

A New TwIST: Two-Step Iterative Shrinkage/Thresholding Algorithms for Image Restoration

José M. Bioucas-Dias, *Member, IEEE*, and Mário A. T. Figueiredo, *Senior Member, IEEE*

Abstract—Iterative shrinkage/thresholding (IST) algorithms have been recently proposed to handle a class of convex unconstrained optimization problems arising in image restoration and other linear inverse problems. This class of problems results from combining a linear observation model with a nonquadratic regularizer (e.g., total variation or wavelet-based regularization). It happens that the convergence rate of these IST algorithms depends heavily on the linear observation operator, becoming very slow when this operator is ill-conditioned or ill-posed. In this paper, we introduce two-step IST (TwIST) algorithms, exhibiting much faster convergence rate than IST for ill-conditioned problems. For a vast class of nonquadratic convex regularizers (\mathcal{L}^p norms, some Besov norms, and total variation), we show that TwIST converges to a minimizer of the objective function, for a given range of values of its parameters. For noninvertible observation operators, we introduce a monotonic version of TwIST (MTwIST); although the convergence proof does not apply to this scenario, we give experimental evidence that MTwIST exhibits similar speed gains over IST. The effectiveness of the new methods are experimentally confirmed on problems of image deconvolution and of restoration with missing samples.

Index Terms—Convex analysis, image deconvolution, image restoration, non-smooth optimization, optimization, regularization, total variation, wavelets.

I. INTRODUCTION

A. Problem Formulation

INVERSE problems abound in many application areas of signal/image processing: remote sensing, radar imaging, tomographic imaging, microscopic imaging, astronomic imaging, digital photography, etc. [1], [5], [35]. Image restoration is one of the earliest and most classical linear inverse problems in imaging, dating back to the 1960s [1].

In an inverse problem, the goal is to estimate an unknown original signal/image \mathbf{x} from a (possibly noisy) observation \mathbf{y} , produced by an operator \mathbf{K} applied to \mathbf{x} . When \mathbf{K} is linear, we have a *linear inverse problem* (LIP). Although we only report image restoration experiments, all the results herein presented are of general applicability in LIPs.

Manuscript received February 26, 2007; revised August 6, 2007. A preliminary, much shorter version of this work is described in [8]. This work was supported in part by the Fundação para a Ciência e Tecnologia (FCT), Portuguese Ministry of Science and Higher Education, under project POSC/EEA-CPS/61271/2004. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Peyman Milanfar.

The authors are with the Instituto de Telecomunicações and the Instituto Superior Técnico, Technical University of Lisbon, 1049-001 Lisboa, Portugal (e-mail: mario.figueiredo@lx.it.pt; jose.bioucas@lx.it.pt).

Digital Object Identifier 10.1109/TIP.2007.909319

Many approaches to LIPs define a solution $\hat{\mathbf{x}}$ (e.g., a restored image/signal) as a minimizer of a convex objective function $f : \mathcal{X} \rightarrow \bar{\mathbb{R}} = [-\infty, +\infty]$, given by

$$f(\mathbf{x}) = \frac{1}{2} \|\mathbf{y} - \mathbf{K}\mathbf{x}\|^2 + \lambda\Phi(\mathbf{x}) \quad (1)$$

where $\mathbf{K} : \mathcal{X} \rightarrow \mathcal{Y}$ is the (linear) *direct operator*, \mathcal{X} and \mathcal{Y} are real Hilbert spaces (both with norm denoted as $\|\cdot\|$), $\Phi : \mathcal{X} \rightarrow \bar{\mathbb{R}}$ is a function (whose meaning and role will be discussed in the next paragraphs), $\lambda \in [0, +\infty[$ is a parameter.

In a regularization framework, minimizing f is seen as a way of overcoming the ill-conditioned, or singular, nature of \mathbf{K} , which precludes inverting it. In this context, Φ is called the *regularizer* and λ the *regularization parameter* [5].

In a (finite-dimensional) Bayesian setting, the reasoning behind (1) is as follows. Assume that $\mathbf{y} = \mathbf{K}\mathbf{x} + \mathbf{w}$, where \mathbf{w} is a sample of a white zero-mean Gaussian random vector/field, of variance σ^2 ; let $p(\mathbf{x})$ be the adopted prior; thus, the logarithm of the *a posteriori* density is $\log p(\mathbf{x}|\mathbf{y}) = -f(\mathbf{x})$ (up to a constant), with $\lambda = \sigma^2$ and $\Phi(\mathbf{x}) = -\log p(\mathbf{x})$; maximum *a posteriori* (MAP) estimates are, thus, minimizers of f (see [2] and references therein). Despite this possible interpretation of (1), we will refer to Φ simply as the regularizer.

Regardless of the adopted formal framework, the intuitive meaning of f is simple: minimizing it corresponds to looking for a compromise between the lack of fitness of a candidate estimate \mathbf{x} to the observed data, measured by $\|\mathbf{y} - \mathbf{K}\mathbf{x}\|^2$, and its degree of undesirability, given by $\Phi(\mathbf{x})$. The regularization parameter λ controls the relative weight of the two terms.

A more detailed discussion of Φ will be postponed to Section II; suffice it to say here that the current state-of-the-art regularizers for image restoration are nondifferentiable. Examples of such choices are total-variation (TV) regularization [11], [15], [42] and wavelet-based regularization [13], [22], [23], [39]. The nondifferentiable nature of f , together with the huge dimension of its argument (for a typical 512×512 image, $\mathcal{X} = \mathbb{R}^{262144}$), place its minimization beyond the reach of standard off-the-shelf optimization methods.

Of course not all approaches to LIPs lead to convex optimization problems such as (1). For example, some wavelet-based deconvolution approaches do not lead to an optimization problem [31], [40]. However, this paper is strictly concerned with algorithms for minimizing (1), and will not discuss its relative merits with respect to other criteria, nor the relative merits of different choices of Φ .

B. Previous Algorithms

In recent years, *iterative shrinkage/thresholding* (IST) algorithms (described in Section IV), tailored for objective

functions with the form (1), were independently proposed by several authors in different frameworks. IST algorithms for wavelet-based deconvolution were first derived in [41] (see also [28]) under the expectation-maximization (EM) framework and, later [29], using a *majorization-minimization* (MM, see [33]) approach. In [21], IST algorithms were placed on solid mathematical grounds, with a rigorous convergence proof in an infinite-dimensional setting. A proof for the finite-dimensional case was independently presented in [4]. IST algorithms have been independently proposed in [24], [25], [45], and [46]. Recently, [19] brought important contributions to the understanding of a class of objective functions which contains f , as well as of a class of algorithms (termed *forward-backward splitting*) which includes IST.

A different (not IST) algorithm, proposed in [6], [7], under a *generalized* EM framework [49], was recently shown to also be an MM scheme [27]. That algorithm, which we will refer to as IRS (iterative reweighted shrinkage), was shown to be much faster than IST when \mathbf{K} is strongly ill-conditioned. Conversely, for mildly ill-conditioned \mathbf{K} and medium to strong noise, IST is faster than IRS [27].

C. Contributions

This paper introduces a new class of iterative schemes, bringing together the best of IRS and IST. Algorithms in this class have a two-step IST (TwIST) structure, i.e., each iterate depends on the two previous iterates, rather than only on the previous one. For ill-conditioned (but invertible) linear observation operators, we prove (linear) convergence of TwIST to minima of the objective function f , for a certain range of the algorithm parameters, and derive bounds for the convergence factor. As a byproduct of this analysis, we provide a bound for the convergence factor of IST in the case of invertible operators which, to best of our knowledge, was not available in the literature.

Experimental results (in wavelet-based and TV-based deconvolution) confirm that TwIST algorithms can be tuned to converge much faster than the original IST versions, specially in severely ill-conditioned problems. Similarly to the IRS algorithm [7], [27], the speed gains can reach up two orders of magnitude in a typical benchmark problem (e.g., 9×9 uniform blur). Conversely, in well-conditioned LIPs, TwIST is still faster than IST (although not as much as in severely ill-conditioned LIPs), thus faster than IRS [27].

The convergence proof mentioned in the previous paragraph applies only to invertible linear operators. For the noninvertible case, we introduce a monotonic variant of TwIST, termed MTwIST. Although we do not have a proof of convergence, we give experimental evidence that, with a noninvertible operator, MTwIST also exhibits a large speed advantage over IST.

D. Summary of the Paper

In Section II, we review several choices of Φ in the context of denoising problems, the solution of which plays a central role in IST and TwIST. Section III studies the existence and uniqueness of minimizers of f . The IST and IRS algorithms are reviewed in Section IV, together with previous results on the convergence of IST. The TwIST algorithm is introduced in Section V, which

also contains the central theorem of the paper. Finally, experimental results are reported in Section VI. Appendices contain brief reviews of basic results from convex analysis and other mathematical tools, as well as the proofs of the new results presented.

II. REGULARIZERS AND DENOISING

A. Denoising With Convex Regularizers

Denoising problems are LIPs in which \mathbf{K} is the identity, $\mathbf{K}\mathbf{x} = \mathbf{x}$. In this case, the objective function (1) simplifies to

$$f_{\text{den}} = (1/2)d_{\mathbf{y}}^2 + \lambda\Phi$$

where $d_{\mathbf{y}} : \mathcal{X} \rightarrow \mathbb{R}$

$$d_{\mathbf{y}}(\mathbf{x}) = \|\mathbf{x} - \mathbf{y}\|. \quad (2)$$

We adopt the following standard assumptions about the regularizer $\Phi : \mathcal{X} \rightarrow \overline{\mathbb{R}}$: it is convex, lower semi-continuous (lsc), and proper (see Appendix I for definitions and implications of these properties).

The fact that Φ is lsc and proper and $d_{\mathbf{y}}^2$ is a continuous, real-valued, coercive function ($\lim_{\|\mathbf{x}\| \rightarrow \infty} d_{\mathbf{y}}^2(\mathbf{x}) = \infty$), guarantees that f_{den} is lsc, proper, and coercive. Consequently, the set of minimizers of f_{den} is not empty (Theorem 5, Appendix I). Finally, the strict convexity of $d_{\mathbf{y}}^2$ implies strict convexity of f_{den} (Theorem 7, Appendix I); thus, its minimizer is unique; this allows defining the *denoising function* (also known as the *Moreau proximal mapping* [19], [37], [44]) $\Psi_{\lambda} : \mathcal{X} \rightarrow \mathcal{X}$ as

$$\Psi_{\lambda}(\mathbf{y}) = \arg \min_{\mathbf{x}} \left\{ \frac{d_{\mathbf{y}}^2(\mathbf{x})}{2} + \lambda\Phi(\mathbf{x}) \right\}. \quad (3)$$

In the following sections, we describe in detail the classes or regularizers considered in this work, as well as the corresponding denoising functions.

B. Denoising With 1-Homogeneous Regularizers

A function Φ that satisfies $\Phi(\zeta\mathbf{x}) = \zeta\Phi(\mathbf{x})$, for all $\zeta \geq 0$ and $\mathbf{x} \in \mathcal{X}$, is called *positively homogeneous of degree 1* (phd-1). Let $\Upsilon(\mathcal{X})$ denote the set of functions $\Phi : \mathcal{X} \rightarrow \overline{\mathbb{R}}$ that are convex, lsc, proper, and phd-1.

