \overline{\mathcal{O}_h(\boldsymbol{x}^*)} = \mathcal{O}_{\sup}(\boldsymbol{x}^*) \cap \operatorname{lev}_{>h} \Theta.$$

So $(\Theta(\boldsymbol{y}^{(p)}))_p$ cannot converge to $\Theta(\boldsymbol{y})$, a contradiction.

(ii) If $\mathcal{O}_{\sup}(\boldsymbol{x}^*) \neq \mathbb{R}^N$, then property (ii) follows directly from (i) and (4.10). If $\mathcal{O}_{\sup}(\boldsymbol{x}^*) = \mathbb{R}^N$, then $\partial \mathcal{O}_{\sup}(\boldsymbol{x}^*)$ is empty and Θ has no global maximizer (since a global maximizer is a stationary point and the only stationary point is the strict local minimizer $\boldsymbol{x}^*$).

(iii) Let $\boldsymbol{x} \in \overline{\mathcal{O}_{\sup}(\boldsymbol{x}^*)} \setminus \{\boldsymbol{x}^*\}$. If $\boldsymbol{x} \in \mathcal{O}_{\sup}(\boldsymbol{x}^*)$, then $\boldsymbol{x}$ is in a cup with bottom $\boldsymbol{x}^*$, and so $\Theta(\boldsymbol{x}) > \Theta(\boldsymbol{x}^*)$ by Lemma 4.8(ii). If $\boldsymbol{x} \in \partial \mathcal{O}_{\sup}(\boldsymbol{x}^*)$, then $\Theta(\boldsymbol{x}) = H(\boldsymbol{x}^*)$ by (i), and so $\Theta(\boldsymbol{x}) > \Theta(\boldsymbol{x}^*)$ by (4.10). ∎

5. Experiments. We consider the problem of restoring the 128×128 images displayed in Figures 3(a) and 4(a) from synthetic data generated using the observation model (1.3) described in the introduction. The pixel intensity in the original "text" image is either 50 (background) or 150 (text regions). The data shown in Figure 3(b) were obtained by blurring with a 1×9 uniform kernel and adding white Gaussian noise at 25 dB SNR (the signal-to-noise ratio defines the standard deviation ς of the noise term via the relation $\mathrm{SNR}_{\mathrm{dB}} := 20 \log_{10}(\varsigma^\sharp/\varsigma)$, where $\varsigma^\sharp$ is the sample standard deviation of the noise-free observation $\boldsymbol{D}\boldsymbol{x}^\sharp$). The data displayed in Figure 3(c) were generated by first blurring with a 3×3 uniform kernel, then adding 30 dB white Gaussian noise, and finally degrading the result by random-valued

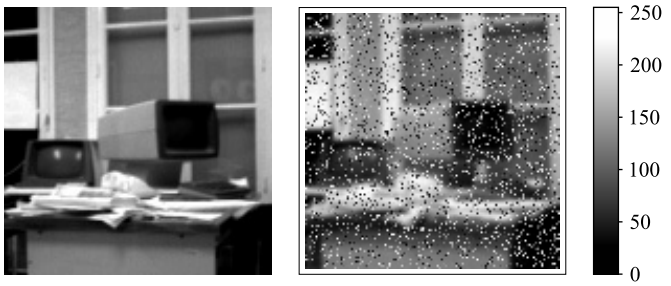

Figure 4. *"Office" image:* (a) *original;* (b) *degraded by a* 7×7*,* 1*-pixel standard deviation Gaussian filter,* 30 dB *white Gaussian noise, and* 20% *random-valued impulse noise.*

impulse noise with range $[0, 255]$ and 15% corruption rate—that is, the data pixels are of the form

$$\boldsymbol{d}_m = \begin{cases} (\boldsymbol{D}\boldsymbol{x}^\sharp + \boldsymbol{\nu})_m & \text{with probability } 0.85, \\ \xi_m & \text{with probability } 0.15, \end{cases}$$

where ξ_m is sampled uniformly at random from $[0, 255]$. The "office" image data shown in Figure 4(b) were obtained in a similar way by blurring with a 7×7 rotationally symmetric Gaussian kernel with a standard deviation of one pixel and corrupting the result with 30 dB white Gaussian noise and 20% random-valued impulse noise. In each example, the blurring kernel has odd dimensions and blurring is performed by convolution without zero-padded edges (as in [2]); so the pixel grid of the data is a subset of the pixel grid $\mathsf{G}^\sharp$ of the original image (the observations in Figures 3(b), 3(c), and 4(b) are of sizes 128×120, 126×126, and 122×122, respectively, and the square boxes around each of them delimit $\mathsf{G}^\sharp$).

In all our experiments, the restorations are obtained by minimizing objective functionals of the form of (1.1) using Algorithm 1 with 100 iterations. We assess restoration quality with the structural similarity (SSIM) index [39] and the improvement in SNR (ISNR). The SSIM index is computed using the MATLAB implementation available at http://ece.uwaterloo.ca/~z70wang/research/ssim with default settings. The ISNR of an estimate $\hat{\boldsymbol{x}}$ of $\boldsymbol{x}^\sharp$ is defined by

$$\text{ISNR}_{\text{dB}} := 20 \log_{10} \left(\frac{\| \boldsymbol{x}^\sharp|_{\mathsf{G}} - \boldsymbol{d} \|}{\| (\boldsymbol{x}^\sharp - \hat{\boldsymbol{x}})|_{\mathsf{G}} \|} \right), \tag{5.1}$$

where $\boldsymbol{x}|_{\mathsf{G}}$ denotes the restriction of $\boldsymbol{x}$ to the pixel grid G of the data. Note that, for simplicity, we make no distinction between images and their vector representations.

5.1. Overview. In section 5.2, we present our experiments on the restoration of the "text" image degraded by 1×9 uniform blur and Gaussian noise. For this example problem, the objective functional is the sum of the squared ℓ_2-norm of the residual and a bounded gradient-based regularizer. The regularizer is chosen so that condition (C4) fails, a difficulty circumvented by introducing a continuation sequence at the beginning of the optimization process. Next, in section 5.3, we consider the restoration of the "text" image degraded by 3×3 uniform blur and Gaussian plus impulse noise. To deal with the impulse noise, we replace the ℓ_2

data-fidelity term by a strictly convex approximation to the ℓ_1-norm of the residual. We emphasize the instability caused by the regularizer, and we discuss the stabilizing effect of adding either a Tikhonov penalty or a wavelet-sparsity penalty to the objective functional. Finally, in section 5.4, we compare different regularization schemes for restoring the "office" image, including convex, nonconvex, and bounded gradient-based regularizers and their association with Tikhonov, wavelet-sparsity, and Hessian-based regularizers.

It should be stressed that our goal here is to show that Algorithm 1 behaves well for various instances of the objective functional (1.1). A thorough comparison of existing regularization strategies is beyond the scope of this paper. Likewise, we do not address the choice of the parameters that control the strengths of the regularization terms.

5.2. "Text" image corrupted by Gaussian noise. Consider the problem of restoring the "text" image from the data shown in Figure 3(b). The blurred observation is corrupted by additive white Gaussian noise, and the original image is piecewise constant. This suggests combining ℓ_2 data-fidelity with gradient regularization. We focus on the difficult case where the regularizer is bounded and has flat regions: the solutions are the minimizers of the functional $\Theta_{\ell_2} : \mathbb{R}^N \longrightarrow \mathbb{R}$, $N = 128^2$, defined by

$$\Theta_{\ell_2}(\boldsymbol{x}) := \|\boldsymbol{D}\boldsymbol{x} - \boldsymbol{d}\|^2 + \gamma \sum_l \vartheta_{\mathrm{TB}}\big(\|\boldsymbol{G}_l\boldsymbol{x}\|/\delta\big), \qquad \gamma, \delta \in (0, +\infty),$$

where ϑ_{TB} is the Tukey biweight function given in (1.10), and $\{\boldsymbol{G}_l\}_l$ is the discrete gradient operator defined as follows. Let (l_1, l_2) and $x(l_1, l_2)$, respectively, denote the coordinates and the value of the lth pixel in $\boldsymbol{x}$ (the indices l_1 and l_2 range from 1 to 128, and $l = 128(l_1 - 1) + l_2$). For every l,

$$\boldsymbol{G}_l\boldsymbol{x} := \begin{pmatrix} x_{\mathrm{h}}(l_1, l_2) \\ x_{\mathrm{v}}(l_1, l_2) \end{pmatrix} \in \mathbb{R}^2, \tag{5.2}$$

where $x_{\mathrm{h}}(l_1, l_2)$ and $x_{\mathrm{v}}(l_1, l_2)$ are the horizontal and vertical differences at the lth pixel location:

$$x_{\mathrm{h}}(l_1, l_2) := \begin{cases} x(l_1, l_2) - x(l_1, l_2 - 1) & \text{if } l_2 \geqslant 2, \\ 0 & \text{otherwise,} \end{cases}$$

$$x_{\mathrm{v}}(l_1, l_2) := \begin{cases} x(l_1, l_2) - x(l_1 - 1, l_2) & \text{if } l_1 \geqslant 2, \\ 0 & \text{otherwise.} \end{cases}$$

With this definition, the gradient $\boldsymbol{G}_1\boldsymbol{x}$ at the upper left pixel is always zero; so the sum in the regularizer is over $[2\,..\,N]$.

For every $t \geqslant 0$, we define

$$\mathcal{G}_+(t) := \big\{\boldsymbol{x} \in \mathbb{R}^N : \|\boldsymbol{G}_l\boldsymbol{x}\| > t \text{ for all } l \geqslant 2\big\}. \tag{5.3}$$

Let $\tau := \delta\sqrt{6}$. The open set $\mathcal{G}_+(\tau)$ is a flat region of the regularizer, as $\vartheta_{\mathrm{TB}}(\|\boldsymbol{G}_l\boldsymbol{x}\|/\delta) = 1$ if the gradient magnitude $\|\boldsymbol{G}_l\boldsymbol{x}\|$ is above the threshold τ. We set $\delta = 20/\sqrt{6}$ to preserve

edges with magnitude greater than $\tau = 20$, and we call γ the *spatial smoothing strength*, for the regularizer penalizes spatial gradients to favor piecewise-smooth solutions.

