Sparsity-inducing Nonconvex Nonseparable Regularization for Convex Image Processing

Sparsity-inducing Non-convex Non-separable Regularization for Convex Image Processing

A. Lanza S. Morigi I. Selesnick F. Sgallari

Last edit: 10:44am, December 14, 2021

Abstract

A popular strategy for determining solutions to linear least-squares problems relies on using sparsity-promoting regularizers and is widely exploited in image processing applications such as, e.g., image denoising, deblurring and inpainting. It is well known that, in general, non-convex regularizers hold the potential for promoting sparsity more effectively than convex regularizers such as, e.g., those involving the ℓ_1 norm. To avoid the intrinsic difficulties related to non-convex optimization, the Convex Non-Convex (CNC) strategy has been proposed, which allows the use of non-convex regularization while maintaining convexity of the total objective function. In this paper, a new CNC variational model is proposed, based on a more general parametric non-convex non-separable regularizer. The proposed model is applicable to a greater variety of image processing problems than prior CNC methods. We derive the convexity conditions and related theoretical properties of the presented CNC model, and we analyze existence and uniqueness of its solutions. A primal-dual forward-backward splitting algorithm is proposed for solving the related saddle-point problem. The convergence of the algorithm is demonstrated theoretically and validated empirically. Several numerical experiments are presented which prove the effectiveness of the proposed approach.

1 Introduction

A widely used technique for determining solutions to linear least-squares problems relies on solving sparsity-regularized variational models which, in general, take the following form

$$x^* \in \arg\min_{x\in\mathbb{R}^n} \mathcal{J}(x), \qquad \mathcal{J}(x) := \frac{1}{2}\,\|Ax-b\|_2^2 \,+\lambda\,\mathcal{R}(x), \tag{1}$$

where $b \in \mathbb{R}^m$ is the vector of observed data, $A \in \mathbb{R}^{m\times n}$ is a matrix representing the linear observation model in some basis and $\lambda > 0$ is the classical regularization parameter which controls the trade-off between fidelity to the observations and regularity of the solution of (1). The sparsity-inducing regularization function $\mathcal{R} : \mathbb{R}^n \to \mathbb{R}$ takes the following general form

$$\mathcal{R}(x) := \Phi(y)\,, \quad y := G(Lx)\,, \tag{2}$$

with $L \in \mathbb{R}^{r\times n}$, $G\colon \mathbb{R}^r \to \mathbb{R}^s$ a possibly nonlinear vector-valued function with $g_i\colon \mathbb{R}^r \to \mathbb{R}$, $i = 1,\dots,s$, representing its scalar-valued components and $\Phi\colon \mathbb{R}^s \to \mathbb{R}$ a sparsity-promoting penalty function [34, 35]. The aim of model (1)-(2) is to determine solutions x^* which are close to the data b according to the observation model and, at the same time, for which the transformed vector $y^* = G(Lx^*)$ is sparse. We notice that the penalty Φ in (2) can be (additively) separable or non-separable with respect to its argument $y = G(Lx)$.

The focus of this paper is on the choice of the sparsity-promoting regularizer $\mathcal{R}$ in (1). In the literature, the most widely used sparsity-promoting penalties belong to the class of additively separable functions, that is they have the form $\Phi(y_1, y_2, \ldots, y_s) = \sum_{i=1}^{s} \phi_i(y_i)$, with $\phi_i : \mathbb{R} \to \mathbb{R}$.

In particular, the most natural sparsity-inducing penalty is the ℓ_0 pseudo-norm, namely

$$\Phi(y) = \|y\|_0 = \#\{i : \, y_i \neq 0\}, \tag{3}$$

which counts the number of non-zero elements in y. However, this choice yields (1)-(2) to be an NP-hard and non-convex optimization problem [26], with all the associated computational and numerical difficulties. The most popular surrogate of the ℓ_0 pseudo-norm is the ℓ_1 norm, defined as

$$\Phi(y) = \|y\|_1 = \sum_{i=1}^{s} |y_i| \, . \tag{4}$$

In fact, it is well known that the ℓ_1 norm is the convex regularizer which induces sparsity most effectively [4]; however, it tends to underestimate high-amplitude components of the vector to which it is applied, in our case $y = G(Lx)$. Consequently, the ℓ_1 norm is in many cases a sub-optimal choice for inducing sparsity, and non-convex penalties are known which outperform the ℓ_1 norm in this respect, in particular yielding more accurate estimation of high-amplitude components. A representative non-convex alternative to the ℓ_1 norm is the ℓ_p quasi-norm [21], defined as

$$\Phi(y) = \frac{1}{p} \|y\|_p^p = \frac{1}{p} \sum_{i=1}^{s} |y_i|^p \, , \qquad 0 < p < 1 \, . \tag{5}$$

However, by using such non-convex regularizers, the cost function $\mathcal{J}$ is non-convex and can present extraneous suboptimal local minimizers [27]. A solution to this problem is the CNC strategy, consisting in constructing and then optimizing convex functionals containing non-convex (sparsity-promoting) regularization terms. In particular, a quite recent alternative to (5) is represented by the class of additively separable penalties

$$\Phi(y) = \sum_{i=1}^{s} \phi(y_i; a) \, , \tag{6}$$

with $\phi \colon \mathbb{R} \to \mathbb{R}$ a non-convex function parametrized by the scalar parameter a, which controls the degree of non-convexity of ϕ. Examples of ϕ functions can be found in [22, 27, 34, 35]. Recently, CNC variational models containing non-convex separable regularizers of the form (6) have been successfully used to solve several data processing problems [9, 18, 22, 34, 35].

In the particular case that the observation matrix A in (1) has full column rank, it has been demonstrated [19] that it is possible to choose a in such a way that the total cost functional $\mathcal{J}$ in (1) is convex even if the penalty Φ is a non-convex function of the form (6). Such penalties can promote sparsity better than the convex ℓ_1 norm penalty. However, when the matrix A does not have full-column rank, a CNC formulation of this kind is not possible when Φ is separable as in (6), see [19].

For what concerns non-separable sparsity-inducing penalties in (1)-(2), pioneering work has been conducted in [37,42]; however, such penalties were not designed to maintain cost function convexity. More recently, families of non-convex non-separable penalty functions have been proposed in [31,33] that do maintain convexity of the cost functional $\mathcal{J}$ for any matrix A, but only in the special case where both G and L in (2) are identity operators.

In other related work, non-convex regularizers have been constructed via convexification of particular functionals involving the ℓ_0 pseudo-norm and rank functions [7, 36]. An aim in these

works is to reduce the number of non-optimal local minima as compared with ℓ_0 psuedo-norm regularized problems, so these non-convex regularizers do not maintain convexity of the objective function in the general case.

In this paper we extend the applicability of the CNC strategy: we propose a family of non-convex penalties that maintain the convexity of the cost functional $\mathcal{J}$ for any matrices A and L and quite general functions G. The proposed family of penalties generalizes earlier work [19,31,32]. More precisely, we extend the results in [31] wherein L is the identity, and the results in [32] wherein A is the identity and L is specifically the first-order finite-difference operator. Moreover, in [31] and [32] the authors consider only one-dimensional signals, whereas in this paper we consider the application of CNC to image processing. We note that while [31] and [32] are special cases of our proposal when L is the identity operator, and when both L and G are the identity operators, respectively, the model in [19] can not be obtained as a special case of our proposal. Some popular variational image processing models which rely on different combinations of the operators A, L G and of the penalty function ϕ will be briefly outlined in Section 2.

The new CNC non-separable (CNC-NS) regularization approach proposed in this work consists of a general strategy for constructing non-convex non-separable regularizers starting from any convex regularizer of the form (2). More precisely, our proposal is as follows. For any convex regularizer $\mathcal{R}$ of the form (2), we propose the parameterized non-convex non-separable regularizer

$$\mathcal{R}_B(x) := \mathcal{R}(x) - \left(\mathcal{R} \,\square\, \tfrac{1}{2}\|B\cdot\|_2^2\right)(x), \tag{7}$$

where $\square$ denotes the infimal convolution operator and matrix $B \in \mathbb{R}^{q\times n}$ is a parameter. The cost function to be minimized, instead of $\mathcal{J}$ in (1), is then given by

$$\mathcal{J}_B(x) := \frac{1}{2}\|Ax-b\|_2^2 + \lambda\,\mathcal{R}_B(x). \tag{8}$$

The rationale of our proposal, the properties of the regularizer R_B, and guidelines for choosing B will be outlined in the next sections. The proposed regularizer exhibits two noteworthy properties:

- Under mild assumptions on A, L, and G, there always exists a parameter matrix B such that the total cost function $\mathcal{J}_B$ is convex and $\mathcal{R}_B$ is non-convex. The assumptions do not require that $A^\mathsf{T}A$ be invertible. Hence the assumptions are milder than for non-convex separable regularizers (6).

- The proposed regularizer $\mathcal{R}_B$ induces sparsity more effectively than both the convex ℓ_1 norm penalty (4) and the class of non-convex separable penalties (6).

In this paper, we propose two strategies to set the parameter matrix B, we analyze some theoretical aspects of the CNC-NS model, and we prove the existence and the uniqueness of its solution under mild assumptions. Regarding the numerical solution of the proposed optimization problem, we cast the problem as a saddle-point problem for which the forward-backward splitting algorithm can be applied. The convergence of the algorithm is demonstrated and the effectiveness of the proposed approach is assessed by several examples.

The paper is organized as follows. In Section 2 we briefly outline some popular related variational models. In Section 3 we present in detail the proposed strategy for constructing non-separable regularizers and in Section 4 we highlight their properties. Convexity, coerciveness and existence/uniqueness of solution(s) to the introduced model are discussed in Section 5. In Section 6 we describe in detail the forward-backward splitting optimization algorithm used to solve the CNC model. Numerical experiments are reported in Section 7 and conclusions are drawn in Section 8.

2 Related Models

The general least-squares model (1) with sparsity-promoting convex regularizer $\mathcal{R}$ of the form (2) encompasses a large class of specific convex models which are of great importance in many different research areas, including numerical linear algebra [2, 43], image restoration [30], pattern recognition [13, 23], and compressed sensing [5, 14]. Various choices of the operators A, L, G and of the penalty function Φ yield a variety of popular models that have been successfully used in many research and application fields.

In general, model (1)-(2) with $A \in \mathbb{R}^{m\times n}$, $m < n$, G and L the identity operators and Φ the ℓ_1 norm, can be used in numerical linear algebra to compute sparse solutions of undetermined linear systems.

In the field of image processing, which is the one of interest in this paper, many important sparsity-inducing convex regularizers can be casted in the form (2).

For instance, the function $\mathcal{R}$ in (2) represents the popular Total Variation (TV) regularizer when $L = [D_h^\mathsf{T}, D_v^\mathsf{T}]^\mathsf{T}$ with $D_h, D_v \in \mathbb{R}^{n\times n}$ representing finite difference approximations of the first-order partial derivatives along the horizontal and vertical directions, respectively, Φ is the ℓ_1 norm penalty, and G characterizes anisotropic or isotropic TV [25]. In formulas:

$$\mathcal{R}(x) = \|G(Lx)\|_1 = \begin{cases} \displaystyle\sum_{i=1}^{2n} |g_i(Lx)| \quad \text{with} \quad g_i(Lx) = (Lx)_i & \text{(anisotropic TV)} \\ \displaystyle\sum_{i=1}^{n} |g_i(Lx)| \quad \text{with} \quad g_i(Lx) = \sqrt{(D_h x)_i^2 + (D_v x)_i^2} & \text{(isotropic TV)}\,. \end{cases} \tag{9}$$

Other popular regularizers employed to derive regularized solutions of linear inverse imaging problems are based on the Schatten p-norm of the Hessian matrix computed at every pixel of the image [24], where, we recall, the Shatten p-norm $\|M\|_{\mathcal{S}_p}$ of a matrix $M \in \mathbb{R}^{z\times z}$ is defined by

$$\|M\|_{\mathcal{S}_p} := \left(\sum_{i=1}^{z} \sigma_i^p(M)\right)^{\frac{1}{p}}, \quad p > 0\,, \tag{10}$$

with $\sigma_i(M)$ indicating the i-th singular value of matrix M. These kind of regularizers can be considered as second-order extensions of the TV regularizer, in the sense that they promote sparsity of the Hessian Schatten norms instead of the gradient norms. Their main feature is that they favor piecewise-affine solutions - as opposed to TV which favors piecewise-constant solutions - and thus avoid the staircase effect, a common artifact of TV-based reconstructions. Let $L = [D_{hh}^\mathsf{T}, D_{vv}^\mathsf{T}, D_{hv}^\mathsf{T}]^\mathsf{T}$ with $D_{hh}, D_{vv}, D_{hv} \in \mathbb{R}^{n\times n}$ representing finite difference approximations of second-order derivatives along horizontal, vertical and mixed horizontal/vertical directions, respectively. Then the discrete Hessian Schatten p-norm regularizer is defined by

$$\mathcal{R}(x) = \|G(Lx)\|_1 = \sum_{i=1}^{n} |g_i(Lx)| \quad \text{with} \quad g_i(Lx) = \left\| \begin{bmatrix} (D_{hh}x)_i & (D_{hv}x)_i \\ (D_{vh}x)_i & (D_{vv}x)_i \end{bmatrix} \right\|_{\mathcal{S}_p}, \tag{11}$$

where, commonly, $p \in (0, 1]$ is used [24]. We recall that the Schatten p-norm reduces to the nuclear norm when $p = 1$.

The nuclear norm is also used in the general context of matrix completion (MC) as a convex approximation of the matrix rank operator. In image processing, such an approximation is used as

a convex regularizer aimed to enforce the low-rank property of a $d_1 \times d_2$ image (or patches of an image) regarded as a matrix. More precisely, the regularization term takes the form

$$\mathcal{R}(x) = \|G(x)\|_1 = \sum_{i=1}^{\bar{d}} |g_i(x)| \quad \text{with} \quad g_i(x) = \sigma_i(x)\,, \tag{12}$$

where L in (12) is clearly the identity matrix, $\bar{d} = \min\{d_1, d_2\}$ and x in the right-most expression of (12) must be regarded as a matrix.

For the above considered popular regularization functions, the external sparsity-promoting penalty Φ - see the definition of $\mathcal{R}(x)$ in (2) - is the ℓ_1 norm. However, the procedure proposed in (7) for constructing non-convex non-separable regularizers is not limited to the ℓ_1 norm. Indeed, the proposed procedure is applicable for any convex external penalty Φ. Alternative sparsity-inducing convex penalties are, e.g., the Huber penalty [8], the ℓ_1-ℓ_2 group-sparsity penalty [41], and the block- or joint-sparsity penalties [17]. For what concerns the matrix L, besides the identity or some differential operator, it can be the analysis operator of a dictionary - eventually a redundant dictionary, or a frame - in the context of the analysis sparsity priors. Popular examples of redundant dictionaries used in sparsity-based image processing are the oversampled Discrete Fourier Transform (DFT), the Gabor frames, the curvelet frames and the (undecimated) wavelet frames; see [6] and the references therein.

In the considered variational model (1)-(2), the above introduced regularizers are coupled with a quadratic fidelity term involving the measurement matrix A. Depending on A, model (1)-(2) can be applied to various image processing applications [8, 24], such as image deblurring when A is a blurring operator, image inpainting or compressed sensing based image reconstruction when A is a masking operator, and image super-resolution when A is the composition of a sub-sampling operator and a blurring operator.

3 Construction of non-convex non-separable regularizers

In this section, we propose a general strategy for constructing non-convex non-separable regularizers starting from any regularizer $\mathcal{R}$ of the form (2) under the following assumptions:

A1) $\mathcal{R}(\cdot) = \Phi(G(L\,\cdot))$ is convex and bounded from below by zero with $\mathcal{R}(0) = 0$;

A2) $\Phi(G(\cdot))$ is a proper, lower semicontinuous and coercive function.

In the following, first we introduce a function S_B which will then be involved in the construction of the proposed regularizer.

Definition 1. *Given a function $\mathcal{R}\colon \mathbb{R}^n \to \mathbb{R}$ as defined in (2), namely $\mathcal{R}(x) = \Phi(G(Lx))$, such that assumptions A1)–A2) are satisfied, and a matrix $B \in \mathbb{R}^{q\times n}$, we define the function $S_B\colon \mathbb{R}^n \to \mathbb{R}$ as*

$$S_B(x) := \inf_{v\in\mathbb{R}^n} \big\{F(v) := \mathcal{R}(v) + \tfrac{1}{2}\|B(x-v)\|_2^2\big\}. \tag{13}$$

In the notation of infimal convolution, we have

$$S_B(\,\cdot\,) = \mathcal{R}(\,\cdot\,) \,\square\, \tfrac{1}{2}\|B\,\cdot\|_2^2. \tag{14}$$

Note that if $C^{\mathsf{T}}C = B^{\mathsf{T}}B$, then $S_B(x) = S_C(x)$ for all $x \in \mathbb{R}^n$. That is, the function S_B depends only on $B^{\mathsf{T}}B$ and not B itself. Therefore, without loss of generality, we may assume B has full row

rank. In fact, if a given matrix B does not have full row rank, then there is another matrix C with full row rank such that $C^{\mathsf{T}}C = B^{\mathsf{T}}B$ which yields the same function S_B.