An important recent result states that denoising with regularizers from $\Upsilon(\mathcal{X})$ corresponds to the residual of the projection onto a convex set, as formalized in the following theorem (see [11], [19], and [36] for proofs).

Theorem 1: If $\Phi \in \Upsilon(\mathcal{X})$, then the denoising function Ψ_{λ} defined in (3) is given by

$$\Psi_{\lambda}(\mathbf{y}) = \mathbf{y} - \mathbf{P}_{\lambda C}(\mathbf{y}) \quad (4)$$

where $C \subset \mathcal{X}$ is a closed convex set depending on the regularizer Φ , and $P_A : \mathcal{X} \rightarrow \mathcal{X}$ denotes the orthogonal projection operator onto the convex set $A \subset \mathcal{X}$.

C. Total Variation

In the original formulation of total-variation (TV) [11], [15], [42], \mathcal{X} is an infinite-dimensional Hilbert space $L^2(I)$, where I is a bounded open domain of \mathbb{R}^2 , e.g., $\Omega =]0, 1[^2$. With digital images, \mathcal{X} is simply a finite-dimensional space of pixel values

on a 2-D lattice, say $\mathcal{X} = \mathbb{R}^m$, equipped with the Euclidean norm; thus, discrete TV regularizers have to be used [11], [12]. Standard choices are the “isotropic” and “nonisotropic” discrete TV regularizers, given, respectively, by

$$\Phi_{\text{ITV}}(\mathbf{x}) = \sum_i \sqrt{(\Delta_i^h \mathbf{x})^2 + (\Delta_i^v \mathbf{x})^2} \quad (5)$$

$$\Phi_{\text{niTV}}(\mathbf{x}) = \sum_i (|\Delta_i^h \mathbf{x}| + |\Delta_i^v \mathbf{x}|) \quad (6)$$

where Δ_i^h and Δ_i^v denote horizontal and vertical (on the 2-D lattice) first-order local difference operators (omitting boundary corrections). It is clear from (5) and (6) that $\Phi_{\text{ITV}}, \Phi_{\text{niTV}} \in \Upsilon(\mathbb{R}^m)$. Although there is no closed form for the projection onto C , i.e., to solve the TV denoising problem, fast iterative methods have been recently introduced [11], [12], [20], [30].

D. Weighted ℓ^p Norm

Weighted ℓ^p norms, for $p \geq 1$, are defined as

$$\Phi_{\mathbf{w}}^{\ell^p}(\mathbf{x}) = \|\mathbf{x}\|_{p,\mathbf{w}} = \left(\sum_i w_i |x_i|^p \right)^{1/p} \quad (7)$$

where $\mathbf{w} = [w_1, w_2, \dots, w_i, \dots]$, with $w_i \geq 0$ and $p \geq 1$. The underlying Hilbert space is simply $\mathcal{X} = \mathbb{R}^m$, in the finite-dimensional case [with the sum in (7) extending from $i = 1$ to $i = m$], or $\mathcal{X} = \ell^2(\mathbb{N})$, in the infinite-dimensional case [where the sum in (7) is for $i \in \mathbb{N}$]. Being a norm, $\Phi_{\mathbf{w}}^{\ell^p}$ clearly belongs to Υ .

The denoising function Ψ_{λ} under a $\Phi_{\mathbf{w}}^{\ell^p}$ regularizer cannot be obtained in closed form, except in some particular cases, the most notable of which is $p = 1$; in this case, Ψ_{λ} is the well-known soft-thresholding function [23], that is $\Psi_{\lambda}(\mathbf{z}) = \hat{\mathbf{x}} = [\hat{x}_1, \dots, \hat{x}_i, \dots]$, with

$$\hat{x}_i = \text{soft}(z_i, \lambda w_i) = \text{sign}(z_i) \max\{0, |z_i| - \lambda w_i\}. \quad (8)$$

Next, we discuss some approaches involving $\Phi_{\mathbf{w}}^{\ell^p}$ regularizers.

1) *Orthogonal Representations*: A classical approach consists in modeling images as elements of some Besov space $B_p^a(L^p(I))$, where $I = [0, 1]^2$ is the unit square. The adopted regularizer is then the corresponding Besov norm $\|\cdot\|_{B_p^a(L^p(I))}$, which has an equivalent weighted ℓ^p sequence norm of wavelet coefficients on an orthogonal wavelet basis (see [13] for details). To exploit this equivalence, the problem is formulated w.r.t. the coefficients, rather than the image itself. Letting $\mathbf{W} : \mathcal{X} \rightarrow B_p^a(L^p(I))$ denote the linear operator that produces an image from its wavelet coefficients, the objective function becomes

$$f(\mathbf{x}) = \frac{1}{2} d_{\mathbf{y}}^2(\mathbf{H}\mathbf{W}\mathbf{x}) + \lambda \Phi_{\mathbf{w}}^{\ell^p}(\mathbf{x}) \quad (9)$$

where the weights depend on the scale of each coefficient and on the parameters (p and a) of the Besov norm (see [13]), and \mathbf{H} is the observation operator. In practice, for digital images, \mathcal{X} is finite dimensional, say $\mathcal{X} = \mathbb{R}^m$, $\Phi_{\mathbf{w}}^{\ell^p}$ is a weighted ℓ^p norm on \mathbb{R}^m , and $\mathbf{W} \in \mathbb{R}^{m \times m}$ is an unitary $m \times m$ matrix. Notice that (9) has the same form as (1), with $\mathbf{K} = \mathbf{H}\mathbf{W}$ and $\Phi = \Phi_{\mathbf{w}}^{\ell^p}$.

2) *Frames and Redundant Representations*: Another formulation (in a finite-dimensional setting) leading to an objective function with the same form as (9) is the following. Let the columns of \mathbf{W} contain a redundant dictionary (e.g., a frame) with respect to which we seek a representation of the unknown image. If the image is directly observed, \mathbf{H} is the identity; in this case, minimizing (9) corresponds to finding a regularized representation of the observed image on the dictionary \mathbf{W} [24], [25]. For $p = 1$, this is the well-known *basis-pursuit denoising* criterion [17]. If the original image is not directly observed (\mathbf{H} is not identity), minimizing (9) corresponds to reconstructing/restoring the original image by looking for a (regularized) representation on an over-complete dictionary. This formulation has been used for shift-invariant wavelet-based deconvolution [7], [28], [29].

E. p th Power of a Weighted ℓ^p Norm

This class of regularizers, defined as

$$\Phi_{\mathbf{w}}^p(\mathbf{x}) = \|\mathbf{x}\|_{p,\mathbf{w}}^p = \sum_i w_i |x_i|^p \quad (10)$$

appears in many wavelet-based approaches [7], [21], [28]–[30], [43]. This regularizer can also be motivated as being equivalent to the p th power of a Besov norm, $\|\cdot\|_{B_p^a(L^p(I))}^p$ [21], [36].

For $p = 1$, $\Phi_{\mathbf{w}}^1 = \Phi_{\mathbf{w}}^{\ell^1}$; thus, the denoising operator (3) is given by (8). For $p > 1$, $\Phi_{\mathbf{w}}^p$ is not phd-1, and the denoising operator does not have the form (4). In this case, however, we can write $\Psi_{\lambda}(\mathbf{z}) = \hat{\mathbf{x}} = [\hat{x}_1, \dots, \hat{x}_i, \dots]$, with

$$\hat{x}_i = S_{\lambda w_i, p}(z_i) \quad (11)$$

where $S_{\tau, p} = F_{\tau, p}^{-1}$ is the inverse function of

$$F_{\tau, p}(x) = x + \tau p \text{sign}(x) |x|^{p-1}. \quad (12)$$

Notice that, for $p > 1$, $F_{\tau, p} : \mathbb{R} \rightarrow \mathbb{R}$ is one-to-one; thus, $S_{\tau, p} = F_{\tau, p}^{-1}$ is well defined. The function $S_{\tau, p}$, called the *shrinkage* function, has simple closed forms when $p = 4/3$, $p = 3/2$, or $p = 2$ [16]. For example, the function $S_{\tau, 2}$ is a simple linear shrinkage, $S_{\tau, 2}(z) = z/(1+2\tau)$. Important features of $S_{\tau, p}$ (for $p > 1$) are: it is strictly monotonic, continuously¹ differentiable, and its derivative is upper bounded by 1 (since the derivative of its inverse $F_{\tau, p}$ is uniformly lower bounded by 1) [21].

III. EXISTENCE AND UNIQUENESS OF SOLUTIONS

The existence and uniqueness of minimizers of (1) are addressed in the following proposition, the proof of which can be found in [19, Propositions 3.1 and 5.3].

Proposition 1: Let $f : \mathcal{X} \rightarrow \overline{\mathbb{R}}$ be defined as in (1), where operator $\mathbf{K} : \mathcal{X} \rightarrow \mathcal{Y}$ is linear and bounded, and $\Phi : \mathcal{X} \rightarrow \overline{\mathbb{R}}$ is a proper, lsc, convex function. Let G denote the set of minimizers of f . Then:

- i) if Φ is coercive, then G is nonempty;

¹Continuous differentiability is not claimed in [21], only its differentiability. However, the continuity (for $p > 1$) of the derivative of $S_{\tau, p}$, denoted $S'_{\tau, p}$, is easily shown. First, it is trivial to check that $\lim_{x \rightarrow 0} S'_{\tau, p}(x) = 0$, where $S'_{\tau, p}(x) = 1/F'_{\tau, p}(S_{\tau, p}(x))$. Second, it is also easy to show, via the definition of derivative, that $S'_{\tau, p}(0) = 0$.

- ii) if Φ is strictly convex or \mathbf{K} is injective, then G contains at most one element;
- iii) if \mathbf{K} is bounded below, that is, if there exists $\kappa \in]0, +\infty]$, such that, for any $\mathbf{x} \in \mathcal{X}$, $\|\mathbf{K}\mathbf{x}\| \geq \kappa\|\mathbf{x}\|$, then G contains exactly one element.

We will now comment on the application of Proposition 1 to the several regularization functions above considered. If all the weights are strictly positive ($w_i > 0, \forall_i$), both the weighted ℓ^p norm and its p th power (for $p \geq 1$) are coercive (see [19, Proposition 5.15 and Problem 5.18]); thus, Proposition 1 i) ensures existence of minimizers of f . Under these regularizers, if \mathbf{K} is injective, the minimizer is unique; otherwise, the minimizer is unique with $\Phi_{\ell^p}^p$, with $p > 1$ (which is strictly convex).

In the finite-dimensional case ($\mathcal{X} = \mathbb{R}^m$), injectivity of \mathbf{K} is sufficient to guarantee existence and uniqueness of the solution (under any convex regularizer, strictly or not, coercive or not). This results from Proposition 1 iii), because any finite-dimensional injective operator is bounded below.

When Φ is a TV regularizer (e.g., Φ_{ITV} or Φ_{niTV}) and \mathbf{K} is not bounded below, Proposition 1 can not be used to guarantee existence of minimizers of (1). The reason is that TV regularizers are not coercive since they equal zero when the argument is a constant image. However, under the additional condition that constant images do not belong to the null space of \mathbf{K} , it can still be shown that G is not empty [14].

