

Structure Tensor Total Variation*

Stamatios Lefkimmiatis[†], Anastasios Roussos[‡], Petros Maragos[§], and Michael Unser[¶]

Abstract. We introduce a novel generic energy functional that we employ to solve inverse imaging problems within a variational framework. The proposed regularization family, termed as structure tensor total variation (STV), penalizes the eigenvalues of the structure tensor and is suitable for both grayscale and vector-valued images. It generalizes several existing variational penalties, including the total variation seminorm and vectorial extensions of it. Meanwhile, thanks to the structure tensor's ability to capture first-order information around a local neighborhood, the STV functionals can provide more robust measures of image variation. Further, we prove that the STV regularizers are convex while they also satisfy several invariance properties w.r.t. image transformations. These properties qualify them as ideal candidates for imaging applications. In addition, for the discrete version of the STV functionals we derive an equivalent definition that is based on the patch-based Jacobian operator, a novel linear operator which extends the Jacobian matrix. This alternative definition allow us to derive a dual problem formulation. The duality of the problem paves the way for employing robust tools from convex optimization and enables us to design an efficient and parallelizable optimization algorithm. Finally, we present extensive experiments on various inverse imaging problems, where we compare our regularizers with other competing regularization approaches. Our results are shown to be systematically superior, both quantitatively and visually.

Key words. structure tensor, patch-based Jacobian, image reconstruction, convex optimization, total variation, inverse problems

AMS subject classifications. 65F22, 65K10, 90C06, 90C25, 49M29

DOI. 10.1137/14098154X

1. Introduction. Inverse problems are ubiquitous in science and engineering and they have been an important topic of interest for many years. Nowadays, with the advent of modern imaging techniques, inverse problems appear increasingly in a host of imaging applications ranging from microscopy and medical imaging to remote sensing and astronomical imaging [6]. They also appear in a plethora of computer vision applications, including motion estimation, image registration, stereo, and dense three-dimensional (3D) reconstruction [42]. These problems pertain to estimating unknown features of interest from partial or indirect measurements

*Received by the editors August 8, 2014; accepted for publication (in revised form) February 5, 2015; published electronically May 7, 2015. The first and second authors have a joint first authorship.

<http://www.siam.org/journals/siims/8-2/98154.html>

[†]Department of Mathematics, University of California, Los Angeles, Los Angeles, CA 90095 (stamatis@math.ucla.edu). The research of this author was supported by the Swiss National Science Foundation (SNF) under grant P300P2_151325.

[‡]Department of Computing, Imperial College London, London NW1 2FD, UK (troussos@cs.ucl.ac.uk). The research of this author was supported by the European Research Council Starting Grant 204871-HUMANIS.

[§]School of ECE, National Technical University of Athens (maragos@cs.ntua.gr). The research of this author was supported in part by the grant COGNIMUSE under the ARISTEIA Action co-funded by the European Social Fund and Greek National Resources.

[¶]School of Engineering, Swiss Federal Institute of Technology, Lausanne (michael.unser@epfl.ch). The research of this author was supported in part by the European Research Council under grant ERC-2010-AdG 267439-FUN-SP.

of image attributes. In practice, inverse imaging problems are typically ill-posed. This is to say that the equations relating the object of interest to the measurements are not enough by themselves to uniquely characterize the solution; a model of known properties of the underlying image is necessary for the recovered information to be physically and/or statistically meaningful.

A common strategy for dealing with ill-posed inverse problems is the variational approach. The main element of this framework is the regularization functional, whose role is to favor solutions that exhibit desirable properties. The proper selection of a regularizer is of paramount importance since it significantly influences the reconstruction quality. During the past years, this has triggered an increasing research interest on the design of regularizers able to encode important properties of natural images.

One of the most popular regularizers for imaging applications is the total variation (TV) seminorm [47]. TV is a convex functional which allows the use of powerful optimization techniques. Usually it is intended for grayscale images but it has also been extended to apply to vector-valued images; see, for instance, [7, 30, 48, 55, 58]. The main reason for its success is its ability to reconstruct images with sharp, well-preserved edges. However, it has the drawbacks of oversmoothing homogenous regions and creating staircase artifacts [16].

To overcome the limitations of TV, several regularization functionals have been proposed in the literature. Among them there are extensions that adapt the penalization of image variation, by locally varying the regularization weight [22, 29, 32] or by incorporating an anisotropy that is steered by the local structures of the image [5, 49, 53]. However, these methods extract the local information either from the input image in a preprocessing step or as an additional unknown function of the optimization problem, not directly depending on the underlying image. On the contrary, the so-called *anisotropic TV* (ATV) [33, 38] adapts the penalization of image variation by introducing a “diffusion” tensor that depends on the structure tensor of the unknown image itself. Nevertheless, in this case the adaptivity on image structures is heuristically designed, similarly to the design of the coherence-enhancing diffusion [57]. Another related regularizer is the Beltrami functional [52] of vector-valued images, which corresponds to the area of the generalized graph of the image embedded in a high dimensional space, using the induced metric. This framework has been recently generalized and unified with the Mumford–Shah functional [51], while in [46] an extension of the Beltrami framework that uses image patches has been introduced. A different regularization approach that also aims to eliminate the staircase effect of TV, involves functionals that can encode higher-order information. These regularizers promote piecewise-linear instead of piecewise-constant reconstructions and are either combined with TV as in [13, 41] or used in a standalone fashion as the total generalized variation (TGV) [10, 50] and the Hessian Schatten-norm regularizers [37]. Finally, there also exist nonlocal variational techniques that exploit the nonlocal self-similarity property of natural images. This is made possible by employing functionals that allow nonlocal interactions between image points [24, 28]. Among the most efficient regularizer of this category is the nonlocal TV functional (NLTV) [28]. Discrete-domain extensions of NLTV for vector-valued images are studied in [17, 18].

1.1. Contributions. In this work, we propose a family of convex regularizers that penalize the image variation at every point of the domain by taking into account the information in a local neighborhood. This is accomplished by adopting a penalty that depends on the

eigenvalues of the structure tensor of the underlying image. In detail, our key contributions are the following:

1. The introduction of a family of energy functionals that generalize TV and vectorial extensions of it. Our regularizers preserve all the favorable properties of TV, such as convexity and invariance w.r.t. image transformations. Moreover, they provide a more robust and richer measure of image variation by exploiting additional information from the neighborhood of every point. We use these functionals in a variational framework to obtain regularized solutions of ill-posed inverse imaging problems.
2. The proof that our regularization criteria, which involve the eigenvalues of the nonlinear structure tensor operator, are convex and invariant to image transformations.
3. The concept of the patch-based Jacobian as a novel neighborhood-aware extension of the discrete Jacobian operator. We employ this operator to obtain an equivalent definition for the discrete version of our proposed energy functionals. This further allows us to derive a dual problem formulation.
4. The development of a practical first-order algorithm for evaluating the proximal map of the introduced regularizers. Our optimization algorithm copes with the nonsmoothness of the energy functionals and leads to efficient numerical solutions.
5. An extensive quantitative experimental evaluation of our method on several inverse problems for different degradation conditions. The comparisons are performed on a large dataset of both graylevel and color images, where we compare the proposed method with TV and vectorial extensions of it, as well as with the state-of-the-art TGV [10]. The results show that in all the studied problems, our method systematically outperforms all other evaluated techniques.

We note that in the conference papers [36, 46], which contain preliminary parts of this work, we established a connection between the minimization of a class of structure tensor-based regularizers, which also includes a relaxation of the proposed family of regularizers, and a novel generic type of nonlinear anisotropic diffusion. This type of diffusion is driven by a diffusion tensor that depends on the solution itself and is spatially regularized (in the sense that it involves convolutions with a kernel, as, e.g., in [55, 57]). This connection, in conjunction with the efficient optimization strategy that we present here, implies that we can employ powerful convex optimization techniques for implementing this type of anisotropic diffusion schemes. We intend to present in depth the relations of the proposed regularizers and anisotropic diffusion in a separate article.

An important extension of the present work with respect to the conference paper in [36] is that we introduce a continuous-domain formulation of the problem, which is guided by fundamental invariance properties (translation and rotation invariance, 1-homogeneity), and further provide a proof of convexity. These new results are presented in Theorem 2.1. Additionally, we establish a bound for the Lipschitz constant of the dual objective function that arises in the evaluation of the proximal operator of our functionals (Proposition 4.2). This allows us to use a fixed step-size in our gradient-based optimization algorithm. Further, we provide extended comparisons on several inverse imaging problems, including two new applications, namely, image reconstruction from sparse Fourier measurements and image magnification. Finally, for the deblurring problem that we study, we consider an improved forward model than the one we used in [36]. This formulation avoids the need of making any specific assumption about

the image boundaries and holds the regularizer responsible for reconstructing the image in a way that best explains the observation.

This paper is organized as follows: in section 2, we introduce the proposed family of regularizers and motivate their use. In section 3, we focus on the discrete domain and introduce the patch-based Jacobian operator. In section 4, we design our optimization algorithm, while in section 5 we assess the reconstruction performance of our approach with extensive experiments on several linear inverse imaging problems. Conclusions are drawn in section 6, while all the mathematical proofs are provided in Appendix A.

2. Regularization for inverse problems.

2.1. Problem formulation and variational recovery. The relationship between the image of interest and the measurements is given by the physical forward model that provides the basis of the acquisition process. For many imaging applications, the acquisition is well-described by a linear process, which can be mathematically formulated as

$$(2.1) \quad \mathbf{v}(\mathbf{x}) \sim \mathcal{N}(\mathcal{A}\mathbf{u}(\mathbf{x})).$$

In this case, $\mathbf{u}(\mathbf{x}) = [u_1(\mathbf{x}) \dots u_M(\mathbf{x})] : \mathbb{R}^2 \mapsto \mathbb{R}^M$ represents the generic vector-valued image with M channels that we wish to reconstruct and \mathcal{A} is a linear operator that provides a mapping from the space of underlying images to the space of measurements. The symbol \mathcal{N} represents the measurement noise which accounts for all possible types of errors during the acquisition, including stochastic noise.