We recall the result in Lemma 2.7.1 in [11], where $\operatorname{null} M$ denotes the null space of matrix M.

Proposition 1. *Assume that the following mappings are given:*
i) Two matrices / linear mappings $L\colon \mathbb{R}^n \to \mathbb{R}^r$; $B\colon \mathbb{R}^n \to \mathbb{R}^q$ with

$$\operatorname{null} L \cap \operatorname{null} B \;=\; \{0\}. \tag{15}$$

ii) Two proper, lower semicontinuous and coercive mappings $f_1, f_2\colon \mathbb{R}^n \to [-\infty, +\infty]$. Then the mapping $h\colon \mathbb{R}^n \to [-\infty, +\infty]$, given by

$$h(v) = f_1(Lv) + f_2(Bv), \tag{16}$$

is lower semicontinuous and coercive. In particular, the mapping h attains its infimum $\inf h \in [-\infty, +\infty]$ at some point in $\mathbb{R}^n$.

Proposition 2. *Given two matrices L, B which satisfy assumption i) in Prop. 1, the infimal convolution (14) is exact, denoted by $\boxdot$, i.e.,*

$$S_B(x) \;=\; \mathcal{R}(\,\cdot\,) \boxdot \tfrac{1}{2}\|B\,\cdot\|_2^2 \;=\; \min_{v\in\mathbb{R}^n} \left\{\mathcal{R}(v) + \tfrac{1}{2}\|B(x-v)\|_2^2\right\}. \tag{17}$$

Proof. Let $f_1 := \Phi(G(\cdot))$ and $f_2 := \frac{1}{2}\|\,\cdot - c\|_2^2$ where c is a constant. They satisfy the assumption ii) in Prop. 1 due to assumption A2) on $\Phi(G(\cdot))$. Then $F(v)$ in (13) can be written as $F(v) = f_1(Lv) + f_2(Bv)$, and following Prop. 1, it is lower semicontinuous and coercive. Moreover, F is also convex, being the sum of convex functions. Thus, $F(v)$ takes its minimum at some point(s) in $\mathbb{R}^n$, that is, the infimal convolution is exact. □

Proposition 3. *For any matrix $B \in \mathbb{R}^{q\times n}$, the function S_B in (17) is a proper lower semicontinuous convex function.*

Proof. Set $f = \mathcal{R}(\,\cdot\,)$ and $g = \|B\,\cdot\|_2^2$. Both f and g are proper, convex and lower semicontinuous; hence $f \,\square\, g$ is convex and lower semicontinuous by Prop. 12.11 in [1]. Moreover, since the infimal convolution is exact by Prop. 2, the function S_B is proper. □

In the following, we highlight some properties of the function S_B in (17).

Proposition 4. *For any matrix $B \in \mathbb{R}^{q\times n}$, the function S_B defined in* (17) *satisfies*

$$0 \leqslant S_B(x) \leqslant \mathcal{R}(x), \quad \forall x \in \mathbb{R}^n. \tag{18}$$

The proof of Proposition 4 mirrors that of Proposition 8 in [31].

Proposition 5. *Let $B \in \mathbb{R}^{q\times n}$ be an arbitrary matrix and let $\alpha := \|B\|_2$. Then, the function S_B in* (17) *satisfies*

$$S_B(x) \leqslant S_{\alpha I}(x), \quad \forall x \in \mathbb{R}^n. \tag{19}$$

The proof of Proposition 5 mirrors that of Proposition 10 in [31].

The function S_B is closely related to the Moreau envelope of the function $\mathcal{R}$. The Moreau envelope [1] of $\mathcal{R}\colon \mathbb{R}^n \to \mathbb{R}$ is given by

$$\mathcal{R}^{\mathsf{M}}(x) = \inf_{v\in\mathbb{R}^n} \left\{\mathcal{R}(v) + \tfrac{1}{2}\|x-v\|_2^2\right\} = \left(\mathcal{R} \,\square\, \tfrac{1}{2}\|\,\cdot\,\|_2^2\right)(x). \tag{20}$$

The exact relationship of S_B to the Moreau envelope is described by the following two results.

Lemma 1. *For any invertible matrix $B \in \mathbb{R}^{n\times n}$, the function S_B in (17) can be expressed in terms of a Moreau envelope as*

$$S_B = \big(\mathcal{R}(\,\cdot\,) \circ B^{-1}\big)^{\mathsf{M}} \circ B. \tag{21}$$

Proof. Using (13), we have

$$\begin{aligned} S_B &= \mathcal{R}(\,\cdot\,) \,\square\, \big(\tfrac{1}{2}\|\cdot\|_2^2 \circ B\big) \\ &= \Big(\mathcal{R}(\,\cdot\,) \,\square\, \big(\tfrac{1}{2}\|\cdot\|_2^2 \circ B\big)\Big) \circ B^{-1} \circ B \\ &= \Big(\big(\mathcal{R}(\,\cdot\,) \circ B^{-1}\big) \,\square\, \big(\tfrac{1}{2}\|\cdot\|_2^2\big)\Big) \circ B \\ &= \big(\mathcal{R}(\,\cdot\,) \circ B^{-1}\big)^{\mathsf{M}} \circ B. \end{aligned}$$

□

Lemma 2. *For any full row rank matrix $B \in \mathbb{R}^{q\times n}$, the function S_B in (17) can be expressed in terms of a Moreau envelope as*

$$S_B = \big(d \circ B^{+}\big)^{\mathsf{M}} \circ B \tag{22}$$

where $d\colon \mathbb{R}^n \to \mathbb{R}$ is the convex function

$$d(x) = \min_{w\in \operatorname{null} B} \mathcal{R}(x-w)\,. \tag{23}$$

Proof. Using (17), we have

$$\begin{aligned} S_B(x) &= \min_{v\in\mathbb{R}^n} \big\{\mathcal{R}(v) + \tfrac{1}{2}\|B(x-v)\|_2^2\big\} \\ &= f(Bx) \end{aligned}$$

where $f\colon \mathbb{R}^q \to \mathbb{R}$ is given by

$$\begin{aligned} f(z) &= \min_{v\in\mathbb{R}^n} \big\{\mathcal{R}(v) + \tfrac{1}{2}\|z - Bv\|_2^2\,\big\} \\ &= \min_{u\in(\operatorname{null} B)^\perp}\ \min_{w\in\operatorname{null} B} \big\{\mathcal{R}(u+w) + \tfrac{1}{2}\|z - B(u+w)\|_2^2\,\big\} \\ &= \min_{u\in(\operatorname{null} B)^\perp}\ \min_{w\in\operatorname{null} B} \big\{\mathcal{R}(u+w) + \tfrac{1}{2}\|z - Bu\|_2^2\,\big\} \\ &= \min_{u\in(\operatorname{null} B)^\perp} \big\{d(u) + \tfrac{1}{2}\|z - Bu\|_2^2\,\big\} \end{aligned}$$

where d is the convex function defined in (23). The fact that d is convex follows from Prop. 8.26 of [1] and Example 3.16 of [3], Section 3.2.5. Since $(\operatorname{null} B)^\perp = \operatorname{range} B^{\mathsf{T}}$,

$$\begin{aligned} f(z) &= \min_{u\in\operatorname{range} B^{\mathsf{T}}} \big\{d(u) + \tfrac{1}{2}\|z - Bu\|_2^2\big\} \\ &= \min_{v\in\mathbb{R}^q} \big\{d(B^{\mathsf{T}}v) + \tfrac{1}{2}\|z - BB^{\mathsf{T}}v\|_2^2\big\} \\ &= \min_{v\in\mathbb{R}^q} \big\{d(B^{\mathsf{T}}(BB^{\mathsf{T}})^{-1}v) + \tfrac{1}{2}\|z - BB^{\mathsf{T}}(BB^{\mathsf{T}})^{-1}v\|_2^2\big\} \\ &= \min_{v\in\mathbb{R}^q} \big\{d(B^{+}v) + \tfrac{1}{2}\|z - v\|_2^2\big\} \\ &= \big(d(B^{+}\,\cdot)\big)^{\mathsf{M}}(z). \end{aligned}$$

Hence, $S_B(x) = \big(d(B^{+}\,\cdot)\big)^{\mathsf{M}}(Bx)$ which completes the proof. □

Note that (22) reduces to (21) when B is invertible. (Suppose B is invertible. Then $\operatorname{null} B = \{0\}$; hence $d(x) = \mathcal{R}(x)$ in (23). Additionally, $B^+ = B^{-1}$.)

Proposition 6. *For any full row rank matrix $B \in \mathbb{R}^{q\times n}$, the function S_B in (17) is differentiable.*

The proof of Proposition 6 mirrors that of Proposition 11 in [31].

The following result gives an expression for the gradient of S_B. When B is the identity operator, we recover a classic result (Proposition 12.29 in Ref. [1]).

Lemma 3. *The gradient of S_B in (17) is given by*

$$\nabla S_B(x) = B^{\mathsf{T}} B\Big(x - \arg\min_{v\in\mathbb{R}^n}\Big\{\frac{1}{2}\|B(x-v)\|_2^2 + \mathcal{R}(v)\Big\}\Big). \tag{24}$$

Proof. Using Lemma 2 we write S_B as

$$S_B = g \circ B \tag{25}$$

where

$$g = \big(d \circ B^+\big)^{\mathsf{M}} \tag{26}$$

where d is given by (23). Hence, by the chain rule,

$$\nabla S_B = B^{\mathsf{T}} \circ \nabla g \circ B, \tag{27}$$

i.e., $\nabla S_B(x) = B^{\mathsf{T}} \nabla g(Bx)$. Next, we write g as

$$g = h^{\mathsf{M}} \tag{28}$$

where

$$h = d \circ B^+, \tag{29}$$

i.e.,

$$h(x) = \min_{w\in\operatorname{null} B} \mathcal{R}(B^+x - w) \;= \min_{w\in\operatorname{null} B} \mathcal{R}(B^{\mathsf{T}}(BB^{\mathsf{T}})^{-1}x - w)\,, \tag{30}$$

where we assume B has full row-rank.

From Proposition 12.29 in Ref. [1], we have

$$\nabla g(x) = \nabla h^{\mathsf{M}}(x) = x - \operatorname{prox}_h(x) \tag{31}$$

where prox_h is given by

$$\begin{aligned}
\operatorname{prox}_h(x) &= \arg\min_{y\in\mathbb{R}^q}\big\{h(y) + \tfrac{1}{2}\|x-y\|_2^2\big\} && (32)\\
&= \arg\min_{y\in\mathbb{R}^q}\big\{\min_{w\in\operatorname{null} B} \mathcal{R}(B^{\mathsf{T}}(BB^{\mathsf{T}})^{-1}y - w) + \tfrac{1}{2}\|x-y\|_2^2\big\} && (33)\\
&= BB^{\mathsf{T}} \arg\min_{u\in\mathbb{R}^q}\big\{\min_{w\in\operatorname{null} B} \mathcal{R}(B^{\mathsf{T}}u - w) + \tfrac{1}{2}\|x - BB^{\mathsf{T}}u\|_2^2\big\} && (34)\\
&= B \arg\min_{z\in\operatorname{range} B^{\mathsf{T}}}\big\{\min_{w\in\operatorname{null} B} \mathcal{R}(z-w) + \tfrac{1}{2}\|x-Bz\|_2^2\big\} && (35)\\
&= B \arg\min_{z\in(\operatorname{null} B)^{\perp}}\big\{\min_{w\in\operatorname{null} B} \mathcal{R}(z-w) + \tfrac{1}{2}\|x-Bz\|_2^2\big\} && (36)\\
&= B \arg\min_{v\in\mathbb{R}^n}\big\{\mathcal{R}(v) + \tfrac{1}{2}\|x-Bv\|_2^2\big\}. && (37)
\end{aligned}$$

In the last line we used the following: if $v = z - w$ with $z \in (\operatorname{null} B)^{\perp}$ and $w \in \operatorname{null} B$, then $Bv = Bz - Bw = Bz$.

Using (37) in (31), we have

$$\nabla g(x) = x - B \arg\min_{v\in\mathbb{R}^n} \big\{\mathcal{R}(v) + \tfrac{1}{2}\|x - Bv\|_2^2\big\}. \tag{38}$$

Using (38) in (27), we have

$$\nabla S_B(x) = B^{\mathsf{T}} \nabla g(Bx) \tag{39}$$

$$= B^{\mathsf{T}} \Big(Bx - B \arg\min_{v\in\mathbb{R}^n} \big\{\mathcal{R}(v) + \tfrac{1}{2}\|Bx - Bv\|_2^2\big\}\Big) \tag{40}$$

which completes the proof. □

Definition 2. *Given a function $\mathcal{R}\colon \mathbb{R}^n \to \mathbb{R}$ as defined in* (2) *and satisfying assumptions A1)–A2), and a matrix $B \in \mathbb{R}^{q\times n}$, we define the non-separable penalty function $\mathcal{R}_B\colon \mathbb{R}^n \to \mathbb{R}$ as*

$$\mathcal{R}_B(x) := \mathcal{R}(x) - S_B(x) = \mathcal{R}(x) - \big(\mathcal{R}(\,\cdot\,) \boxdot \tfrac{1}{2}\|B\,\cdot\|_2^2\big)\,(x)\,, \tag{41}$$

where the second equality comes from the definition of S_B in (17).

The following corollary follows directly from Proposition 4 and Definition 2.

Corollary 1. *For any matrix $B \in \mathbb{R}^{q\times n}$, the penalty function $\mathcal{R}_B$ defined in (41) satisfies*

$$0 \leqslant \mathcal{R}_B(x) \leqslant \mathcal{R}(x) \quad \textit{for all } x \in \mathbb{R}^n. \tag{42}$$

As mentioned in the introduction, our proposal is to replace the original linear least squares model (1) with the following one:

$$x^* \in \arg\min_{x\in\mathbb{R}^n} \mathcal{J}_B(x), \qquad \mathcal{J}_B(x) := \frac{1}{2}\,\|Ax - b\|_2^2 \,+ \lambda\,\mathcal{R}_B(x), \tag{43}$$

where $\mathcal{R}_B(x)$ is the regularization function (41). We refer to (43) as the Convex-NonConvex Non-Separable least-squares model, abbreviated CNC-NS-ℓ_2.

4 Characterizing the properties of the regularizer

In this section, we carry out a theoretical analysis of the proposed regularizer $\mathcal{R}_B$ in (41), with focus on its sparsity-promoting properties. To this aim, we consider the special case of a regularizer $\mathcal{R} = \Phi(G(Lx))$ with $\Phi = \|\,\cdot\,\|_1$ and G the identity operator, that is

$$\mathcal{R}(x) = \|Lx\|_1 = \sum_{i=1}^{r} |L_i x| = \sum_{i=1}^{r} f_i(x)\,, \tag{44}$$

where L_i is the i-th row of matrix L and $f_i(x) := |L_i x|$. According to Definition 2, the proposed regularizer $\mathcal{R}_B$ associated with $\mathcal{R}$ above reads as follows:

$$R_B(x) = \mathcal{R}(x) - S_B(x) = \|Lx\|_1 - \min_{v\in\mathbb{R}^n} \Big\{ \frac{1}{2}\|B(x - v)\|_2^2 + \|Lv\|_1 \Big\}. \tag{45}$$

In Lemma 4 and Corollary 2 below, we analyze the gradient of the function S_B which, we recall, is differentiable according to Proposition 6. In the following, we use the set-valued signum function, defined as

$$\operatorname{sign}(t) := \begin{cases} \{1\}, & t > 0 \\ [-1,\ 1], & t = 0 \\ \{-1\}, & t < 0 \end{cases}$$

for $t \in \mathbb{R}$. For $x \in \mathbb{R}^n$, the signum function is defined in a component-wise fashion.