IV. PREVIOUS ALGORITHMS

This section reviews algorithms previously proposed for finding minimizers of f . From this point on, we focus on the finite-dimensional case, $\mathcal{X} = \mathbb{R}^m$, $\mathcal{Y} = \mathbb{R}^n$, and denote the standard Euclidean vector norm as $\|\cdot\|_2$.

A. Iterative Shrinkage/Thresholding (IST)

IST algorithms has the form

$$\mathbf{x}_{t+1} = (1 - \beta)\mathbf{x}_t + \beta\Psi_\lambda(\mathbf{x}_t + \mathbf{K}^T(\mathbf{y} - \mathbf{K}\mathbf{x}_t)) \quad (13)$$

where $\beta > 0$. The original IST algorithm has the form (13), with $\beta = 1$ [21], [28], [29]. Schemes with $\beta \neq 1$ can be seen as under ($\beta < 1$) or over ($\beta > 1$) relaxed versions of the original IST algorithm.

Each iteration of the IST algorithm only involves sums, matrix-vector products by \mathbf{K} and \mathbf{K}^T , and the application of the denoising operation Ψ_λ . In wavelet-based methods, Ψ_λ is a coefficient-wise nonlinearity, thus very computationally efficient. When \mathbf{K} represents the convolution with some kernel \mathbf{k} , the corresponding product can be computed efficiently using the fast Fourier transform (FFT).

Convergence of IST, with $\beta = 1$, was first shown in [21]. Later, convergence of a more general version of the algorithm (including $\beta \neq 1$), was shown in [19]. The following theorem is a simplified version of Theorems 3.4 and 5.5 from [19]; the simplifications result from considering finite-dimensional spaces (no difference between strong and weak convergence) and from (13) being a particular case of the somewhat more general version studied in [19].

Theorem 2: Let f be given by (1), where $\Phi : \mathcal{X} \rightarrow \mathbb{R}$ is convex² and $\|\mathbf{K}\|_2^2 < 2$. Let G , the set of minimizers of f , be nonempty. Fix some \mathbf{x}_1 and let the sequence $\{\mathbf{x}_t, t \in \mathbb{N}\}$ be produced by (13), with $\beta \in]0, 1]$. Then, $\{\mathbf{x}_t, t \in \mathbb{N}\}$ converges to a point $\mathbf{x} \in G$.

B. Iterative Re-Weighted Shrinkage (IRS)

The IRS algorithm was specifically designed for wavelet-based problems of the form (9), where \mathbf{W} contains an orthogonal or redundant wavelet basis and the regularizer is not necessarily a weighted ℓ^p norm [7]. The iterations of the IRS algorithm are given by

$$\mathbf{x}_{t+1} = \text{solution}\{\mathbf{A}_t\mathbf{x} = \mathbf{b}\} \quad (14)$$

with $\mathbf{b} = \mathbf{K}^T\mathbf{y}$ and $\mathbf{A}_t = \lambda\mathbf{D}_t + \mathbf{K}^T\mathbf{K}$, where \mathbf{D}_t is a diagonal matrix (of non-negative elements) that depends on \mathbf{x}_t and Φ . Observe that matrix \mathbf{D}_t shrinks the components of \mathbf{x}_{t+1} , thus the term *iterative reweighted shrinkage*. Each iteration of IRS resembles a weighted ridge regression problem, with design matrix \mathbf{K} ; algorithms with a similar structure have been used for sparse regression [26], [32].

The huge size of \mathbf{A}_t forces the use of iterative methods to implement (14). In [7], this is done with a two-step (or second-order) stationary iterative method [3], which we will next briefly review.

C. Two-Step Methods for Linear Systems

Consider the linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$, with \mathbf{A} positive definite; define a so-called *splitting* of \mathbf{A} as $\mathbf{A} = \mathbf{C} - \mathbf{R}$, such that \mathbf{C} is positive definite and easy to invert (e.g., a diagonal matrix). A stationary *two-step iterative method* (TwSIM) for solving $\mathbf{A}\mathbf{x} = \mathbf{b}$ is defined as

$$\begin{aligned} \mathbf{x}_1 &= \mathbf{x}_0 + \beta_0\mathbf{C}^{-1}(\mathbf{b} - \mathbf{A}\mathbf{x}_0) \\ \mathbf{x}_{t+1} &= (1 - \alpha)\mathbf{x}_{t-1} + \alpha\mathbf{x}_t + \beta\mathbf{C}^{-1}(\mathbf{b} - \mathbf{A}\mathbf{x}_t) \end{aligned} \quad (15)$$

for $t \geq 1$, where \mathbf{x}_0 is the initial vector, and α, β, β_0 are the parameters of the algorithm (more on this below). The designation “two-step” stems from the fact that \mathbf{x}_{t+1} depends on both \mathbf{x}_t and \mathbf{x}_{t-1} , rather than only on \mathbf{x}_t .

The main result concerning TwSIM is given in following theorem [3, Theorem 5.9].

Theorem 3: Let $\{\mathbf{x}_t, t \in \mathbb{N}\}$ be the sequence produced by (15), with arbitrary \mathbf{x}_0 . Let λ_1 and λ_m denote the smallest and largest eigenvalues of matrix $\mathbf{C}^{-1}\mathbf{A}$, and $\kappa = \lambda_1/\lambda_m$ be its inverse condition number. Then, $\{\mathbf{x}_t, t \in \mathbb{N}\}$ converges to the solution of $\mathbf{A}\mathbf{x} = \mathbf{b}$ if and only if $0 < \alpha < 2$ and $0 < \beta < 2\alpha/\lambda_m$. The optimal asymptotic convergence factor³ is $\rho \equiv (1 - \sqrt{\kappa})/(1 + \sqrt{\kappa})$, obtained for $\alpha = \rho^2 + 1$ and $\beta = 2\alpha/(\lambda_1 + \lambda_m)$. With $\alpha = 1$, the two-step method (15) becomes a one-step method for which the best asymptotic converge factor is $\bar{\rho} \equiv (1 - \kappa)/(1 + \kappa)$.

²In a finite-dimensional space, every real convex function is continuous, so we can drop the lsc condition.

³See Appendix II for a brief review of convergence factors.

D. Comparing IST With IRS

It was shown in [7] that, for ill conditioned systems, IRS is much faster than IST. This fact can be traced to the use of the TwSIM in each step of IRS. On the other hand, when noise is the main factor, and the observation operator is not too ill-conditioned, IST outperforms IRS because it uses a closed-form (usually nonlinear) denoising step in each iteration [27]. In fact, in a pure denoising problem ($\mathbf{K} = \mathbf{I}$ or \mathbf{K} orthogonal), IST (with $\beta = 1$ and initialized with a zero image) converges in one step, while IRS does not.

V. TWO-STEP IST (TWIST)

A. Motivation and Definition

The TwIST method proposed in this paper aims at keeping the good denoising performance of the IST scheme, while still being able to handle ill-posed problems as efficiently as the IRS algorithm.

Taking $\mathbf{C} = \mathbf{I} + \lambda \mathbf{D}_t$ and $\mathbf{R} = \mathbf{I} - \mathbf{K}^T \mathbf{K}$ in the splitting $\mathbf{A} = \mathbf{C} - \mathbf{R}$ of matrix $\mathbf{A} = \lambda \mathbf{D}_t + \mathbf{K}^T \mathbf{K}$, the two-step iteration (15) for the linear system $\mathbf{A} \mathbf{x} = \mathbf{K}^T \mathbf{y}$ becomes

$$\mathbf{x}_{t+1} = (1 - \alpha) \mathbf{x}_{t-1} + (\alpha - \beta) \mathbf{x}_t + \beta \mathbf{C}^{-1} (\mathbf{x}_t + \mathbf{K}^T (\mathbf{y} - \mathbf{K} \mathbf{x}_t)). \quad (16)$$

Observe the relationship between (13) and (16): the former can be obtained from the latter by setting $\alpha = 1$ and replacing the multiplication by matrix \mathbf{C}^{-1} by the denoising operator Ψ_λ . This similarity suggests a two-step version of IST (TwIST) as

$$\mathbf{x}_1 = \Gamma_\lambda(\mathbf{x}_0) \quad (17)$$

$$\mathbf{x}_{t+1} = (1 - \alpha) \mathbf{x}_{t-1} + (\alpha - \beta) \mathbf{x}_t + \beta \Gamma_\lambda(\mathbf{x}_t) \quad (18)$$

for $t \geq 1$, where $\Gamma_\lambda : \mathbb{R}^m \rightarrow \mathbb{R}^m$ is defined as

$$\Gamma_\lambda(\mathbf{x}) = \Psi_\lambda(\mathbf{x} + \mathbf{K}^T (\mathbf{y} - \mathbf{K} \mathbf{x})). \quad (19)$$

A key observation is that TwIST, IST, and the original IST with $\beta = 1$ all have the same fixed points. In fact, elementary manipulation allows showing that the three following equations are equivalent:

$$\mathbf{x} = (1 - \alpha) \mathbf{x} + (\alpha - \beta) \mathbf{x} + \beta \Gamma_\lambda(\mathbf{x})$$

$$\mathbf{x} = (1 - \beta) \mathbf{x} + \beta \Gamma_\lambda(\mathbf{x})$$

$$\mathbf{x} = \bar{\Gamma}_\lambda(\mathbf{x}).$$

B. Convergence of TwIST

Fundamental questions concerning TwIST are: for what values of α and β does it converge? How does the convergence rate depend of α and β ? The main theorem of this paper partially characterizes the convergence of the TwIST algorithm, when f has a unique minimizer.

Theorem 4: Let f be given by (1), where Φ is a convex regularizer. Let ξ_1 and ξ_m be two real numbers such that $0 < \xi_1 \leq \lambda_i(\mathbf{K}^T \mathbf{K}) \leq \xi_m$, where $\lambda_i(\cdot)$ is the i th eigenvalue of its argument, let $\kappa = \xi_1 / \xi_m$, where $\xi_m \equiv \max(1, \xi_m)$, and

$$\hat{\rho} \equiv \frac{1 - \sqrt{\kappa}}{1 + \sqrt{\kappa}} < 1. \quad (20)$$

Let $\hat{\mathbf{x}}$ be the unique (because \mathbf{K} is injective) minimizer of f and define the ‘‘error vector’’ as $\mathbf{e}_t = \mathbf{x}_t - \hat{\mathbf{x}}$ and the ‘‘stacked error vector’’ as

$$\mathbf{w}_t = \begin{bmatrix} \mathbf{e}_{t+1} \\ \mathbf{e}_t \end{bmatrix}. \quad (21)$$

- i) There exists a matrix set \mathcal{Q} such that \mathbf{w}_{t+1} can be written as $\mathbf{w}_{t+1} = \mathbf{Q}_t \mathbf{w}_t$, where $\mathbf{Q}_t \in \mathcal{Q}$, for $t \in \mathbb{N}$ (\mathbf{Q}_t may depend on \mathbf{w}_t); moreover, if $0 < \alpha < 2$ and $0 < \beta < 2\alpha / \xi_m$, then $\rho(\mathbf{Q}_t) < 1$, for any $\mathbf{Q}_t \in \mathcal{Q}$, where $\rho(\mathbf{Q}_t)$ is the spectral radius of \mathbf{Q}_t (see Appendix II).
- ii) Setting

$$\alpha = \hat{\alpha} \equiv \hat{\rho}^2 + 1 \quad (22)$$

$$\beta = \hat{\beta} \equiv 2\hat{\alpha} / (\xi_m + \xi_1) \quad (23)$$

guarantees that $\rho(\mathbf{Q}_t) = \hat{\rho}$.