The recovery of \mathbf{u} from the measurements \mathbf{v} in (2.1) belongs to the category of linear inverse problems. For most cases of practical interest, the operator \mathcal{A} is either ill-conditioned or singular. This kind of *ill-posedness* [56] is dealt with within a variational framework, where the reconstruction of \mathbf{u} is cast as a minimization problem of an objective function of the form

$$(2.2) \quad \mathcal{E}(\mathbf{u}) = \varphi(\mathcal{A}\mathbf{u}) + \tau\psi(\mathbf{u}).$$

Such a cost function consists of (a) the *data fidelity* term $\varphi(\mathcal{A}\mathbf{u})$, which measures how well a candidate solution explains the observed data, and (b) the *regularizer* $\psi(\mathbf{u})$ which encodes any available prior information about the underlying image. The exact form of the data fidelity term depends on the assumed noise model perturbing the measurements. From a Bayesian point of view, the overall reconstruction approach corresponds to either a penalized maximum likelihood or a maximum a posteriori estimation problem [26].

2.2. TV regularization. As stated earlier, TV [47] applies to grayscale images $u(M = 1)$ and for smooth images is defined as

$$(2.3) \quad \text{TV}(u) = \int_{\mathbb{R}^2} \|\nabla u\|_2 \, d\mathbf{x}.$$

Apart from the summation over the domain, another key ingredient of TV is the gradient magnitude $\|\nabla u\|_2$ of the underlying image. Their combination leads to an L_1 weighting that does not overpenalize high variations of u . Therefore, the reconstructed images feature sharp edges, which are visually appealing. Due to this favorable property, TV has been

extended to vector-valued images \mathbf{u} . A popular extension that retains TV's properties and at the same time introduces a coupling between the different channels of \mathbf{u} is the *vectorial TV* (VTV) [7, 48]. VTV has a definition similar to (2.3), where the gradient magnitude is replaced by its straight-forward vectorial extension $\|\nabla\mathbf{u}\|_2 = \{\sum_{i=1}^M \|\nabla u_i\|_2^2\}^{\frac{1}{2}}$.

A known drawback of TV is that it favors piecewise-constant solutions and, thus, it can create strong staircase artifacts in smooth regions of the image. Furthermore, a basic shortcoming of both TV and VTV is that the gradient magnitude, employed to penalize the image variation at every point \mathbf{x} , is too simple as an image descriptor; it relies only on \mathbf{x} without taking into account the available information from its neighborhood. In fact, most of the existing extensions of TV [30, 33, 55, 58] as well as related regularizers [52] share the same drawback: they integrate a penalty of image variation that is completely localized.

To overcome these limitations, we next adopt more sophisticated descriptors of image variations that take into account information from the neighborhood of every point. Consequently, the resulting regularizers are able to provide a richer and more robust measure of image regularity.

2.3. Directional derivatives and the structure tensor. Our goal is to develop neighborhood-aware measures of vectorial image variations that will be incorporated in a novel variational framework. To do so, we first revisit the theory¹ behind the structure tensor [35, 57]. Hereafter, we assume that the vector-valued image \mathbf{u} belongs to the Sobolev space $W^{1,2}(\mathbb{R}^2, \mathbb{R}^M)$. Let \mathbf{n} be an arbitrary two-dimensional (2D) direction ($\|\mathbf{n}\|_2 = 1$). The *vectorial directional derivative* of the vector-valued image \mathbf{u} in the direction \mathbf{n} and at any specific image point \mathbf{x} is given by [21]

$$(2.4) \quad \frac{\partial \mathbf{u}}{\partial \mathbf{n}}(\mathbf{x}) = (J\mathbf{u}(\mathbf{x}))\mathbf{n},$$

where $J\mathbf{u}$ is the *Jacobian matrix* of \mathbf{u} defined as

$$(2.5) \quad J\mathbf{u}(\mathbf{x}) = [\nabla u_1(\mathbf{x}) \quad \dots \quad \nabla u_M(\mathbf{x})]^T.$$

The magnitude of the directional derivative $\|\partial\mathbf{u}/\partial\mathbf{n}\|_2$ yields a measure of the amount of change of the image \mathbf{u} at the direction \mathbf{n} for any point \mathbf{x} . This measure is computed by concentrating completely at the point \mathbf{x} . In order to be more robust and also capture the behavior of the image \mathbf{u} in the neighborhood of \mathbf{x} , we consider instead the weighted *root mean square* (RMS) of $\|\partial\mathbf{u}/\partial\mathbf{n}\|_2$, which we call (local) *directional variation*:

$$(2.6) \quad \text{RMS}_K \{\|\partial\mathbf{u}/\partial\mathbf{n}\|_2\} = \sqrt{K * \|\partial\mathbf{u}/\partial\mathbf{n}\|_2^2} = \sqrt{\mathbf{n}^T (S_K \mathbf{u}) \mathbf{n}}.$$

In the above equation $*$ denotes the convolution operation, $K(\mathbf{x})$ is a nonnegative, rotationally symmetric convolution kernel, $K(\mathbf{x}) = K(|\mathbf{x}|)$, that performs the weighted averaging (e.g., a 2D Gaussian) and $S_K \mathbf{u}$ is the so-called *structure tensor* [27] of the image \mathbf{u} at point \mathbf{x} defined as

$$(2.7) \quad S_K \mathbf{u}(\mathbf{x}) = K * [J\mathbf{u}^T J\mathbf{u}](\mathbf{x}).$$

¹Instead of the usual focus on the eigenvalues of the structure tensor [35, 57], our study focuses on their square roots.

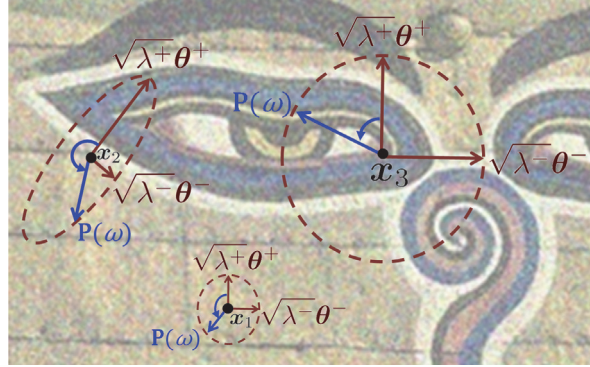


Figure 1. Structure tensor of an image at sample points $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3$. The structure tensor is visualized as an ellipse and its unit eigenvectors $\boldsymbol{\theta}^+, \boldsymbol{\theta}^-$ and rooted eigenvalues $\sqrt{\lambda^+}, \sqrt{\lambda^-}$ are also depicted. For any arbitrary direction \mathbf{n} , which is characterized by its angle ω with the eigenvector $\boldsymbol{\theta}^+$, there is a corresponding point $\mathbf{P}(\omega)$ on the ellipse. The distance of this point from the ellipse center yields the directional variation $V(\omega)$.

Note that in the above definition, similarly to [55], we do not consider any presmoothing of the image before computing its Jacobian. The reason is that the single convolution with K seems sufficient for the needs of image regularization while it also avoids the introduction of additional complexity in the computations.

Let $\lambda^+ = \lambda^+(S_K \mathbf{u}(\mathbf{x}))$, $\lambda^- = \lambda^-(S_K \mathbf{u}(\mathbf{x}))$ be the *eigenvalues* of $S_K \mathbf{u}(\mathbf{x})$, ordered so that $\lambda^+ \geq \lambda^-$, and $\boldsymbol{\theta}^+, \boldsymbol{\theta}^-$ are the corresponding unit *eigenvectors*. Also, let $\omega \in (-\pi, \pi]$ denote the angle between the direction vector \mathbf{n} and the eigenvector $\boldsymbol{\theta}^+$. Using the eigendecomposition of $S_K(\mathbf{u})$, the directional variation (2.6) can be expressed as a function of the angle ω :

$$(2.8) \quad V(\omega) \triangleq \text{RMS}_K \{ \|\partial \mathbf{u} / \partial \mathbf{n}\|_2 \} = \sqrt{\lambda^+ \cos^2 \omega + \lambda^- \sin^2 \omega}.$$

Now, let us consider the ellipse given by the parametric equation

$$(2.9) \quad \mathbf{P}(\omega) = \cos \omega \sqrt{\lambda^+} \boldsymbol{\theta}^+ + \sin \omega \sqrt{\lambda^-} \boldsymbol{\theta}^-, \omega \in [0, 2\pi).$$

The major and minor radius of this ellipse are $\sqrt{\lambda^+}$ and $\sqrt{\lambda^-}$, respectively, whereas the direction of the major and minor axis are given by $\boldsymbol{\theta}^+$ and $\boldsymbol{\theta}^-$, respectively. Therefore, this type of ellipse offers a way to visualize the structure tensor $S_K \mathbf{u}$ as shown in Figure 1. More importantly, we observe that $V(\omega) = \|\mathbf{P}(\omega)\|_2$. This implies that the directional variation $V(\omega)$ can be interpreted as the distance of any point $\mathbf{P}(\omega)$ from the center of the ellipse.