Lemma 4. *The gradient of S_B satisfies*

$$\nabla S_B(x) \in L^{\mathsf{T}}\text{sign}(Lv) \quad \textit{for some } v \in \mathbb{R}^n. \tag{46}$$

Proof. Using (24), we write

$$\nabla S_B(x) = B^{\mathsf{T}}B(x - v^{\mathsf{opt}}) \tag{47}$$

where

$$v^{\mathsf{opt}} \in \arg\min_{v\in\mathbb{R}^n} \big\{ \tfrac{1}{2}\|B(x-v)\|_2^2 + \|Lv\|_1 \big\}. \tag{48}$$

Since any minimizer v^{opt} of a convex function S_B satisfies $0 \in \partial S_B(v^{\mathsf{opt}})$, we have

$$0 \in \partial\big[\tfrac{1}{2}\|B(x - \,\cdot\,)\|_2^2 + \|L\cdot\|_1\big](v^{\mathsf{opt}}). \tag{49}$$

That is

$$0 \in B^{\mathsf{T}}B(v^{\mathsf{opt}} - x) + L^{\mathsf{T}}\{u : u \in \text{sign}(Lv^{\mathsf{opt}})\} \tag{50}$$

or

$$B^{\mathsf{T}}B(x - v^{\mathsf{opt}}) \in L^{\mathsf{T}}\{u : u \in \text{sign}(Lv^{\mathsf{opt}})\}. \tag{51}$$

In light of (47), it follows that

$$\nabla S_B(x) \in L^{\mathsf{T}}\{u : u \in \text{sign}(Lv^{\mathsf{opt}})\} \tag{52}$$

where v^{opt} is defined in (48). □

Corollary 2. *The gradient of S_B satisfies*

$$\nabla S_B(x) = L^{\mathsf{T}}u \tag{53}$$

for some $u \in \mathbb{R}^r$ with $|u_i| \leqslant 1$ for $i = 1, \ldots, r$.

We now aim to show that changes in x that cause $f_i(x)$ to increase also causes R in (44) to increase. Since x is a vector, we use directional derivatives. Actually, since our functions are not differentiable, we use directional subgradients. The directional subgradient of a convex function f at x in direction v is denoted $\partial_v f(x)$. If f is differentiable at x in direction v, then $\partial_v f(x)$ is single-valued. If f is differentiable at x, then $\partial_v f(x) = [\nabla f(x)]^{\mathsf{T}}v$. We have

$$\partial_v R(x) = [\partial R(x)]^{\mathsf{T}}v = [L^{\mathsf{T}}\text{sign}(Lx)]^{\mathsf{T}}v = \text{sign}(Lx)^{\mathsf{T}}Lv \tag{54}$$

and

$$\partial_v f_i(x) = \text{sign}(L_i x)^{\mathsf{T}}L_i v = \text{sign}(L_i x)L_i v \tag{55}$$

where we can drop the transpose because $L_i x$ is scalar. In the following proposition, we show that for any direction v in which all f_i are increasing at x, the total regularizer R_B is increasing too.

Proposition 7. *Let $x \in \mathbb{R}^n$ and $v \in \mathbb{R}^n$. Suppose $\partial_v f_i(x) \geqslant 0$ for $i = 1, \ldots, r$. Then $\partial_v R_B(x) \geqslant 0$ also.*

Proof. It is given that $\partial_v f_i(x) \geqslant 0$ for $i = 1, \ldots, r$. That is,

$$\text{sign}(L_i x)L_i v \geqslant 0, \quad i = 1, \ldots, r \tag{56}$$

where the left-hand side of (56) is a set. Note that if $L_i x = 0$, then $\text{sign}(L_i x)$ is the interval $[-1, 1]$. Hence, if $L_i x = 0$, then $\text{sign}(L_i x) L_i v \geqslant 0$ (which is given) implies that $L_i v = 0$. Hence, it follows from (56) that

$$L_i x = 0 \implies L_i v = 0 \tag{57}$$
$$L_i x > 0 \implies L_i v \geqslant 0 \tag{58}$$
$$L_i x < 0 \implies L_i v \leqslant 0. \tag{59}$$

We now consider the function $R_B(x) = R(x) - S_B(x)$. The directional derivative of R_B at x in direction v is as follows

$$\partial_v R_B(x) = \partial_v R(x) - \partial_v S_B(x). \tag{60}$$

Using (53), the directional derivative of S_B at x in direction v is given by

$$\partial_v S_B(x) = [\nabla S_B(x)]^{\mathsf{T}} v = [L^{\mathsf{T}} u]^{\mathsf{T}} v = u^{\mathsf{T}} L v \tag{61}$$

for some u with $|u_i| \leqslant 1$. Thus,

$$\begin{aligned}
\partial_v R_B(x) &= \partial_v R(x) - \partial_v S_B(x) && (62)\\
&= \text{sign}(Lx)^{\mathsf{T}} L v - u^{\mathsf{T}} L v && (63)\\
&= [\text{sign}(Lx) - u]^{\mathsf{T}} L v && (64)\\
&= \sum_i [\text{sign}(L_i x) - u_i] L_i v && (65)\\
&= \sum_{i,\ L_i x \neq 0} [\text{sign}(L_i x) - u_i] L_i v. && (66)
\end{aligned}$$

We can omit the terms for which $L_i x = 0$ because for these terms we have $L_i v = 0$ as noted after (56). Note that if $L_i x > 0$, then $\text{sign}(L_i x) - u_i \geqslant 0$ because $|u_i| \leqslant 1$. Similarly, if $L_i x < 0$, then $\text{sign}(L_i x) - u_i \leqslant 0$. Hence, using (58) and (59), it follows that each term in the sum (66) is non-negative. Hence, $\partial_v R_B(x) \geqslant 0$. □

The result in Prop. 7 is a generalization of the result for the Generalized MC penalty in Prop.13 in [31].

In the following simple example, we provide some visual insights on the proposed regularizer $\mathcal{R}_B(x)$ in (41).

Example 1. By the way of illustration, to show the properties/effects of the proposed non-separable non-convex regularizer $\mathcal{R}_B(x)$ in (41), we consider the two bivariate models defined by

$$\arg\min_{x \in \mathbb{R}^2} \mathcal{J}(x), \qquad \mathcal{J}(x) := \frac{1}{2} \|Ax - b\|_2^2 + \lambda \mathcal{R}(x), \tag{67}$$

$$\arg\min_{x \in \mathbb{R}^2} \mathcal{J}_{\mathcal{B}}(x), \qquad \mathcal{J}_{\mathcal{B}}(x) := \frac{1}{2} \|Ax - b\|_2^2 + \lambda \mathcal{R}_{\mathcal{B}}(x), \tag{68}$$

with

$$\mathcal{R}(x) = \|Lx\|_1, \qquad \mathcal{R}_B(x) = \mathcal{R}(x) - S_B(x) = \|Lx\|_1 - \left(\|L \cdot\|_1 \boxdot \tfrac{1}{2}\|B \cdot\|_2^2\right)(x) \tag{69}$$

and

$$\lambda = 1.5, \quad b = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \quad A = \begin{bmatrix} 1 & 2 \\ -1 & 0.5 \end{bmatrix}, \quad L = \begin{bmatrix} -0.5 & 1 \\ 2 & -0.5 \end{bmatrix}, \quad B = \sqrt{\frac{\gamma}{\lambda}}\, A,\ \gamma = 0.99\,. \tag{70}$$

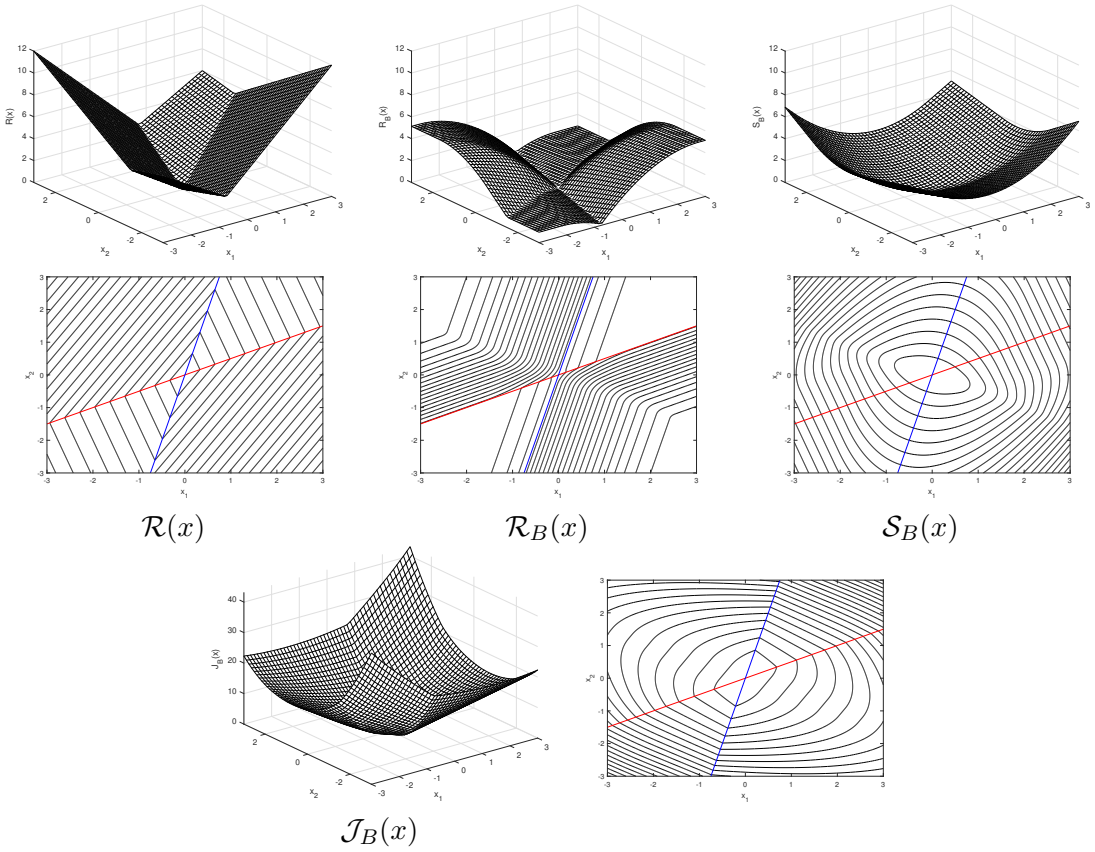

Figure 1: Example 1: plots of functions $\mathcal{R}(x)$, $\mathcal{R}_B(x)$, $\mathcal{S}_B(x)$ (first row) and of the associated contour plots (second row); total cost function $\mathcal{J}_B(x)$ and associated contour plot (third row).

As it will be demonstrated in Section 5, the definition of matrix B in (70) with $\gamma \in [0, 1]$ yields the total cost functional $\mathcal{J}_B$ in (68) to be convex in spite of the regularizer $\mathcal{R}_B$ in (69) being non-convex. In the first and second row of Fig. 1 we show the functions $\mathcal{R}(x)$, $\mathcal{R}_B(x)$, $\mathcal{S}_B(x)$ and their associated contour plots, respectively, whereas the total cost functional $\mathcal{J}_B$ together with its contour plot are reported in the third row. The solid red and blue lines in the contour plots represent the hyperplanes H_1 and H_2, respectively, with $H_i = \{x \in \mathbb{R}^2 \colon\ L_i x = 0\},\ i \in \{1, 2\}$ and L_i the i-th row of matrix L.

It can be noticed from Fig. 1 that the function $\mathcal{S}_B$ is convex (see Proposition 3) such that our regularizer $\mathcal{R}_B$, obtained by subtracting $\mathcal{S}_B$ from $\mathcal{R}$, is non-convex. In particular, the shape of the regularizer $\mathcal{R}_B$ (and of its contour lines) suggests that $\mathcal{R}_B$ holds the potential for promoting sparsity of the vector $Lx = (L_1 x, L_2 x)^{\mathsf{T}}$ more effectively than the regularizer $\mathcal{R}$. Finally, the plots in the last row of Fig. 1 confirm that the total cost functional $\mathcal{J}_B$ is convex.

5 Convexity Conditions

In this section, we investigate if and how the matrix B of parameters of the proposed regularizer $\mathcal{R}_B$ in (41) can be set such that the total objective function $\mathcal{J}_B$ in (43) is convex and then discuss existence and uniqueness of its minimizer(s), that is of the solution(s) of the proposed CNC-NS-ℓ_2 variational model in (43). In the following, $M \succcurlyeq 0\,(\succ 0)$ indicates that the matrix M is positive semidefinite (positive definite) and I_n denotes the identity matrix of order n. First, we rewrite $\mathcal{J}_B$

in (43) in the following equivalent form

$$
\begin{aligned}
\mathcal{J}_B(x) &= \tfrac{1}{2}\|Ax-b\|_2^2 + \lambda\big(\mathcal{R}(x) - S_B(x)\big) \\
&= \tfrac{1}{2}\|Ax-b\|_2^2 + \lambda\Big(\mathcal{R}(x) - \min_{v\in\mathbb{R}^n}\big\{\mathcal{R}(v) + \tfrac{1}{2}\|B(x-v)\|_2^2\big\}\Big) \\
&= \tfrac{1}{2}\|Ax-b\|_2^2 + \lambda\Big(\mathcal{R}(x) - \min_{v\in\mathbb{R}^n}\big\{\mathcal{R}(v) + \tfrac{1}{2}\|Bx\|_2^2 + \tfrac{1}{2}\|Bv\|_2^2 - v^{\mathsf{T}}B^{\mathsf{T}}Bx\big\}\Big) \\
&= \tfrac{1}{2}\|Ax-b\|_2^2 - \tfrac{\lambda}{2}\|Bx\|_2^2 + \lambda\Big(\mathcal{R}(x) - \min_{v\in\mathbb{R}^n}\big\{\mathcal{R}(v) + \tfrac{1}{2}\|Bv\|_2^2 - v^{\mathsf{T}}B^{\mathsf{T}}Bx\big\}\Big) \\
&\quad \overbrace{\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad}^{\mathcal{Q}(x)} \\
&= \tfrac{1}{2}\,x^{\mathsf{T}}\big(\overbrace{A^{\mathsf{T}}A - \lambda\,B^{\mathsf{T}}B}^{Q}\big)x - b^{\mathsf{T}}Ax + \tfrac{1}{2}\,\|b\|_2^2 \qquad (71) \\
&\qquad + \lambda\Big(\mathcal{R}(x) + \underbrace{\max_{v\in\mathbb{R}^n}\big\{\underbrace{-\mathcal{R}(v) - \tfrac{1}{2}\|Bv\|_2^2 + v^{\mathsf{T}}B^{\mathsf{T}}Bx}_{g(x,v)}\big\}}_{G(x)}\Big) \qquad (72) \\
&\qquad\qquad \underbrace{\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad}_{\mathcal{G}(x)}
\end{aligned}
$$

The following Lemma 5 and Definition 3 are useful for the analysis of convexity and coerciveness of our cost functional $\mathcal{J}_B$ carried out in Proposition 8 .

Lemma 5. *In case the (symmetric) matrix $Q \in \mathbb{R}^{n\times n}$ defined in (71) is positive semidefinite, such that* $\operatorname{null} Q \neq \{0\}$*, we have*

$$\operatorname{null} Q \supseteq \operatorname{null} A \cap \operatorname{null} B = \operatorname{null} A\,. \tag{73}$$

Definition 3. *Let $Z\colon \mathbb{R}^n \to \mathbb{R}$ be a (not necessarily smooth) function. Then, Z is said to be δ-strongly convex if and only if there exists a real constant $\delta > 0$, called the modulus of strong convexity of Z, such that the function $Z(x) - \frac{\delta}{2}\|x\|_2^2$ is convex [1].*

Proposition 8. *Let $\mathcal{R}\colon \mathbb{R}^n \to \mathbb{R}$ be a function having the form in* (2) *and satisfying assumptions A1)–A2), and let $\mathcal{R}_B\colon \mathbb{R}^n \to \mathbb{R}$ be the function defined in (41). Then, the functional $\mathcal{J}_B\colon \mathbb{R}^n \to \mathbb{R}$ in (43) is proper, lower semicontinuous, and bounded from below by zero. Moreover, a sufficient condition for $\mathcal{J}_B$ to be convex (strictly convex) is that the matrix of parameters $B \in \mathbb{R}^{q\times n}$ satisfies*

$$Q = A^{\mathsf{T}}A - \lambda\,B^{\mathsf{T}}B \succcurlyeq 0 \ \ (\succ 0)\,. \tag{74}$$

In particular, we have:

- $Q \succ 0 \quad \Longrightarrow \mathcal{J}_B(x)$ *is $\underline{e}_Q$–strongly convex* $\Longrightarrow$ *admits a unique minimizer,*
- $\begin{cases} Q \succcurlyeq 0 \\ \operatorname{null} A \cap \operatorname{null} L = \{0\} \end{cases} \Longrightarrow \mathcal{J}_B(x)$ *is coercive* $\Longrightarrow$ *admits minimizer(s),*

where $\underline{e}_Q > 0$ denotes the minimum eigenvalue of the real, symmetric, positive definite matrix Q.

Proof. The functional $\mathcal{J}_B$ in (43) is the sum of a nonnegative convex quadratic fidelity term - which is clearly proper and everywhere continuous - and the regularization term $\mathcal{R}_B$ defined in (41), which is bounded from below by zero according to Corollary 1 and proper lower semicontinuous. Hence, $\mathcal{J}_B$ is proper lower semicontinuous and bounded from below by zero. To derive conditions for convexity and coerciveness of $\mathcal{J}_B$, first we analyze separately the terms $\mathcal{Q}(x)$ and $\mathcal{G}(x)$ defined in (71) and (72), respectively.