- iii) Setting $\alpha = 1$ (i.e., the IST algorithm) and

$$\beta = \bar{\beta} \equiv 2 / (\xi_m + \xi_1) \quad (24)$$

guarantees that

$$\rho(\mathbf{Q}_t) \leq \bar{\rho} \equiv \frac{1 - \kappa}{1 + \kappa} < 1. \quad (25)$$

- iv) If $\xi_m < 1$, $0 < \alpha \leq 1$, and $0 < \beta < 2\alpha$, then $\lim_{t \rightarrow \infty} \mathbf{w}_t = \mathbf{0}$.

Theorem 4 extends the results about the convergence of the linear TwSIM (see Section IV-C and [3]) to the nonlinear/non-differentiable case. While the proof in [3] uses linear algebra tools, the possible nonlinear/nondifferentiable nature of Ψ_λ demands nonsmooth analysis techniques [18], [44]. The proof of Theorem 4 can be found in Appendix III.

If matrix \mathbf{Q}_t is not time dependent, i.e., $\mathbf{Q}_t = \mathbf{Q}$, the condition $\rho(\mathbf{Q}) < 1$ would be sufficient for convergence to zero of \mathbf{w}_t . However, in TwIST, \mathbf{Q}_t is in general not constant, thus $\rho(\mathbf{Q}_t) < 1, \forall t$, is not a sufficient condition for convergence to zero of \mathbf{w}_t . Convergence of a nonstationary linear iteration $\mathbf{w}_{t+1} = \mathbf{Q}_t \mathbf{w}_t$, where \mathbf{Q}_t belongs to some set of matrices \mathcal{Q} , depends on the so-called *joint spectral radius* (JSR) of \mathcal{Q} [48, Proposition 3.2]. Computing (or bounding) the JSR of (even very small) matrix sets is a hard problem, currently under active research (see [48] and the many references therein). The convergence stated in Theorem 4 iv) results from the following fact: for $\alpha \leq 1$, there exists a matrix norm, say $\|\cdot\|_a$, for which $\|\mathbf{Q}_t\|_a \leq \varepsilon < 1$, for any $\mathbf{Q}_t \in \mathcal{Q}$, which is a sufficient condition for convergence to zero of the iteration $\mathbf{w}_{t+1} = \mathbf{Q}_t \mathbf{w}_t$.

Although, when $\alpha > 1$, Theorem 4 does not guarantee convergence, we have observed, in a large number of image deconvolution experiments, that the algorithm always converges for a wide range of choices of parameters α and β . In Section VI, we will discuss practical rules for choosing these parameters.

As in linear stationary algorithms, we have experimentally verified that $\hat{\rho}$ and $\bar{\rho}$, respectively, are good indicators of the relative speed of TwIST and IST. Seeing the algorithms as linear stationary, quantities $-1 / \log_{10} \hat{\rho}$ and $-1 / \log_{10} \bar{\rho}$ are approximately the numbers of iterations needed to reduce the error norm by a factor of 10 (see Appendix II). For example, with $\kappa \simeq 10^{-4}$ (common in image restoration), $-1 / \log_{10} \hat{\rho} \simeq 10^2$

and $-1/\log_{10} \bar{\rho} \simeq 10^4$; i.e., in this case, TwIST is expected to be roughly two orders of magnitude faster than IST, as confirmed in our experiments. To the best of our knowledge, the bound on the convergence factor of IST given by (25) has not appeared previously in the literature.

C. Monotonic Version: MTwIST

Monotonicity underlies the derivation of many algorithms and is instrumental in several convergence proofs; e.g., the proof of convergence of IST (which is monotonic for $\beta = 1$) in [21]. Monotonicity is not used in our convergence proof of TwIST (which is not necessarily monotonic), but the proof requires the condition that the observation operator is invertible. To handle noninvertible operators, we introduce a monotonic version of TwIST (MTwIST); the rationale is that, even though we cannot guarantee convergence of the estimate, monotonicity combined with the fact that the objective function is bounded below guarantees convergence of the objective function values $f(\mathbf{x}_t)$. Although this is a weaker result, we have observed in many experiments that MTwIST always converges and still does so much faster than IST.

The structure of MTwIST is very similar to that of TwIST, with a single difference. Formally, for $t > 1$, let \mathbf{z} be given by (18); then

$$\mathbf{x}_{t+1} = \begin{cases} \mathbf{z}, & \Leftarrow f(\mathbf{z}) \leq f(\mathbf{x}_t) \\ \Gamma_\lambda(\mathbf{x}_t), & \Leftarrow f(\mathbf{z}) > f(\mathbf{x}_t). \end{cases}$$

Notice that setting $\mathbf{x}_{t+1} = \Gamma_\lambda(\mathbf{x}_t)$ corresponds to taking a TwIST step with $\alpha = \beta = 1$, that is, a (monotonic) IST step.

VI. EXPERIMENTAL RESULTS

In this section, we present image restoration results illustrating the convergence speed of TwIST in comparison with IST. Our goal is not to assess the performance of image restoration criteria of the form (1); this has been carried out in several other publications, in comparison with other state of the art criteria (see [7], [25], [28], [31], and [34]). It is clear that the performance of such criteria (e.g., in terms of SNR improvement) does not depend on the optimization algorithms used to implement them, but only on the choice of the type of regularizer Φ . On the other hand, the relative convergence speed of the algorithms is essentially dependent on their structure.

We consider two classes of regularizers: i) $\Phi(\mathbf{x}) = \|\mathbf{x}\|_1$, where \mathbf{x} denotes wavelet coefficients of the image to be inferred, and ii) $\Phi(\mathbf{x}) = \Phi_{\text{ITV}}(\mathbf{x})$, i.e., isotropic TV. See Sections II-C and D for further details. In the case i) we use the simplest possible choice of wavelet: Haar discrete wavelet transform (DWT). We are well aware that this does not lead to state-of-the-art performance in terms of SNR improvement; however, the conclusions obtained concerning the relative speed of the algorithms are valid for other choices of wavelets and penalty functions. To implement Ψ_λ corresponding to the regularizer $\Phi_{\text{ITV}}(\mathbf{x})$, we use the algorithm introduced in [11].

Table I shows the setting used in each of the five experiments conducted. Experiments 1 and 2 correspond to a strong blur with low noise, whereas experiments 3 and 4 correspond to mild blur with medium noise. Our aim in choosing these two scenarios is

TABLE I
EXPERIMENTAL SETTING (\mathbf{W} IS THE INVERSE
DISCRETE WAVELET TRANSFORM)

Exp	Image	Linear Operator \mathbf{K}	$\Phi(\mathbf{x})$	BSNR
1	Camera	\mathbf{H}_1 (9×9 uniform)	Φ_{ITV}	40dB
2	Camera	$\mathbf{H}_1 \mathbf{W}$	$\ \mathbf{x}\ _1$	40dB
3	Lena	$\mathbf{H}_2 \left(\frac{[1,4,6,4,1]^T [1,4,6,4,1]}{256} \right)$	Φ_{ITV}	17dB
4	Lena	$\mathbf{H}_2 \mathbf{W}$	$\ \mathbf{x}\ _1$	17dB
5	Camera	40% missing samples	Φ_{ITV}	40 dB

to illustrate that TwIST converges much faster than IST in severely ill-conditioned LIPs and still faster than IST in mildly ill-conditioned LIPs. In all the experiments, the operator \mathbf{K} is normalized to have $\xi_m = 1$, thus $\kappa = \xi_1$, which we simply denote as ξ . Finally, experiment 5 considers a problem in which matrix \mathbf{K} models the loss of 40% of the image pixels (at random locations); matrix \mathbf{K} is, thus, $0.6 m \times m$, thus noninvertible (40% of its singular values are zero). This experiment illustrates the behavior of MTwIST on an severely ill-posed ($\xi_1 = 0$) problem.

Inspired by Theorem 4 ii), the TwIST parameters are initially set to

$$\alpha = \hat{\rho}^2 + 1 \quad (26)$$

$$\beta = 2\alpha/(1 + \xi) \quad (27)$$

where $\hat{\rho}$ is given by (20) and ξ is set according to a qualitative classification: $\xi = 10^{-1}$ or $\xi = 10^{-3}$ for, respectively, mildly or severely ill-conditioned LIPs. This choice may be, of course, far from optimal; we have observed, however, that it leads to speeds of convergence very close to the best ones obtained by hand tuning (α, β) . The reason for this is that, as illustrated below, TwIST is very robust with respect to the parameters (α, β) , namely for severely ill-conditioned LIPs. Another alternative is to run a few TwIST iterations, say t_0 , for each $\xi = 10^{-i}$, with $i = 1, 2, \dots$ and choose the value that leads to lowest value of $f(\mathbf{x}_{t_0})$. In the experiments reported below, we use this procedure with $t_0 = 5$.

The parameter β of IST is set according to Theorem 4 iii); i.e., $\beta = 2/(1 + \xi)$. This setting, yields the optimal spectral radius $\rho(\mathbf{Q}_t)$ associated to the one-step iterative method. We will refer to this particular *optimal* version of IST as IST_{opt} and to the original IST (i.e., $\beta = 1$) as IST_1 . Notice that since, in ill-conditioned problems, $\xi_1 \ll 1$, the optimal β is very close to the maximum allowed value that guarantees convergence; for example, for $\xi_1 = 10^{-3}$, we have $\beta = 1.998$.

In all the examples, the algorithms are initialized with \mathbf{x}_0 given by a Wiener filter and the parameter λ is hand tuned for the best SNR improvement.

Experiments 1 and 2: Fig. 1 shows the evolution of the objective function along the iterations⁴ confirming that TwIST converges much faster than IST_{opt} and IST_1 , which take, respectively, 2400 and 5800 iterations to reach the value of f

⁴Arguably, the horizontal axes should represent CPU time instead of number of iterations; however, we have verified that the CPU time per iteration differs by less than 1% between TwIST and IST, so this change would only imply a change of scale of these horizontal axes.




Fig. 1. TV-based deconvolution in a severely ill-conditioned problem (experiment 1). Evolution of the objective function $f(\mathbf{x}_t)$ produced by TwIST, IST_{opt} , and IST_1 .




Fig. 2. TV-based deconvolution in a severely ill-conditioned problem (experiment 1). Evolution of the SNR improvement (ISNR) produced by TwIST, IST_{opt} , and IST_1 .