From (2.8), we observe that the maximum of the directional variation $V(\omega)$ is $\sqrt{\lambda^+}$ and is achieved for $\omega = 0, \pi$, whereas the minimum of $V(\omega)$ is $\sqrt{\lambda^-}$ and is achieved for $\omega = \pi \pm \pi/2$. In conclusion, the structure tensor measures the geometry of image structures in the neighborhood of each point: its eigenvectors $\boldsymbol{\theta}^+$ and $\boldsymbol{\theta}^-$ describe the orientation of maximum and minimum vectorial variation of \mathbf{u} and the square roots of its eigenvalues, $\sqrt{\lambda^+}$ and $\sqrt{\lambda^-}$, describe measures of these variations.

The eigenvalues of the structure tensor offer a rich and discriminative description of the local geometry of the image. When both λ^+ and λ^- are relatively small (e.g., point \mathbf{x}_1 in Figure 1) there are small vectorial variations around the current point, indicating that the

region is homogeneous. When λ^+ is large and λ^- is small (e.g., point \mathbf{x}_2 in Figure 1) there are strong variations, but only in a dominant direction. Therefore, the current point is located close to an image edge. When both λ^+ and λ^- are large (e.g., point \mathbf{x}_3 in Figure 1) there are high variations in all directions, which implies that the current point is close to an image corner.

2.4. The structure tensor TV functional. To design regularizers that integrate a scalar penalty of the local image variation at every image point, we need to consider measures that provide a synopsis of the function $V(\omega)$. Such measures, which are also computed from the eigenvalues of the structure tensor, are the following:

Case 1. RMS value of $V(\omega)$: $((2\pi)^{-1} \int_0^{2\pi} V^2(\omega) d\omega)^{1/2} = \sqrt{\lambda^+ + \lambda^-} / \sqrt{2}$, as it can be easily verified using equation (2.8).

Case 2. Maximum value of $V(\omega)$: the maximum of $V(\omega)$ w.r.t. ω is $\sqrt{\lambda^+}$.

Case 3. Midrange of $V(\omega)$: this is defined as the average of minimum and maximum values of $V(\omega)$, therefore, it is given by $(\sqrt{\lambda^+} + \sqrt{\lambda^-})/2$.

If for every image point \mathbf{x} we define the 2D vector

$$\sqrt{\boldsymbol{\lambda}} = \sqrt{\boldsymbol{\lambda}(S_K \mathbf{u}(\mathbf{x}))} = \left(\sqrt{\lambda^+(S_K \mathbf{u}(\mathbf{x}))}, \sqrt{\lambda^-(S_K \mathbf{u}(\mathbf{x}))} \right),$$

we observe that the above measures of image variation correspond (up to appropriate scale factors) to the following norms of $\sqrt{\boldsymbol{\lambda}}$: $\|\sqrt{\boldsymbol{\lambda}}\|_2$ (*Case 1*), $\|\sqrt{\boldsymbol{\lambda}}\|_\infty$ (*Case 2*), $\|\sqrt{\boldsymbol{\lambda}}\|_1$ (*Case 3*).

This connection further motivates us to consider the more general case of the ℓ_p norm $\|\sqrt{\boldsymbol{\lambda}}\|_p$ ($p \geq 1$) as a measure of image variation. These norms measure the image variation more coherently and robustly than the gradient magnitude used in TV, as they take into account the variations in its neighborhood. At the same time, they incorporate richer information, since they depend both on the maximum and minimum of the directional variation. In this way, their response is in general better adapted to the image geometry.

The fact that the ℓ_p norms of $\sqrt{\boldsymbol{\lambda}}$ are able to measure the local image variation suggests generalizing the TV (2.3) via replacing $\|\nabla u\|_2$ by $\|\sqrt{\boldsymbol{\lambda}}\|_p$. This leads us to define the following novel class of regularizers, with $\mathbf{u} \in W^{1,2}(\mathbb{R}^2, \mathbb{R}^M)$ and $p \geq 1$:

$$(2.10) \quad \text{STV}_p(\mathbf{u}) = \int_{\mathbb{R}^2} \left\| \left(\sqrt{\lambda^+}, \sqrt{\lambda^-} \right) \right\|_p d\mathbf{x}.$$

We call the resulting family of regularizers *structure tensor TV* (STV), due to its dependence on the structure tensor of the image and to the fact that it generalizes TV.

Next, we provide a result about the properties that the proposed family of regularizers satisfies. Specifically, the following theorem, whose proof is given in Appendix A.2, verifies that all the favorable properties of TV are indeed preserved by STV.

Theorem 2.1. *The STV functional, $\text{STV}_p(\mathbf{u})$, is translation and rotation invariant, 1-homogeneous, and convex.*

Connections to TV-based functionals. Several existing variational methods emerge as special cases of the proposed regularizers when $K(\mathbf{x})$ is chosen to be the Dirac delta $\delta(\mathbf{x})$ (degenerated case where no convolution takes place at the computation of the structure tensor). Specifically,

if we consider grayscale images ($M = 1$) and $K(\mathbf{x}) = \delta(\mathbf{x})$, then all the regularizers of the type (2.10) are equivalent to TV. The reason is that in this case, $\lambda^+ = \|\nabla u\|_2^2$ and λ^- is always 0, therefore $\|\sqrt{\lambda}\|_p = \|\nabla u\|_2$ for any $p \geq 1$. Furthermore, if we consider vector-valued images ($M > 1$) and $K(\mathbf{x}) = \delta(\mathbf{x})$, then the choice of $p = 2$ corresponds to the VTV, while the choice of $p = \infty$ corresponds to the regularizer of [30], which the authors call *natural vectorial TV* (TV_J). Finally, by choosing $K(\mathbf{x}) = \delta(\mathbf{x})$ and $p = 1$ we obtain a novel vectorial extension of TV that we initially introduced in [36]. This regularizer penalizes the nuclear norm of the Jacobian. We denote this regularizer as STV₁(loc), since it corresponds to a completely localized version of STV₁, where the corresponding structure tensor is formed by only combining the information from the different image channels without any spatial smoothing.

3. Discrete STV. In practice, for all inverse problems of interest we need to deal with discrete data. Therefore, from now on we solely focus on the discrete problem formulation of (2.1) and its manipulation. In this case, \mathbf{A} will denote the system matrix that models the spatial response of the imaging device, while \mathbf{v} and \mathbf{u} will denote the vectorized versions of the observed image and the image to be estimated, respectively.

3.1. Notation and definitions. Hereafter, lower-case bold-faced letters are used for vectors while upper-case bold-faced letters are reserved for matrices and multidimensional arrays. The set of unitary matrices is denoted as $\mathbb{U}^N = \{\mathbf{X} \in \mathbb{C}^{N \times N} : \mathbf{X}^{-1} = \mathbf{X}^H\}$, where \mathbb{C} is the set of complex numbers and $(\cdot)^H$ is the Hermitian transpose. Further, the set of positive semidefinite diagonal matrices is denoted as $\mathbb{D}^{N_1 \times N_2} = \{\mathbf{X} \in \mathbb{R}_+^{N_1 \times N_2} : \mathbf{X}(i, j) = 0 \forall i \neq j\}$.

We assume that the discretized vector-valued image $\mathbf{u} = (\mathbf{u}_1, \dots, \mathbf{u}_M)$ is defined on a rectangular grid with unary steps and contains N pixels. Each channel m of \mathbf{u} ($m = 1, \dots, M$) is rasterized in the vector $\mathbf{u}_m \in \mathbb{R}^N$ and by combining all the image channels, we have that $\mathbf{u} \in \mathbb{R}^{NM}$. We use the index $n = 1, \dots, N$ to refer to a specific pixel of the grid and we denote by \mathbf{x}_n its coordinates. Further, the convolution kernel K is also discretized and truncated in order to have compact support $\mathcal{P} = \{-L_K, \dots, L_K\}^2$, where L_K is a nonnegative integer much smaller than the dimensions of the discrete image. We denote by $[\mathbf{S}_K \mathbf{u}]_n$ the discrete structure tensor of \mathbf{u} evaluated at pixel n , which is defined by adopting a discrete gradient in (2.5) and discrete convolution in (2.7). To properly handle values outside the image boundaries in the definition of this convolution, we assume a type of image extension beyond the boundaries such as mirroring, periodic extension, or zero padding. In our implementation we have used the mirroring boundary extension but the framework that we describe next can easily accommodate all the boundary extensions. Now, we can formally define the discrete STV as follows:

$$(3.1) \quad \text{STV}_p(\mathbf{u}) = \sum_{n=1}^N \left\| \left(\sqrt{\lambda_n^+}, \sqrt{\lambda_n^-} \right) \right\|_p,$$

where $p \geq 1$ and λ_n^+, λ_n^- denote the two nonnegative eigenvalues of the discrete structure tensor $[\mathbf{S}_K \mathbf{u}]_n$ at pixel n .

Finally, the following definition will be of help in what follows.

Definition 3.1 (Schatten norms). Let $\mathbf{X} \in \mathbb{C}^{N_1 \times N_2}$ be a matrix with the singular-value decomposition (SVD) $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^H$, where $\mathbf{U} \in \mathbb{U}^{N_1}$ and $\mathbf{V} \in \mathbb{U}^{N_2}$ consist of the singular vectors

of \mathbf{X} , and $\Sigma \in \mathbb{D}^{N_1 \times N_2}$ consists of the singular values of \mathbf{X} . The Schatten norm of order p (\mathcal{S}_p norm) of \mathbf{X} , is defined as

$$(3.2) \quad \|\mathbf{X}\|_{\mathcal{S}_p} = \left(\sum_{n=1}^{\min(N_1, N_2)} \sigma_n^p \right)^{\frac{1}{p}},$$

where $p \geq 1$, and σ_n is the n th singular value of \mathbf{X} , which corresponds to the (n, n) entry of Σ .