Analysis of $\mathcal{Q}(x)$. After recalling the definition of the quadratic function $\mathcal{Q}(x)$ in (71), it is clear how the matrix $Q \in \mathbb{R}^{n\times n}$ in (74) represents the (constant) Hessian matrix of $\mathcal{Q}(x)$. It follows that $\mathcal{Q}(x)$ is convex (strictly convex) if and only if $Q \succcurlyeq 0$ ($Q \succ 0$). In particular, we notice that Q is a real symmetric matrix, hence it admits the eigenvalue decomposition

$$Q = V_Q E_Q V_Q^{\mathsf{T}}, \quad E_Q, V_Q \in \mathbb{R}^{n\times n}, \quad E_Q = \mathrm{diag}(e_{Q,1}, e_{Q,2}, \dots, e_{Q,n}), \quad V_Q^{\mathsf{T}} V_Q = V_Q V_Q^{\mathsf{T}} = I_n . \tag{75}$$

According to Definition 3, $\mathcal{Q}(x)$ is δ-strongly convex - with $\delta > 0$ - if and only if the quadratic function $\widetilde{\mathcal{Q}}(x) := \mathcal{Q}(x) - (\delta/2)\|x\|_2^2$ is convex, that is if and only if its Hessian matrix $\widetilde{Q} = Q - \delta\, I_n$ is positive semidefinite. Using (75), we can write:

$$\begin{aligned} &\widetilde{Q} = Q - \delta\, I_n \succcurlyeq 0 \iff V_Q E_Q V_Q^{\mathsf{T}} - \delta\, I_n \succcurlyeq 0 \iff V_Q\,(E_Q - \delta\, I_n)\, V_Q^{\mathsf{T}} \succcurlyeq 0 \\ \iff\ &E_Q - \delta\, I_n \succcurlyeq 0 \iff \mathrm{diag}(e_{Q,1} - \delta, \dots, e_{Q,n} - \delta) \succcurlyeq 0 \iff 0 < \delta \leqslant \min_i e_{Q,i} . \end{aligned} \tag{76}$$

It follows from (76) that, when the matrix Q in (74) is positive definite ($Q \succ 0$) and, hence, $\min_i e_{Q,i} > 0$, then the function $\mathcal{Q}(x)$ in (71) is δ-strongly convex with $\delta = \min_i e_{Q,i}$. Instead, when the matrix Q is only positive semidefinite ($Q \succcurlyeq 0$), then the function $\mathcal{Q}(x)$ in (71) is only convex and nothing can be said about its coerciveness or even about its lower boundedness.

Analysis of $\mathcal{G}(x)$. First, we notice that, for any given v, the function $g(x, v)$ in (72) is affine - hence, convex - in x. It follows that the function $G(x)$ in (72) is convex in x as it is the pointwise maximum of a set of convex functions (Proposition 8.14 in [1]). This implies that the function $\mathcal{G}(x)$ in (72) is convex, as $\mathcal{R}(x)$ is convex by assumption. Moreover, $\mathcal{G}(x)$ is bounded from below by zero as we can write

$$\mathcal{G}(x) = \mathcal{R}(x) + \max_{v\in\mathbb{R}^n} g(x, v) \geqslant \mathcal{R}(x) + g(x, v = 0) = \mathcal{R}(x) - \mathcal{R}(0) = \mathcal{R}(x) \geqslant 0 \quad \forall\, x \in \mathbb{R}^n, \tag{77}$$

where the last equality and inequality in (77) come from assumption A1) on the regularizer $\mathcal{R}(x)$.

Analysis of $\mathcal{J}_B(x)$. Recalling that, according to (72), we can write $\mathcal{J}_B(x) = \mathcal{Q}(x) + \lambda\,\mathcal{G}(x)$, in light of previous analysis we have that, when the matrix Q in (74) is positive definite ($Q \succ 0$), then our total cost functional $\mathcal{J}_B(x)$ is δ-strongly convex with $\delta = \min_i e_{Q,i} > 0$, hence it admits a unique minimizer. Instead, when the matrix Q is only positive semidefinite ($Q \succcurlyeq 0$, such that $\mathrm{null}\, Q \neq \{0\}$), then $\mathcal{J}_B(x)$ is only convex as it is the sum of two functions which are both only convex. Hence, the existence of minimizers of $\mathcal{J}_B(x)$ in this case is not guaranteed unless $\mathcal{J}_B(x)$ is coercive. To investigate coerciveness of $\mathcal{J}_B(x)$, we partition its domain $\mathbb{R}^n$ as follows:

$$\mathbb{R}^n = P_1 \cup P_2, \quad P_1 = \mathbb{R}^n \setminus \mathrm{null}\, A, \quad P_2 = \mathrm{null}\, A\,, \tag{78}$$

and we analyze coerciveness of the restrictions of $\mathcal{J}_B(x)$ to the disjoint sub-domains P_1 and P_2. Before considering the two cases separately, it is useful to derive the two following lower bounds for $\mathcal{J}_B(x)$:

$$\mathcal{J}_B(x) \ \geqslant\ \tfrac{1}{2}\|Ax - b\|_2^2 \quad \forall\, x \in \mathbb{R}^n, \tag{79}$$

$$\mathcal{J}_B(x) \ \geqslant\ \mathcal{Q}(x) + \lambda\,\mathcal{R}(x) = \tfrac{1}{2}\, x^{\mathsf{T}} Q x - b^{\mathsf{T}} A x + \tfrac{1}{2}\,\|b\|_2^2 + \lambda\,\mathcal{R}(x) \quad \forall\, x \in \mathbb{R}^n\,, \tag{80}$$

where (79) comes from (42)–(43) and (80) follows from (71)–(72) and (77). We also recall that, given an unbounded set $S \subseteq \mathbb{R}^n$ and two functions $f_1, f_2 : S \to \mathbb{R}$ such that $f_1(x) \geqslant f_2(x)\, \forall\, x \in S$, coerciveness of $f_2(x)$ (over S) implies coerciveness of $f_1(x)$ (over S).

Case $x \in P_1$. The restriction of the quadratic function $\frac{1}{2}\|Ax - b\|_2^2$ to P_1 defined in (78) is clearly strongly convex, hence coercive. It thus follows from inequality (79) that $\mathcal{J}_B(x)$ restricted to P_1 is coercive.

Case $x \in P_2$. According to statement (73) of Lemma 5, $x \in P_2 = \operatorname{null} A$ implies $x \in \operatorname{null} Q$. Hence, when restricted to P_2 inequality (80) reduces to

$$\mathcal{J}_B(x) \;\geqslant\; \tfrac{1}{2}\,\|b\|_2^2 + \lambda\,\mathcal{R}(x) \;=\; \tfrac{1}{2}\,\|b\|_2^2 + \lambda\,\Phi(G(Lx)) \quad \forall\, x \in P_2 = \operatorname{null} A\,, \tag{81}$$

where the second equality in (81) comes from definition of the regularizer $\mathcal{R}(x)$ in (2). In case that $(P_2 = \operatorname{null} A) \cap \operatorname{null} L = \{0\}$, then the right-hand side function in (81) is coercive over P_2 due to assumption A2) on the regularizer $\mathcal{R}(x)$. It follows from inequality (81) that $\mathcal{J}_B(x)$ is also coercive over P_2. This concludes the proof after recalling that a proper, l.s.c., convex and coercive function admits minimizer(s). □

In the next subsection, we discuss some simple strategies for choosing the matrix of parameters B among all those satisfying the convexity condition in (74).

5.1 Setting the parameter matrix B

The convexity condition in (74) sets a constraint on $B^{\mathsf{T}}B$, hence on the matrix of parameters B. However, infinitely many $B^{\mathsf{T}}B$ matrices – and hence B matrices – satisfy this constraint; therefore, some guidelines must be given. In the following, we illustrate two simple strategies for choosing B.

The first and simplest strategy consists in setting $B = \sqrt{\gamma/\lambda}\,A$, that is

$$B^{\mathsf{T}}B = (\gamma/\lambda)A^{\mathsf{T}}A\,, \quad \gamma \in\, [0,1]\,. \tag{82}$$

It is immediate to verify that this choice of B fulfills condition (74) for convexity of our cost functional $\mathcal{J}_B$. Moreover, it is intuitive that the scalar parameter γ controls the degree of non-convexity of the regularization term $\mathcal{R}_B$ in $\mathcal{J}_B$, namely the greater γ the more non-convex $\mathcal{R}_B$. If $\gamma = 0$, then $B = 0$ and our penalty $\mathcal{R}_B$ reduces to the convex regularizer $\mathcal{R}$ in (2) under assumptions A1) and A2). If $\gamma = 1$, then (74) is satisfied with equality and our regularizer $\mathcal{R}_B$ is 'maximally' non-convex under the constraint that $\mathcal{J}_B$ is convex.

We also propose the following second, more sophisticated strategy for constructing a matrix $B^{\mathsf{T}}B \in \mathbb{R}^{n\times n}$ satisfying the convexity condition in (74). Since the matrix $A^{\mathsf{T}}A \in \mathbb{R}^{n\times n}$ is symmetric and positive semidefinite, it admits the eigenvalue decomposition

$$A^{\mathsf{T}}A = VEV^{\mathsf{T}}, \quad E, V \in \mathbb{R}^{n\times n}, \quad E = \operatorname{diag}(e_1, e_2, \ldots, e_n)\,, \quad V^{\mathsf{T}}V = VV^{\mathsf{T}} = I_n\,, \tag{83}$$

with e_i, $i = 1, \ldots, n$, indicating the real non-negative eigenvalues of $A^{\mathsf{T}}A$. We set

$$B^{\mathsf{T}}B = \frac{1}{\lambda}\,V\,\Gamma E V^{\mathsf{T}}, \quad \Gamma := \operatorname{diag}(\gamma_1, \gamma_2, \ldots, \gamma_n)\,, \quad \gamma_i \in\, [0,1]\;\forall\, i \in \{1, 2, \ldots, n\}\,, \tag{84}$$

such that, replacing (84) into convexity condition (74), we have

$$Q = V\,(E - \Gamma E)\,V^{\mathsf{T}} \;\succcurlyeq\; 0 \;\; (\succ\; 0) \iff E\,(I_n - \Gamma) \;\succcurlyeq\; 0 \;\; (\succ\; 0)\,, \tag{85}$$

which is clearly satisfied given the definition of matrix Γ in (84). We notice that, when one chooses $\gamma_1 = \gamma_2 = \cdots = \gamma_n = \gamma \in [0,1]$, then (84) reduces to (82), that is strategy (82) is included in the more general strategy (84). A suitable choice of the γ_i will be discussed in the experimental section for the case of image restoration.

In the following proposition we adapt the results in Proposition 8 to the special case of a matrix $B^{\mathsf{T}}B$ chosen according to the strategy illustrated in (83)–(84).

Proposition 9. *Let the matrix* $B^{\mathsf{T}}B \in \mathbb{R}^{n\times n}$ *be set as in (83)–(84). Then, we have*

- $\begin{cases} \text{null}\, A = \{0\} \\ \gamma_i \neq 1 \;\forall\, i \in \{1, \ldots, n\} \end{cases} \implies \mathcal{J}_B(x)$ *is* $\underline{e}_Q$*–strongly convex* $\implies$ *admits a unique minimizer,*

- $\begin{cases} \text{null}\, A \neq \{0\} \vee \exists\, i \in \{1, \ldots, n\}\text{:}\ \gamma_i = 1 \\ \text{null}\, A \cap \text{null}\, L = \{0\} \end{cases} \implies \mathcal{J}_B(x)$ *is coercive* $\implies$ *admits minimizer(s),*

where

$$\underline{e}_Q \;=\; \min_{i=1,\ldots,n} \{e_i(1-\gamma_i)\}\;. \tag{86}$$

Proof. The proof of this proposition comes in a straightforward manner from the results in Proposition 8 once the two conditions $Q \succ 0$ and $Q \succcurlyeq 0$ are made explicit for the considered case of a matrix $B^{\mathsf{T}}B$ chosen as in (83)–(84). Starting from (85) and using (83)–(84), we directly obtain the eigenvalue decomposition of the real symmetric matrix Q in (74):

$$Q \;=\; V\; \text{diag}\big(\; \underbrace{e_1(1-\gamma_1)}_{e_{Q,1}},\, \underbrace{e_2(1-\gamma_2)}_{e_{Q,2}},\, \ldots\,,\, \underbrace{e_n(1-\gamma_n)}_{e_{Q,n}} \;\big)\; V^{\mathsf{T}}\,, \tag{87}$$

where $e_{Q,i}$, $i = 1, \ldots, n$, represent the eigenvalues of Q and, hence, $\underline{e}_Q$ in (86) represents the smallest eigenvalue of Q. The matrix Q is positive definite ($Q \succ 0$) if and only if all its eigenvalues are positive that is - see (87) - if and only if $e_i > 0$ and $\gamma_i \neq 1$ for any $i \in \{1, \ldots, n\}$. Since e_i, $i = 1, \ldots, n$, denote the eigenvalues of matrix $A^{\mathsf{T}}A$, the condition $e_i > 0\; \forall\, i \in \{1, \ldots, n\}$ is equivalent to $\text{null}\, A^{\mathsf{T}}A = \text{null}\, A = \{0\}$. The matrix Q is only positive semidefinite ($Q \succcurlyeq 0$) if and only if at least one of its eigenvalues is null that is - see (87) - if and only if there exists $i \in \{1, \ldots . n\}$ such that $e_i = 0$ or $\gamma_i = 1$. The condition $\exists\, i \in \{1, \ldots, n\}\text{:}\ e_i = 0$ is equivalent to $\text{null}\, A^{\mathsf{T}}A = \text{null}\, A \neq \{0\}$. □

6 Optimization Algorithm

Even though the proposed class of non-separable regularization functions $\mathcal{R}_B$ in (41) does not have a simple explicit formula, a global minimizer of the sparse-regularized cost function $\mathcal{J}_B$ in (43) can be readily calculated using proximal algorithms.

To minimize the cost function $\mathcal{J}_B$ in (43) using proximal algorithms, we rewrite it as a saddle-point problem:

$$(x^*,\, v^*) = \arg \min_{x \in \mathbb{R}^n} \max_{v \in \mathbb{R}^n} \mathcal{F}(x, v) \tag{88}$$

where

$$\mathcal{F}(x, v) = \frac{1}{2}\, \|Ax - b\|_2^2 + \lambda\, \mathcal{R}(x) - \lambda\, \mathcal{R}(v) - \frac{\lambda}{2} \|B(x - v)\|_2^2. \tag{89}$$

The solution of the saddle-point problem can be calculated using a general form of the forward-backward (FB) algorithm (Theorem 25.8 in Ref. [1]). This form of the FB algorithm is formulated to solve monotone inclusion problems, of which the saddle-point problem (89) is a special case. The FB algorithm involves operators A, A^{T}, B, and B^{T}, and the proximity operator of $\mathcal{R}$.

Proposition 10. *Let $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, $\lambda > 0$, and let $\mathcal{R}_B \colon \mathbb{R}^n \to \mathbb{R}$ be the regularization function in (41), with the matrix $B^{\mathsf{T}}B \in \mathbb{R}^{n \times n}$ set as in (83)–(84). Then the saddle-point $(x^*,\, v^*)$ of the function $\mathcal{F}$ in (89) can be obtained as the limit point of the sequence of iterates $\{x_k, v_k\}_{k=1}^{\infty}$ generated by the following forward-backward iterative algorithm:*

$$\textbf{\textit{set}} \quad \rho = \max_k \left\{ \frac{1 - 2\gamma_k + 2\gamma_k^2}{1 - \gamma_k}\; e_k \right\}$$

set $\mu \in \,]\,0\,,\,2/\rho\,[$

for $k = 0, 1, 2, \ldots$

$$w_k = x_k - \mu\left[A^{\mathsf{T}}(Ax_k - b) + \lambda\, B^{\mathsf{T}}B\,(v_k - x_k)\right]$$
$$u_k = v_k - \mu\lambda B^{\mathsf{T}}B(v_k - x_k)$$
$$x_{k+1} = \arg\min_{x\in\mathbb{R}^n}\left\{\mathcal{R}(x) + \frac{1}{2\mu\lambda}\|x - w_k\|_2^2\right\}$$
$$v_{k+1} = \arg\min_{v\in\mathbb{R}^n}\left\{\mathcal{R}(v) + \frac{1}{2\mu\lambda}\|v - u_k\|_2^2\right\}$$

end

where e_k and γ_k are defined in (83)–(84), and k is the iteration counter.