Fig. 3. TV-based deconvolution in a severely ill-conditioned problem (experiment 1). Evolution of the objective function $f(\mathbf{x}_t)$ produced by TwIST, for different parameters $(\alpha(\xi), \beta(\xi))$, and by IST_1 . Notice the low sensitivity of TwIST with respect to $(\alpha(\xi), \beta(\xi))$.

obtained by TwIST just after 100 iterations. Notice also that IST_{opt} converges approximately twice as fast as IST_1 . This pattern of behavior was systematically observed in severely ill-conditioned LIPs. Fig. 2 shows the evolution of the SNR

TABLE II
AVERAGE NUMBER OF ITERATIONS REQUIRED BY IST_1 AND IST_{opt} TO REACH THE SAME VALUE OF f OBTAINED BY 100 ITERATIONS OF TwIST

Initialization	IST_{opt}	IST_1
Wiener filter	2.4×10^3	5.8×10^3
Zeros	1.1×10^3	2.1×10^3
Observed image	1.1×10^3	2.0×10^3




Fig. 4. Wavelet-based deconvolution in a severely ill-conditioned problem (experiment 2). Evolution of the objective function $f(\mathbf{x}_t)$ produced by TwIST, IST_{opt} , and IST_1 .

improvement (ISNR) produced by TwIST, IST_{opt} and IST_1 . As expected,⁵ $ISNR(t)$ also converges much faster with TwIST than with IST_1 and IST_{opt} . Fig. 3 shows the evolution of the objective function $f(\mathbf{x}_t)$ produced by TwIST, for different parameters $(\alpha(\xi), \beta(\xi))$, and by IST_1 . Notice the low sensitivity of TwIST with respect to $(\alpha(\xi), \beta(\xi))$. This is a relevant characteristic of TwIST, because the optimal setting for (α, β) is rarely know.

In order to assess the impact of the initialization on the relative performance of the algorithms, we considered two other initialization methods: an all zeros image and the observed image. Table II shows the average (over ten runs) number of iterations required by IST_1 and IST_{opt} to reach the value of the objective obtained by 100 iterations of TwIST. Initialization with zeros or the observed image decreases the advantage of TwIST by roughly 50%; however, the Wiener initialization leads (after 100 iterations) to a final value of f and an ISNR which are a little better than the other two methods.

Figs. 4 and 5 plot wavelet-based deconvolution results obtained with the setting of experiment 2. The comments to Figs. 4 and 5 are similar to those made for Figs. 1 and 3: TwIST converges much faster than IST_1 and IST_{opt} ; TwIST has low sensitivity with respect to ξ .

Experiments 3 and 4: Figs. 6 and 7 plot results obtained in mildly ill-conditioned LIP. The first aspect to note is that (as expected) all the algorithms converge much faster than in the severely ill-conditioned case. The limit situation is a denoising LIP (i.e., $\mathbf{K} = \mathbf{I}$ or unitary) in which the solution is obtained in just one step (with $\alpha = \beta = 1$ and $\mathbf{x}_0 = \mathbf{0}$). The other aspect to

⁵ $ISNR(t) = 10 \log_{10}(\|\mathbf{y} - \mathbf{x}\|/\|\mathbf{x}_t - \mathbf{x}\|)$, where \mathbf{x} is the original image.




Fig. 5. Wavelet-based deconvolution in a severely ill-conditioned problem (experiment 2). Evolution of the objective function $f(\mathbf{x}_t)$ produced by TwIST, for different parameters $(\alpha(\xi), \beta(\xi))$, and by IST_1 . Notice the low sensitivity of TwIST with respect to $(\alpha(\xi), \beta(\xi))$.




Fig. 8. TV-based image restoration from 40% missing samples (experiment 5). Evolution of the objective function $f(\mathbf{x}_t)$ produced by MTwIST, IST_{opt} , and IST_1 .




Fig. 6. TV-based deconvolution in a mildly ill-conditioned problem (experiment 3). Evolution of the objective function $f(\mathbf{x}_t)$ produced by TwIST, IST_{opt} , and IST_1 .




Fig. 7. Wavelet-based deconvolution in a mildly ill-conditioned problem (experiment 3). Evolution of the objective function $f(\mathbf{x}_t)$ produced by TwIST, IST_{opt} , and IST_1 .



Fig. 9. TV-based image restoration from 40% missing samples (experiment 5); top: observed image; bottom: restored image.

note is that although the behavior of all the algorithms is almost identical, TwIST is still slightly faster than IST.

Experiment 5: In this example, the goal is not to present a state-of-the-art method for restoration from missing samples,

but simply to illustrate the behavior of the algorithms with a noninvertible observation model. The evolution of the objective function in Fig. 8 shows that MTwIST converges considerably faster than IST_1 and IST with $\beta = 1.998$. In line with the results reported in 3 and 5, MTwIST is, again, rather insensitive to the choice of ξ_1 (which in this case can no longer be related to the minimum singular value of \mathbf{K} , which is zero) Fig. 9 shows the observed image (the missing samples are set to the mid level gray value) and the restored image produced by MTwIST.

VII. CONCLUDING REMARKS

In this paper we have introduced a new class of iterative methods, called TwIST, which have the form of **two**-step iter-

ative shrinkage/thresholding (TwIST) algorithms. The update equation depends on the two previous estimates (thus, the term *two-step*), rather than only on the previous one. This class contains and extends the iterative shrinkage/thresholding (IST) methods recently introduced.

We proved convergence of TwIST to minima of the objective function (for a certain range of the algorithm parameters) and derived bounds for the convergence factor as a function of the parameters defining the algorithm. Experimental results (in wavelet-based and TV-based deconvolution) have shown that TwIST can in fact be tuned to converge much faster than the original IST, specially in severely ill-conditioned problems, where the speed up can reach two orders of magnitude in a typical deblurring problem. We have also introduced MTwIST, a monotonic variant of TwIST, conceived for noninvertible observation operators; the performance of MTwIST was illustrated on a problem of image restoration from missing samples.

APPENDIX I CONVEX ANALYSIS

We very briefly review some basic convex analysis results used in this paper. For more details see [44] and [50].

Consider a function $f : \mathcal{X} \rightarrow [-\infty, +\infty] = \overline{\mathbb{R}}$, where $\overline{\mathbb{R}}$ is the extended real line, and \mathcal{X} is a real Hilbert space.

The function f is *convex* if $f(\alpha \mathbf{u} + (1 - \alpha)\mathbf{v}) \leq \alpha f(\mathbf{u}) + (1 - \alpha)f(\mathbf{v})$, for any $\mathbf{u}, \mathbf{v} \in \mathcal{X}$ and any $\alpha \in [0, 1]$. Convexity is said to be *strict* if the inequality holds strictly ($<$) for any $\mathbf{u}, \mathbf{v} \in \mathcal{X}$ and any $\alpha \in]0, 1[$.

The function f is *proper* if $f(\mathbf{x}) < \infty$, for at least one $\mathbf{x} \in \mathcal{X}$, and $f(\mathbf{x}) > -\infty$, for all $\mathbf{x} \in \mathcal{X}$.

The function f is *lower semi-continuous* (lsc) at \mathbf{v} if

$$\lim_{\delta \searrow 0} \inf_{\mathbf{x} \in B(\mathbf{v}, \delta)} f(\mathbf{x}) \geq f(\mathbf{v})$$

where $B(\mathbf{v}, \delta) = \{\mathbf{x} : \|\mathbf{x} - \mathbf{v}\| \leq \delta\}$ is the δ -ball around \mathbf{v} , and $\|\cdot\|$ is the norm in the Hilbert space \mathcal{X} .

A function f is called *coercive* if it verifies $\lim_{\|\mathbf{x}\| \rightarrow \infty} f(\mathbf{x}) = +\infty$.

Proper, lsc, coercive functions play a key role in optimization because of the following theorem (see [44]).

Theorem 5: If f is a proper, lsc, coercive, convex function, then $\inf_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ is finite and the set $\arg \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ is nonempty.

The next theorems concern strictly convex functions.

Theorem 6: If f is a strictly convex function, the set $\arg \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ possesses at most one element.

Theorem 7: If f_1 is a convex function, f_2 is a strictly convex function, and $0 < \lambda < \infty$, then λf_2 and $f_1 + \lambda f_2$ are strictly convex.

APPENDIX II

MATRIX NORMS, SPECTRAL RADIUS, CONVERGENCE, CONVERGENCE FACTORS, AND CONVERGENCE RATES

Given a vector norm $\|\cdot\|$, $\|\mathbf{A}\| = \max_{\|\mathbf{x}\|=1} \|\mathbf{A}\mathbf{x}\|$ is the matrix norm of \mathbf{A} induced by this vector norm. A vector norm and the corresponding induced matrix norm are *consistent*, i.e., they satisfy $\|\mathbf{A}\mathbf{v}\| \leq \|\mathbf{A}\| \|\mathbf{v}\|$.

When the vector norm is the Euclidean norm (denoted $\|\cdot\|_2$), the induced matrix norm (also denoted $\|\cdot\|_2$) is called *spectral*

norm, and is given by $\|\mathbf{A}\|_2 = \sqrt{\rho(\mathbf{A}^T \mathbf{A})}$, where $\rho(\mathbf{B}) = \max_i |\lambda_i(\mathbf{B})|$ is the largest absolute eigenvalue of \mathbf{B} , called the *spectral radius*. Key results involving $\rho(\cdot)$ are

$$\lim_{k \rightarrow \infty} \mathbf{A}^k = 0 \Leftrightarrow \rho(\mathbf{A}) < 1 \quad (28)$$

$$\lim_{k \rightarrow \infty} \|\mathbf{A}^k\|^{1/k} = \rho(\mathbf{A}) \quad (29)$$

$$\forall \mathbf{A}, \varepsilon \exists \|\cdot\|_a : \|\mathbf{A}\|_a \leq \rho(\mathbf{A}) + \varepsilon. \quad (30)$$

If \mathbf{A} is Hermitian, $\|\mathbf{A}\|_2 = \max_i |\lambda_i(\mathbf{A})| = \rho(\mathbf{A})$.

Consider the linear system $\mathbf{B}\mathbf{x} = \mathbf{b}$, with solution $\hat{\mathbf{x}}$ and an iterative scheme yielding a sequence of iterates $\{\mathbf{x}_t, t \in \mathbb{N}\}$. For a linear stationary iterative algorithm, the error $\mathbf{e}_t = \mathbf{x}_t - \hat{\mathbf{x}}$ evolves according to $\mathbf{e}_t = \mathbf{A}\mathbf{e}_{t-1}$, thus $\mathbf{e}_t = \mathbf{A}^t \mathbf{e}_0$. From (28), the error goes to zero if and only if $\rho(\mathbf{A}) < 1$. Because of (29), $\rho(\mathbf{A})$ is also called the *asymptotic convergence factor*. The *asymptotic convergence rate*, given by $r = -\log_{10} \rho(\mathbf{A})$, is roughly the number of new correct decimal places obtained per iteration, while its inverse approximates the number of iterations required to reduce the error by a factor of 10.

APPENDIX III PROOF OF THEOREM 4

Before proving Theorem 4, we introduce several results on which the proof is built, one of them being Clarke's mean value theorem for nondifferentiable functions [18]. Other preliminary results are presented and proved in Appendix III-C. Finally, Appendix III-C-F contain the proofs of parts i)–iv) of Theorem 4, respectively.

A. Nonsmooth Mean Value Theorem

Definition: Let $\mathbf{F} : \mathbb{R}^m \rightarrow \mathbb{R}^m$ be such that each of its components is Lipschitz and Ω_F the set of points at which F is non-differentiable. Let $JF(\mathbf{x})$ denote the $m \times m$ Jacobian matrix of F at \mathbf{x} , when $\mathbf{x} \notin \Omega_F$. The (Clarke's [18]) *generalized Jacobian* of F at \mathbf{x} is given by

$$\partial F(\mathbf{x}) = \text{co} \left\{ \lim_{\mathbf{x}_i \rightarrow \mathbf{x}, \mathbf{x}_i \notin \Omega_F} JF(\mathbf{x}_i) \right\} \quad (31)$$

where $\text{co}(\mathbf{A})$ denotes the convex hull of \mathbf{A} . If F is continuously differentiable at \mathbf{x} , then $\partial F(\mathbf{x}) = \{JF(\mathbf{x})\}$ [18].