Condition (C4) reduces here to $\mathrm{null}(\boldsymbol{D}) = \{\boldsymbol{0}\}$; but since $\boldsymbol{D}$ is a low-pass filtering operator, it has a nontrivial null space composed of high frequency band images. So, according to Proposition 2.4, Algorithm 1 is not well-defined. A simple way to get around this problem is to run the first iterations with ϑ_{TB} replaced by strictly increasing potential functions $\vartheta^{(p)}$ such that the optimization difficulty increases gradually with p. The expected effect of this continuation technique is to reach a cup of Θ_{ℓ_2} satisfying condition (4.4) prior to running iterations with the target potential function ϑ_{TB}. (By Example 3 following Corollary 4.17, condition (4.4) holds for a cup with bottom $\boldsymbol{x}^*$ if every image $\boldsymbol{y} \in \mathrm{null}(\boldsymbol{D}) \setminus \{\boldsymbol{0}\}$ has a pixel l where $\boldsymbol{G}_l\boldsymbol{y} \neq \boldsymbol{0}$ and $\|\boldsymbol{G}_l\boldsymbol{x}^*\| < \tau$, which is the case if $\boldsymbol{x}^*$ is piecewise smooth.) We use a continuation sequence of the form

$$\vartheta^{(p)} := \kappa_p \vartheta_{\mathrm{TB}} + (1 - \kappa_p)\vartheta_{\mathrm{MS}}, \tag{5.4}$$

where ϑ_{MS} is the function of minimal surfaces defined in (1.6), and κ_p increases linearly from 0 to 1 during the first half of the iterations (after which $\vartheta^{(p)} = \vartheta_{\mathrm{TB}}$):

$$\kappa_p := \begin{cases} (p-1)/50 & \text{if } p \in [1\,..\,50], \\ 1 & \text{if } p > 50. \end{cases}$$

The initial solution $\boldsymbol{x}^{(0)}$ is taken in $\mathrm{null}(\boldsymbol{D}) \cap \mathcal{G}_+(\tau)$ so that the algorithm cannot be started without continuation. (We obtain $\boldsymbol{x}^{(0)}$ by first computing the null space component $\boldsymbol{v}$ of a uniform noise image using the method described in [40] and then setting $\boldsymbol{x}^{(0)} = 2\tau\boldsymbol{v}/\min_{l \geqslant 2} \|\boldsymbol{G}_l\boldsymbol{v}\|$ so that $\|\boldsymbol{G}_l\boldsymbol{x}^{(0)}\| \geqslant 2\tau$ for all $l \geqslant 2$.)

Figures 5(a) and 5(c) display the SSIM and the ISNR of the restorations obtained by varying γ in a range covering both underregularization ($\gamma < 3$) and overregularization ($\gamma > 30$). The maximum SSIM and ISNR, namely 0.9885 and 20.06 dB, are attained simultaneously for $\gamma \approx 6.3$ (the corresponding estimate is shown in the upper left of Figure 6). At first glance, the sharp drop after the maximum ISNR may be interpreted as instability of the algorithm, but the enlargements in Figures 5(b) and 5(d) reveal a smooth behavior. The local SSIM and ISNR variations, whether of small or large amplitude, actually reflect the changes in the objective landscape as γ varies—not instability. This assertion is further supported by Figure 6: the restored image varies smoothly as γ increases in the interval where the drop occurs, which suggests the distortion of a basin of attraction rather than a change in the topology of the objective landscape (both behaviors are exemplified in Appendix E). We conclude that a simple continuation sequence suffices to properly initialize the algorithm in the extreme case where the regularizer has flats.

5.3. "Text" image corrupted by Gaussian plus impulse noise. Since ℓ_2 data-fidelity is not adapted to impulse noise (see, e.g., [8, 12, 41]), the functional Θ_{ℓ_2} used above is not appropriate for the restoration of the "text" image data shown in Figure 3(c). So we now consider an objective functional Θ_{ℓ_1} similar to Θ_{ℓ_2} but with the data-fidelity term replaced by a strictly convex approximation to the ℓ_1-norm of the residual:

$$\Theta_{\ell_1}(\boldsymbol{x}) := \|\boldsymbol{D}\boldsymbol{x} - \boldsymbol{d}\|_{1,\epsilon} + \gamma \sum_l \vartheta_{\mathrm{TB}}\big(\|\boldsymbol{G}_l\boldsymbol{x}\|/\delta\big),$$

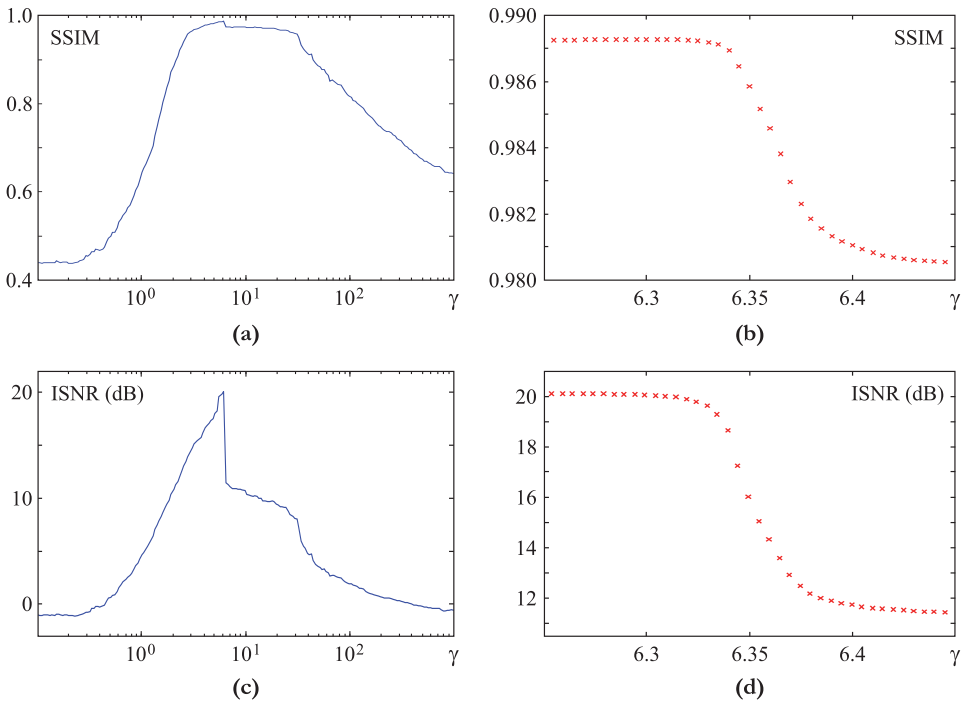

Figure 5. *Restoration of the "text" image degraded by* 1×9 *uniform blur and Gaussian noise:* (a) *SSIM as a function of the spatial smoothing strength* γ*;* (b) *zoom in on of the small drop following the maximum SSIM;* (c) *ISNR as a function of* γ*;* (d) *zoom in on the sharp drop following the maximum ISNR.*

where $\|\cdot\|_{1,\epsilon} : \mathbb{R}^M \longrightarrow \mathbb{R}$ is defined by

$$\|\boldsymbol{u}\|_{1,\epsilon} := \sum_{m=1}^{M} \mathrm{id}_\epsilon(u_m), \qquad \mathrm{id}_\epsilon(t) := (\epsilon^2 + t^2)^{1/2} - \epsilon.$$

In the present context, $\|\cdot\|_{1,\epsilon}$ behaves almost like the ℓ_1-norm when ϵ is two or more orders of magnitude smaller than the width of the original intensity range—we set $\epsilon = 0.1$ throughout the rest of the experiments.

5.3.1. Restoration with gradient regularization alone. Figure 7(a) shows the best solution obtained by minimizing Θ_{ℓ_1} with the same continuation strategy and initial solution as in section 5.2. This estimate, say $\boldsymbol{x}^*$, is close to the original image $\boldsymbol{x}^\sharp$, but the restoration process is unstable because the regularizer yields deep spurious minimizers. For example, the minimizer $\boldsymbol{y}^*$ displayed in Figure 7(b) is a solution of the *forward system* $\boldsymbol{D}\boldsymbol{x} = \boldsymbol{d}$ and belongs to $\mathcal{G}_+(\tau)$. Therefore, $\Theta_{\ell_1}(\boldsymbol{y}^*) = \gamma(N-1)$, which is about 5.7×10^4 for the value of γ that gives $\boldsymbol{x}^*$ (namely, $\gamma \approx 3.5$). By contrast, $\Theta_{\ell_1}(\boldsymbol{x}^*) \approx 2.0 \times 10^5$, and so $\boldsymbol{y}^*$ is much deeper than $\boldsymbol{x}^*$, although it has nothing in common with $\boldsymbol{x}^\sharp$. The existence of such spurious minimizers is explained as follows. Since $\boldsymbol{D}$ has full rank, the forward system is consistent and its general solution is the sum of a particular solution $\boldsymbol{x}_\mathrm{p}$ and a vector in $\mathrm{null}(\boldsymbol{D})$. Furthermore, since $\mathrm{null}(\boldsymbol{D}) \setminus \{\boldsymbol{0}\}$ is composed of high frequency band images,

$$\mathcal{N}_+ := \mathrm{null}(\boldsymbol{D}) \cap \mathcal{G}_+(0) \neq \emptyset.$$

Figure 6. *Restorations of the "text" image degraded by* 1×9 *uniform blur and Gaussian noise for increasing values of* γ *in the interval considered in Figures* 5(b) *and* 5(d).

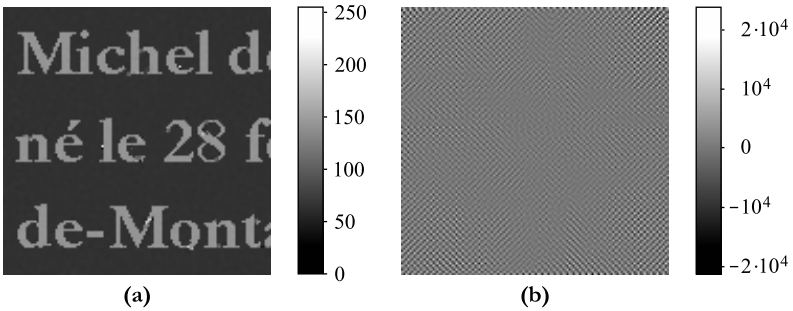

Figure 7. *Restoration of the "text" image degraded by* 3×3 *uniform blur and Gaussian plus impulse noise:* (a) *best result achieved by varying the spatial smoothing strength (*SSIM $= 0.9796$, ISNR $= 13.46$ dB*);* (b) *deep spurious minimizer.*

For every $\boldsymbol{v} \in \mathcal{N}_+$, we can choose $\alpha > 0$ large enough so that $\boldsymbol{x}_{\rm p} + \alpha\boldsymbol{v} \in \mathcal{G}_+(\tau)$, and thus

$$\mathcal{S}_{\boldsymbol{v}} := \big\{\boldsymbol{x}_{\rm p} + \alpha\boldsymbol{v} : \alpha \in \mathbb{R}\big\} \cap \mathcal{G}_+(\tau) \neq \emptyset.$$

The set $\mathcal{S}_{\boldsymbol{v}}$ is contained in the solution set of the forward system, and the regularizer is constant in a neighborhood of each point in $\mathcal{S}_{\boldsymbol{v}}$. Consequently,

$$\mathcal{S}_+ := \bigcup_{\boldsymbol{v}\in\mathcal{N}_+} \mathcal{S}_{\boldsymbol{v}} \tag{5.5}$$

is a set of local minimizers of Θ_{ℓ_1}.