Proof. The point (x^*, v^*) is a saddle-point of $\mathcal{F}$ in (89) if $0 \in \partial\mathcal{F}(x^*, v^*)$ where $\partial\mathcal{F}$ is the subdifferential of $\mathcal{F}$. From (89), we have

$$\partial_x\mathcal{F}(x, v) = A^{\mathsf{T}}(Ax - b) - \lambda B^{\mathsf{T}}B(x - v) + \lambda\,\partial\mathcal{R}(x),$$
$$\partial_v\mathcal{F}(x, v) = \lambda B^{\mathsf{T}}B(x - v) - \lambda\,\partial\mathcal{R}(v).$$

Hence, $0 \in \partial\mathcal{F}(x, v)$ if $0 \in \mathcal{P}(x, v) + \mathcal{Q}(x, v)$ where

$$\mathcal{P}(x, v) = \underbrace{\begin{bmatrix} A^{\mathsf{T}}A - \lambda\, B^{\mathsf{T}}B & \lambda\, B^{\mathsf{T}}B \\ -\lambda\, B^{\mathsf{T}}B & \lambda\, B^{\mathsf{T}}B \end{bmatrix}}_{P}\begin{bmatrix} x \\ v \end{bmatrix} - \begin{bmatrix} A^{\mathsf{T}}b \\ 0 \end{bmatrix},$$
$$\mathcal{Q}(x, v) = \begin{bmatrix} \lambda\,\partial\mathcal{R}(x) \\ \lambda\,\partial\mathcal{R}(v) \end{bmatrix}.$$

The operators $\mathcal{P}$ and $\mathcal{Q}$ are maximally monotone, since they represent the subdifferentials of proper convex l.s.c. functions. Hence, the pair (x, v) such that $0 \in \mathcal{P}(x, v) + \mathcal{Q}(x, v)$ is a zero of a sum of maximally monotone operators. Since $\mathcal{P}$ is single-valued and β-cocoercive with $\beta > 0$, the pair (x, v) can be calculated via the forward-backward algorithm (Theorem 25.8 in [1]). In the current notation, the forward-backward algorithm is

$$\begin{bmatrix} w_k \\ u_k \end{bmatrix} = \begin{bmatrix} x_k \\ v_k \end{bmatrix} - \mu\mathcal{P}(x_k,\, v_k)$$
$$\begin{bmatrix} x_{k+1} \\ v_{k+1} \end{bmatrix} = (I + \mu\mathcal{Q})^{-1}(w_k,\, u_k)\,.$$

The constant μ should be chosen so that $0 < \mu < 2\beta$ where $\mathcal{P}$ is β-cocoercive (Definition 4.4 in [1]), i.e., $\beta\mathcal{P}$ is firmly non-expansive. We now address the value β. By Corollary 4.3(v) in [1], this condition is equivalent to

$$\underbrace{\tfrac{1}{2}P + \tfrac{1}{2}P^{\mathsf{T}} - \beta P^{\mathsf{T}}P}_{M} \succcurlyeq 0\,, \tag{90}$$

with matrix $P \in \mathbb{R}^{2n\times 2n}$ defined above. Recalling (83)-(84), we can rewrite the previously defined matrix P as follows:

$$P = \begin{bmatrix} VEV^{\mathsf{T}} - V\Gamma EV^{\mathsf{T}} & V\Gamma EV^{\mathsf{T}} \\ -V\Gamma EV^{\mathsf{T}} & V\Gamma EV^{\mathsf{T}} \end{bmatrix}$$

$$= \underbrace{\begin{bmatrix} V & 0 \\ 0 & V \end{bmatrix}}_{\widetilde{V}} \begin{bmatrix} E-\Gamma E & \Gamma E \\ -\Gamma E & \Gamma E \end{bmatrix} \underbrace{\begin{bmatrix} V^{\mathsf{T}} & 0 \\ 0 & V^{\mathsf{T}} \end{bmatrix}}_{\widetilde{V}^{\mathsf{T}}} \tag{91}$$

where the matrix $\widetilde{V} \in \mathbb{R}^{2n\times 2n}$ is clearly orthogonal as V is orthogonal. Substituting (91) for P in (90), we have:

$$\begin{aligned} M = \widetilde{V} &\begin{bmatrix} E-\Gamma E & 0 \\ 0 & \Gamma E \end{bmatrix} \widetilde{V}^{\mathsf{T}} \\ &- \beta \widetilde{V} \begin{bmatrix} E-\Gamma E & -\Gamma E \\ \Gamma E & \Gamma E \end{bmatrix} \widetilde{V}^{\mathsf{T}} \widetilde{V} \begin{bmatrix} E-\Gamma E & \Gamma E \\ -\Gamma E & \Gamma E \end{bmatrix} \widetilde{V}^{\mathsf{T}} \\ = \widetilde{V} \widetilde{M} \widetilde{V}^{\mathsf{T}} &\succcurlyeq 0 \iff \widetilde{M} \succcurlyeq 0\,, \end{aligned} \tag{92}$$

where

$$\begin{aligned} \widetilde{M} = &\begin{bmatrix} E-\Gamma E & 0 \\ 0 & \Gamma E \end{bmatrix} \\ &- \beta \begin{bmatrix} E-\Gamma E & -\Gamma E \\ \Gamma E & \Gamma E \end{bmatrix} \begin{bmatrix} E-\Gamma E & \Gamma E \\ -\Gamma E & \Gamma E \end{bmatrix} \\ = &\begin{bmatrix} E-\Gamma E & 0 \\ 0 & \Gamma E \end{bmatrix} \\ &- \beta \begin{bmatrix} (E-\Gamma E)^2 + \Gamma^2 E^2 & (E-2\,\Gamma E)\,\Gamma E \\ (E-2\,\Gamma E)\,\Gamma E & 2\,\Gamma^2 E^2 \end{bmatrix} \\ = &\begin{bmatrix} E-\Gamma E & 0 \\ 0 & \Gamma E \end{bmatrix} \\ &- \beta \begin{bmatrix} E^2 - 2\,\Gamma E^2 + 2\,\Gamma^2 E^2 & \Gamma E^2 - 2\,\Gamma^2 E^2 \\ \Gamma E^2 - 2\,\Gamma^2 E^2 & 2\,\Gamma^2 E^2 \end{bmatrix} \\ = &\begin{bmatrix} \widetilde{M}_{1,1} & \widetilde{M}_{1,2} \\ \widetilde{M}_{2,1} & \widetilde{M}_{2,2} \end{bmatrix}, \end{aligned} \tag{93}$$

and where

$$\widetilde{M}_{1,1} = E\left(I_n - \Gamma - \beta E\left(I_n - 2\,\Gamma + 2\,\Gamma^2\right)\right), \tag{94}$$

$$\widetilde{M}_{1,2} = \widetilde{M}_{2,1} = -\beta\,\Gamma E^2\left(I_n - 2\,\Gamma\right), \tag{95}$$

$$\widetilde{M}_{2,2} = \Gamma E\left(I_n - 2\,\beta\,\Gamma E\right). \tag{96}$$

According to [40], there exists a permutation matrix $\Pi \in \mathbb{R}^{2n\times 2n}$ such that we can use the similarity transformation:

$$\Pi\,\widetilde{M}\,\Pi^{\mathsf{T}} = \widetilde{M}'_1 \oplus \widetilde{M}'_2 \oplus \cdots \oplus \widetilde{M}'_n = \widetilde{M}', \tag{97}$$

with

$$\left[\widetilde{M}'_k\right]_{i,j} = \widetilde{M}^k_{i,j}\,, \qquad k \in \{1,\ldots,n\}, \quad i,j \in \{1,2\}\,. \tag{98}$$

Thus, the original 2×2 block matrix $\widetilde{M}$ in (93)–(96) with diagonal blocks $\widetilde{M}_{i,j}$ of order n has been transformed into the block-diagonal matrix $\widetilde{M}'$ with n full blocks $\widetilde{M}'_k$ of order 2. Hence, the eigenvalues of $\widetilde{M}$ coincide with those of all the 2×2 blocks $\widetilde{M}'_k$.

The generic 2×2 block $\widetilde{M}'_k$ takes the following form:

$$\widetilde{M}'_k = \begin{bmatrix} e_k \left(1 - \gamma_k - \beta\, e_k \left(1 - 2\,\gamma_k + 2\,\gamma_k^2\right)\right) & -\beta\,\gamma_k\, e_k^2\,(1 - 2\,\gamma_k) \\ -\beta\,\gamma_k\, e_k^2\,(1 - 2\,\gamma_k) & \gamma_k\, e_k\,(1 - 2\,\beta\,\gamma_k\, e_k) \end{bmatrix}, \quad k = 1, \ldots, n\,. \tag{99}$$

Hence, $\widetilde{M}$ in (93)–(96) is positive semidefinite if and only if all the blocks $\widetilde{M}'_k$ in (99) are positive semidefinite. According to Sylvester's criterion, $\widetilde{M}'_k$ is positive semi-definite if and only if all its principal minors are non-negative, that is if and only if:

$$\left\{ \begin{array}{rcl} e_k \left(1 - \gamma_k - \beta\, e_k \left(1 - 2\,\gamma_k + 2\,\gamma_k^2\right)\right) & \geqslant & 0 \\ \gamma_k\, e_k\,(1 - 2\,\beta\,\gamma_k\, e_k) & \geqslant & 0 \\ \gamma_k\, e_k^2 \left(1 - \gamma_k - \beta\, e_k \left(1 - 2\,\gamma_k + 2\,\gamma_k^2\right)\right)(1 - 2\,\beta\,\gamma_k\, e_k) & & \\ -\beta^2\,\gamma_k^2\, e_k^4\,(1 - 2\,\gamma_k)^2 & \geqslant & 0. \end{array} \right. \tag{100}$$

First, we notice that if $e_k = 0$ then $\widetilde{M}'_k$ is a matrix of all zeros, hence it is positive semidefinite. We thus consider blocks for which $e_k > 0$. In case $\gamma_k = 0$, then the second and third inequalities in (100) are satisfied for any β and the first inequality reduces to

$$\beta \leqslant \frac{1}{e_k} \quad \forall\, k \in \{0, 1, \ldots, n\} \iff \beta \leqslant \frac{1}{\max_k e_k}\,. \tag{101}$$

We finally consider blocks for which $e_k, \gamma_k > 0$. The inequalities in (100) reduce to

$$\left\{ \begin{array}{rcl} 1 - \gamma_k - \beta\, e_k \left(1 - 2\,\gamma_k + 2\,\gamma_k^2\right) & \geqslant & 0 \\ 1 - 2\,\beta\,\gamma_k\, e_k & \geqslant & 0 \\ \left(1 - \gamma_k - \beta\, e_k \left(1 - 2\,\gamma_k + 2\,\gamma_k^2\right)\right)(1 - 2\,\beta\,\gamma_k\, e_k) & & \\ -\beta^2\,\gamma_k\, e_k^2\,(1 - 2\,\gamma_k)^2 & \geqslant & 0 \end{array} \right. \tag{102}$$

that is

$$\left\{ \begin{array}{rcl} \beta & \leqslant & \dfrac{1 - \gamma_k}{1 - 2\gamma_k + 2\gamma_k^2}\,\dfrac{1}{e_k} \\ \beta & \leqslant & \dfrac{1}{2\gamma_k}\,\dfrac{1}{e_k} \\ \beta^2 e_k^2 \gamma_k \left[2\left(1 - 2\,\gamma_k + 2\,\gamma_k^2\right) - (1 - 2\,\gamma_k)^2\right] & & \\ -\beta\, e_k \left[1 - 2\,\gamma_k + 2\,\gamma_k^2 + 2\,\gamma_k\,(1 - \gamma_k)\right] + 1 - \gamma_k & \geqslant & 0 \end{array} \right. \tag{103}$$

that is

$$\left\{ \begin{array}{rcl} \beta & \leqslant & \dfrac{1 - \gamma_k}{1 - 2\gamma_k + 2\gamma_k^2}\,\dfrac{1}{e_k} \\ \beta & \leqslant & \dfrac{1}{2\gamma_k}\,\dfrac{1}{e_k} \\ 0 & \leqslant & \beta^2 e_k^2 \gamma_k - \beta\, e_k + 1 - \gamma_k. \end{array} \right. \tag{104}$$

Simple algebraic manipulations allow to prove that the first inequality is more stringent than the second one and that the third inequality is always satisfied. Since (104) must hold for any k, we have:

$$\beta \leqslant \frac{1 - \gamma_k}{1 - 2\gamma_k + 2\gamma_k^2}\,\frac{1}{e_k} \quad \forall\, k \in \{0, 1, \ldots, n\} \iff \beta \leqslant \frac{1}{\max\limits_k \left\{\dfrac{1 - 2\gamma_k + 2\gamma_k^2}{1 - \gamma_k}\; e_k\right\}}\,. \tag{105}$$

This concludes the proof, after noticing that (105) contains (101) as a special case. □

We note that the FB algorithm requires that $\mathcal{P}$ be β-cocoercive with $\beta > 0$; hence, the choice $\gamma_k = 1$ for some k is precluded. Note also that if one chooses $\gamma_k = 0$ for all k and $\mathcal{R}$ is the ℓ_1 norm, then the algorithm in Proposition 10 reduces to the classic iterative shrinkage/thresholding algorithm (ISTA) [12,16]. The Douglas-Rachford algorithm (Theorem 25.6 in [1]) may also be used to find the saddle-point of F in (89).

Finally, we observe that, based on definitions (88)–(89), at each iteration the proposed algorithm in Prop. 10 performs a step of minimization in the primal variable x of the function

$$\mathcal{F}^{(x)}(x; v) := \frac{1}{2} \|Kx - b\|_2^2 + \lambda \, \mathrm{TV}(x) - \frac{\lambda}{2} \|Bx\|_2^2 + \lambda \, x^{\mathsf{T}} B^{\mathsf{T}} B v \tag{106}$$

and a step of maximization in the dual variable v of the function

$$\mathcal{F}^{(v)}(v; x) := -\lambda \, \mathrm{TV}(v) - \frac{\lambda}{2} \|Bv\|_2^2 + \lambda \, x^{\mathsf{T}} B^{\mathsf{T}} B v \, . \tag{107}$$

Hence the algorithm in Prop. 10 can be defined as belonging to the wide class of Primal-Dual Forward-Backward (PDFB) algorithms, with the only exception being the updating of the dual variable which is computed in terms of x_k instead of x_{k+1}.

7 Numerical Examples

In this section, we evaluate the performance of the CNC-NS-ℓ_2 convex variational model in (43) which is defined in terms of the proposed non-convex non-separable regularizer $\mathcal{R}_B$ defined in (41). More broadly, we investigate the advantages of using convex models defined using non-convex non-separable regularizers instead of either non-convex separable regularizers or convex separable regularizers, for restoring images characterized by some sparsity properties.

The presented numerical experiments involve different image processing problems, namely image denoising, deblurring and inpainting, with different regularizers. In particular, we consider gradient-based regularizers for the first three examples, while regularizers based on the nuclear norm and the Hessian Schatten norm will be investigated in Example 4 and Example 5, respectively.

We conduct experiments on the four gray-scale test images `geometric` (256×256 pixels), `QRcode` (320×320 pixels), `rectangles` (200×200 pixels) and `roof` (200×200 pixels), shown in Figure 2. The first three test images are characterized by constant regions and sharp edges – that is, they belong to the class of piecewise constant images – so as to highlight the sparsity-inducing properties of the CNC-NS-ℓ_2 model in the case of sparse gradient magnitudes. The fourth image considered, shown in Figure 2(d), belongs instead to the class of piecewise affine images and will be used in Example 5 to test the proposed CNC-NS-ℓ_2 model in case of images characterized by sparse Hessian magnitudes.

In the first three examples, we consider the proposed CNC-NS-ℓ_2 model (43) where $\mathcal{R}$ is the isotropic TV defined in (9), and we compare it with the popular convex Rudin-Osher-Fatemi (ROF) model [30], based on the minimization of the isotropic TV-ℓ_2 functional defined in (1).

In the third example, which deals with image denoising, since the matrix A in the quadratic fidelity term is reduced to the identity matrix, we can extend the comparison to a separable TV-based CNC denoising competitor, namely the one proposed in [20, 22], referred to as CNC-S-ℓ_2. Moreover, by appropriately setting the convexity parameter, the CNC-S-ℓ_2 model works in a non-convex regime. This makes our comparison with non-convex methods, namely the CNC-S-ℓ_2 model with suitable convexity parameter, and the TV_p model with $0 < p < 1$, more challenging. In the

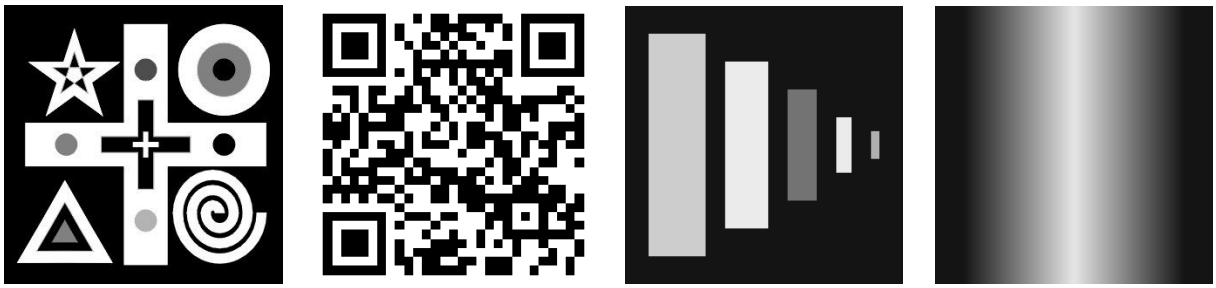

(a) (b) (c) (d)

Figure 2: Test images `geometric` (a), `qrcode` (b), `rectangles` (c) and `roof` (d).

last two examples, we investigate regularizers based on the convex nuclear and Hessian Schatten norms defined in (11)-(12).