Theorem 8: (Mean value theorem [18]) Let F be as in Definition 1 and $\mathbf{u}, \mathbf{v} \in \mathbb{R}^m$ be any two points. Then

$$F(\mathbf{u}) - F(\mathbf{v}) \in \text{co} \partial F([\mathbf{u}, \mathbf{v}]) (\mathbf{u} - \mathbf{v}) \quad (32)$$

where $\text{co} \partial F([\mathbf{u}, \mathbf{v}])$ denotes the convex hull of the set $\{\mathbf{A} : \mathbf{A} \in \partial F(\mathbf{r}), \mathbf{r} \in [\mathbf{u}, \mathbf{v}]\}$, with $[\mathbf{u}, \mathbf{v}]$ denoting the line segment between \mathbf{u} and \mathbf{v} . Expression (32) means that there exists a matrix $\mathbf{B} \in \text{co} \partial F([\mathbf{u}, \mathbf{v}])$, such that $F(\mathbf{u}) - F(\mathbf{v}) = \mathbf{B}(\mathbf{u} - \mathbf{v})$.

B. Preliminary Results

The two following propositions characterize the elements of the generalized Jacobian of denoising functions, $\partial \Psi_\lambda$, and of $\text{co} \partial \Psi_\lambda([\mathbf{u}, \mathbf{v}])$.

Proposition 2: For any $\mathbf{x} \in \mathbb{R}^m$, any $\mathbf{D} \in \partial \Psi_\lambda(\mathbf{x})$ is symmetric, positive semi-definite (psd), and $\|\mathbf{D}\|_2 \leq 1$.

Proof: The proof distinguishes two classes of regularizers. Consider first that Ψ_λ results from a regularizer in class $\Upsilon(\mathbb{R}^m)$ (see Section II-B); e.g., iTV, niTV, or $\Phi_{\ell_w^p}$. From Theorem 1 in Section II-B, $\Psi_\lambda(\mathbf{x}) = \mathbf{x} - \mathbf{P}_{\lambda C}(\mathbf{x})$. Thus, $\partial\Psi_\lambda(\mathbf{x}) = \mathbf{I} - \partial\mathbf{P}_{\lambda C}(\mathbf{x})$, that is, any element, say \mathbf{D} , of $\partial\Psi_\lambda(\mathbf{x})$ can be written as $\mathbf{D} = \mathbf{I} - \mathbf{A}$, where $\mathbf{A} \in \partial\mathbf{P}_{\lambda C}(\mathbf{x})$. Theorem 2.3 in [47] guarantees that \mathbf{A} is symmetric, psd, and $\|\mathbf{A}\|_2 \leq 1$. Thus, $\mathbf{D} = \mathbf{I} - \mathbf{A}$ is also symmetric, psd, and $\|\mathbf{D}\|_2 \leq 1$.

Consider now that Ψ_λ results from a $\Phi_{\ell_w^p}$ regularizer, with $p > 1$ (see Section II-E). Due to the component-wise structure of Ψ_λ , shown in (11), and since $S_{\tau,p}$ is continuously differentiable (see footnote 1), $\partial\Psi_\lambda(\mathbf{x})$ contains a single diagonal (thus, symmetric) matrix, say \mathbf{D} . As shown in [21], $S_{\tau,p}$ (for $p > 1$) is strictly monotonic and its derivative is upper bounded by 1, which implies that each entry of \mathbf{D} belongs to $]0,1[$. This implies that \mathbf{D} is psd and $\|\mathbf{D}\|_2 \leq 1$. ■

Proposition 3: For any pair of points $\mathbf{u}, \mathbf{v} \in \mathbb{R}^m$, any $\mathbf{B} \in \text{co}\partial\Psi_\lambda([\mathbf{u}, \mathbf{v}])$ is symmetric, psd, and $\|\mathbf{B}\|_2 \leq 1$.

Proof: From Proposition 2, for any $\mathbf{r} \in \mathbb{R}^m$, any $\mathbf{A} \in \partial\Psi_\lambda(\mathbf{r})$ is symmetric, psd, and has $\|\mathbf{A}\|_2 \leq 1$. Thus, $\text{co}\partial\Psi_\lambda([\mathbf{u}, \mathbf{v}])$ is the convex hull of a set matrices which are all symmetric, psd, and have norm no larger than 1. Therefore, any matrix $\mathbf{B} \in \text{co}\partial\Psi_\lambda([\mathbf{u}, \mathbf{v}])$ is also symmetric, psd, and has $\|\mathbf{B}\|_2 \leq 1$. ■

C. Proof of Theorem 4 i)

Recalling that $\mathbf{e}_t = \mathbf{x}_t - \hat{\mathbf{x}}$ and using (18), we write

$$\mathbf{e}_{t+1} = (1-\alpha)\mathbf{e}_{t-1} + (\alpha-\beta)\mathbf{e}_t + \beta[\mathbf{\Gamma}_\lambda(\mathbf{x}_t) - \mathbf{\Gamma}_\lambda(\hat{\mathbf{x}})]. \quad (33)$$

Using the definition of $\mathbf{\Gamma}_\lambda$ given in (19) and the mean value theorem (Theorem 8), we may write

$$\begin{aligned} \mathbf{\Gamma}_\lambda(\mathbf{x}_t) - \mathbf{\Gamma}_\lambda(\hat{\mathbf{x}}) &= \Psi_\lambda \left(\underbrace{\mathbf{x}_t + \mathbf{K}^T(\mathbf{y} - \mathbf{K}\mathbf{x}_t)}_{\mathbf{z}_t} \right) \\ &\quad - \Psi_\lambda \left(\underbrace{\hat{\mathbf{x}} + \mathbf{K}^T(\mathbf{y} - \mathbf{K}\hat{\mathbf{x}})}_{\hat{\mathbf{z}}} \right) \\ &= \mathbf{B}_t [\mathbf{x}_t - \hat{\mathbf{x}} + \mathbf{K}^T\mathbf{K}(\mathbf{x}_t - \hat{\mathbf{x}})] \\ &= \mathbf{B}_t[\mathbf{I} - \mathbf{K}^T\mathbf{K}]\mathbf{e}_t \end{aligned} \quad (34)$$

where $\mathbf{B}_t \in \text{co}\partial\Psi_\lambda([\mathbf{z}_t, \hat{\mathbf{z}}])$. Recall that Proposition 3 states that \mathbf{B}_t is symmetric, psd, and has $\|\mathbf{B}_t\|_2 \leq 1$. Inserting (34) into (33)

$$\begin{aligned} \mathbf{e}_{t+1} &= (1-\alpha)\mathbf{e}_{t-1} + \alpha\mathbf{e}_t - \beta[\mathbf{I} - \mathbf{B}_t[\mathbf{I} - \mathbf{K}^T\mathbf{K}]]\mathbf{e}_t \\ &= (1-\alpha)\mathbf{e}_{t-1} + [\alpha\mathbf{I} - \beta\mathbf{M}_t]\mathbf{e}_t \end{aligned} \quad (35)$$

where

$$\mathbf{M}_t = \mathbf{I} - \mathbf{B}_t[\mathbf{I} - \mathbf{K}^T\mathbf{K}]. \quad (36)$$

Recalling that the *stacked error vector* $\mathbf{w}_t \in \mathbb{R}^{2m}$ is

$$\mathbf{w}_t = \begin{bmatrix} \mathbf{e}_{t+1} \\ \mathbf{e}_t \end{bmatrix}$$

we can use (35) to write $\mathbf{w}_t = \mathbf{Q}_t\mathbf{w}_{t-1}$, where

$$\mathbf{Q}_t = \begin{bmatrix} (\alpha\mathbf{I} - \beta\mathbf{M}_t) & (1-\alpha)\mathbf{I} \\ \mathbf{I} & \mathbf{0} \end{bmatrix}. \quad (37)$$

Thus, \mathcal{Q} is the set of matrices with the form (37), where \mathbf{M}_t is given by (36) and \mathbf{B}_t is symmetric, psd, and has $\|\mathbf{B}_t\|_2 \leq 1$.

To prove the second statement in Theorem 4 i), we need to study how the choice of α and β affects $\rho(\mathbf{Q}_t) = \max_i |\lambda_i(\mathbf{Q}_t)|$, for any possible \mathbf{M}_t . We begin by considering the following facts: a) $\mathbf{I} - \mathbf{K}^T\mathbf{K}$ is symmetric and $1 - \xi_m \leq \lambda_i(\mathbf{I} - \mathbf{K}^T\mathbf{K}) \leq 1 - \xi_1$ (because $\xi_1 \leq \lambda_i(\mathbf{K}^T\mathbf{K}) \leq \xi_m$); b) according to Proposition 3, \mathbf{B}_t is symmetric, psd, and $\|\mathbf{B}_t\|_2 \leq 1$, thus $0 \leq \lambda_i(\mathbf{B}_t) \leq 1$. Consequently, using results on bounds of eigenvalues of products of symmetric matrices, one of which is psd, [38, Theorem 2.2]

$$\min(0, 1 - \xi_m) \leq \lambda_i(\mathbf{B}_t[\mathbf{I} - \mathbf{K}^T\mathbf{K}]) \leq 1 - \xi_1 \quad (38)$$

finally, since $\mathbf{M}_t = \mathbf{I} - \mathbf{B}_t[\mathbf{I} - \mathbf{K}^T\mathbf{K}]$

$$0 < \xi_1 \leq \lambda_i(\mathbf{M}_t) \leq \max(1, \xi_m) \equiv \bar{\xi}_m. \quad (39)$$

Following [3], let (μ, \mathbf{z}) denotes any eigenpair of \mathbf{Q}_t , i.e., $\mathbf{Q}_t\mathbf{z} = \mu\mathbf{z}$; writing $\mathbf{z} = [\mathbf{z}_a^T, \mathbf{z}_b^T]^T$, we have

$$\begin{bmatrix} \alpha\mathbf{I} - \beta\mathbf{M}_t & (1-\alpha)\mathbf{I} \\ \mathbf{I} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{z}_a \\ \mathbf{z}_b \end{bmatrix} = \mu \begin{bmatrix} \mathbf{z}_a \\ \mathbf{z}_b \end{bmatrix}. \quad (40)$$

The bottom m rows of (40) give $\mathbf{z}_a = \mu\mathbf{z}_b$; inserting this equality into the top half of (40), we obtain

$$[\mu(\alpha\mathbf{I} - \beta\mathbf{M}_t) + (1-\alpha)\mathbf{I}]\mathbf{z}_b = \mu^2\mathbf{z}_b. \quad (41)$$

Since the matrix in the left-hand side (l.h.s.) of (41) can be written as $(\mu\alpha + 1 - \alpha)\mathbf{I} - \mu\beta\mathbf{M}_t$, its eigenvectors coincide with those of \mathbf{M}_t . Thus, with λ denoting some eigenvalue of \mathbf{M}_t , μ has to be a solution of the following second degree equation

$$(\mu\alpha + 1 - \alpha) - \mu\beta\lambda = \mu^2. \quad (42)$$

Let

$$\tilde{\rho}(\alpha, \beta, \lambda) = \max\{|\mu_1|, |\mu_2|\} \quad (43)$$

where μ_1 and μ_2 are the two solutions of (42). We, thus, need to study how $\tilde{\rho}(\alpha, \beta, \lambda)$ behaves for $\lambda \in [\lambda_{\min}(\mathbf{M}_t), \lambda_{\max}(\mathbf{M}_t)] \equiv [\tau_1, \tau_m]$, for each choice of α and β . Notice that (39) does not provide τ_1 and τ_m (all it guarantees is that $[\tau_1, \tau_m] \subseteq [\xi_1, \bar{\xi}_m]$). It is shown in [3, Lemma 5.8] that $\tilde{\rho}(\alpha, \beta, \lambda) < 1$, for any $\lambda \in [\tau_1, \tau_m]$, if $0 < \alpha < 2$ and $0 < \beta < 2\alpha/\tau_m$. Since $\tau_m \leq \bar{\xi}_m$, any β satisfying $\beta < 2\alpha/\bar{\xi}_m$ also satisfies $\beta < 2\alpha/\tau_m$.