5.3.2. Association with Tikhonov and wavelet-sparsity regularization. We discuss here two techniques which allow us to exclude spurious minimizers and which guarantee global convergence by ensuring coercivity—indeed, coercivity implies condition (C4) (see Proposition 2.5) as well as boundedness of the sublevel sets, and hence it implies the conditions common to the global convergence theorems (Theorems 3.7, 3.9, and 3.12).

The first technique was introduced in Remark 1 after Proposition 2.5. It consists in adding a Tikhonov penalty to the objective functional: the solutions are minimizers of

$$\Theta^{\mathrm{T}}_{\ell_1}(\boldsymbol{x}) \;:=\; \Theta_{\ell_1}(\boldsymbol{x}) + \rho\|\boldsymbol{x}\|^2, \qquad \rho \in (0, +\infty).$$

We call ρ the *Tikhonov strength*; this parameter must be small enough so that $\Theta^{\mathrm{T}}_{\ell_1}(\boldsymbol{x}) \approx \Theta_{\ell_1}(\boldsymbol{x})$ when $\boldsymbol{x}$ is piecewise smooth, for otherwise the solutions are pulled toward $\mathbf{0}$. A rule of thumb is to begin with a value one or two orders of magnitude smaller than the ratio $\Theta_{\ell_1}(x^\sharp_{\max}\mathbf{1})/\|x^\sharp_{\max}\mathbf{1}\|^2$, where $x^\sharp_{\max}$ is the presumed maximum absolute value of the pixels in $\boldsymbol{x}^\sharp$, and $\mathbf{1}$ denotes the constant image with value 1. We may then increase or decrease ρ depending on whether the restoration has unexpected high-frequency content or is oversmoothed.

The second technique is to add a penalty $\|\boldsymbol{T}\boldsymbol{x}\|_{1,\epsilon}$ defined by a sparsifying transform $\boldsymbol{T}$ such that the resulting functional $\boldsymbol{x} \longmapsto \Theta_{\ell_1}(\boldsymbol{x}) + \|\boldsymbol{T}\boldsymbol{x}\|_{1,\epsilon}$ is coercive, which is the case if $\mathrm{null}(\boldsymbol{D}) \cap \mathrm{null}(\boldsymbol{T}) = \{\mathbf{0}\}$. A natural choice is the wavelet transform, as piecewise smoothness translates to sparse wavelet coefficients. Let $\boldsymbol{W} \in \mathbb{R}^{N\times N}$ be a nonredundant, J-level wavelet transform: given an image $\boldsymbol{x}$, the vector $\boldsymbol{W}\boldsymbol{x}$ is the concatenation of $\boldsymbol{V}_J\boldsymbol{x}, \boldsymbol{W}_1\boldsymbol{x}, \ldots, \boldsymbol{W}_J\boldsymbol{x}$, where $\boldsymbol{V}_J\boldsymbol{x}$ is the approximation of $\boldsymbol{x}$ at resolution 2^{-J}, and for every decomposition level $j \in [1\,..\,J]$, the components of $\boldsymbol{W}_j\boldsymbol{x}$ are the wavelet coefficients of $\boldsymbol{x}$ at resolution 2^{-j} (that is, $\boldsymbol{W}_j\boldsymbol{x}$ contains the information in $\boldsymbol{x}$ that is lost when going from resolution 2^{-j+1} to resolution 2^{-j}). Restorations with sparse wavelet representations are obtained by minimizing

$$\Theta^{\mathrm{W}}_{\ell_1,J}(\boldsymbol{x}) \;:=\; \Theta_{\ell_1}(\boldsymbol{x}) \;+\; \widetilde{\gamma}\sum_{j=1}^{J}\|2^{-j+1}\boldsymbol{W}_j\boldsymbol{x}\|_{1,\epsilon}, \qquad \widetilde{\gamma} \in (0, +\infty),$$

where the *wavelet-sparsity strength* $\widetilde{\gamma}$ controls the degree of sparsity of the minimizers, and the weights 2^{-j+1} compensate for the decay of the wavelet-coefficient magnitudes as the resolution increases [8]. For this objective functional, condition (2.8) in Proposition 2.5 becomes

$$\mathrm{null}(\boldsymbol{D}) \cap \mathcal{N}_J = \{\mathbf{0}\}, \qquad \mathcal{N}_J := \bigcap_{j=1}^{J}\mathrm{null}(\boldsymbol{W}_j). \tag{5.6}$$

The vector space $\mathcal{N}_J$ is the set of low-pass filtered images obtained by interpolating every possible approximation at resolution 2^{-J} back to full resolution (that is, resolution 1). Therefore, since the images in $\mathrm{null}(\boldsymbol{D})$ have strong high-frequency content, condition (5.6) holds, and hence $\Theta^{\mathrm{W}}_{\ell_1,J}$ is coercive. Note that large values of $\widetilde{\gamma}$ pull the solutions toward $\mathcal{N}_J$ rather than $\mathbf{0}$. In this respect, the minimizers of $\Theta^{\mathrm{W}}_{\ell_1,J}$ are less sensitive to $\widetilde{\gamma}$ than those of $\Theta^{\mathrm{T}}_{\ell_1}$ are to ρ. In our experiments below, we use a single decomposition level (that is, we set $J = 1$) and biorthogonal spline wavelets with two vanishing moments [42].

The Tikhonov and wavelet-sparsity regularizers penalize the images in $\mathcal{G}_+(\tau)$ much more strongly than for piecewise-smooth images (e.g., for the "good" and "bad" minimizers $\boldsymbol{x}^*$ and

$\boldsymbol{y}^*$ shown in Figure 7, we have $\|\boldsymbol{y}^*\|^2/\|\boldsymbol{x}^*\|^2 \approx 5.1\times10^3$ and $\|\boldsymbol{W}_1\boldsymbol{y}^*\|_{1,\epsilon}/\|\boldsymbol{W}_1\boldsymbol{x}^*\|_{1,\epsilon} \approx 6.4\times10^2$). Therefore, both regularizers are appropriate to exclude the spurious minimizers in the set $\mathcal{S}_+$ defined in (5.5). Besides, since $\Theta^{\mathrm{T}}_{\ell_1}$ and $\Theta^{\mathrm{W}}_{\ell_1,1}$ are coercive, there is theoretically no need for the continuation strategy used so far—yet we continue to use it because it provides a good starting point for the algorithm.

Figure 8 displays the SSIM and the ISNR of the restorations obtained by minimizing $\Theta^{\mathrm{T}}_{\ell_1}$ and $\Theta^{\mathrm{W}}_{\ell_1,1}$ when varying the Tikhonov and wavelet-sparsity strengths (the spatial smoothing strength is kept constant at the value that gives $\boldsymbol{x}^*$, the best achieved minimizer of Θ_{ℓ_1}). Both additional regularizers improve restoration quality if their strengths are not too high. When the Tikhonov strength increases above its optimum value, the SSIM and the ISNR decrease rapidly to the values corresponding to the zero image (namely, 1.2×10^{-3} and $-4.94\,\mathrm{dB}$, respectively). On the other hand, when increasing the wavelet-sparsity strength, the SSIM and the ISNR stabilize around 0.86 and 8 dB. Indeed, increasing $\widetilde{\gamma}$ brings the minimizers closer to null($\boldsymbol{W}_1$) (the vector space of the images that contain no detail information between resolutions 1 and $\frac{1}{2}$), and in the limit when $\widetilde{\gamma} \longrightarrow +\infty$, minimizing $\Theta^{\mathrm{W}}_{\ell_1,1}$ amounts to searching for minimizers of Θ_{ℓ_1} on null($\boldsymbol{W}_1$). The best restorations achieved by adding Tikhonov and wavelet-sparsity regularization are shown in Figures 9(a) and 9(b), respectively. They are similar, and their ISNR values are more than 6 dB higher than that of $\boldsymbol{x}^*$. The increase in SSIM is less marked because the perceptual improvement compared to $\boldsymbol{x}^*$ lies in the smoothing of a few outliers.

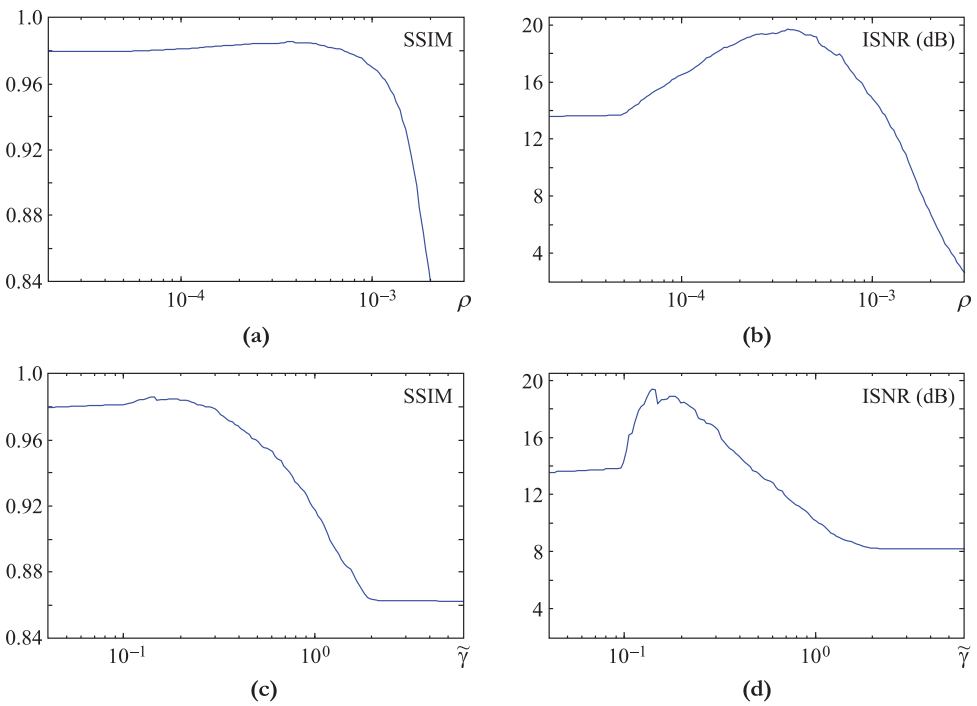

Figure 8. *Restoration of the "text" image degraded by* 3×3 *uniform blur and Gaussian plus impulse noise:* (a)–(b) *SSIM and ISNR as functions of the Tikhonov strength when minimizing* $\Theta^{\mathrm{T}}_{\ell_1}$*;* (c)–(d) *SSIM and ISNR as functions of the wavelet-sparsity strength when minimizing* $\Theta^{\mathrm{W}}_{\ell_1,1}$.