In words, the Schatten matrix norm of order p corresponds to the ℓ_p norm of the vector that contains the singular values of the matrix. Three of the most popular matrix norms of this family are the \mathcal{S}_1 (nuclear), \mathcal{S}_2 (Frobenius), and \mathcal{S}_∞ (spectral) norms, corresponding to the ℓ_1 , ℓ_2 , and ℓ_∞ norms of the matrix singular values, respectively.

3.2. Patch-based Jacobian operator. As described earlier, the proposed regularizers depend on the eigenvalues of the structure tensor. Their current form is challenging to work with, mainly because of the nonlinearity of the operator and the presence of the convolution kernel K . These two features pose significant difficulties in the efficient minimization of our functionals.

To overcome these difficulties, in this section we introduce an alternative formulation of the proposed regularizers. This equivalent formulation depends on a novel generalization of the Jacobian of an image based on local weighted patches. This new operator, which we call *patch-based Jacobian*, contains weighted shifted versions of the Jacobian of \mathbf{u} , whose weights are determined by the smoothing kernel K . As we will show next, the alternative definition of STV provides more intuition and paves the way for employing tools from convex optimization, which in their turn facilitate the development of an efficient optimization strategy.

We define the *patch-based Jacobian* of an image \mathbf{u} as the linear mapping $\mathbf{J}_K : \mathbb{R}^{NM} \mapsto \mathcal{X}$, where $\mathcal{X} \triangleq \mathbb{R}^{N \times (LM) \times 2}$ and $L = (2L_K + 1)^2$. This definition implies that if we apply the patch-based Jacobian on the n th pixel of \mathbf{u} , then the result is a matrix of size $(LM \times 2)$, which we denote by $[\mathbf{J}_K \mathbf{u}]_n$. This matrix is constructed by (1) evaluating the discrete versions of the $M \times 2$ Jacobian matrices of \mathbf{u} for all the pixels $\{\mathbf{x}_n - \mathbf{s} : \mathbf{s} \in \mathcal{P}\}$ in the \mathcal{P} -neighborhood of the pixel \mathbf{x}_n , (2) weighting these matrices with the window function $\omega[\mathbf{s}] = \sqrt{K[\mathbf{s}]}$, and (3) stacking all of them vertically to form the final matrix. Formally, the patch-based Jacobian can be defined as

$$(3.3) \quad [\mathbf{J}_K \mathbf{u}]_n = \left([\tilde{\nabla} \mathbf{u}_1]_n, \dots, [\tilde{\nabla} \mathbf{u}_M]_n \right)^T,$$

where

$$(3.4) \quad [\tilde{\nabla} \mathbf{u}_m]_n = \left([T_{\mathbf{s}_1, \omega} \circ \nabla \mathbf{u}_m]_n, \dots, [T_{\mathbf{s}_L, \omega} \circ \nabla \mathbf{u}_m]_n \right),$$

∇ is the discrete gradient, $(\cdot)^T$ is the transpose operator, \circ denotes the composition of operators, the shift vectors \mathbf{s}_l ($l = 1, \dots, L$) are the elements of the lattice \mathcal{P} , and $T_{\mathbf{s}_l, \omega}$ is a weighted translation operator. The latter takes into account the mirror boundary conditions and is defined as

$$(3.5) \quad [T_{\mathbf{s}_l, \omega} \circ \nabla \mathbf{u}_m]_n = \omega[\mathbf{s}_l] \nabla \mathbf{u}_m[\mathbf{x}_n - \mathbf{s}_l].$$

Next, we equip the space \mathcal{X} , which is the target space of \mathbf{J}_K , with the inner product $\langle \cdot, \cdot \rangle_{\mathcal{X}}$ and norm $\|\cdot\|_{\mathcal{X}}$. To define them, let $\mathbf{X}, \mathbf{Y} \in \mathcal{X}$, with $\mathbf{X}_n, \mathbf{Y}_n \in \mathbb{R}^{(LM) \times 2} \forall n = 1, 2, \dots, N$. Then, we have

$$(3.6) \quad \langle \mathbf{X}, \mathbf{Y} \rangle_{\mathcal{X}} = \sum_{n=1}^N \text{trace}(\mathbf{Y}_n^T \mathbf{X}_n)$$

and

$$(3.7) \quad \|\mathbf{X}\|_{\mathcal{X}} = \sqrt{\langle \mathbf{X}, \mathbf{X} \rangle_{\mathcal{X}}} = \left(\sum_{n=1}^N \|\mathbf{X}_n\|_F^2 \right)^{\frac{1}{2}},$$

where $\text{trace}(\cdot)$ is the trace operator of a matrix and $\|\cdot\|_F$ is the Frobenius (\mathcal{S}_2) norm. For the Euclidean space \mathbb{R}^{NM} we use the standard inner product $\langle \cdot, \cdot \rangle_2$ and norm $\|\cdot\|_2$.

Since the patch-based Jacobian is a bounded linear operator, we can further define its adjoint, which performs a reverse mapping from the space \mathcal{X} to \mathbb{R}^{NM} . Indeed, the adjoint of the patch-based Jacobian is the discrete linear operator $\mathbf{J}_K^* : \mathcal{X} \mapsto \mathbb{R}^{NM}$ such that

$$(3.8) \quad \langle \mathbf{Y}, \mathbf{J}_K \mathbf{u} \rangle_{\mathcal{X}} = \langle \mathbf{J}_K^* \mathbf{Y}, \mathbf{u} \rangle_2$$

should hold for every $\mathbf{Y} \in \mathcal{X}$. The following proposition expresses \mathbf{J}_K^* in a more suitable form that facilitates its numerical computation. The proof is provided in Appendix A.3.

Proposition 3.2. *The adjoint operator \mathbf{J}_K^* of the patch-based Jacobian is given by*

$$(3.9) \quad [\mathbf{J}_K^* \mathbf{Y}]_k = \sum_{l=1}^L \left[-\mathbf{div} \left(T_{s_l, \omega}^* \circ \mathbf{Y}^{((m-1)L+l, :)} \right) \right]_n,$$

where $k = (m-1)N + n$ with $1 \leq n \leq N$ and $1 \leq m \leq M$, \mathbf{div} is the discrete divergence,² $T_{s_l, \omega}^*$ is the adjoint of the weighted translation operator, and $\mathbf{Y}_n^{(i, :)}$ is a two-element vector extracted from the i th row of the n th matrix component, $\mathbf{Y}_n \in \mathbb{R}^{(LM) \times 2}$, of \mathbf{Y} .

3.3. Equivalent formulation of the discrete STV. Having introduced the discrete patch-based Jacobian, we can now employ it to express the discrete structure tensor in a novel way. This is done in the next proposition.

Proposition 3.3. *The discrete structure tensor of \mathbf{u} evaluated at the pixel location n can be written in terms of the patch-based Jacobian as*

$$(3.10) \quad [\mathbf{S}_K \mathbf{u}]_n = [\mathbf{J}_K \mathbf{u}]_n^T [\mathbf{J}_K \mathbf{u}]_n.$$

The proof of Proposition 3.3 can be derived by a direct computation of the right-hand side (r.h.s.) of (3.10) and, thus, is omitted. However, it is worth mentioning that by exploiting

²The exact formula for the discrete divergence depends on the discretization scheme one uses for the gradient. In our implementation we consider the discretization that uses forward differences, as in [12]. In this case, the adjoint operator is the discrete divergence that is defined using backward differences (see [12] for details).

the special structure of the patch-based Jacobian we manage to reproduce the convolution operation that takes place in the formulation of the structure tensor.

Now, from (3.10) we have that the 2×2 matrix $[\mathbf{S}_K \mathbf{u}]_n$ can be decomposed with the help of the $LM \times 2$ matrix $[\mathbf{J}_K \mathbf{u}]_n$. Therefore, the singular values of $[\mathbf{J}_K \mathbf{u}]_n$ are equal to $\sqrt{\lambda_n^+}$ and $\sqrt{\lambda_n^-}$. This connection highlights that by using the \mathcal{S}_p matrix norms given in (3.2), we can equivalently write the discrete version of the proposed regularizers (3.1) as

$$(3.11) \quad \text{STV}_p(\mathbf{u}) = \sum_{n=1}^N \|[\mathbf{J}_K \mathbf{u}]_n\|_{\mathcal{S}_p}$$

with $p \geq 1$. This equivalent formulation of $\text{STV}_p(\mathbf{u})$ in (3.11) provides more intuition about the fact that the proposed regularizers are effective generalizations of TV. In particular, $[\mathbf{J}_K \mathbf{u}]_n$ encodes the vectorial variation of the image \mathbf{u} in the vicinity of the pixel n . Therefore, the Schatten norms of this matrix provide different measures of the local variation of \mathbf{u} , by taking into account its neighborhood in a weighted manner. In addition, a significant advantage that we gain from the above result, is that the expression (3.1), which involves the eigenvalues of the nonlinear structure tensor, has been transformed into the expression (3.11). The latter expression is much easier to handle in an optimization setting while it also allows us to derive the following result.

Proposition 3.4. *The discrete $\text{STV}_p(\mathbf{u})$ regularizer defined in (3.1) is convex w.r.t. $\mathbf{u} \forall p \geq 1$.*

The validity of the above proposition stems from the fact that the discrete STV can be formed as the composition of a mixed vector/matrix norm ($\ell_1\text{-}\mathcal{S}_p$) and the linear operator \mathbf{J}_K . Therefore, our discrete regularizers retain the important property of convexity that we established with Theorem 2.1 for their continuous counterparts.