Finally, notice that $\rho(\mathbf{Q}_t) = \max_i \{\tilde{\rho}(\alpha, \beta, \lambda_i(\mathbf{M}_t))\}$; thus, $\tilde{\rho}(\alpha, \beta, \lambda) < 1$ implies that $\rho(\mathbf{Q}_t) < 1$, concluding the proof of Theorem 4 i). ■

D. Proof of Theorem 4 ii)

We begin by re-writing (42), for $\alpha = \hat{\alpha}$ and $\beta = \hat{\beta}$, as

$$\mu^2 + (\lambda\hat{\beta} - \hat{\alpha})\mu + (\hat{\alpha} - 1) = 0 \quad (44)$$

and proving that the solutions of (44) are complex conjugate for any $\lambda \in [\tau_1, \tau_m]$. From the classical formula for the solutions of a second degree equation, it is clear that the two roots of (44) are complex conjugate if and only if $(\lambda\hat{\beta} - \hat{\alpha})^2 \leq 4(\hat{\alpha} - 1)$, for any $\lambda \in [\tau_1, \tau_m]$; this inequality is equivalent to

$$\hat{\beta}^2 \lambda^2 - 2\hat{\alpha}\hat{\beta}\lambda + \hat{\alpha}^2 - 4(\hat{\alpha} - 1) \leq 0. \quad (45)$$

It is easy to show that the two roots of l.h.s. of (45) are ξ_1 and $\bar{\xi}_m$; thus, since $\hat{\beta}^2 > 0$, inequality (45) is satisfied when λ is between these two roots. Therefore, when $\lambda \in [\tau_1, \tau_m] \subseteq [\xi_1, \bar{\xi}_m]$, the roots of (44) are indeed complex conjugate.

Recall that the product of the two roots of a second order polynomial equals its independent term; applying this fact to (44) yields $\mu_1\mu_2 = (\hat{\alpha} - 1)$. For $\lambda \in [\tau_1, \tau_m]$, we have $\mu_1 = \mu_2^*$, thus $\mu_1\mu_2 = |\mu_1|^2 = |\mu_2|^2 = (\hat{\alpha} - 1)$; thus $\tilde{\rho}(\hat{\alpha}, \hat{\beta}, \lambda) = \max\{|\mu_2|, |\mu_1|\} = \sqrt{\hat{\alpha} - 1} = \hat{\rho}$, for any $\lambda \in [\tau_1, \tau_m]$, as stated in Theorem 4 ii). ■

E. Proof of Theorem 4 iii)

Inserting $\alpha = 1$ and $\beta = \bar{\beta}$ in (42) leads to the equation

$$\mu(1 - \bar{\beta}\lambda) = \mu^2 \quad (46)$$

which has solutions $\mu_1 = 0$ and $\mu_2 = (1 - \bar{\beta}\lambda)$. Consequently, $\tilde{\rho}(1, \bar{\beta}, \lambda) = \max\{|\mu_1|, |\mu_2|\} = |1 - \bar{\beta}\lambda|$. To show part iii) of the theorem, we need to show that

$$\max_{\lambda \in [\tau_1, \tau_m]} \tilde{\rho}(1, \bar{\beta}, \lambda) \leq \frac{1 - \kappa}{1 + \kappa}.$$

Because $\tilde{\rho}(1, \bar{\beta}, \lambda)$ and $(1 - \kappa)/(1 + \kappa)$ are positive, both sides of the previous inequality can be squared. Simple manipulation allows showing that

$$\tilde{\rho}^2(1, \bar{\beta}, \xi_1) = \tilde{\rho}^2(1, \bar{\beta}, \bar{\xi}_m) = \left(\frac{1 - \kappa}{1 + \kappa}\right)^2.$$

Finally, since $\tilde{\rho}^2(1, \bar{\beta}, \lambda) = (1 - \beta\lambda)^2$ is a convex function of λ , and $[\tau_1, \tau_m] \subseteq [\xi_1, \bar{\xi}_m]$

$$\max_{\lambda \in [\tau_1, \tau_m]} \tilde{\rho}^2(1, \bar{\beta}, \lambda) \leq \max_{\lambda \in [\xi_1, \bar{\xi}_m]} \tilde{\rho}^2(1, \bar{\beta}, \lambda) \leq \left(\frac{1 - \kappa}{1 + \kappa}\right)^2$$

concluding the proof of Theorem 4 iii). ■

F. Proof of Theorem 4 iv)

A sufficient condition for convergence to zero of the switched linear system $\mathbf{z}_{t+1} = \mathbf{T}_t \mathbf{z}_t$, where $\mathbf{T}_t \in \mathcal{T}$, and \mathcal{T} is a bounded set of matrices, is the existence of a matrix norm $\|\cdot\|$, such that $\|\mathbf{T}_t\| \leq \varepsilon < 1$, for any $\mathbf{T}_t \in \mathcal{T}$. Our proof uses the matrix norm $\|\cdot\|_{\mathbf{A}}$, defined as

$$\|\mathbf{B}\|_{\mathbf{A}} = \|\mathbf{A}\mathbf{B}\mathbf{A}^{-1}\|_2, \quad (47)$$

where \mathbf{A} is a symmetric positive definite matrix, which is induced by the vector norm $\|\cdot\|_{\mathbf{A}} = \|\mathbf{A}\mathbf{v}\|_2$ [10].

We split the proof into two cases.

a) With $\alpha = 1$, the error \mathbf{e}_t evolves according to the one-step iteration

$$\mathbf{e}_{t+1} = (\mathbf{I} - \beta\mathbf{M}_t)\mathbf{e}_t \quad (48)$$

matrix \mathbf{M}_t [see (36)] can be written as $\mathbf{M}_t = \mathbf{I} - \mathbf{B}_t\mathbf{U}$, where $\mathbf{U} = \mathbf{I} - \mathbf{K}^T\mathbf{K}$ is a symmetric positive definite matrix; thus, so is $\mathbf{U}^{1/2}$. Computing the $\|\cdot\|_{\mathbf{U}^{1/2}}$ norm of $(\mathbf{I} - \beta\mathbf{M}_t)$

$$\begin{aligned} \|\mathbf{I} - \beta\mathbf{M}_t\|_{\mathbf{U}^{1/2}} &= \left\| \mathbf{U}^{1/2}(\mathbf{I} - \beta\mathbf{M}_t)\mathbf{U}^{-1/2} \right\|_2 \\ &= \left\| \mathbf{I} - \beta(\mathbf{I} - \mathbf{U}^{1/2}\mathbf{B}_t\mathbf{U}^{1/2}) \right\|_2 \\ &= \rho(\mathbf{I} - \beta\mathbf{M}_t) \end{aligned} \quad (49)$$

where we have used the following facts: for a real symmetric matrix \mathbf{A} , $\|\mathbf{A}\|_2 = \rho(\mathbf{A})$ and, for any pair of square matrices \mathbf{A} and \mathbf{B} , $\rho(\mathbf{A}\mathbf{B}) = \rho(\mathbf{B}\mathbf{A})$.

Finally, notice that, as shown in Appendix III-E

$$\rho(\mathbf{I} - \beta\mathbf{M}_t) = \max_{\lambda \in [\tau_1, \tau_m]} \tilde{\rho}(1, \beta, \lambda) \leq \frac{1 - \kappa}{1 + \kappa} < 1$$

concluding the convergence proof, for $\alpha = 1$.

b) With $\alpha < 1$, let us define the matrix

$$\mathbf{V} = \begin{bmatrix} \mathbf{U} & \mathbf{0} \\ \mathbf{0} & (1 - \alpha)\mathbf{U} \end{bmatrix}. \quad (50)$$

With \mathbf{Q}_t given by (37), it is simple to conclude that

$$\mathbf{V}^{1/2}\mathbf{Q}_t\mathbf{V}^{-1/2} = \begin{bmatrix} (\alpha - \beta)\mathbf{I} + \beta\mathbf{U}^{1/2}\mathbf{B}_t\mathbf{U}^{1/2} & \sqrt{1 - \alpha}\mathbf{I} \\ \sqrt{1 - \alpha}\mathbf{I} & \mathbf{0} \end{bmatrix} \quad (51)$$

which is a real symmetric matrix. This allows writing

$$\begin{aligned} \|\mathbf{Q}_t\|_{\mathbf{V}^{1/2}} &= \|\mathbf{V}^{1/2}\mathbf{Q}_t\mathbf{V}^{-1/2}\|_2 \\ &= \rho(\mathbf{V}^{1/2}\mathbf{Q}_t\mathbf{V}^{-1/2}) \\ &= \rho(\mathbf{Q}_t) \\ &= \max_i \{\tilde{\rho}(\alpha, \beta, \lambda_i(\mathbf{M}_t))\} \end{aligned} \quad (52)$$

$$\leq \max_{\lambda \in [\tau_1, \tau_m]} \{\tilde{\rho}(\alpha, \beta, \lambda)\} \quad (53)$$

$$< 1 \quad (54)$$

where the equality in (52) and the inequalities (53) and (54) were shown in Section III-C. ■

REFERENCES

- [1] H. Andrews and B. Hunt, *Digital Image Restoration*. Englewood Cliffs, NJ: Prentice-Hall, 1977.
- [2] G. Archer and D. Titterton, "On Bayesian/regularization methods for image restoration," *IEEE Trans. Image Process.*, vol. 4, no. 3, pp. 989–995, Jul. 1995.
- [3] O. Axelsson, *Iterative Solution Methods*. New York: Cambridge Univ. Press, 1996.