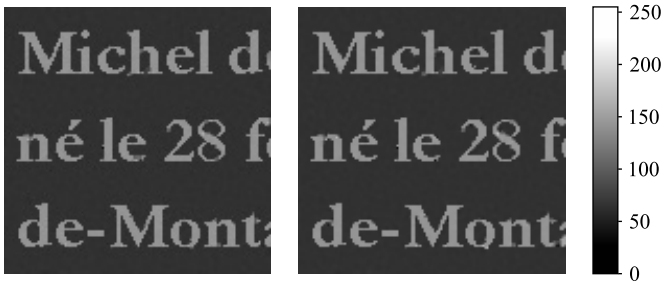

Figure 9. *Restoration of the "text" image degraded by* 3×3 *uniform blur and Gaussian plus impulse noise: best results obtained by adding* (a) *Tikhonov regularization (*SSIM $= 0.9852$, ISNR $= 19.72$ dB*) and* (b) *wavelet-sparsity regularization (*SSIM $= 0.9858$, ISNR $= 19.42$ dB*).*

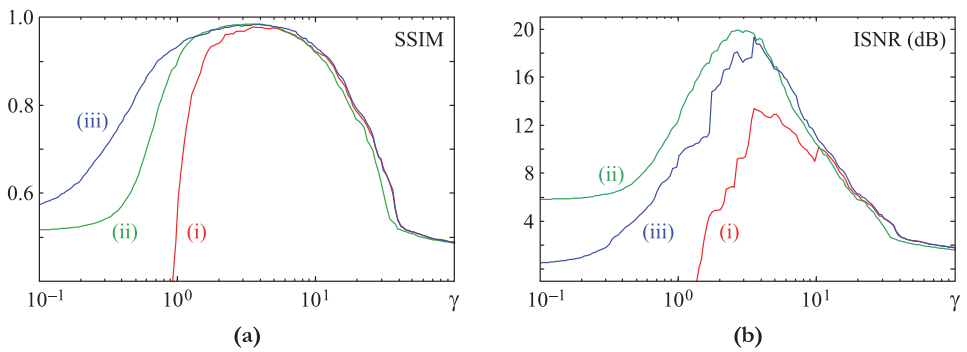

Figure 10. *Restoration of the "text" image degraded by* 3×3 *uniform blur and Gaussian plus impulse noise. SSIM and ISNR as functions of the spatial smoothing strength:* (i) *without Tikhonov or wavelet-sparsity regularization;* (ii) *with Tikhonov regularization;* (iii) *with wavelet-sparsity regularization.*

We can also observe the impact of Tikhonov and wavelet-sparsity regularization by keeping their strengths constant while varying that of the gradient regularizer. Figure 10 compares the SSIM and ISNR versus γ curves obtained by minimizing Θ_{ℓ_1} (in red), $\Theta^{\mathrm{T}}_{\ell_1}$ (in green), and $\Theta^{\mathrm{W}}_{\ell_1,1}$ (in blue). The Tikhonov and wavelet-sparsity regularizers improve restoration quality up to a certain value of γ (about 10) above which their effects are negligible. Furthermore, both reduce the sensitivity to γ in two senses. First, they maintain a minimum quality when gradient regularization is too weak: when γ goes to zero, the SSIM decreases to 0.5109 with Tikhonov regularization and to 0.5325 with wavelet sparsification, whereas it goes to zero when neither is used. Second, they enlarge the range of values of γ that yield estimates with SSIM (or ISNR) above a given value: for example, the length of the interval of values of γ yielding an SSIM greater than 0.9 is 0.91 decades when minimizing Θ_{ℓ_1}, 1.05 decades when adding Tikhonov regularization, and 1.27 decades with wavelet sparsification.

5.4. "Office" image. We conclude this paper with the restoration of the "office" image from the data shown in Figure 4(b). All the results presented in this section are the best obtained by varying the regularization strength parameters.

As an alternative to the ℓ_1-norm of the residual, we deal with Gaussian plus impulse noise by using the Huber function ϑ_{Hu} defined in (1.7), a choice rooted in robust statistics [21]. The

instances of (1.1) considered here are of the form

$$\Theta(\boldsymbol{x}) \;=\; \sum_{m=1}^{M} \vartheta_{\mathrm{Hu}}\big(|(\boldsymbol{D}\boldsymbol{x}-\boldsymbol{d})_m|/\varsigma\big) \;+\; \Theta_{\mathrm{reg},1}(\boldsymbol{x}) \;+\; \Theta_{\mathrm{reg},2}(\boldsymbol{x}), \tag{5.7}$$

where $\varsigma \approx 2.12$ is the standard deviation of the Gaussian noise component, and the regularizers $\Theta_{\mathrm{reg},1}$ and $\Theta_{\mathrm{reg},2}$ impose different types of constraints. The first regularizer is gradient- or Hessian-based:

$$\Theta_{\mathrm{reg},1}(\boldsymbol{x}) \;=\; \gamma_1 \sum_{l} \vartheta\big(\|\boldsymbol{R}_l\boldsymbol{x}\|/\delta\big),$$

where $\{\boldsymbol{R}_l\}_l$ is either the gradient operator $\{\boldsymbol{G}_l \in \mathbb{R}^{2\times N}\}_l$ given in (5.2) or the Hessian operator $\{\boldsymbol{H}_l \in \mathbb{R}^{3\times N}\}_l$ that defines Frobenius-norm regularization [14] (Hessian regularization was not introduced in the previous experiments because it is not adapted to piecewise-constant images). In either case, we consider convex, nonconvex, and bounded regularization by setting $\vartheta = \mathrm{id}_\epsilon$, $\vartheta = \vartheta_{\mathrm{LE}}$, and $\vartheta = \vartheta_{\mathrm{GM}}$, respectively ($\vartheta_{\mathrm{LE}}$ and ϑ_{GM} are the Lorentzian error and the Geman and McClure functions defined in (1.8) and (1.9)). When $\vartheta = \mathrm{id}_\epsilon$, the scaling parameter δ is set to 1, and thus $\Theta_{\mathrm{reg},1}(\boldsymbol{x}) \approx \gamma_1 \sum_l \|\boldsymbol{R}_l\boldsymbol{x}\|$; in particular, if $\{\boldsymbol{R}_l\}_l$ is the gradient operator, then $\Theta_{\mathrm{reg},1}(\boldsymbol{x})/\gamma_1$ is a smooth approximation to the discrete total variation (TV) of $\boldsymbol{x}$. When $\vartheta = \vartheta_{\mathrm{LE}}$ or ϑ_{GM}, we set $\delta = 10$ for gradient regularization and $\delta = 10\sqrt{2}$ for Hessian regularization, where the factor $\sqrt{2}$ accounts for the fact that $\|\boldsymbol{H}_l\boldsymbol{x}\| = \sqrt{2}\|\boldsymbol{G}_l\boldsymbol{x}\|$ when the lth pixel in $\boldsymbol{x}$ is adjacent to a horizontal or vertical step-edge. The second regularizer is either zero or one of the following:

$$\Theta_{\mathrm{reg},2}(\boldsymbol{x}) \;=\; \begin{cases} \gamma_2\|\boldsymbol{x}\|^2 & \text{(Tikhonov regularizer)},\\ \gamma_2\|\boldsymbol{W}_1\boldsymbol{x}\|_{1,\epsilon} & \text{(wavelet-sparsity regularizer)},\\ \gamma_2\sum_l \mathrm{id}_\epsilon\big(\|\boldsymbol{H}_l\boldsymbol{x}\|\big) & \text{(convex Hessian regularizer)}. \end{cases}$$

If $\Theta_{\mathrm{reg},2}$ is the Tikhonov regularizer or the wavelet sparsifier, Θ is coercive independently of the choice of ϑ in $\Theta_{\mathrm{reg},1}$ (the reasons are the same as those given in section 5.3.2). The objective functional is also coercive when $\Theta_{\mathrm{reg},2}$ is the convex Hessian regularizer. Indeed, the null space of the vertical concatenation $\boldsymbol{H}$ of the matrices $\boldsymbol{H}_l$ is the set of planar images (we call $\boldsymbol{x}$ planar if there are real numbers a_1, a_2, a_3 such that $x(l_1,l_2) = a_1l_1 + a_2l_2 + a_3$ for all (l_1,l_2)). Therefore, since $\boldsymbol{D}$ is a low-pass filtering operator, we have $\mathrm{null}(\boldsymbol{D}) \cap \mathrm{null}(\boldsymbol{H}) = \{\boldsymbol{0}\}$, and coercivity follows from Proposition 2.5.

The only case when Θ is not coercive is when $\vartheta = \vartheta_{\mathrm{GM}}$ and $\Theta_{\mathrm{reg},2} = 0$, but since condition (C4) holds, we can introduce a continuation sequence to reach a bounded sublevel set of Θ and hence guarantee global convergence. Continuation is actually helpful when Θ is nonconvex, whether coercive or not; so we use this technique whenever $\vartheta = \vartheta_{\mathrm{LE}}$ or $\vartheta = \vartheta_{\mathrm{GM}}$ (the corresponding sequences of potential functions are similar to that defined in (5.4) but with ϑ_{TB} replaced by either ϑ_{LE} or ϑ_{GM} depending on which is considered).

Figure 11 shows the restorations obtained using gradient regularization alone—the corresponding SSIM and ISNR values are listed in Table 1 (in the row labeled "Gradient"), which summarizes our experiments on the "office" image. The solution associated with the

convex potential function is fairly good; it contains no outlier and is approximately piecewise smooth, and so we expect little if no improvement from the introduction of Tikhonov or wavelet-sparsity regularization. We indeed observe that Tikhonov regularization is useless and that wavelet sparsification improves restoration quality only slightly. The restoration obtained by adding the wavelet sparsifier is shown in Figure 12(a). It looks similar to that achieved with convex gradient regularization alone, but closer examination reveals a reduced staircase effect and smoother contour lines (by "smooth contour lines" we mean piecewise-smooth level curves, not smooth edges).