The proposed CNC-NS-ℓ_2 method requires the prescription of a matrix B satisfying the convexity condition (74), i.e., $B^{\mathsf{T}}B \preccurlyeq (1/\lambda)A^{\mathsf{T}}A$. Many such matrices B exist. A simple option is to set B to be proportional to A as in (82). In the case of image denoising where $A = I$, this choice leads to B being a scalar multiple of identity. However, we found experimentally that using the more sophisticated strategy in (84), with Γ a suitably designed two-dimensional dc-notch filter, gave better results [28]. More precisely, let us first recall that only $B^{\mathsf{T}}B$ influences the regularizer, not B itself. Correspondingly, in the algorithm of Prop. 10, the matrix B arises only in the form $B^{\mathsf{T}}B$. Hence, it is sufficient to prescribe the matrix (operator) $B^{\mathsf{T}}B$; it is not necessary to prescribe B itself. In the following experiments, we set Γ to be a two-dimensional dc-notch filter defined by $\Gamma = I - H$ where H is a two-dimensional low-pass filter with a dc-gain of unity and $H \preccurlyeq I$. In our experiments, we set $\gamma = 0.98$ and $H = H_0^{\mathsf{T}}H_0$ where H_0 is the most basic two-dimensional low-pass filter: the moving-average filter with square support. Hence, H is a row-column separable two-dimensional filter given by convolution with a triangle sequence.

Regarding the optimization algorithms, the proposed CNC-NS-ℓ_2 model is numerically solved by the primal-dual forward-backward algorithm illustrated in Prop. 10. In particular, the computation of both v_{k+1} and x_{k+1} in the algorithm is carried out by an ADMM-based iterative algorithm. For the minimization of the CNC-S-ℓ_2 model in Example 2, we use the ADMM-based procedure proposed in [22] with parameters $\beta_t = \beta_z = 50$, and $\tau_c = 0.99$ such that for all the experiments the convexity condition is satisfied. For the TV-ℓ_2 model, we use a similar ADMM-based algorithm (see [10]) using the same parameters $\beta_t = \beta_z = 50$.

In all the experiments and for all the algorithms, we use the observed corrupted image as the initial iterate, i.e. $x_0 = b$, and we terminated the iterations as soon as two successive iterates satisfy

$$\delta_k^{(x)} := \frac{\|x_k - x_{k-1}\|_2}{\|x_{k-1}\|_2} < 10^{-5}\,. \tag{108}$$

The quality of the obtained restorations is evaluated by the Signal-to-Noise Ratio (SNR), defined as

$$\mathrm{SNR}(x^*, \bar{x}) := 10\log_{10}\big(\,\|\bar{x} - E[\bar{x}]\|_2^2 \;/\; \|x^* - \bar{x}\|_2^2\,\big)\,,$$

where $x^* \in \mathbb{R}^n$ is the computed estimate of the uncorrupted image $\bar{x} \in \mathbb{R}^n$ and $E[\bar{x}]$ denotes the mean value of $\bar{x}$. A high SNR value indicates that x^* is an accurate approximation of $\bar{x}$. In all the experiments and for all the algorithms, we hand-tuned the regularization parameter λ so as to obtain the highest possible SNR value.

Example 1. In this example, we test the performance of the proposed non-convex non-separable regularizer – or, better, of the overall CNC-NS-ℓ_2 proposed model in (43) – when applied to the

inverse problem of restoring images corrupted by blur and additive white Gaussian noise (AWGN). The degradation model we aim to 'invert' as accurately as possible reads as

$$b \;=\; K\bar{x} \;+\; \varepsilon\,, \tag{109}$$

where $\bar{x}, \varepsilon, b \in \mathbb{R}^n$ represent vectorized forms of the unknown uncorrupted image, unknown noise realization and observed corrupted image, respectively, and where $K \in \mathbb{R}^{n\times n}$ is a known linear blurring operator. The matrix K is typically so ill-conditioned (if not numerically singular) that recovering $\bar{u}$ given b and K by means of a naive (not regularized) least-squares procedure leads to meaningless results. Some sort of regularization is required. We consider here the restoration of the three test images `geometric`, `QRcode` and `rectangles` depicted in Fig. 2. These images belong to the class of piecewise constant images, which are characterized by sparse derivatives. One of the most popular and effective convex regularizers capable of promoting sparsity of image derivatives is the TV regularizer, the two common forms being defined in (9). Due to its well-known superior performance, we consider here the isotropic TV regularizer, which induces sparsity of the image gradient magnitudes. We thus assume the isotropic TV-ℓ_2 model as the baseline convex model and we seek to evaluate the relative performance of the proposed associated CNC-NS-ℓ_2 model as defined in (43), which, we note, is also convex in spite of its non-convex regularizer. For the sake of clarity, we recall that the two compared models obtain the restored images as the minimizers of the following two functionals:

$$\mathcal{J}(x) \;=\; \frac{1}{2}\,\|Kx-b\|_2^2 \;+\; \lambda\;\mathrm{TV}(x) \qquad [\;\mathrm{TV}{-}\ell_2\;] \tag{110}$$

$$\mathcal{J}_B(x) \;=\; \frac{1}{2}\,\|Kx-b\|_2^2 \;+\; \lambda\Big(\mathrm{TV}(x) - \Big(\mathrm{TV}\boxdot \tfrac{1}{2}\|B\cdot\|_2^2\Big)(x)\Big) \qquad [\;\mathrm{CNC}-\mathrm{NS}{-}\ell_2\;] \tag{111}$$

The experimental setting is as follows. The original test images are first synthetically corrupted by space-invariant Gaussian blur under the assumption of periodic boundary conditions. The blur matrix $K \in \mathbb{R}^{n\times n}$ in (109)–(111) is thus block-circulant with circulant blocks and is constructed starting from the Gaussian convolution kernel, or point-spread function, generated by the Matlab command `fspecial`('gaussian',`band`,`sigma`). The parameters `band` and `sigma` determine the bandwidth and the values of each circulant block in K, respectively. In particular, `band` represents the side length (in pixels) of the square support of the kernel, whereas `sigma` is the standard deviation of the circular, zero-mean, bivariate Gaussian probability density function representing the Gaussian point-spread function in the continuous setting. The blurred image $K\bar{x} \in \mathbb{R}^n$ is then corrupted by AWGN with standard deviation $\sigma = 40$ to obtain the image $b \in \mathbb{R}^n$. Given K and b, as accurate as possible estimates x^* of the original uncorrupted image $\bar{x}$ are computed for the two models in (110)–(111) by hand-tuning the regularization parameter λ.

Quantitative and qualitative (visual) measures of the accuracy of the obtained results have been produced. In Table 1 we report quantitative accuracy results measured by the SNR values associated with the achieved restorations for different blurring corruptions characterized by the parameters `band` and `sigma`. The best SNR results are marked in boldface. Figure 3 shows both the corrupted images and the restored images computed by the two compared methods in the case of blur with (`band`,`sigma`) = (5,1.5) - see the associated SNR values in the last column of Table 1.

From the SNR values in Table 1 and the visual inspection of the restored images in Fig. 3, the improvement in accuracy provided by the proposed non-convex non-separable regularizer, with respect to the convex separable TV regularizer, is evident. It is worth remarking that such improvement is obtained without renouncing any of the well-known advantages of (strongly) convex optimization, namely the existence of a unique (global) minimizer and of numerical algorithms

		(band,sigma)	
image	model	(3,1.5)	(5,1.5)
geometric	TV-ℓ_2	19.09	18.12
	CNC-NS-ℓ_2	**20.31**	**19.30**
QRcode	TV-ℓ_2	14.94	12.88
	CNC-NS-ℓ_2	**17.39**	**14.26**
rectangles	TV-ℓ_2	21.57	20.08
	CNC-NS-ℓ_2	**26.17**	**22.92**

Table 1: Example 1: SNR values obtained by deblurring the piecewise-constant test images `geometric`, `QRcode` and `rectangles` corrupted by space-invariant Gaussian blur with parameters (`band`,`sigma`) and AWGN with standard deviation $\sigma = 40$.

with proved convergence towards such minimizer. In this first example, we provide some empirical evidence for the numerical convergence of the proposed PDFB algorithm outlined in Prop. 10 towards the global minimizer of the CNC-NS-ℓ_2 objective function. In particular, we consider the `QRcode` deblurring test with blur parameters (`band`, `sigma`) = (5, 1.5) and in Fig. 4 we report the obtained convergence graphs for some relevant quantities. In particular, in the first row of Fig. 4 we show the graphs of the quantities $\delta_k^{(x)}$ and $\delta_k^{(v)}$ [see definition (108)] which indicate the good convergence behavior of both the sequences $\{x_k\}$ and $\{v_k\}$ of primal and dual iterates. In regards to the function values, in the second row of Fig. 4 we report the graphs of the quantities $F^{(x)}(x_k; v_k)$ and $F^{(v)}(v_k; x_k)$ defined in (106)-(107) which not only demonstrate the good convergence behavior of the sequences of primal and dual function values but also strongly indicate monotonicity of such sequences. Finally, the plots of the SNR value for each primal iterate x_k, shown in the third row of Fig. 4, suggest that few iterations (in the order of tens) of the PDFB algorithm are sufficient to achieve high quality results; see the two values $\mathrm{SNR}(x_{60}) = 14.24$ and $\mathrm{SNR}(x_{750}) = 14.26$ highlighted in the graphs. In Fig. 4(left column) we plot the values of these quantities for the first 800 iterations of the PDFB algorithm, whereas in Fig. 4(right column) we show only the first 80 iterations for detail.

Example 2. Another popular and successful application of sparsity-promoting regularization is image inpainting, where it is desired to reconstruct image pixels that are lost - i.e., not sensed correctly - or corrupted due, e.g., to scratches, texts or logos superposed on the image. The associated degradation model is

$$b \;=\; M\bar{x} \;+\; \varepsilon\,, \tag{112}$$

where $\bar{x} \in \mathbb{R}^n$ and $\varepsilon, b \in \mathbb{R}^m$ represent vectorized forms of the unknown uncorrupted image, unknown noise realization, and observed corrupted image, respectively, and where $M \in \mathbb{R}^{m\times n}$ with $m < n$ is a known selection (i.e., masking) operator which in practice selects the uncorrupted pixels. The matrix M is a binary wide matrix (hence, it does not have full column rank) obtained from the $n \times n$ identity matrix by eliminating $n - m$ different rows. The matrix $M^{\mathsf{T}}M \in \mathbb{R}^{n\times n}$ is a binary diagonal matrix with $1s$ in the positions on the main diagonal corresponding to the indexes of the rows of the identity matrix contained in M. The matrix $M^{\mathsf{T}}M$ is thus singular, in particular it has rank equal to m.

In this example, we focus on inpainting of the piecewise constant test image `QRcode` - see Fig. 2(b) - synthetically corrupted by AWGN with standard deviations $\sigma \in \{10, 20, 40\}$ and with two different types of inpainting masks, namely `scratch` and `random` - see the first column of

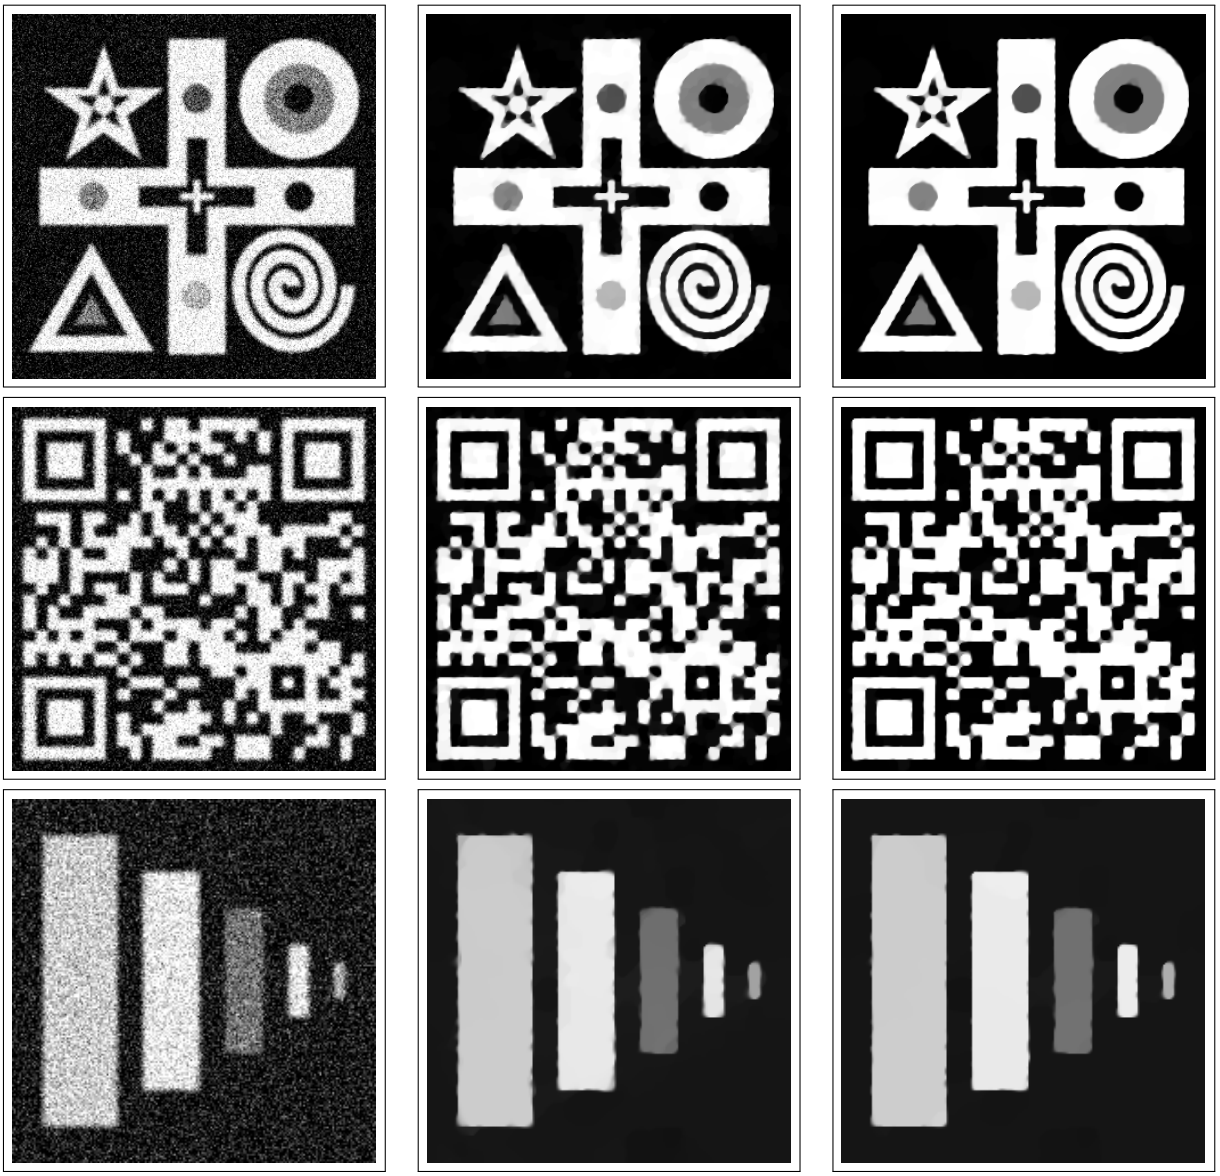

Figure 3: Example 1: restoration results obtained by the TV-ℓ_2 (second column) and the associated CNC-NS-ℓ_2 (third column) models applied to the piecewise constant test images `geometric`, `QRcode` and `rectangles` corrupted by Gaussian blur with parameters (`band`,`sigma`) $= (5, 1.5)$ and by AWGN with standard deviation $\sigma = 40$ (first column). The SNR values associated to the depicted restored images are reported in the last column of Table 1.

Fig. 5. We thus apply the convex isotropic TV-ℓ_2 and the associated CNC-NS-ℓ_2 models, which rely, respectively, on the minimization of the two functionals in (110) and (111) with the blurring matrix K replaced by the selection matrix M.