4. Discrete STV minimization.

4.1. Proximal map of STV_p . Having introduced the discrete version of the STV functionals, the next step is to develop a scheme for their efficient minimization. In this section, we focus on the evaluation of the proximal map associated with them. The proximal map or the *Moreau proximity operator* of a function ψ is denoted as $\text{prox}_\psi(\cdot) : \mathbb{R}^N \mapsto \mathbb{R}^N$ and is formally defined as [19]

$$(4.1) \quad \text{prox}_\psi(\mathbf{z}) = \arg \min_{\mathbf{u}} \frac{1}{2} \|\mathbf{u} - \mathbf{z}\|_2^2 + \psi(\mathbf{u}).$$

If ψ is a proper, closed, convex function, then a minimizer of (4.1) exists and the solution is unique. For certain choices of the function ψ the solution of (4.1) can be expressed in closed form. However, in several other cases the solution cannot be obtained explicitly, but instead an iterative approach must be followed. This is also the case for our proposed regularizers.

In our case, $\psi(\mathbf{u}) = \tau \text{STV}_p(\mathbf{u}) + \iota_{\mathcal{C}}$ and the evaluation of the proximal map associated with the STV regularizers corresponds to finding a numerical solution to the following problem, for any $p \geq 1$:

$$(4.2) \quad \arg \min_{\mathbf{u}} \frac{1}{2} \|\mathbf{u} - \mathbf{z}\|_2^2 + \tau \text{STV}_p(\mathbf{u}) + \iota_{\mathcal{C}}(\mathbf{u}).$$

In the above formulation, \mathcal{C} is a convex set that represents additional constraints on the solution, such as nonnegativity or box constraints, and $\iota_{\mathcal{C}}$ is its indicator function: $\iota_{\mathcal{C}}(\mathbf{u})$ takes the value 0 for $\mathbf{u} \in \mathcal{C}$ and ∞ otherwise. The unconstrained minimization case simply corresponds to choosing $\mathcal{C} = \mathbb{R}^N$. Note that from a signal processing standpoint the solution of (4.2) can also be interpreted as a denoising problem under independently and identically distributed (i.i.d.) Gaussian noise, where the STV functionals are used to regularize the solution.

The reason for focusing on the solution of (4.2) is that the proximal map of the STV penalties can serve as the building block for several optimization algorithms, including the fast iterative shrinkage-thresholding algorithm (FISTA) [4], the alternating direction method of multipliers [1, 8, 25], the split-Bregman algorithm [31], and the primal-dual algorithms introduced in [14, 43]. In their turn, these methods are commonly employed to solve a variety of inverse problems in imaging with different data fidelity terms. Therefore, by being able to accurately and efficiently obtain the solution of (4.2), one can also deal with more general inverse problems than merely Gaussian denoising.

Next, we propose a first-order method for solving the problem under study. To proceed, we write our regularizers in the compact form $\text{STV}_p(\mathbf{u}) = \|\mathbf{J}_K \mathbf{u}\|_{1,p}$. Here, $\|\cdot\|_{1,p}$ corresponds to the mixed ℓ_1 - \mathcal{S}_p norm which, for an argument $\mathbf{X} = [\mathbf{X}_1^T, \dots, \mathbf{X}_N^T]^T \in \mathcal{X}$, is defined as

$$(4.3) \quad \|\mathbf{X}\|_{1,p} = \sum_{n=1}^N \|\mathbf{X}_n\|_{\mathcal{S}_p}.$$

At this point special attention must be paid to the nonsmoothness of the regularizers. The fact that the proposed STV functionals are not differentiable everywhere implies that the minimizer of (4.2) cannot be obtained simply by employing a gradient-based scheme. To circumvent this difficulty, we follow an alternative strategy which involves the derivation of a dual formulation for our problem. To do so, we make use of Lemma 4.1 that follows.

Lemma 4.1 (see [37]). *Let $p \geq 1$, and let q be the conjugate exponent of p , i.e., $\frac{1}{p} + \frac{1}{q} = 1$. Then, the mixed vector-matrix norm $\|\cdot\|_{\infty,q}$ is dual to the mixed vector-matrix norm $\|\cdot\|_{1,p}$.*

Using Lemma 4.1 and the fact that the dual of the dual norm is the original norm [45], we equivalently write (4.3) as

$$(4.4) \quad \|\mathbf{X}\|_{1,p} = \max_{\mathbf{\Omega} \in \mathcal{B}_{\infty,q}} \langle \mathbf{\Omega}, \mathbf{X} \rangle_{\mathcal{X}},$$

where $\mathcal{B}_{\infty,q}$ denotes the ℓ_{∞} - \mathcal{S}_q unit-norm ball, defined as

$$(4.5) \quad \mathcal{B}_{\infty,q} = \{\mathbf{\Omega} \in \mathcal{X} : \|\mathbf{\Omega}_n\|_{\mathcal{S}_q} \leq 1, \forall n = 1, \dots, N\}.$$

From the definition of the convex set $\mathcal{B}_{\infty,q}$, it turns out that the orthogonal projection of $\mathbf{\Omega}$ can be performed by projecting separately each submatrix $\mathbf{\Omega}_n$ on a \mathcal{S}_q unit-norm ball ($\mathcal{B}_{\mathcal{S}_q}$). This property will be used later for designing our first-order minimization method.

Now, combining (3.8) and (4.4) we rewrite (4.2) as

$$(4.6) \quad \hat{\mathbf{u}} = \arg \min_{\mathbf{u} \in \mathcal{C}} \frac{1}{2} \|\mathbf{u} - \mathbf{z}\|_2^2 + \tau \max_{\mathbf{\Omega} \in \mathcal{B}_{\infty,q}} \langle \mathbf{J}_K^* \mathbf{\Omega}, \mathbf{u} \rangle_2.$$

This formulation naturally leads us to the minimax problem

$$(4.7) \quad \min_{\mathbf{u} \in \mathcal{C}} \max_{\Omega \in \mathcal{B}_{\infty, q}} \mathcal{L}(\mathbf{u}, \Omega),$$

where $\mathcal{L}(\mathbf{u}, \Omega) = \frac{1}{2} \|\mathbf{u} - \mathbf{z}\|_2^2 + \tau \langle \mathbf{J}_K^* \Omega, \mathbf{u} \rangle_2$.

Since \mathcal{L} is strictly convex w.r.t. \mathbf{u} and concave w.r.t. Ω , we have the guarantee that a saddle value of \mathcal{L} is attained [45] and, thus, the order of the minimum and the maximum in (4.7) does not affect the solution. This means that there exists a common saddle point $(\hat{\mathbf{u}}, \hat{\Omega})$ when the minimum and the maximum are interchanged, i.e.,

$$(4.8) \quad \min_{\mathbf{u} \in \mathcal{C}} \max_{\Omega \in \mathcal{B}_{\infty, q}} \mathcal{L}(\mathbf{u}, \Omega) = \mathcal{L}(\hat{\mathbf{u}}, \hat{\Omega}) = \max_{\Omega \in \mathcal{B}_{\infty, q}} \min_{\mathbf{u} \in \mathcal{C}} \mathcal{L}(\mathbf{u}, \Omega).$$

According to the above, it is now possible to define two optimization problems, the primal and the dual one. This can be accomplished by identifying the primal and dual objective functions, respectively. In particular, the left-hand side (l.h.s.) of (4.8) corresponds to the minimization of the primal objective function

$$(4.9) \quad p(\mathbf{u}) = \max_{\Omega \in \mathcal{B}_{\infty, q}} \mathcal{L}(\mathbf{u}, \Omega) = \frac{1}{2} \|\mathbf{u} - \mathbf{z}\|_2^2 + \tau \|\mathbf{J}_K \mathbf{u}\|_{1, p},$$

which matches the problem of interest in (4.2), and the r.h.s. corresponds to the maximization of the dual objective function

$$(4.10) \quad d(\Omega) = \min_{\mathbf{u} \in \mathcal{C}} \mathcal{L}(\mathbf{u}, \Omega) = \frac{1}{2} \|\mathbf{w} - \Pi_C(\mathbf{w})\|_2^2 + \frac{1}{2} \left(\|\mathbf{z}\|_2^2 - \|\mathbf{w}\|_2^2 \right),$$

where $\mathbf{w} = \mathbf{z} - \tau \mathbf{J}_K^* \Omega$ and Π_C is the orthogonal projection operator on the convex set \mathcal{C} . In addition, if we consider the r.h.s. of (4.8), then we can find the minimizer $\hat{\mathbf{u}}$ of the primal objective from the maximizer $\hat{\Omega}$ of the dual objective, through the following closed-form relation

$$(4.11) \quad \hat{\mathbf{u}} = \Pi_C \left(\mathbf{z} - \tau \mathbf{J}_K^* \hat{\Omega} \right).$$

The advantage of this last formulation is that the solution to our problem boils down to computing the maximizer of $d(\Omega)$. As opposed to $p(\mathbf{u})$, this objective function is smooth and has a well-defined gradient. Thus, we can exploit gradient information to derive the solution. This would not have been an option if we had considered the original formulation of the problem in (4.2). Next, we provide the details of our approach.

4.2. Solving the dual problem. From the r.h.s. of (4.8) and from (4.10) it follows that $\hat{\Omega}$ can be derived as the solution of the constrained optimization problem

$$(4.12) \quad \hat{\Omega} = \arg \max_{\Omega \in \mathcal{B}_{\infty, q}} \frac{1}{2} \|\mathbf{w} - \Pi_C(\mathbf{w})\|_2^2 + \frac{1}{2} \left(\|\mathbf{z}\|_2^2 - \|\mathbf{w}\|_2^2 \right).$$

Since (4.12) does not admit a closed-form solution (\mathbf{J}_K does not have an inverse operator) we have to resort to a numerical iterative scheme. In this work, we choose to pursue a first-order

projected gradient approach. To this end, the essential tools that we need are the computation of the gradient and the orthogonal projection on the convex set $\mathcal{B}_{\infty,q}$, defined in (4.5).