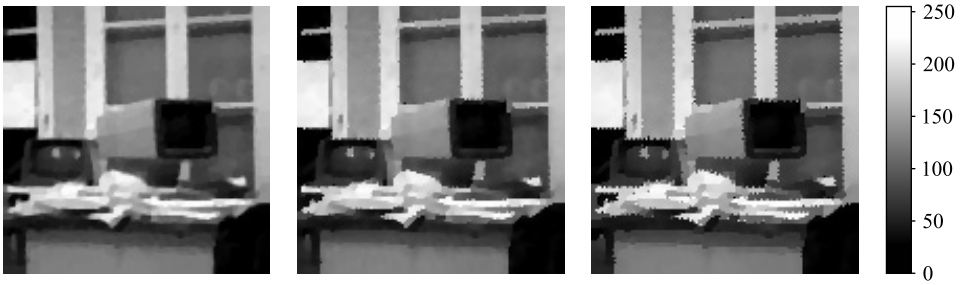

Figure 11. *Restoration of the "office" image with different gradient regularizers:* (a) *smooth approximation to TV (*SSIM = 0.9292*);* (b) *Lorentzian error (*SSIM = 0.9069*);* (c) *Geman and McClure (*SSIM = 0.8650*).*

Table 1

Best SSIM *and* ISNR *(*dB*) obtained by minimizing the objective functional* (5.7) *for various regularization schemes (WS stands for "wavelet sparsity").*

		Potential function in the first regularizer		
First regularizer (+ second regularizer)	Quality measure	id_ϵ (convex)	ϑ_{LE} (nonconvex)	ϑ_{GM} (bounded)
Gradient	SSIM	0.9292	0.9069	0.8650
	ISNR	14.56	11.42	8.94
Gradient + Tikhonov	SSIM	0.9292	0.9095	0.8833
	ISNR	14.56	12.12	10.66
Gradient + WS	SSIM	0.9356	0.9308	0.9264
	ISNR	15.04	14.22	13.90
Gradient + Hessian	SSIM	0.9415	0.9421	0.9436
	ISNR	15.94	15.95	15.78
Hessian	SSIM	0.9368	0.9345	0.9216
	ISNR	15.19	14.60	12.07
Hessian + WS	SSIM	0.9368	0.9356	0.9306
	ISNR	15.19	14.65	13.80

In contrast with the convex case, nonconvex gradient regularization produces noisy contour lines, an effect even more pronounced if the potential function is bounded. This deterioration

Figure 12. *Restoration of the "office" image with wavelet sparsification added to different gradient regularizers:* (a) *smooth approximation to TV (*SSIM = 0.9356*);* (b) *Lorentzian error (*SSIM = 0.9308*);* (c) *Geman and McClure (*SSIM = 0.9264*).*

in image quality is clearly visible in Figures 11(b) and 11(c) and translates to smaller SSIM and ISNR. Since the Tikhonov penalty does not distinguish between smooth and noisy contour lines, it is not surprising that additional Tikhonov regularization does not improve the situation much: although it slightly increases the SSIM and the ISNR, the restorations obtained with or without it have similar noisy boundaries. The wavelet sparsifier, on the other hand, brings real improvements. The solutions obtained using nonconvex gradient regularization together with wavelet sparsification are shown in Figures 12(b) and 12(c). Compared to the results achieved with nonconvex gradient regularization alone, the contour lines are smooth (as in the original image) and the increases in SSIM and ISNR are significant.

Figure 13. *Restoration of the "office" image with convex Hessian regularization added to different gradient regularizers:* (a) *smooth approximation to TV (*SSIM = 0.9415*);* (b) *Lorentzian error (*SSIM = 0.9421*);* (c) *Geman and McClure (*SSIM = 0.9436*).*

We can further improve restoration quality by combining gradient and Hessian regularization. The resulting estimates are shown in Figure 13; they are qualitatively and quantitatively very close to the original image, especially considering the poor quality of the data. This suggests that gradient plus Hessian regularization is the best adapted to piecewise-smooth images. To support this claim, we performed restorations using Hessian regularization either alone or with wavelet sparsification. The corresponding SSIM and ISNR values are reported in the last two rows of Table 1; they are all smaller than those achieved with gradient plus Hessian regularization. (We also observe that wavelet sparsification brings negligible improvement to

Hessian regularization. The reason is that Hessian regularization alone produces smooth contour lines even when the potential function is nonconvex.) More importantly, gradient plus Hessian regularization is the most stable: contrary to the other regularization schemes, there is almost no difference between the restorations obtained with the three potential functions considered. It is also noteworthy that the bounded potential function offers a slight advantage in terms of SSIM. This again confirms that Algorithm 1 behaves well in difficult situations where the objective functional contains a bounded regularizer.

Appendix A. List of symbols and acronyms. Below is a summary of the symbols and acronyms used in this paper. As a general rule, we denote matrices by bold upper-case roman letters, vectors by bold lower-case roman letters, and sets by calligraphic upper-case letters. The interior, the closure, and the boundary of a set $\mathcal{G}$ are denoted by $\mathcal{G}^\circ$, $\overline{\mathcal{G}}$, and $\partial\mathcal{G}$, respectively.

$\boldsymbol{A}_k$, $\boldsymbol{a}_k$	Matrices and vectors in the definition (1.1) of the objective functional
$\boldsymbol{A}$, $\boldsymbol{a}$	Vertical concatenations of the matrices $\boldsymbol{A}_k$ and of the vectors $\boldsymbol{a}_k$
$\mathcal{A}(\boldsymbol{x}^*)$	Basin of attraction of a stationary point $\boldsymbol{x}^*$ (Definition 4.14)
$\mathcal{B}(\boldsymbol{x}, \alpha)$	Open ball with center $\boldsymbol{x}$ and radius α
$\mathcal{C}_\mathcal{X}$	Set of cluster points of the sequence $\mathcal{X}$
$\boldsymbol{D}$, $\boldsymbol{d}$	Observation matrix and data vector, (1.3)
$\boldsymbol{E}(\boldsymbol{x})$	Diagonal matrix of the weighting coefficients at point $\boldsymbol{x}$, (2.3)
$\{\boldsymbol{G}_l\}_l$	Discrete gradient operator, (5.2)
$\mathcal{G}_+(t)$	Set of images with gradient norm everywhere greater than t, (5.3)
$H(\mathcal{O})$	Height of a cup $\mathcal{O}$ (Definition 4.4)
$H(\boldsymbol{x}^*)$	Supremum of the heights of the cups with bottom $\boldsymbol{x}^*$, (4.7)
ISNR	Improvement in SNR, (5.1)
$\mathcal{J}_+$	Index set of the strictly increasing potential functions, (C4)
$\mathcal{L}_0$	Sublevel set at the initial height $\Theta(\boldsymbol{x}^{(0)})$, (1.13)
lev	Level, sublevel, and superlevel sets, (4.2)
null	Null-space operator
$\mathcal{O}$	Cup (Definition 4.4)
$\mathcal{O}_h(\boldsymbol{x}^*)$	Unique cup with bottom $\boldsymbol{x}^*$ and height h
$\mathcal{O}_{\sup}(\boldsymbol{x}^*)$	Supremum cup with bottom $\boldsymbol{x}^*$ (Definition 4.15)
$\{\boldsymbol{R}_l\}_l$	Discrete regularization operator, (1.4)
$\mathcal{S}$	Set of stationary points of the objective functional, (1.12)
$\mathcal{S}_\Xi$	Set of stationary points of a differentiable functional Ξ
SNR	Signal-to-noise ratio
SSIM	Structural similarity
TV	Total variation
$\boldsymbol{x}^\sharp$	Original signal, (1.3)
$\mathcal{X}$	Sequence $(\boldsymbol{x}^{(p)})_p$ generated by Algorithm 1

$\varepsilon_k(\boldsymbol{x})$	kth weighting coefficient at point $\boldsymbol{x}$, (2.2)
θ, ϑ	Potential functions
$\theta^\dagger$	Weighting function $t \in (0,+\infty) \longmapsto t^{-1}\theta'(t)$
$\theta^\ddagger$	Continuous extension of $\theta^\dagger$ to $\mathbb{R}_+$, (2.1)
Θ	Objective functional of the form of (1.1)
$\Theta^\ddagger$	Half-quadratic functional associated with Θ, (2.6a)
$\overline{\Theta}$	Surrogate generator of the majorization-minimization interpretation of Algorithm 1, (2.15)
Ξ	Continuous functional on some subset of $\mathbb{R}^N$ (Theorem 3.4) or differentiable functional on $\mathbb{R}^N$ (section 4.1)
$\boldsymbol{\Phi}$	Iteration map of Algorithm 1, (2.11b)
$\Omega(\boldsymbol{x}^*)$	Set of all the cups with bottom $\boldsymbol{x}^*$ (Definition 4.15)
∇	Gradient operator in $\mathbb{R}^N$

Appendix B. Sufficient conditions for strict convexity. Propositions B.1 and B.2 provide sufficient conditions for strict convexity of functionals of the form of (1.1)—the former when the potential functions $\theta_1, \dots, \theta_K$ are all convex and the latter when some of the potential functions are nonconvex. We conclude this appendix with the example of image restoration using ℓ_2 data-fidelity and nonconvex regularization under periodic boundary conditions.

Notation. Given a nonempty set $\mathcal{J} \subseteq [1\,..\,K]$, we denote by $\boldsymbol{A}_\mathcal{J}$ any vertical concatenation of the matrices $\boldsymbol{A}_k$ indexed by $\mathcal{J}$, that is,

$$\boldsymbol{A}_\mathcal{J} := \begin{pmatrix} \boldsymbol{A}_{k_1}^T & \cdots & \boldsymbol{A}_{k_j}^T \end{pmatrix}^T \quad \text{with} \quad \{k_1, \dots, k_j\} = \mathcal{J}.$$

Proposition B.1. *Let $\Theta : \mathbb{R}^N \longrightarrow \mathbb{R}$ be a functional of the form of (1.1) with potential functions $\theta_1, \dots, \theta_K$ increasing and convex. If the set*

$$\mathcal{K}_+ := \big\{k \in [1\,..\,K] : \theta_k \text{ is strictly convex}\big\}$$

is nonempty and

$$\operatorname{null}(\boldsymbol{A}_{\mathcal{K}_+}) = \{\boldsymbol{0}\},$$

then Θ is strictly convex.