Some inpainting results with inpainting masks `scratch`, `random` 20%, and `random` 60% – see the red pixels in the leftmost column - are illustrated in Fig. 5. Table 2 reports the quantitative accuracy results measured by the SNR values associated with the best inpainted images obtained by the TV-ℓ_2 and CNC-NS-ℓ_2 approaches on the `QRcode` image with inpainting masks `scratch`, `random` 20% and `random` 60% and AWGN of standard deviations $\sigma \in \{10, 20, 40\}$. The best SNR results are indicated in boldface. As in Example 1, we observe the improved performance yielded

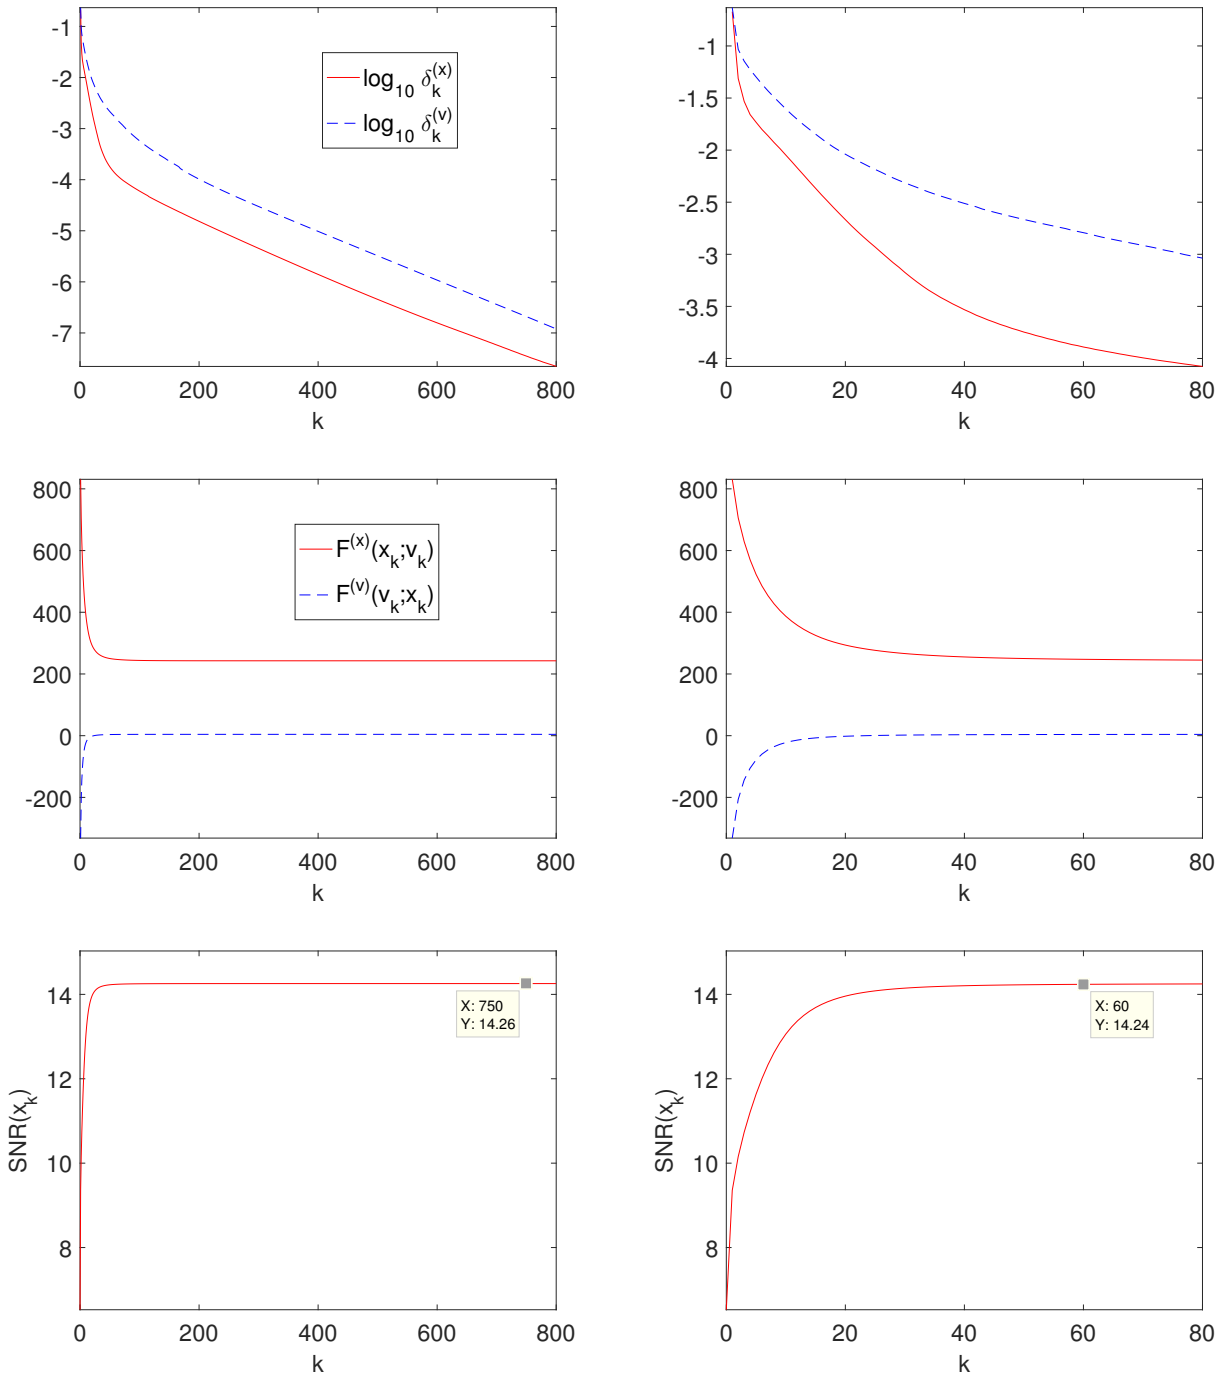

Figure 4: Example 1: empirical convergence graphs.

by the proposed model, especially at higher noise levels.

Example 3. In this example, we test the performance of the proposed CNC-NS-ℓ_2 model for denoising the three piecewise constant test images considered in Example 1, which call for regularizers that force the sparsity of image gradient magnitudes. Hence the degradation model and the functionals to be minimized are the same as in (109) and (110)-(111), with the blur matrix K replaced by the identity matrix.

		σ		
image	model	10	20	40
scratch	TV-ℓ_2	15.92	15.35	13.73
	CNC-NS-ℓ_2	**16.03**	**15.67**	**14.40**
random 20%	TV-ℓ_2	21.33	19.57	16.11
	CNC-NS-ℓ_2	**21.88**	**21.07**	**18.67**
random 60%	TV-ℓ_2	12.50	12.01	10.74
	CNC-NS-ℓ_2	**12.61**	**12.41**	**11.61**

Table 2: Example 2: SNR values obtained by inpainting the piecewise-constant test image `QRcode` with inpainting masks `scratch`, `random` 20% and `random` 60% (see first column of Fig.5) and corrupted by zero-mean AWGN with standard deviation σ.

Furthermore, in this example we can also extend the comparison to CNC models with non-convex but *separable* regularizers. To the best of our knowledge, the only TV-based CNC variational model for image denoising proposed so far is the one presented in [20, 22], from now on referred to as CNC-S-ℓ_2, which was demonstrated to outperform the classical TV-ℓ_2 denoising model in inducing sparsity of the gradient magnitudes in the denoised images. Moreover, the challenging comparison comes with purely non-convex models, namely, the TV_p in (5) with $0 < p < 1$ and the non-convex model based on separable regularizers of the form (6), which is achieved by using values of the convexity parameter a that violate the convexity limit in the CNC-S-ℓ_2 model. For both non-convex models, we hand-tuned their non-convexity parameters p and a so as to achieve the best SNR values for each test.

In Table 3 we report the SNR values achieved by the five compared methods for various noise levels, i.e., for various values of the standard deviation σ of the synthetic AWGN degradation. The best SNR results are indicated in boldface. The results in Table 3 indicate that the proposed non-convex non-separable regularizer can promote sparsity of the gradient magnitudes more effectively than both the standard convex separable TV regularizer and the non-convex separable regularizer used in the CNC-S-ℓ_2 model in convex regime [20, 22]. For what concerns the CNC-S-ℓ_2 model, we observe that when used in its non-convex regime it is always better than the convex counterpart. However, our proposal outperforms both the non-convex CNC-S-ℓ_2 model and the non-convex $\mathrm{TV}_p-\ell_2$ model on most of the tests. Overall, the results shown in Table 3 highlight the potential of the purely non-convex models but, at the same time, indicate, by the variability of their performance, their intrinsic weakness related to the existence of local minimizers. Instead, the proposed CNC approach yields high-quality solutions which can be robustly computed as they represent global minimizers of convex functionals.

Example 4. As discussed in Section 2, estimating a low-rank approximation of an unknown matrix from few of its entries is known as Matrix Completion (MC). In image processing, the MC procedure can be used for the inpainting problem as a valuable alternative to the method considered - and tested - in Example 2 or, more in general, to variational methods based on regularizers promoting sparsity of local differential quantities. In particular, MC is suitable for inpainting images $x \in \mathbb{R}^{d_1 \times d_2}$ satisfying a low-rank assumption: $\mathrm{rank}(x) \ll \bar{d} = \min\{d_1, d_2\}$.

The low-rank prior is typically implemented via convex relaxation, i.e., the rank of the image is substituted by its nuclear norm - the Shatten p-norm defined in (10) with $p = 1$ - and the regularizer in (12) can be used. We recall that the nuclear norm is proven to be the tightest convex surrogate of the rank operator [15], and the analogy between using the nuclear norm for low-rank

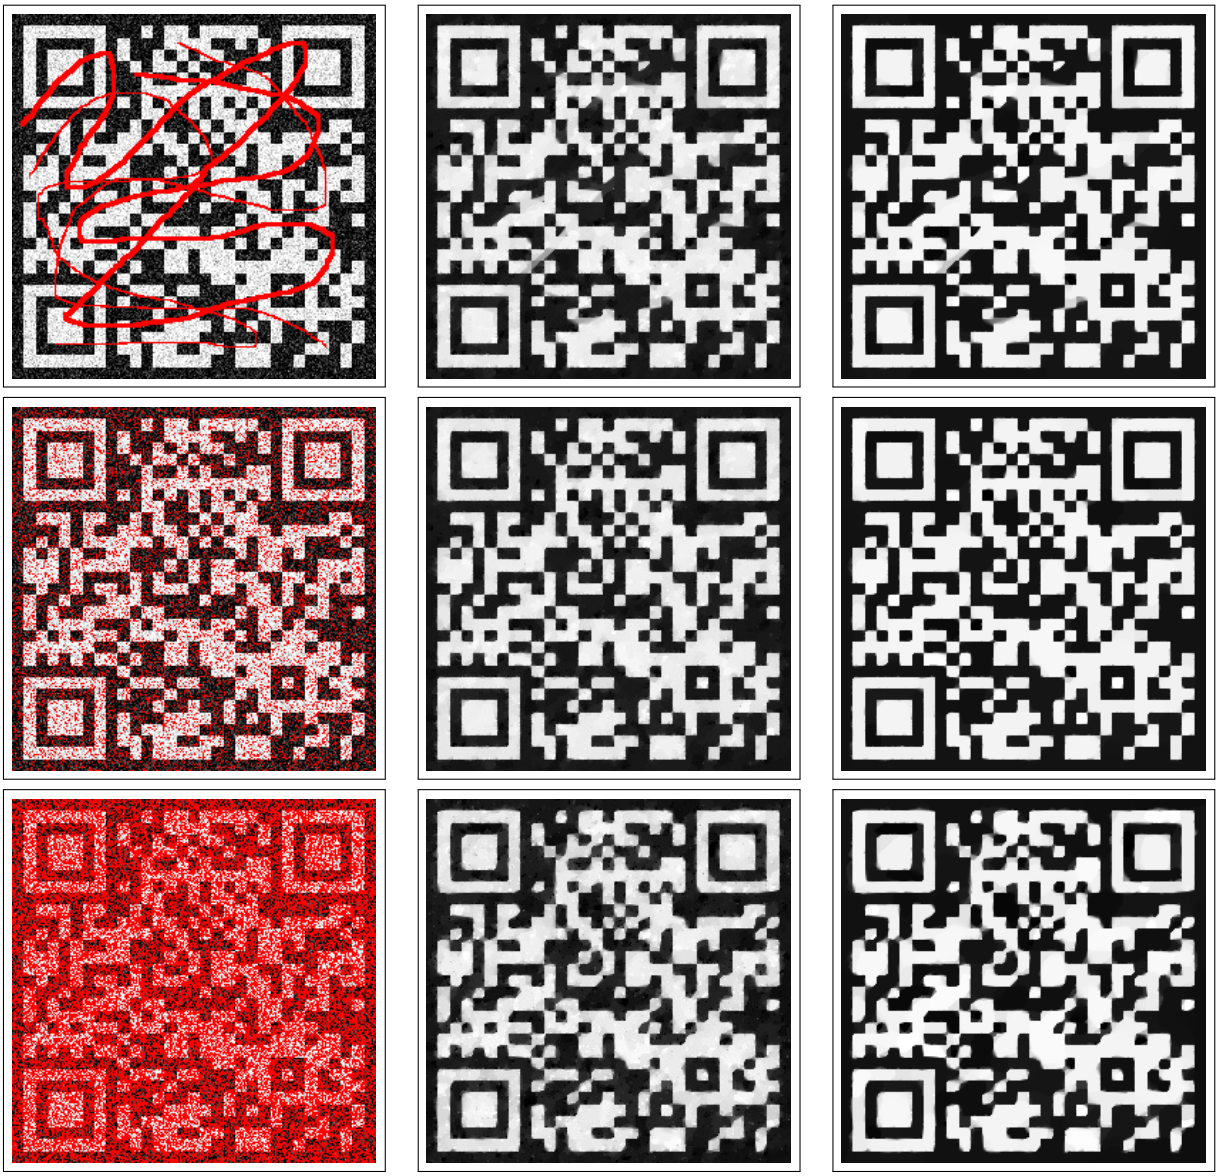

Figure 5: Example 2: inpainting results obtained by the TV-ℓ_2 (second column) and the associated CNC-NS-ℓ_2 (third column) models applied to the test image `QRcode` corrupted by zero-mean AWGN with standard deviation $\sigma = 20$.

matrix recovery and using the ℓ_1-norm for sparse signal recovery has been well established [29].

With all this said, in this example we consider the image inpainting problem - such that the image degradation model we aim to invert is the same as in (112) - solved by using both the convex least-squares variational model with the nuclear norm-based regularizer defined in (12), referred to as $\mathcal{S}_1 - \ell_2$ model, and the associated CNC-NS-ℓ_2 model. The two compared models rely on minimizing the following cost functionals:

$$\mathcal{J}(x) \;=\; \frac{1}{2}\,\|Mx - b\|_2^2 \;+\; \lambda\;\mathcal{S}_1(x) \qquad [\,\mathcal{S}_1 - \ell_2\,]\,, \tag{113}$$

$$\mathcal{J}_B(x) \;=\; \frac{1}{2}\,\|Mx - b\|_2^2 \;+\; \lambda\Big(\mathcal{S}_1(x) - \Big(\mathcal{S}_1 \boxdot \tfrac{1}{2}\|B\cdot\|_2^2\Big)(x)\Big) \qquad [\,\mathrm{CNC}-\mathrm{NS}-\ell_2\,]\,, \tag{114}$$

with M the selection matrix associated with the inpainting mask.

		σ				
image	model	10	20	40	60	80
geometric	TV-ℓ_2	30.15	25.76	20.87	18.13	16.26
	CNC-S-ℓ_2 (convex)	30.89	26.57	21.98	19.10	17.01
	CNC-S-ℓ_2 (non-convex)	**30.99**	**26.73**	22.04	19.15	17.08
	TV$_p$-ℓ_2 ($p < 1$)	30.82	25.98	21.10	18.28	16.16
	CNC-NS-ℓ_2	30.87	26.47	**22.68**	**20.16**	**18.31**
QRcode	TV-ℓ_2	30.79	24.95	19.27	15.41	13.81
	CNC-S-ℓ_2 (convex)	40.08	34.53	23.90	18.81	15.39
	CNC-S-ℓ_2 (non-convex)	42.03	36.03	25.05	20.53	17.01
	TV$_p$-ℓ_2 ($p < 1$)	**45.13**	**37.13**	**30.38**	21.77	16.18
	CNC-NS-ℓ_2	44.06	36.69	29.58	**23.31**	**17.64**
rectangles	TV-ℓ_2	34.97	29.32	23.77	20.62	18.44
	CNC-S-ℓ_2 (convex)	45.30	36.26	28.05	23.26	20.10
	CNC-S-ℓ_2 (non-convex)	48.30	39.31	31.72	24.41	20.75
	TV$_p$-ℓ_2 ($p < 1$)	49.65	40.84	29.00	21.38	18.76
	CNC-NS-ℓ_2	**51.00**	**42.61**	**33.53**	**27.58**	**23.16**

Table 3: Example 3: SNR values obtained by denoising the test images `geometric`, `QRcode` and `rectangles` corrupted by zero-mean AWGN with different values of the standard deviation σ.