First, we proceed by deriving the gradient. To do so, we use that the gradient of a function $h(\mathbf{u}) = \|\mathbf{u} - \Pi_C(\mathbf{u})\|_2^2$ is well-defined and is equal to $\nabla h(\mathbf{u}) = 2(\mathbf{u} - \Pi_C(\mathbf{u}))$ [3, Lemma 4.1]. Based on this, we obtain

$$(4.13) \quad \nabla d(\mathbf{\Omega}) = \tau \mathbf{J}_K \Pi_C(\mathbf{z} - \tau \mathbf{J}_K^* \mathbf{\Omega}).$$

Having the gradient of the dual objective at hand, we also need an efficient way to perform the projection. We accomplish this by projecting independently each of the N components $\mathbf{\Omega}_n$ of $\mathbf{\Omega}$ on the set $\mathcal{B}_{\mathcal{S}_q} = \{\mathbf{X} \in \mathbb{R}^{LM \times 2} : \|\mathbf{X}\|_{\mathcal{S}_q} \leq 1\}$.

Now, let $\mathbf{X} \in \mathbb{R}^{N_1 \times N_2}$ with an SVD $\mathbf{X} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$ and $\mathbf{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_N)$ with $N = \min(N_1, N_2)$. According to [37, Proposition 1], the projection of \mathbf{X} onto the unit-norm \mathcal{S}_q ball is computed as

$$(4.14) \quad \Pi_{\mathcal{S}_q}(\mathbf{X}) = \mathbf{U} \mathbf{\Sigma}_q \mathbf{V}^T,$$

where $\mathbf{\Sigma}_q = \text{diag}(\boldsymbol{\sigma}_q)$ and the $\boldsymbol{\sigma}_q$ are the projected singular values of $\mathbf{\Sigma}$ onto the ℓ_q unit-norm ball $\mathcal{B}_q = \{\boldsymbol{\sigma} \in \mathbb{R}_+^N : \|\boldsymbol{\sigma}\|_q \leq 1\}$. We note that the projection in (4.14) requires knowledge of the singular vectors and singular values of \mathbf{X} . In our case $N_2 = 2 < N_1$, and we compute the projection in an efficient way as described next. First, we observe that the matrix $\mathbf{X}^T \mathbf{X}$ is $N_2 \times N_2$ symmetric with an eigenvalue decomposition $\mathbf{V} \mathbf{\Sigma}^2 \mathbf{V}^T$. Therefore, for $N_2 = 2$ both \mathbf{V} and $\mathbf{\Sigma}$ can be computed in closed form. Moreover, if $\mathbf{\Sigma}^+$ is the pseudoinverse matrix of $\mathbf{\Sigma}$, defined as $\mathbf{\Sigma}^+ = \text{diag}(\sigma_1^{-1}, \dots, \sigma_k^{-1}, 0, \dots, 0)$, with σ_k the smallest nonzero singular value, then $\mathbf{U} = \mathbf{X} \mathbf{V} \mathbf{\Sigma}^+$. Using this result we write (4.14) as

$$(4.15) \quad \Pi_{\mathcal{S}_q}(\mathbf{X}) = \mathbf{X} \mathbf{V} \mathbf{\Sigma}^+ \mathbf{\Sigma}_q \mathbf{V}^T,$$

which actually avoids the computation of \mathbf{U} .

Thanks to (4.15) the projection of \mathbf{X} reduces to the projection of the singular values. In our case, where $N_2 = 2$, for $q = 1, 2, \infty$ these projections can be computed analytically. Indeed, for $q = 1$ the projection corresponds to the soft-thresholding operation $S_\gamma(\boldsymbol{\sigma}) = \max(\boldsymbol{\sigma} - \gamma, 0)$ [23], where the max operator is applied componentwise and γ is given by

$$(4.16) \quad \gamma = \begin{cases} 0 & \text{if } \sigma_1 \leq 1 - \sigma_2, \\ \frac{\sigma_1 + \sigma_2 - 1}{2} & \text{if } 1 - \sigma_2 < \sigma_1 \leq 1 + \sigma_2, \\ \sigma_1 - 1 & \text{if } \sigma_1 > 1 + \sigma_2 \end{cases}$$

with the singular values sorted in decreasing order, i.e., $\sigma_1 \geq \sigma_2$. For $q = 2$ the projection of \mathbf{X} is computed at a reduced computational cost since it does not require the knowledge of the singular values and vectors. In particular, we have $\Pi_{\mathcal{S}_2}(\mathbf{X}) = \mathbf{X} / \max(1, \|\mathbf{X}\|_F)$. Finally, for $q = \infty$ the projection is performed by setting to one the singular values that exceed this value and leaving the rest untouched. For other values of q , we can employ the ℓ_q norm projection algorithm described in [54], but at an increased computational cost. This projection method is based on an iterative proximity algorithm for ℓ_q norms [39].

Algorithm 1 : Proximal algorithm for the STV functionals.

Input: $\mathbf{z}, \tau > 0, p \geq 1, \Pi_{\mathcal{C}}$.

Initialization: $\Psi_1 = \Omega_0 = \mathbf{0} \in \mathcal{X}, t_1 = 1$.

while *stopping criterion is not satisfied* **do**

$$\left\{ \begin{array}{l} \Omega_n \leftarrow \Pi_{\mathcal{B}_{\infty,q}} \left(\Psi_n + \frac{1}{8\sqrt{2}\tau} \mathbf{J}_K \Pi_{\mathcal{C}} (\mathbf{z} - \tau \mathbf{J}_K^* \Psi_n) \right); \\ t_{n+1} \leftarrow \frac{1 + \sqrt{1 + 4t_n^2}}{2}; \\ \Psi_{n+1} \leftarrow \Omega_n + \left(\frac{t_n - 1}{t_{n+1}} \right) (\Omega_n - \Omega_{n-1}); \\ n \leftarrow n + 1; \end{array} \right.$$

end

return $\Pi_{\mathcal{C}} (\mathbf{z} - \tau \mathbf{J}_K^* \Omega_{n-1})$;

4.3. Numerical algorithm. With all the necessary tools at our disposal, we can now invoke a gradient-based scheme for deriving the solution of the dual problem. In this work, we employ Nesterov’s iterative method [40] for smooth functions. This scheme exhibits convergence rates of one order higher than the standard gradient-ascent method. To ensure convergence of the algorithm, we further need to choose an appropriate step size. Since our dual objective is smooth with Lipschitz continuous gradient, we use a constant step size equal to the inverse of the Lipschitz constant of $\nabla d(\Omega)$. We derive an upper bound of this quantity in the following proposition. The proof can be found in Appendix A.4.

Proposition 4.2. *Let $L(d)$ denote the Lipschitz constant of the dual objective function d defined in (4.10). Then,*

$$(4.17) \quad L(d) \leq 8\sqrt{2}\tau^2.$$

To conclude, in Algorithm 1 we provide a detailed description of all the steps involved in the proposed algorithm for computing the proximal associated with the STV functionals.

4.4. Solution of general inverse problems. So far, we have focussed on computing the proximal map of the STV functionals that solves a simple denoising problem. This basic tool, however, can also be deployed for solving more general inverse problems as we discuss next for the case of a quadratic data term. In this case, the solution is derived as

$$(4.18) \quad \hat{\mathbf{u}} = \arg \min_{\mathbf{u} \in \mathcal{C}} \frac{1}{2} \|\mathbf{v} - \mathbf{A}\mathbf{u}\|_2^2 + \tau \text{STV}_p(\mathbf{u}) \quad \forall p \geq 1.$$

We note that it is possible, by using appropriate optimization tools, for one to also deal with nonquadratic data terms. The description of such techniques, however, exceeds the scope of this paper.

An efficient way to cope with the presence of the operator \mathbf{A} in (4.18) is to employ the monotone FISTA (MFISTA) [3] algorithm, which exhibits state-of-the-art convergence rates. The basic idea behind MFISTA is to obtain the minimizer in (4.18) via the successive minimization of a sequence of surrogate functions that upper bound the initial objective function and are easier to minimize. In this context, the solution of (4.18) boils down to iteratively computing the proximal of STV (4.2) for different inputs, which are formed as a combination of solutions from the past iterations of the algorithm. Our main reason for preferring MFISTA

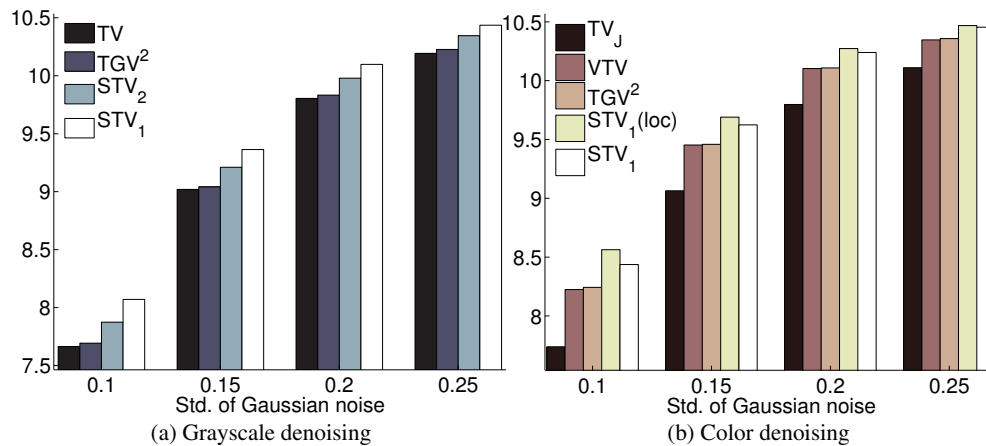


Figure 3. Image denoising comparisons among different regularizers for four noise levels. The performance is measured in terms of the average SNR improvement (in dB) w.r.t. the noisy input, over all 30 images of Figure 2.