Proof. Suppose all the potential functions are increasing and convex. Let $\boldsymbol{x}$ and $\boldsymbol{y}$ be distinct points in $\mathbb{R}^N$, and let $\alpha \in (0,1)$. For every $k \in [1\,..\,K]$,

$$\begin{aligned} &\theta_k\big(\|\boldsymbol{A}_k(\alpha\boldsymbol{x} + (1-\alpha)\boldsymbol{y}) - \boldsymbol{a}_k\|\big) \\ &\qquad \leqslant \alpha\theta_k\big(\|\boldsymbol{A}_k\boldsymbol{x} - \boldsymbol{a}_k\|\big) + (1-\alpha)\theta_k\big(\|\boldsymbol{A}_k\boldsymbol{y} - \boldsymbol{a}_k\|\big). \end{aligned} \tag{B.1}$$

So Θ is strictly convex if this inequality is strict for some k which may depend on $\boldsymbol{x}$ and $\boldsymbol{y}$.

Suppose further that $\operatorname{null}(\boldsymbol{A}_{\mathcal{K}_+}) = \{\boldsymbol{0}\}$, and let $k \in \mathcal{K}_+$ be such that $\boldsymbol{y} - \boldsymbol{x} \notin \operatorname{null}(\boldsymbol{A}_k)$. If $\|\boldsymbol{A}_k\boldsymbol{x} - \boldsymbol{a}_k\| \neq \|\boldsymbol{A}_k\boldsymbol{y} - \boldsymbol{a}_k\|$, then inequality (B.1) is strict by the strict convexity of θ_k. Now consider the case where $\|\boldsymbol{A}_k\boldsymbol{x} - \boldsymbol{a}_k\| = \|\boldsymbol{A}_k\boldsymbol{y} - \boldsymbol{a}_k\|$. Let

$$\boldsymbol{v} := \boldsymbol{A}_k\boldsymbol{x} - \boldsymbol{a}_k \quad \text{and} \quad \boldsymbol{w} := \boldsymbol{A}_k(\boldsymbol{y} - \boldsymbol{x}).$$

Clearly, $\|\boldsymbol{v}+\boldsymbol{w}\| = \|\boldsymbol{v}\|$, and thus $\|\boldsymbol{w}\|^2 = -2\langle \boldsymbol{v}, \boldsymbol{w}\rangle$, where $\langle\cdot,\cdot\rangle$ denotes the Euclidean inner product. It follows that

$$\begin{aligned}\|\boldsymbol{v}+(1-\alpha)\boldsymbol{w}\|^2 &= \|\boldsymbol{v}\|^2 + 2(1-\alpha)\langle \boldsymbol{v},\boldsymbol{w}\rangle + (1-\alpha)^2\|\boldsymbol{w}\|^2 \\ &= \|\boldsymbol{v}\|^2 - \alpha(1-\alpha)\|\boldsymbol{w}\|^2 \\ &< \|\boldsymbol{v}\|^2.\end{aligned}$$

Therefore, since θ_k is strictly increasing (for it is increasing and strictly convex),

$$\theta_k\big(\|\boldsymbol{v}+(1-\alpha)\boldsymbol{w}\|\big) < \theta_k\big(\|\boldsymbol{v}\|\big) = \alpha\theta_k\big(\|\boldsymbol{v}\|\big) + (1-\alpha)\theta_k\big(\|\boldsymbol{v}+\boldsymbol{w}\|\big),$$

which proves that inequality (B.1) is strict. ∎

Notation. Suppose the potential functions $\theta_1,\dots,\theta_K$ are twice differentiable with some strictly convex and others nonconvex. Let $\mathcal{J}$ be a nonempty subset of $\mathcal{K}_+$, and let

$$\mathcal{K}_- := \big\{k \in [1\,..\,K] : \theta_k \text{ is nonconvex}\big\}.$$

We define $\underline{\Gamma}(\mathcal{J})$ to be the minimum convexity over the potential functions indexed by $\mathcal{J}$, and we let $\overline{\Gamma}(\mathcal{K}_-)$ be the maximum concavity over the nonconvex potential functions; that is,

$$\underline{\Gamma}(\mathcal{J}) := \min_{k\in\mathcal{J}} \inf_{\mathbb{R}_+} \theta_k'' \quad \text{and} \quad \overline{\Gamma}(\mathcal{K}_-) := \max_{k\in\mathcal{K}_-} \sup_{\mathbb{R}_+}(-\theta_k'').$$

Proposition B.2. *Let $\Theta : \mathbb{R}^N \longrightarrow \mathbb{R}$ be a functional of the form of (1.1). Suppose that the potential functions $\theta_1,\dots,\theta_K$ are twice differentiable, that they satisfy conditions* (C1) *and* (C3)*, and that at least one of them is nonconvex. If there is a nonempty set $\mathcal{J} \subseteq \mathcal{K}_+$ such that*

$$\underline{\Gamma}(\mathcal{J})\,\sigma_{\min}^2(\boldsymbol{A}_{\mathcal{J}}) > \overline{\Gamma}(\mathcal{K}_-)\,\sigma_{\max}^2(\boldsymbol{A}_{\mathcal{K}_-}),$$

where $\sigma_{\min}(\cdot)$ and $\sigma_{\max}(\cdot)$ respectively denote the smallest and largest singular values of their matrix argument, then Θ is strictly convex.

Proof. We show that the Hessian matrix of Θ is positive definite everywhere under the stated assumptions. Let $\boldsymbol{x} \in \mathbb{R}^N$. The Hessian matrix of Θ at $\boldsymbol{x}$ is

$$\nabla^2\Theta(\boldsymbol{x}) = \sum_{k=1}^{K} \theta_k''\big(\|\boldsymbol{A}_k\boldsymbol{x}-\boldsymbol{a}_k\|\big)\,\boldsymbol{A}_k^T\boldsymbol{A}_k$$

(this expression is derived using Lemmas 2.1(iii) and 2.3, which is why we assume that the potential functions satisfy conditions (C1) and (C3) in addition to twice differentiability). For every $\boldsymbol{y} \in \mathbb{R}^N$,

$$\begin{aligned}\boldsymbol{y}^T\nabla^2\Theta(\boldsymbol{x})\,\boldsymbol{y} &= \sum_{k=1}^{K} \theta_k''\big(\|\boldsymbol{A}_k\boldsymbol{x}-\boldsymbol{a}_k\|\big)\,\|\boldsymbol{A}_k\boldsymbol{y}\|^2 \\ &\geqslant \underline{\Gamma}(\mathcal{J}) \sum_{k\in\mathcal{J}} \|\boldsymbol{A}_k\boldsymbol{y}\|^2 - \overline{\Gamma}(\mathcal{K}_-) \sum_{k\in\mathcal{K}_-} \|\boldsymbol{A}_k\boldsymbol{y}\|^2 \\ &= \boldsymbol{y}^T\big(\underbrace{\underline{\Gamma}(\mathcal{J})(\boldsymbol{A}_{\mathcal{J}})^T\boldsymbol{A}_{\mathcal{J}}}_{=:\,\boldsymbol{B}_{\mathcal{J}}} - \underbrace{\overline{\Gamma}(\mathcal{K}_-)(\boldsymbol{A}_{\mathcal{K}_-})^T\boldsymbol{A}_{\mathcal{K}_-}}_{=:\,\boldsymbol{B}_{\mathcal{K}_-}}\big)\,\boldsymbol{y}.\end{aligned}$$

Let $\lambda_{\min}(\boldsymbol{C})$ and $\lambda_{\max}(\boldsymbol{C})$ denote the smallest and largest eigenvalues of a real symmetric matrix $\boldsymbol{C}$. Using Weyl's theorem (see, e.g., Corollary 4.3.15 in [43]), we have

$$\begin{aligned}\lambda_{\min}(\boldsymbol{B}_{\mathcal{J}} - \boldsymbol{B}_{\mathcal{K}_-}) &\geqslant \lambda_{\min}(\boldsymbol{B}_{\mathcal{J}}) + \lambda_{\min}(-\boldsymbol{B}_{\mathcal{K}_-}) \\ &= \underline{\Gamma}(\mathcal{J})\,\lambda_{\min}\big((\boldsymbol{A}_{\mathcal{J}})^T\boldsymbol{A}_{\mathcal{J}}\big) - \overline{\Gamma}(\mathcal{K}_-)\,\lambda_{\max}\big((\boldsymbol{A}_{\mathcal{K}_-})^T\boldsymbol{A}_{\mathcal{K}_-}\big) \\ &= \underline{\Gamma}(\mathcal{J})\,\sigma^2_{\min}(\boldsymbol{A}_{\mathcal{J}}) - \overline{\Gamma}(\mathcal{K}_-)\,\sigma^2_{\max}(\boldsymbol{A}_{\mathcal{K}_-}) \\ &> 0.\end{aligned}$$

Therefore, $\boldsymbol{B}_{\mathcal{J}} - \boldsymbol{B}_{\mathcal{K}_-}$ is positive definite, and hence so is $\nabla^2\Theta(\boldsymbol{x})$. ■

Example 4. Consider the classical example of image restoration by minimizing

$$\Theta_{\gamma,\delta}(\boldsymbol{x}) := \|\boldsymbol{D}\boldsymbol{x} - \boldsymbol{d}\|^2 + \gamma\sum_{l=1}^{L}\vartheta\big(\|\boldsymbol{R}_l\boldsymbol{x}\|/\delta\big), \qquad \gamma,\delta\in(0,+\infty),$$

where $\boldsymbol{D}$ implements a two-dimensional (2-D) convolution, and the vertical concatenation of $\boldsymbol{R}_1,\dots,\boldsymbol{R}_L$ is any ordering of the first-order horizontal and vertical difference operators, that is,

$$\boldsymbol{R} := \begin{pmatrix}\boldsymbol{R}_1^T & \cdots & \boldsymbol{R}_L^T\end{pmatrix}^T = \boldsymbol{P}\begin{pmatrix}\boldsymbol{H}\\ \boldsymbol{V}\end{pmatrix},$$

where $\boldsymbol{P}$ is a permutation matrix, and $\boldsymbol{H}$ and $\boldsymbol{V}$ implement the 2-D convolutions with the masks $\begin{pmatrix}-1 & 1\end{pmatrix}$ and $\begin{pmatrix}-1 & 1\end{pmatrix}^T$.