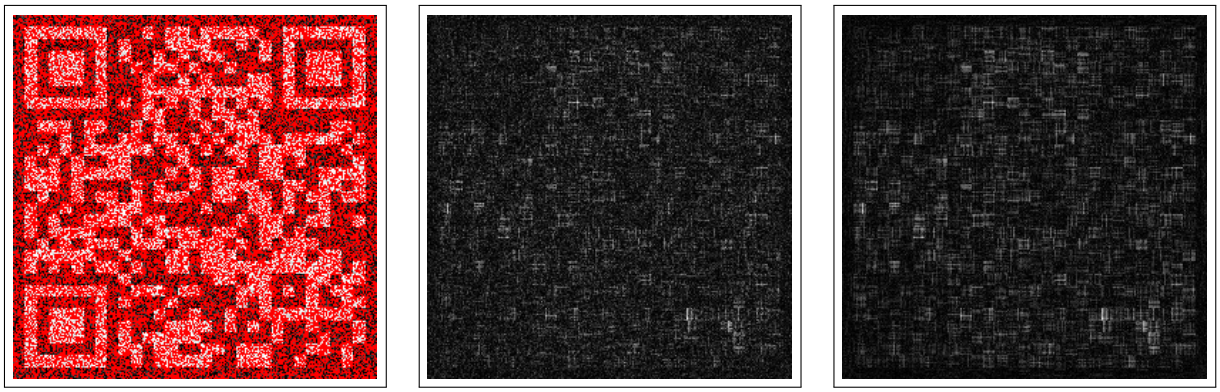

Figure 6: Example 4: Matrix completion for reconstructing the `QRcode` test image corrupted by the inpainting mask `random` 60% and by AWGN of standard deviation $\sigma = 30$ (left); absolute error images associated with the reconstruction results obtained by the $\mathcal{S}_1 - \ell_2$ model (center) - SNR = 8.96 - and by the CNC-NS-ℓ_2 model (right) - SNR = 9.91.

We applied the accelerated proximal gradient singular value thresholding algorithm APGL to solving the $\mathcal{S}_1 - \ell_2$ problem using the MATLAB implementation provided by the authors [39] and available at [38]. For the solution of the CNC-NS-ℓ_2 model, we used the algorithm illustrated in Prop. 10 where the two minimization problems (proximity operators) for x_{k+1} and v_{k+1} are solved using the APGL algorithm where A is the identity matrix.

In Fig. 6 we show some inpainting results for the `QRcode` test image, which satisfies the low-rank assumption and is thus prone to be successfully inpainted by the two models in (113)-(114).

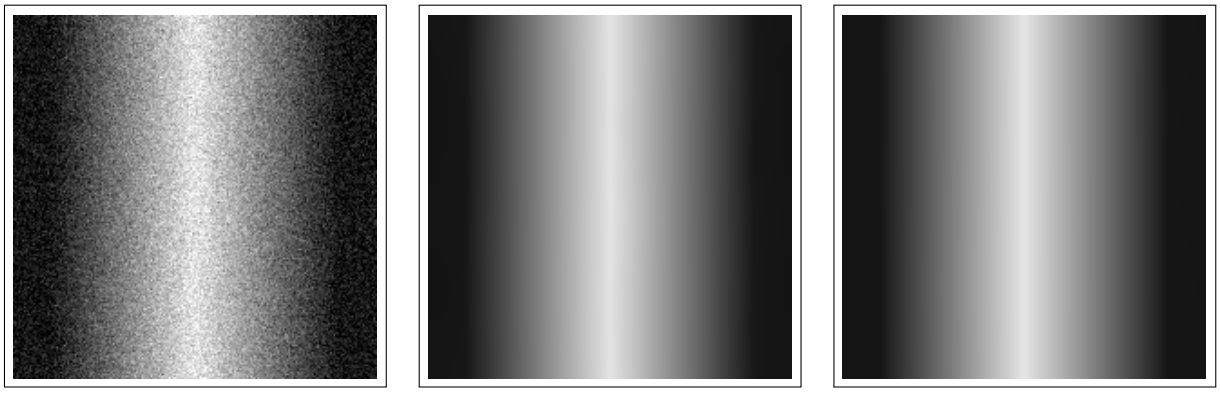

Figure 7: Example 5: restoration results obtained by the S_1H-ℓ_2 model (center), SNR = 37.14, and the associated CNC-NS-ℓ_2 model (right), SNR = 45.97, applied to the piecewise affine test image `roof` corrupted by space-invariant Gaussian blur with parameters (`band`,`sigma`) = (9,2) and by AWGN with standard deviation $\sigma = 20$ (left).

The image has been synthetically corrupted by the inpainting mask `random` 60% and by AWGN of standard deviation $\sigma = 30$ - see Fig. 6(left). In Fig. 6(center) and Fig. 6(right) we report the inpainting absolute error images - defined as the map of the absolute values of the pixel-by-pixel differences between the true image and the inpainted one - associated with the results obtained by the $\mathcal{S}_1 - \ell_2$ and CNC-NS-ℓ_2 models, respectively. From a visual inspection the two results look similar, but the difference in the associated SNR values highlights the more accurate reconstruction obtained by using the proposed CNC-NS-ℓ_2 model.

Example 5. In this example we consider again the image deblurring problem, such that the considered image degradation model is the same as in (109). The difference with respect to Example 1 is that now the penalty function $\mathcal{R}$ is the discrete Hessian Schatten-norm regularizer proposed in [24], defined in (11)-(10) with the particular value $p = 1$, i.e., the nuclear norm of the Hessian. To evaluate the proposed sparse regularization technique, we compare the convex restoration model in [24] with $p = 1$, referred to as S_1H-ℓ_2, with the CNC-NS-ℓ_2 model defined using the corresponding non-convex penalty $\mathcal{R}_B$. the two models minimize the following functionals

$$\mathcal{J}(x) \;=\; \frac{1}{2}\,\|Kx-b\|_2^2 \;+\; \lambda\;\; \mathcal{S}_1H(x) \qquad [\,\mathcal{S}_1H-\ell_2\,] \qquad (115)$$

$$\mathcal{J}_B(x) \;=\; \frac{1}{2}\,\|Kx-b\|_2^2 \;+\; \lambda\Big(\mathcal{S}_1H(x) - \Big(\mathcal{S}_1H \boxdot \tfrac{1}{2}\|B\cdot\|_2^2\Big)(x)\Big) \qquad [\,\mathrm{CNC-NS}{-}\ell_2\,] \qquad (116)$$

We remark that the regularizer $\mathcal{R}$ is known to provide better restorations that avoid staircases artifacts produced by TV regularizers for this class of images. Since such a regularizer is well known to sparsify the Hessian magnitudes, we consider the synthetic piecewise linear test image `roof`, shown in Fig. 2, corrupted by Gaussian blur of parameters `band` = 9, `sigma` = 2, and by AWGN of standard deviation $\sigma = 20$. Figure 7 shows the corrupted image (left) and the restored images obtained by the two models with the associated SNR values (center and right). Figure 8 shows a horizontal 1D scan line (row 100, columns 170–195) of these four images, which more clearly illustrates the improvement in the quality of restoration achieved by the CNC-NS-ℓ_2 model.

8 Conclusion

We proposed a novel strategy for sparsity-inducing regularization of linear least-squares problems. This strategy relies on constructing non-convex non-separable regularizers starting from a quite

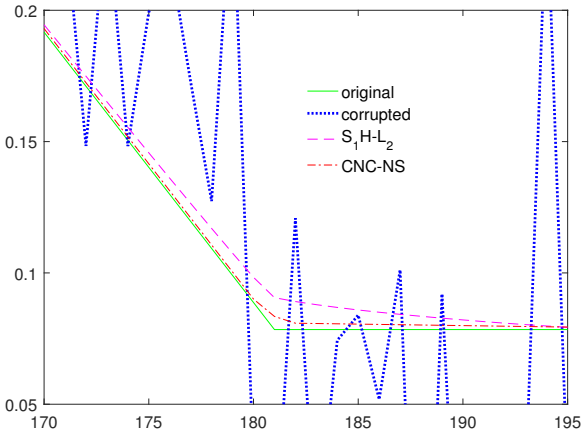

Figure 8: Example 5: horizontal 1D section (row 100, columns 170-195) of images in Fig. 2 and Fig. 7.

wide class of convex regularizers such that the total cost function remains convex but sparsity is promoted more effectively. With respect to previous works on Convex Non-Convex variational models, the proposed approach has a wider spectrum of applicability due to milder constraints on the operators involved in the model. The proposed regularizers depend on a matrix of free parameters which allows to impose the convexity of the cost functional. Two simple strategies for setting this matrix have been proposed; more sophisticated strategies will be matter of future studies together with a deeper investigation on whether and how these strategies depend on the particular image processing application. We presented a primal-dual forward-backward splitting algorithm to solve the saddle-point problem arising from the numerical solution of the proposed variational model. The convergence of the algorithm has been proved. A session of numerical experiments related to image deblurring, denoising and inpainting have been carried out and the reported results strongly indicate that non-separability of the non-convex regularizer holds the potential for achieving higher quality results while remaining in a convex, safe, regime.

Acknowledgements

We would like to thank the referees for comments that lead to improvements of the presentation. This research was supported in part by the National Group for Scientific Computation (GNCS-INDAM), Research Projects 2018; and by the National Science Foundation (NSF) under grant 1525398.

References

[1] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, 2011.

[2] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imag. Sci.*, 2(1):183–202, 2009.

[3] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[4] A. Bruckstein, D. Donoho, and M. Elad. From sparse solutions of systems of equations to sparse modeling of signals and images. *SIAM Review*, 51(1):34–81, 2009.

[5] E. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory*, 52(2):489–509, 2006.

[6] E. J. Candés, Y. C. Eldar, D. Needell, and P. Randall. Compressed sensing with coherent and redundant dictionaries. *Applied and Computational Harmonic Analysis*, 31(1):59–73, 2011.

[7] M. Carlsson. On convexification/optimization of functionals including an l2-misfit term. https://arxiv.org/abs/1609.09378, September 2016.

[8] A. Chambolle and T. Pock. An introduction to continuous optimization for imaging. *Acta Numerica*, 25:161–319, 2016.

[9] R. H. Chan, A. Lanza, S. Morigi, and F. Sgallari. Convex non-convex image segmentation. *Numerische Mathematik*, 138(3):635–680, Mar 2018.

[10] R. H. Chan, M. Tao, and X. M. Yuan. Constrained total variation deblurring models and fast algorithms based on alternating direction method of multipliers. *SIAM J. Imag. Sci.*, 6(1):680–697, 2013.

[11] Renè Ciak. *Coercive functions from a topological viewpoint and properties of minimizing sets of convex functions appearing in image restoration*. PhD thesis, 2015.

[12] I. Daubechies, M. Defrise, and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Commun. Pure Appl. Math*, 57(11):1413–1457, 2004.

[13] T.-M.-T. Do and T. Artières. Regularized bundle methods for convex and non-convex risks. *The Journal of Machine Learning Research*, 13(1):3539–3583, 2012.

[14] D.L. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, April 2006.

[15] M. Fazel. *Matrix rank minimization with applications*. PhD thesis, Stanford University, 2002.

[16] M. Figueiredo and R. Nowak. An EM algorithm for wavelet-based image restoration. *IEEE Trans. Image Process.*, 12(8):906–916, 2003.

[17] R. Giryes, M. Elad, and A. M. Bruckstein. Sparsity based methods for overparameterized variational problems. *SIAM J. Imag. Sci.*, 8(3):2133–2159, 2015.

[18] M. Huska, A. Lanza, S. Morigi, and F. Sgallari. Convex non-convex segmentation of scalar fields over arbitrary triangulated surfaces. *Journal of Computational and Applied Mathematics*, 2018.

[19] A. Lanza, S. Morigi, I. Selesnick, and F. Sgallari. Nonconvex nonsmooth optimization via convex–nonconvex majorization–minimization. *Numerische Mathematik*, 136(2):343–381, Jun 2017.

[20] A. Lanza, S. Morigi, and F. Sgallari. Convex image denoising via non-convex regularization. In J.-F. Aujol, M. Nikolova, and N. Papadakis, editors, *Proc. Scale Space and Variational Methods in Computer Vision (SSVM)*, pages 666–677. Springer International Publishing, 2015.

[21] A. Lanza, S. Morigi, and F. Sgallari. Constrained TV_p–ℓ_2 model for image restoration. *Journal of Scientific Computing*, 68(1):64–91, 2016.

[22] A. Lanza, S. Morigi, and F. Sgallari. Convex image denoising via non-convex regularization with parameter selection. *Journal of Mathematical Imaging and Vision*, 56(2):195–220, Oct 2016.

[23] L. Laporte, R. Flamary, S. Canu, S. Déjean, and J. Mothe. Nonconvex regularizations for feature selection in ranking with sparse SVM. *IEEE Trans. on Neural Networks and Learning Systems*, 25(6):1118–1130, 2014.

[24] S. Lefkimmiatis, J. Ward, and M. Unser. Hessian schatten-norm regularization for linear inverse problems. *Image Processing, IEEE Transactions on*, 22:1873–1888, 05 2013.

[25] Y. Lou, T. Zeng, S. Osher, and J. Xin. A weighted difference of anisotropic and isotropic total variation model for image processing. *SIAM Journal on Imaging Sciences*, 8(3):1798–1823, 2015.

[26] B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM J. Comput.*, 24(2):227–234, April 1995.

[27] M. Nikolova. Energy minimization methods. In O. Scherzer, editor, *Handbook of Mathematical Methods in Imaging*, chapter 5, pages 138–186. Springer, 2011.

[28] T. W. Parks and C. S. Burrus. *Digital Filter Design*. John Wiley and Sons, 1987.

[29] B. Recht, M. Fazel, and P.A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.

[30] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Phys. D*, 60(1-4):259–268, 1992.

[31] I. Selesnick. Sparse regularization via convex analysis. *IEEE Trans. Signal Process.*, 65(17):4481–4494, September 2017.

[32] I. Selesnick. Total variation denoising via the Moreau envelope. *IEEE Signal Processing Letters*, 24(2):216–220, Feb 2017.

[33] I. Selesnick and M. Farshchian. Sparse signal approximation via nonseparable regularization. *IEEE Transactions on Signal Processing*, 65(10):2561–2575, May 2017.

[34] I. W. Selesnick and I. Bayram. Sparse signal estimation by maximally sparse convex optimization. *IEEE Trans. Signal Process.*, 62(5):1078–1092, March 2014.

[35] I. W. Selesnick, A. Parekh, and I. Bayram. Convex 1-D total variation denoising with non-convex regularization. *IEEE Signal Processing Letters*, 22(2):141–144, February 2015.

[36] E. Soubies, L. Blanc-Féraud, and G. Aubert. A continuous exact ℓ_0 penalty (CEL0) for least squares regularized problem. *SIAM J. Imag. Sci.*, 8(3):1607–1639, 2015.

[37] M. E. Tipping. Sparse Bayesian learning and the relevance vector machine. *J. Machine Learning Research*, 1:211–244, 2001.

[38] K.C. Toh and S. Yun. NNLS version 0 – a matlab software for nuclear norm regularized linear least squares problems based on an accelerated proximal gradient method. `http://www.math.nus.edu.sg/~mattohkc/NNLS.html`, 2009.

[39] K.C. Toh and S. Yun. An accelerated proximal gradient algorithm for nuclear norm regularized linear least squares problems. *Pacific Journal of Optimization*, 6:615–640, 2010.

[40] Nikolaos L. Tsitsas. On block matrices associated with discrete trigonometric transforms and their use in the theory of wave propagation. *Journal of Computational Mathematics*, 28(6):864–878, 2010.

[41] Samuel Vaiter, Mohammad Golbabaee, Jalal Fadili, and Gabriel Peyré. Model selection with low complexity priors. *Information and Inference: A Journal of the IMA*, 4(3):230–287, 2015.

[42] D. P. Wipf, B. D. Rao, and S. Nagarajan. Latent variable Bayesian models for promoting sparsity. *IEEE Trans. Inform. Theory*, 57(9):6236–6255, September 2011.

[43] Y.-B. Zhao and D. Li. Reweighted ℓ_1-minimization for sparse solutions to underdetermined linear systems. *SIAM Journal on Optimization*, 22(3):1065–1088, 2012.