- [4] J. Bect, L. Blanc-Féraud, G. Aubert, and A. Chambolle, "A ℓ^1 -unified variational framework for image restoration," in *Proc. Eur. Conf. Computer Vision*, 2004, pp. 1–13, Springer-Verlag.
- [5] M. Bertero and P. Boccacci, *Introduction to Inverse Problems in Imaging*. Bristol, U.K.: IOP, 1998.
- [6] J. Bioucas-Dias, "Fast GEM wavelet-based image deconvolution algorithm," in *Proc. IEEE Int. Conf. Image Processing*, 2003, vol. 2, pp. 961–964.
- [7] J. Bioucas-Dias, "Bayesian wavelet-based image deconvolution: A GEM algorithm exploiting a class of heavy-tailed priors," *IEEE Trans. Image Process.*, vol. 15, no. 4, pp. 937–951, Apr. 2006.
- [8] J. Bioucas-Dias and M. Figueiredo, "Two-step algorithms for linear inverse problems with non-quadratic regularization," presented at the IEEE Int. Conf. Image Processing, San Antonio, TX, 2007.
- [9] J. Bioucas-Dias, M. Figueiredo, and J. Oliveira, "Total variation image deconvolution: A majorization-minimization approach," presented at the IEEE Int. Conf. Acoustics, Speech, and Signal Processing, Toulouse, France, 2006.
- [10] Z.-H. Cao, "On the convergence on nonstationary iterative methods for symmetric positive (semi)definite systems," *Appl. Numer. Math.*, vol. 37, pp. 319–330, 2001.
- [11] A. Chambolle, "An algorithm for total variation minimization and applications," *J. Math. Imag. Vis.*, vol. 20, pp. 89–97, 2004.
- [12] A. Chambolle, "Total variation minimization and a class of binary MRF models," in *Proc. Int. Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, 2005, pp. 136–152, Springer Verlag.
- [13] A. Chambolle, R. De Vore, N. Lee, and B. Lucier, "Nonlinear wavelet image processing: Variational problems, compression, and noise removal through wavelet shrinkage," *IEEE Trans. Image Process.*, vol. 7, no. 3, pp. 319–335, Mar. 1998.
- [14] A. Chambolle and P.-L. Lions, "Image recovery via total variation minimization and related problems," *Numer. Math.*, vol. 76, pp. 167–188, 1997.
- [15] T. Chan, S. Esedoglu, F. Park, and A. Yip, "Recent developments in total variation image restoration," in *Mathematical Models of Computer Vision*, N. Paragios, Y. Chen, and O. Faugeras, Eds. New York: Springer Verlag, 2005.
- [16] C. Chau, P. Combettes, J.-C. Pesquet, and V. Wajs, "Iterative image deconvolution using overcomplete representations," presented at the Eur. Signal Processing Conf., Florence, Italy, 2006.
- [17] S. Chen, D. Donoho, and M. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.*, vol. 20, pp. 33–61, 1998.
- [18] F. Clarke, *Optimization and Nonsmooth Analysis*. New York: Wiley, 1983.
- [19] P. Combettes and V. Wajs, "Signal recovery by proximal forward-backward splitting," *SIAM J. Multiscale Model. Simul.*, vol. 4, pp. 1168–1200, 2005.
- [20] J. Darbon and M. Sigelle, "A fast and exact algorithm for total variation minimization," in *Proc. 2nd Iberian Conf. Pattern Recognition and Image Analysis*, 2005, vol. 3522, pp. 351–359.
- [21] I. Daubechies, M. Defriese, and C. De Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Commun. Pure Appl. Math.*, vol. LVII, pp. 1413–1457, 2004.
- [22] D. Donoho and I. Johnstone, "Adapting to unknown smoothness via wavelet shrinkage," *J. Amer. Statist. Assoc.*, vol. 90, no. 432, pp. 1200–1224, 1995.
- [23] D. Donoho, "De-noising by soft thresholding," *IEEE Trans. Inf. Theory*, vol. 41, no. 3, pp. 613–627, May 1995.
- [24] M. Elad, "Why simple shrinkage is still relevant for redundant representations?," *IEEE Trans. Inf. Theory*, vol. 52, no. 12, pp. 5559–5569, Dec. 2006.
- [25] M. Elad, B. Matalon, and M. Zibulevsky, "Image denoising with shrinkage and redundant representations," presented at the IEEE Comput. Soc. Conf. Computer Vision and Pattern Recognition, New York, 2006.
- [26] M. Figueiredo, "Adaptive sparseness for supervised learning," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, no. 9, pp. 1150–1159, Sep. 2003.
- [27] M. Figueiredo, J. Bioucas-Dias, and R. Nowak, "Majorization-minimization algorithms for wavelet-based image restoration," *IEEE Trans. Image Process.*, to be published.
- [28] M. Figueiredo and R. Nowak, "An EM algorithm for wavelet-based image restoration," *IEEE Trans. Image Process.*, vol. 12, no. 8, pp. 906–916, Aug. 2003.
- [29] M. Figueiredo and R. Nowak, "A bound optimization approach to wavelet-based image deconvolution," presented at the IEEE Int. Conf. Image Processing, Genoa, Italy, 2005.
- [30] M. Figueiredo, J. Bioucas-Dias, J. Oliveira, and R. Nowak, "On total-variation denoising: A new majorization-minimization algorithm and an experimental comparison with wavelet denoising," presented at the IEEE Int. Conf. Image Processing, 2006.
- [31] J. Guerrero-Colon and J. Portilla, "Deblurring-by-denoising using spatially adaptive Gaussian scale mixtures in overcomplete pyramids," presented at the IEEE Int. Conf. Image Processing, Atlanta, GA, 2006.
- [32] Y. Grandvalet and S. Canu, "Outcomes of the equivalence of adaptive ridge with least absolute shrinkage," in *Advances in Neural Information Processing Systems-NIPS-11*. Cambridge, MA: MIT Press, 1999.
- [33] D. Hunter and K. Lange, "A tutorial on MM algorithms," *Amer. Statist.*, vol. 58, pp. 30–37, 2004.
- [34] A. Jalobeanu, N. Kingsbury, and J. Zerubia, "Image deconvolution using hidden Markov tree modeling of complex wavelet packets," presented at the IEEE Int. Conf. Image Processing, Thessaloniki, Greece, 2001.
- [35] A. Katsaggelos, *Digital Image Restoration*, ser. Ser. Inf. Sci.. New York: Springer-Verlag, 1991, vol. 23.
- [36] D. Lorenz, "Wavelet shrinkage in signal and image processing: an investigation of relations and equivalences" Ph.D. dissertation, Zentrum für Technomathematik, Univ. Bremen, Bremen, Germany, 2004 [Online]. Available: www.math.uni-bremen.de/dlorenz
- [37] D. Lorenz, "Variational denoising in besov spaces and interpolation of hard and soft wavelet shrinkage," Tech. Rep. Zentrum für Technomathematik, Univ. Bremen, Germany, 2004 [Online]. Available: www.math.uni-bremen.de/dlorenz
- [38] L. Lu and C. Pearce, "Some new bounds for singular values and eigenvalues of matrix products," *Ann. Oper. Res.*, vol. 98, pp. 141–148, 2000.
- [39] P. Moulin and J. Liu, "Analysis of multiresolution image denoising schemes using generalized-Gaussian and complexity priors," *IEEE Trans. Inf. Theory*, vol. 45, no. 3, pp. 909–919, Apr. 1999.
- [40] R. Neelamani, H. Choi, and R. Baraniuk, "ForWaRD: Fourier-wavelet regularized deconvolution for ill-conditioned systems," *IEEE Trans. Signal Process.*, vol. 52, no. 2, pp. 418–433, Feb. 2004.
- [41] R. Nowak and M. Figueiredo, "Fast wavelet-based image deconvolution using the EM algorithm," in *Proc. 35th Asilomar Conf. Signals, Systems, Computers*, 2001, vol. 1, pp. 371–375.
- [42] S. Osher, L. Rudin, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D*, vol. 60, pp. 259–268, 1992.
- [43] J. Portilla, V. Strela, M. Wainwright, and E. Simoncelli, "Image denoising using scale mixtures of Gaussians in the wavelet domain," *IEEE Trans. Image Process.*, vol. 12, no. 11, pp. 1338–1351, Nov. 2003.
- [44] R. T. Rockafellar and R. Wets, *Variational Analysis*. Berlin, Germany: Springer Verlag, 1998.
- [45] J.-L. Starck, E. Candès, and D. Donoho, "Astronomical image representation by the curvelet transform," *Astron. Astrophys.*, vol. 398, pp. 785–800, 2003.
- [46] J.-L. Starck, M. Nguyen, and F. Murtagh, "Wavelets and curvelets for image deconvolution: A combined approach," *Signal Process.*, vol. 83, pp. 2279–2283, 2003.
- [47] D. Sun and L. Qi, "Solving variational inequality problems via smoothing-nonsmooth reformulations," *J. Comput. Appl. Math.*, vol. 129, pp. 37–62, 2001.
- [48] J. Theys, "Joint spectral radius: Theory and approximations" Ph.D. dissertation, Univ. Catholique de Louvain, Louvain, Belgium, 2005 [Online]. Available: www.inma.ucl.ac.be/~blondel/05thesetheys.pdf
- [49] C. Wu, "On the convergence properties of the EM algorithm," *Ann. Statist.*, vol. 11, pp. 95–103, 1983.
- [50] C. Zalinescu, *Convex Analysis in General Vector Spaces*. Singapore: World Scientific, 2002.



José M. Bioucas-Dias (S'87–M'95) received the E.E., M.Sc., and Ph.D. and "Agregado" degrees in electrical and computer engineering from the Instituto Superior Técnico (IST), Technical University of Lisbon, Lisbon, Portugal, in 1985, 1991, 1995, and 2007, respectively.

He is currently an Assistant Professor with the Department of Electrical and Computer Engineering, IST. He is also a Researcher with the Communication Theory and Pattern Recognition Group, Institute of Telecommunications. His scientific interests

include signal and image processing, pattern recognition, optimization, and remote sensing imaging.

Dr. Bioucas-Dias was an Associate Editor of the IEEE TRANSACTION ON CIRCUITS AND SYSTEMS II. He has been a member of program committees of several international conferences, including CVPR, IGARSS, and ISVC. He is a researcher of several national and international research projects and networks including the Marie Curie Actions "Hyperspectral Imaging Network (HYPER-I-NET)" and the "European Doctoral Program in Signal Processing (SIGNAL)."



Mário A. T. Figueiredo (S'87–M'95–SM'00) received the E.E., M.Sc., Ph.D., and “Agregado” degrees in electrical and computer engineering from the Instituto Superior Técnico (IST), Technical University of Lisbon, Lisbon, Portugal, in 1985, 1990, 1994, and 2004, respectively. Since 1994, he has been with the faculty of the Department of Electrical and Computer Engineering, IST.

He is a Researcher and Area Coordinator at the Institute of Telecommunications, Lisbon. He held visiting positions with the Department of Computer Science and Engineering, Michigan State University, East Lansing, and the Department of Electrical and Computer Engineering, University of Wisconsin, Madison, in 1998 and 2005, respectively. His scientific interests include image processing and analysis, computer vision, statistical pattern recognition, and statistical learning.

Dr. Figueiredo received the Portuguese IBM Scientific Prize in 1995 for work on unsupervised image restoration. He is a member of the IEEE Image and Multidimensional Signal Processing Technical Committee, he is/was an Associate Editor of the IEEE TRANSACTIONS ON IMAGE PROCESSING, the IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, the IEEE TRANSACTIONS ON MOBILE COMPUTING, *Pattern Recognition Letters*, and *Signal Processing*. He was Guest Co-Editor of special issues of the IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE and the IEEE TRANSACTIONS ON SIGNAL PROCESSING. He was a Co-Chair of the 2001 and 2003 Workshops on Energy Minimization Methods in Computer Vision and Pattern Recognition. He has been a member of program committees of several top international conferences, including CVPR, ECCV, ICIAR, ICASSP, ICIP, ICML, ICPR, MLSP, and NIPS.