Suppose the mother potential function ϑ is nonconvex, twice differentiable, and satisfies conditions (C1) and (C3). By Proposition B.2, the functional $\Theta_{\gamma,\delta}$ is strictly convex if

$$\sigma^2_{\min}(\boldsymbol{D}) > \frac{\gamma}{2\delta^2}\sup_{\mathbb{R}_+}(-\vartheta'')\,\sigma^2_{\max}(\boldsymbol{R}). \tag{B.2}$$

Since the singular values of a real matrix $\boldsymbol{C}$ are the square roots of the eigenvalues of $\boldsymbol{C}^T\boldsymbol{C}$, condition (B.2) is of practical interest if we can efficiently estimate the smallest eigenvalue of $\boldsymbol{D}^T\boldsymbol{D}$ and the largest eigenvalue of $\boldsymbol{R}^T\boldsymbol{R} = \boldsymbol{H}^T\boldsymbol{H} + \boldsymbol{V}^T\boldsymbol{V}$. To do so, we further assume periodic boundary conditions for both the data and the regularization models. In this case, $\boldsymbol{D}$, $\boldsymbol{H}$, and $\boldsymbol{V}$ are block-circulant, and thus so are $\boldsymbol{D}^T\boldsymbol{D}$ and $\boldsymbol{R}^T\boldsymbol{R}$. We can then use the fact that the eigenvalues of block-circulant matrices can be computed exactly by using the 2-D discrete Fourier transform (DFT) [44]—the results we need are the following.

Let f be a discrete function defined on $[-M_1\,..\,M_1]\times[-M_2\,..\,M_2]$, and let $\boldsymbol{C}$ be the block-circulant matrix representing the circular convolution of an $N_1\times N_2$ image with f (we assume that $N_1 > 2M_1$ and $N_2 > 2M_2$). The eigenvalues of $\boldsymbol{C}$ are

$$\lambda(n_1,n_2) := \sum_{k_1=-M_1}^{M_1}\sum_{k_2=-M_2}^{M_2} f(k_1,k_2)\exp\left(-2\mathrm{i}\pi\left(\frac{k_1n_1}{N_1}+\frac{k_2n_2}{N_2}\right)\right),$$
$$(n_1,n_2)\in[0\,..\,N_1-1]\times[0\,..\,N_2-1].$$

In other words, the spectrum of $\boldsymbol{C}$ is the 2-D DFT of the $N_1\times N_2$ image obtained by zero-padding f and then performing a circular shift to position $f(0,0)$ in the upper left corner.

Furthermore, if $N_1 > 4M_1$ and $N_2 > 4M_2$, the block-circulant matrix $\boldsymbol{C}^T\boldsymbol{C}$ implements the circular convolution with the autocorrelation function $f \star f$ defined on $[-2M_1 \,..\, 2M_1] \times [-2M_2 \,..\, 2M_2]$ by

$$(f \star f)(k_1, k_2) := \sum_{i \in \mathcal{J}(M_1,k_1)} \sum_{j \in \mathcal{J}(M_2,k_2)} f(i,j)\, f(i-k_1, j-k_2),$$

where $\mathcal{J}(M,k) := [-M + \max\{0,k\} \,..\, M + \min\{0,k\}]$.

Let g be the kernel of the convolution implemented by $\boldsymbol{D}$, and let $\mathfrak{F}(g \star g)$ denote the $N_1 \times N_2$ DFT of the autocorrelation of g. By the above properties of block-circulant matrices, the singular values of $\boldsymbol{D}$ are the square roots of the magnitudes of the Fourier coefficients $\mathfrak{F}(g \star g)(n_1, n_2)$ (since the magnitude of the DFT is invariant to circular shifting in the spatial domain, there is no need to shift the zero-padding of $g \star g$ prior to computing the DFT). It follows that

$$\sigma^2_{\min}(\boldsymbol{D}) \;=\; \min_{(n_1,n_2)} \big|\mathfrak{F}(g \star g)(n_1, n_2)\big|.$$

Besides, $\boldsymbol{H}^T\boldsymbol{H}$ and $\boldsymbol{V}^T\boldsymbol{V}$ implement the convolutions with the masks

$$\begin{pmatrix} -1 & 2 & -1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} -1 & 2 & -1 \end{pmatrix}^T,$$

and thus $\boldsymbol{R}^T\boldsymbol{R}$ represents the convolution with the Laplace kernel

$$\begin{pmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{pmatrix}.$$

Therefore, $\boldsymbol{R}^T\boldsymbol{R}$ has eigenvalues

$$\lambda(n_1, n_2) \;=\; 4 - 2\cos\left(\frac{2\pi n_1}{N_1}\right) - 2\cos\left(\frac{2\pi n_2}{N_2}\right),$$

and hence $\sigma^2_{\max}(\boldsymbol{R}) \leqslant 8$ with equality when N_1 and N_2 are even. So finally, under periodic boundary conditions, $\Theta_{\gamma,\delta}$ is strictly convex if

$$\min_{(n_1,n_2)} \big|\mathfrak{F}(g \star g)(n_1, n_2)\big| \;>\; \frac{4\gamma}{\delta^2} \sup_{\mathbb{R}_+}(-\vartheta'').$$

Appendix C. A nondiscrete, level-discrete set of stationary points. The following example shows the existence of objective functionals that satisfy our assumptions and whose set of stationary points is level-discrete but not discrete.

Consider the function

$$\Theta : x \in \mathbb{R} \longmapsto \theta_1(|x-2|) + \theta_2(|x|),$$

where θ_1 and θ_2 are given by

$$\theta_1(t) \;=\; \begin{cases} t^2(1 - t/3) & \text{if } t \in [0,1], \\ t - 1/3 & \text{if } t > 1 \end{cases}$$

and

$$\theta_2(t) = \begin{cases} \theta_1(t) & \text{if } t \in [0,1], \\ \theta_1(t) + \displaystyle\int_0^{t-1} \frac{u^2 \sin(1/u)}{1+u+u\sin(1/u)}\,\mathrm{d}u & \text{if } t > 1. \end{cases}$$

The potential functions θ_1 and θ_2 satisfy conditions (C1)–(C3); they are also strictly increasing, and so (C4) holds trivially. Leaving details aside, the set of stationary points of Θ is

$$\mathcal{S} = \{1\} \cup \{t_k : k \in \mathbb{N}\}, \qquad t_k := 1 + \frac{2}{(4k+3)\pi}.$$

Furthermore, the derivative of Θ is negative on $(-\infty, 1)$ and positive on $(1, +\infty)\backslash\{t_k : k \in \mathbb{N}\}$. Consequently, $\Theta(1) < \Theta(t_{k+1}) < \Theta(t_k)$ for all $k \in \mathbb{N}$, and thus $\mathcal{S}$ is level-discrete. But $\mathcal{S}$ is not discrete since $\lim_{k\to\infty} t_k = 1 \in \mathcal{S}$.

Appendix D. Supremum cups in the one-dimensional case. Contrary to cups, supremum cups can be unbounded. However, as stated in the following proposition, the one-dimensional case is particular: when defined on $\mathbb{R}$, the objective functions in the class considered do not admit unbounded supremum cups other than $\mathbb{R}$.

Proposition D.1. *Let $\Theta : \mathbb{R} \longrightarrow \mathbb{R}$ be a differentiable function of the form*

$$\Theta(x) = \sum_{k=1}^{K} \theta_k\big(|\alpha_k x - a_k|\big),$$

where, for every k, the potential function $\theta_k : \mathbb{R}_+ \longrightarrow \mathbb{R}$ is increasing and $(\alpha_k, a_k) \in (\mathbb{R}\backslash\{0\}) \times \mathbb{R}$. Then any supremum cup of Θ is either $\mathbb{R}$ or a cup.

Proof. Let Θ be as stated, and let $\boldsymbol{x}^*$ be an isolated minimizer of Θ (assuming one exists). If the supremum cup $\mathcal{O}_{\sup}(\boldsymbol{x}^*)$ is bounded, then it is a cup by Propositions 4.18 and 4.20. So we need only show that if $\mathcal{O}_{\sup}(\boldsymbol{x}^*) \neq \mathbb{R}$, then $\mathcal{O}_{\sup}(\boldsymbol{x}^*)$ is bounded. Seeking a contradiction, suppose $\mathcal{O}_{\sup}(\boldsymbol{x}^*)$ is different from $\mathbb{R}$ and unbounded. Since supremum cups are open and connected by Proposition 4.18, either $\mathcal{O}_{\sup}(\boldsymbol{x}^*) = (-\infty, c)$ with $\boldsymbol{x}^* < c$ or $\mathcal{O}_{\sup}(\boldsymbol{x}^*) = (c, +\infty)$ with $\boldsymbol{x}^* > c$. Without loss of generality, we consider the latter case. By Proposition 4.20(ii), we have $\Theta(x) < \Theta(c)$ for all $x > c$, and thus $\Theta(c) \geqslant \lim_{x\to+\infty} \Theta(x)$. Since the potential functions θ_k are increasing, it follows that $\Theta(c) = \sum_k \sup_{\mathbb{R}_+} \theta_k$. But this is only possible if all the potential functions are eventually constant, which implies that $\Theta(x) = \Theta(c)$ when x is sufficiently large, a contradiction. ∎

Appendix E. Impact of the free parameters on the objective landscape. When the objective functional is nonconvex, the variation of a strength parameter can result in either continuous or discontinuous trajectories of the solutions obtained by Algorithm 1. The following example illustrates the two cases.