For the objective function minimization we employ the MFISTA framework [3] for all methods but TGV. For TGV, we have experimentally observed that this approach does not work well. Therefore, we employ instead the method proposed by its authors in [9]. Finally, for the sake of consistency among comparisons, all the reported results for each regularizer and each degradation condition are derived using the regularization parameter τ that leads to the best reconstruction performance in terms of peak signal to noise ratio (PSNR).

5.2. Image restoration. For the *image denoising* problem we consider additive i.i.d. Gaussian noise and four different noise levels that correspond to a standard deviation of $\sigma_w = \{0.1, 0.15, 0.2, 0.25\}$, respectively. Regarding the stopping criterion for the optimization algorithm of all methods but TGV, this is set to either reaching a relative normed difference of 10^{-5} between two successive estimates, or a maximum of 100 iterations for solving the dual problem, as described in Algorithm 1. For TGV² where we employ the algorithm of [9] we use a maximum of 200 iterations.

In Figure 3 we report the average performance over all tested images of all the methods under comparison. The performance of the regularizers is measured in terms of the average PSNR improvement (ISNR) w.r.t. the noisy input. From these results we observe that, in the grayscale experiments, TV is the worst performing method for all the noise levels. TGV yields an improvement over TV, which reflects the fact that it favors piecewise linear rather than piecewise constant solutions and, thus, it avoids the staircasing artifacts of TV. Both versions of our regularizers, STV₂ and STV₁, systematically outperform TV and TGV². This can be attributed to the increased robustness that the neighborhood information offers. Our best performing method is STV₁ which involves the nuclear norm of the patch-based Jacobian. In the color denoising results of Figure 3, we see that TV_J results in the worst overall performance. This can be explained by the fact that this method penalizes only the maximum singular value of the Jacobian, and thus, it completely ignores the information provided by the minimum singular value. VTV yields better measures than TV_J, since it explores additional



Figure 4. Image denoising examples: close-ups of inputs and optimum results. Top row: grayscale denoising of an input with noise level $\sigma_w = 0.15$. The ground truth corresponds to the grayscale version of the image in row 1, column 4 of Figure 2. Bottom row: color denoising of an input with noise level $\sigma_w = 0.2$. The ground truth corresponds to the image in row 3, column 7 of Figure 2. For each result, the individualized regularization parameter τ and corresponding optimum PSNR are reported.

information about the image variations by imposing an ℓ_2 penalty on the singular values of the Jacobian. Therefore, VTV transfers the characteristics of TV to the vectorial case more effectively than TV_J. Similarly to the improvement of TGV² over TV in the grayscale case, TGV² yields an improvement over VTV in the color experiments. Again, both of the tested versions of the proposed regularizers, STV₁ and STV₁(loc), result in the best error measures for all noise levels. In this case, the completely localized version STV₁(loc) performs better than STV₁, especially in the case of relatively low noise levels. A possible explanation for this outcome is that, in the case of denoising, the coupling of the information from the different channels is more important than the additional spatial information obtained from the neighborhood of each pixel.

For a visual assessment of the reconstruction performance, in Figure 4 we provide representative examples of grayscale and color denoising. In the grayscale example, we observe that the TV's result suffers from strong block artifacts. TGV² manages to reduce these artifacts but cannot completely avoid them while it also excessively blurs some edges. STV₂ and STV₁ succeed in removing the noise and at the same time reconstructing the image in a visually more appealing way than TV and TGV². Comparing all methods we see that STV₁ yields the most effective reconstruction of image structures; compare, e.g., the reconstructions of the numbers in the athlete's budge, the edges and interior of the athlete's left hand. In the color example, we observe that TV_J creates spurious one-dimensional (1D) segments perpendicular to the dominant edges of the image. Staircasing artifacts are also present in the VTV reconstruction. TGV² yields a smoother result, but some of the reconstructed structures are blurred; see, e.g., the edges of the bench back. Finally, STV₁ leads to the best result by avoiding most of the artifacts of the other methods. It sufficiently removes the noise and, in parallel, preserves the sharpness and shape of image edges in the most effective way.

For the *image deblurring* problem the performance of the methods under comparison is assessed for various blurring kernels and different noise levels. In particular, in our set-up we

employ three *point spread functions* (PSFs) to generate blurred versions of the test images. We use a Gaussian PSF of standard deviation $\sigma_b = 4$ and support of 13×13 pixels, a 9×9 uniform (moving-average) PSF, and a real motion blurring kernel³ which has a support of 19×19 pixels. Additional distortion on the measurements is introduced by adding i.i.d. Gaussian noise of four different noise levels, corresponding to a standard deviation of $\sigma_w = \{0.025, 0.05, 0.075, 0.1\}$. Further, to test the deblurring performance of all the studied methods under more realistic conditions, we do not assume the common circulant convolution model, but instead we adopt the approach followed in [44]. Specifically, in the forward model in (2.1) the degradation operator is expressed as $\mathbf{A} = \mathbf{M}\mathbf{H}$, where \mathbf{H} is a circulant convolution operator and \mathbf{M} is a mask that keeps only the valid part of the convolution and truncates the circular wraparound at the boundaries of the measurements. The motivation here is that this observation model allows us to avoid making any assumptions about the image boundaries. Instead, we hold the deblurring algorithm responsible for reconstructing the image in a way that best explains the data. As far as the details of the optimization task are concerned, the stopping criterion is set to the same relative normed difference as previously or a maximum of 100 MFISTA iterations with 10 internal iterations for the proximal evaluation. For TGV², since the algorithm of [9] exhibits a slower convergence than MFISTA, we use a maximum of 2000 iterations.

Similarly to the image denoising case, in Figure 5 we report the average performance of all the tested methods. Once more we verify that for grayscale image deblurring the STV regularizers lead to the best results for all considered degradation conditions. When we consider color images, this is still valid with the STV₁ regularizer achieving systematically the best ISNR scores. It also worths mentioning that in this case, unlikely to color image denoising, STV₁(loc) does not perform better than its patch-based counterpart. This can be attributed to the fact that image deblurring is a more ill-conditioned problem and, thus, the use of a smoothing kernel plays a more important role in the reconstruction. Beyond the ISNR comparisons, the effectiveness of the proposed methods can also be visually appreciated by inspecting the representative deblurring examples provided in Figure 6.

5.3. Image magnification. Image magnification is an inverse problem that is closely related to image deblurring. In particular, the system matrix is expressed as $\mathbf{A} = \mathbf{S}\mathbf{G}$, where \mathbf{S} is a masking operator performing the subsampling, while \mathbf{G} is an antialiasing smoothing operator applied to the underlying image, prior to subsampling. In our experiments, we consider square subsamplings of the type $d \times d$, where d is the zoom factor and we use two different values for the zoom factor, $d = 3$ and $d = 5$ pixels. Also, we consider the antialiasing operator to be a Gaussian kernel with a standard deviation $0.35d$, defined relatively to the zoom factor. Further, we introduce Gaussian noise in the measurements of a standard deviation $\sigma_w = 0.01$. In Figure 7 we provide measurements of the overall performance of all the methods, both for grayscale and color images, while in Figure 8 we present two magnification examples where we juxtapose the results obtained by using different regularization approaches. Similarly to the other two inverse problems we considered so far, once more we verify that STV₁ not only improves the SNR but also leads to reconstructions of enhanced visual quality.

³The blurring kernel was obtained from <http://www.wisdom.weizmann.ac.il/~levina/papers/LevinEtalCVPR09Data.rar>.

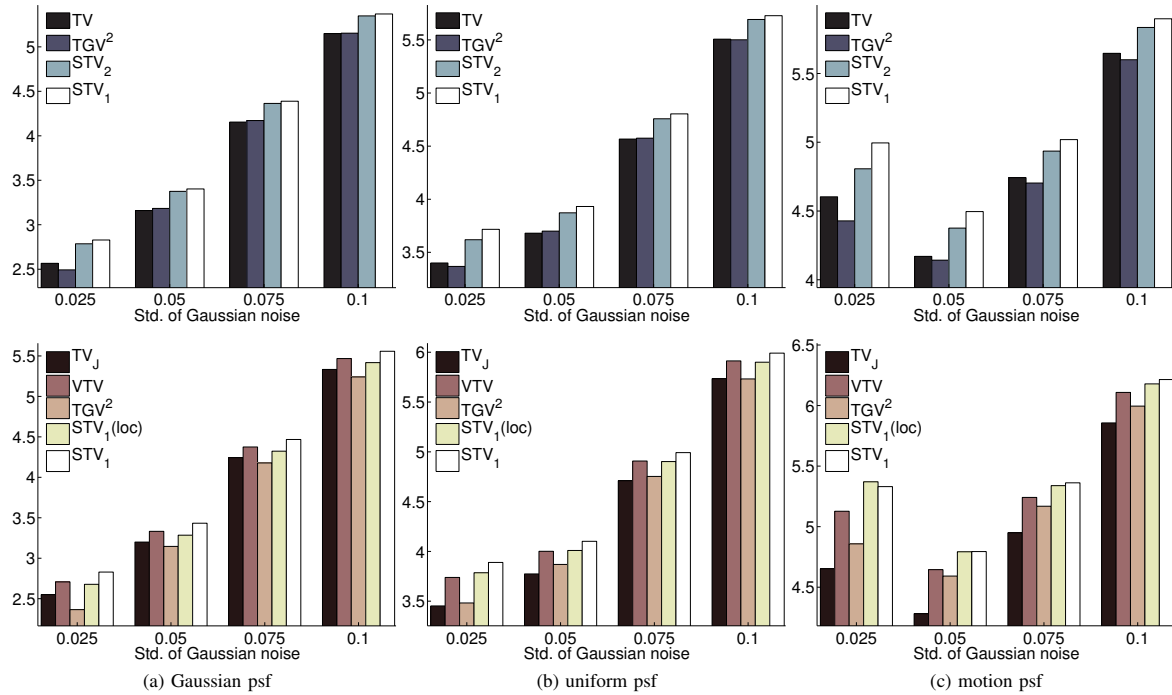


Figure 5. Image deblurring comparisons among different regularizers for three PSFs and four noise levels. The performance is measured in terms of the average PSNR improvement (in dB) w.r.t. the degraded input, over all 30 images of Figure 2. The bars on the top row refer to grayscale results while on the bottom row to color results.