Let $\Theta_{\gamma,\delta} : \mathbb{R} \longrightarrow \mathbb{R}$ be defined by

$$\Theta_{\gamma,\delta}(x) := \vartheta_{\mathrm{GM}}\big(|x-1|/\delta\big) + \gamma\big(\vartheta_{\mathrm{GM}}\big(|x|\big) + \vartheta_{\mathrm{GM}}\big(|x+1|\big)\big), \tag{E.1}$$

where ϑ_{GM} is the Geman and McClure function given in (1.9). Let us first set $\delta = 5/4$. Then, as illustrated by Figure 14(a), the function $\Theta_{\gamma,5/4}$ has a unique minimizer x^*_γ which moves

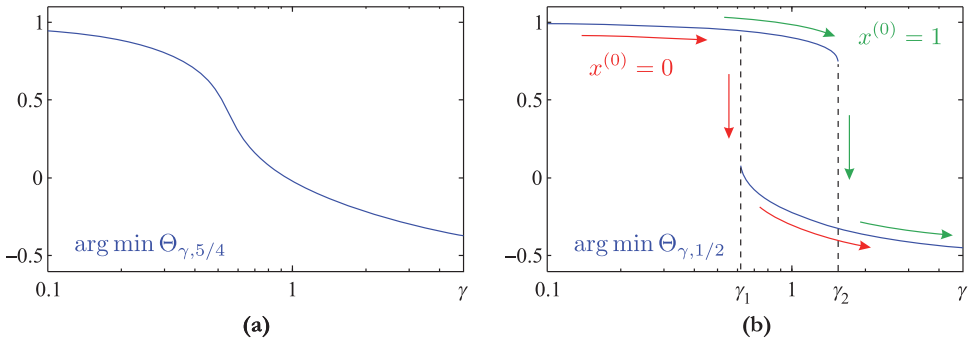

Figure 14. *Minimizers of the function* $\Theta_{\gamma,\delta}$ *defined in* (E.1) *for* (a) $\delta = 5/4$ *and* (b) $\delta = 1/2$. *The red and green arrows in* (b) *indicate the trajectories of the solutions obtained by Algorithm* 1 *when starting from* $x^{(0)} = 0$ *and* $x^{(0)} = 1$, *respectively.*

smoothly from $x_0^* = 1$ to $\lim_{\gamma\to+\infty} x_\gamma^* = -1/2$. Furthermore, since the supremum cup with bottom x_γ^* is $\mathbb{R}$, minimizing $\Theta_{\gamma,\,5/4}$ using Algorithm 1 yields x_γ^* regardless of the initialization. Now let us set $\delta = 1/2$. Then there are two constants $\gamma_1 \approx 0.6189$ and $\gamma_2 \approx 1.5489$ such that $\Theta_{\gamma,\,1/2}$ has a unique minimizer x_γ^* when $\gamma \in [0, \gamma_1) \cup (\gamma_2, +\infty)$ and two distinct minimizers y_γ^* and z_γ^* when $\gamma \in (\gamma_1, \gamma_2)$. In other words, the basin of attraction of x_γ^* splits into two basins with bottoms y_γ^* and z_γ^* when $\gamma > \gamma_1$, and these two basins merge into a single one when $\gamma > \gamma_2$. Consequently, if $\gamma \in [0, \gamma_1) \cup (\gamma_2, +\infty)$, Algorithm 1 converges to x_γ^* from any starting point $x^{(0)}$, whereas if $\gamma \in (\gamma_1, \gamma_2)$, the convergence is toward either y_γ^* or z_γ^* depending on $x^{(0)}$. For instance, the red and green arrows in Figure 14(b) indicate the trajectories of the solutions obtained by starting from $x^{(0)} = 0$ and $x^{(0)} = 1$, respectively. The discontinuity in the first trajectory occurs when $x^{(0)} = 0$ "switches" from a unique basin to a newly emerging basin, and the discontinuity in the second trajectory arises when the basin containing $x^{(0)} = 1$ vanishes.

REFERENCES

[1] M. Nikolova, M. Ng, and C.-P. Tam, *Fast nonconvex nonsmooth minimization methods for image restoration and reconstruction*, IEEE Trans. Image Process., 19 (2010), pp. 3073–3088.

[2] D. Geman and G. Reynolds, *Constrained restoration and the recovery of discontinuities*, IEEE Trans. Pattern Anal. Mach. Intell., 14 (1992), pp. 367–383.

[3] P. Charbonnier, L. Blanc-Féraud, G. Aubert, and M. Barlaud, *Deterministic edge-preserving regularization in computed imaging*, IEEE Trans. Image Process., 6 (1997), pp. 298–311.

[4] A. Delaney and Y. Bresler, *Globally convergent edge-preserving regularized reconstruction: An application to limited-angle tomography*, IEEE Trans. Image Process., 7 (1998), pp. 204–221.

[5] J. Idier, *Convex half-quadratic criteria and interacting auxiliary variables for image restoration*, IEEE Trans. Image Process., 10 (2001), pp. 1001–1009.

[6] M. Nikolova and M. Ng, *Analysis of half-quadratic minimization methods for signal and image recovery*, SIAM J. Sci. Comput., 27 (2005), pp. 937–966.

[7] M. Allain, J. Idier, and Y. Goussard, *On global and local convergence of half-quadratic algorithms*, IEEE Trans. Image Process., 15 (2006), pp. 1130–1142.

[8] M. Robini, Y. Zhu, and J. Luo, *Edge-preserving reconstruction with contour-line smoothing and non-quadratic data-fidelity*, Inverse Probl. Imaging, 7 (2013), pp. 1331–1366.

[9] M. Nikolova and R. Chan, *The equivalence of half-quadratic minimization and the gradient linearization iteration*, IEEE Trans. Image Process., 16 (2007), pp. 1623–1627.
[10] S. Durand and M. Nikolova, *Stability of the minimizers of least squares with a non-convex regularization. Part* I: *Local behavior*, Appl. Math. Optim., 53 (2006), pp. 185–208.
[11] S. Durand and M. Nikolova, *Stability of the minimizers of least squares with a non-convex regularization. Part* II: *Global behavior*, Appl. Math. Optim., 53 (2006), pp. 259–277.
[12] Y.-R. Li, L. Shen, D.-Q. Dai, and B. Suter, *Framelet algorithms for de-blurring images corrupted by impulse plus Gaussian noise*, IEEE Trans. Image Process., 20 (2011), pp. 1822–1837.
[13] Y. Hu and M. Jacob, *Higher degree total variation (HDTV) regularization for image recovery*, IEEE Trans. Image Process., 21 (2012), pp. 2559–2571.
[14] S. Lefkimmiatis, A. Bourquard, and M. Unser, *Hessian-based norm regularization for image restoration with biomedical applications*, IEEE Trans. Image Process., 21 (2012), pp. 983–995.
[15] I. Daubechies, M. Defrise, and C. De Mol, *An iterative thresholding algorithm for linear inverse problems with a sparsity constraint*, Comm. Pure Appl. Math., 57 (2004), pp. 1413–1457.
[16] M. Zibulevsky and M. Elad, *L1−L2 optimization in signal and image processing*, IEEE Signal Proc. Mag., 27 (2010), pp. 76–88.
[17] S. Li, *Markov Random Field Modeling in Computer Vision*, Springer, London, 1995.
[18] M. Nikolova, *Analysis of the recovery of edges in images and signals by minimizing nonconvex regularized least-squares*, Multiscale Model. Simul., 4 (2005), pp. 960–991.
[19] K. Lange, *Convergence of EM image reconstruction algorithms with Gibbs priors*, IEEE Trans. Med. Imag., 9 (1990), pp. 439–446.
[20] S. Li, *On discontinuity-adaptive smoothness priors in computer vision*, IEEE Trans. Pattern Anal. Mach. Intell., 17 (1995), pp. 576–586.
[21] M. Black and A. Rangarajan, *On the unification of line processes, outlier rejection, and robust statistics with applications in early vision*, Int. J. Comput. Vis., 19 (1996), pp. 57–91.
[22] S. Li, *Close-form solution and parameter selection for convex minimization-based edge-preserving smoothing*, IEEE Trans. Pattern Anal. Mach. Intell., 20 (1998), pp. 916–932.
[23] C. Bouman and K. Sauer, *A generalized Gaussian image model for edge-preserving MAP estimation*, IEEE Trans. Image Process., 2 (1993), pp. 296–310.
[24] M. Nikolova, *A variational approach to remove outliers and impulse noise*, J. Math. Imaging Vision, 20 (2004), pp. 99–120.
[25] T. Hohage and F. Werner, *Convergence rates for inverse problems with impulsive noise*, SIAM J. Numer. Anal., 52 (2014), pp. 1203–1221.
[26] R. Acar and C. Vogel, *Analysis of bounded variation penalty method for ill-posed problems*, Inverse Problems, 10 (1994), pp. 1217–1229.
[27] C. R. Vogel and M. E. Oman, *Iterative methods for total variation denoising*, SIAM J. Sci. Comput., 17 (1996), pp. 227–238.
[28] C. Vogel and M. Oman, *Fast, robust total variation-based reconstruction of noisy, blurred images*, IEEE Trans. Image Process., 7 (1998), pp. 813–824.
[29] R. Meyer, *Sufficient conditions for the convergence of monotonic mathematical programming algorithms*, J. Comput. System Sci., 12 (1976), pp. 108–121.
[30] M. Jacobson and J. Fessler, *An expanded theoretical treatment of iteration-dependent majorize-minimize algorithms*, IEEE Trans. Image Process., 16 (2007), pp. 2411–2422.
[31] E. Chouzenoux, A. Jezierska, J.-C. Pesquet, and H. Talbot, *A majorize-minimize subspace approach for $\ell_2-\ell_0$ image regularization*, SIAM J. Imaging Sci., 6 (2013), pp. 563–591.
[32] R. He, W.-S. Zheng, T. Tan, and Z. Sun, *Half-quadratic-based iterative minimization for robust sparse representation*, IEEE Trans. Pattern Anal. Mach. Intell., 36 (2014), pp. 261–275.
[33] D. Hunter and K. Lange, *A tutorial on MM algorithms*, Amer. Statist., 58 (2004), pp. 30–37.
[34] K. Kuratowski, *Topology*, Vol. II, Academic Press, New York, London, 1968.
[35] A. Ostrowski, *Solution of Equations in Euclidean and Banach Spaces*, Academic Press, New York, London, 1973.
[36] P. Ciarlet, *Introduction to Numerical Linear Algebra and Optimisation*, Cambridge University Press, Cambridge, UK, 1989.

[37] M. Nikolova, *Markovian reconstruction using a GNC approach*, IEEE Trans. Image Process., 8 (1999), pp. 1204–1220.
[38] M. Pogu and J. Souza De Cursi, *Global optimization by random perturbation of the gradient method with a fixed parameter*, J. Global Optim., 5 (1994), pp. 159–180.
[39] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, *Image quality assessment: From error visibility to structural similarity*, IEEE Trans. Image Process., 13 (2004), pp. 600–612.
[40] T. Hebert and K. Lu, *Expectation-maximization algorithms, null spaces, and MAP image restoration*, IEEE Trans. Image Process., 4 (1995), pp. 1084–1095.
[41] J.-F. Cai, R. Chan, and M. Nikolova, *Fast two-phase image deblurring under impulse noise*, J. Math. Imaging Vision, 36 (2010), pp. 46–53.
[42] A. Cohen, I. Daubechies, and J.-C. Feauveau, *Biorthogonal bases of compactly supported wavelets*, Comm. Pure Appl. Math., 45 (1992), pp. 485–560.
[43] R. Horn and C. Johnson, *Matrix Analysis*, 2nd ed., Cambridge University Press, Cambridge, UK, 2013.
[44] B. Hunt, *The application of constrained least squares estimation to image restoration by digital computer*, IEEE Trans. Comput., C-22 (1973), pp. 805–812.