Figure 6. Image deblurring examples: close-ups of inputs and optimum results. Top row: grayscale deblurring of an input with motion blurring and noise level $\sigma_w = 0.025$. The ground truth corresponds to the grayscale version of the image in row 3, column 4 of Figure 2. Bottom row: color deblurring of an input with uniform blurring and noise level $\sigma_w = 0.025$. The ground truth corresponds to the image in row 3, column 6 of Figure 2. For each result, the individualized regularization parameter τ and corresponding optimum PSNR are reported.

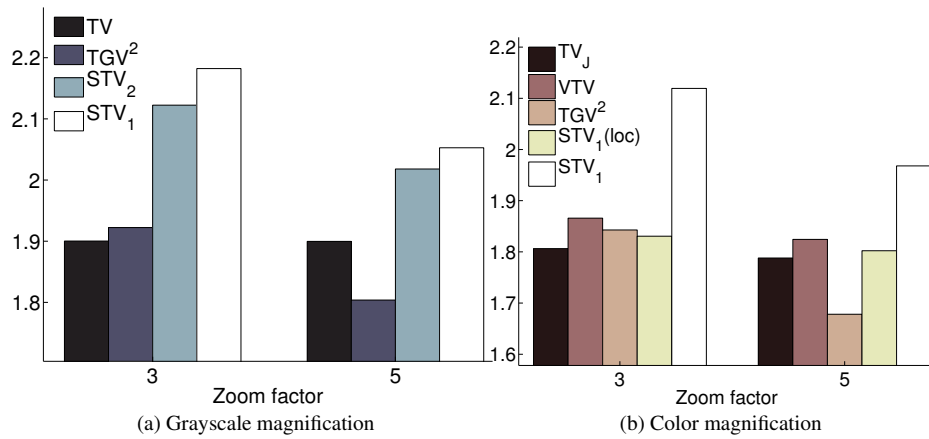


Figure 7. Image magnification comparisons among different regularizers for two zoom factors. The performance is measured in terms of the average SNR improvement (in dB) w.r.t. the input, over all 30 images of Figure 2.

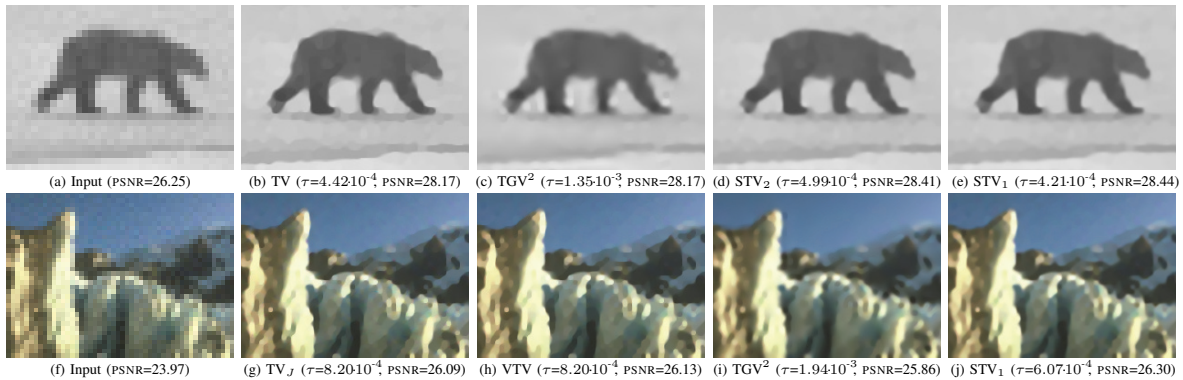


Figure 8. Image magnification examples: close-ups of inputs and optimum results. Inputs are enlarged by simple zero order hold. Top row: grayscale magnification of an input with a zoom factor $d = 5$. The ground truth corresponds to the grayscale version of the image in row 1, column 1 of Figure 2. Bottom row: color magnification of an input with a zoom factor $d = 5$. The ground truth corresponds to the image in row 1, column 3 of Figure 2. For each result, the individualized regularization parameter τ and corresponding optimum PSNR are reported.

5.4. Reconstruction from sparse Fourier measurements. In this section we consider the problem of image reconstruction from a limited number of Fourier measurements. In this setting, the forward operator of (2.1) corresponds to $\mathbf{A} = \mathbf{M}\mathbf{F}$, where \mathbf{F} denotes the Fourier transform and \mathbf{M} is a masking operator that retains only a subset of the available Fourier coefficients. The adopted forward model is closely related to the one of magnetic resonance imaging acquisition, where the k-space of the image is undersampled.

The reported experiments are conducted on the images shown in Figure 9. These are extracted slices from a 3D brain phantom that was obtained from the BrainWeb database (<http://www.bic.mni.mcgill.ca/brainweb/>). To create the measured data we use two different Fourier sampling patterns, namely, a radial mask with 32, 48, and 64 radial lines and a




Figure 9. Sparse Fourier image reconstruction setup. (a)–(c) Slices of a 3D brain phantom, (d) Poisson disk with 20% sampling density (white pixel values indicate sample locations), (e) radial sampling mask with 32 lines.

Table 1

ISNR Comparisons on sparse Fourier image reconstruction for several sampling patterns and noise conditions.

Sampling	Poisson disk – 20%				Poisson disk – 30%				Poisson disk – 40%				Radial – 32 lines				Radial – 48 lines				Radial – 64 lines				
	SNR	10dB	20dB	30dB	∞ dB	10dB	20dB	30dB	∞ dB	10dB	20dB	30dB	∞ dB	10dB	20dB	30dB	∞ dB	10dB	20dB	30dB	∞ dB	10dB	20dB	30dB	∞ dB
Slice 1	TV	2.14	1.92	2.65	2.86	3.27	2.89	3.97	4.62	4.44	3.90	5.79	7.43	3.00	4.15	5.42	5.72	3.30	3.99	6.27	7.23	3.95	3.34	5.91	7.97
	TGV ²	2.74	2.35	2.95	3.22	3.76	3.20	4.23	4.83	4.79	4.16	5.83	7.03	3.48	4.68	5.89	6.25	3.71	4.47	6.77	7.74	4.36	3.82	6.51	8.34
	STV ₁	2.99	2.68	3.29	3.73	4.14	3.74	4.74	5.67	5.34	4.84	6.65	8.82	4.10	5.43	6.90	7.50	4.48	5.55	8.21	10.04	5.15	5.01	8.27	11.88
Slice 2	TV	2.22	1.84	2.51	2.67	3.33	2.66	3.77	4.28	4.28	3.29	5.23	6.43	2.58	3.48	5.02	4.99	3.07	3.25	5.22	6.16	3.83	2.66	5.03	6.94
	TGV ²	2.74	2.30	2.90	3.10	3.67	2.92	3.95	4.49	4.55	3.47	5.05	6.13	3.06	3.86	4.98	5.32	3.42	3.71	5.76	6.60	4.20	3.10	5.56	7.39
	STV ₁	2.99	2.52	3.23	3.53	4.07	3.35	4.40	5.31	5.05	4.08	5.76	7.66	3.60	4.62	6.12	6.64	4.18	4.71	7.19	9.06	4.97	4.29	7.35	10.88
Slice 3	TV	2.48	2.46	3.48	3.91	3.75	3.63	5.38	6.59	4.90	4.47	6.83	9.65	3.45	5.01	6.66	7.20	3.67	3.91	6.51	8.11	4.47	3.49	6.20	9.00
	TGV ²	3.29	3.16	3.92	4.31	4.41	3.94	5.37	6.39	5.40	4.58	6.59	8.55	3.98	5.49	7.12	7.66	4.20	4.41	7.01	8.49	4.92	3.85	6.72	9.34
	STV ₁	3.43	3.38	4.26	4.92	4.71	4.55	6.18	7.87	5.88	5.44	7.74	11.34	4.56	6.38	8.31	9.45	4.86	5.49	8.56	11.60	5.65	5.19	8.61	13.84

Poisson disk that retains 20%, 30%, and 40% of the Fourier coefficients. As an additional degradation factor we consider the presence of complex Gaussian noise in the Fourier domain of four different levels. These correspond to a signal-to-noise-ratio (SNR) of $\{10, 20, 30, \infty\}$ dB, respectively. Note that the last SNR value indicates the absence of noise in the measurements. In this case, for all regularizers under comparison, we do not fine tune the regularization parameter but instead we set it to $\tau = 10^{-6}$. The value of τ is chosen to be small so as to ensure that the reconstructed image in the sampled locations will retain the same frequency content as that of the measurements. Since the very small value of τ can affect the convergence speed of the minimization process, we adopt a simple “continuation” strategy that significantly speeds up the convergence. The idea is that we start with a large value for τ and then during the iterations of the minimization algorithm we gradually decrease it until we reach the target value. We note that for this problem instead of 100 we run 200 MFISTA iterations. The rest of the optimization set-up remains the same.

In Table 1 we report the reconstruction results we obtained using TV, TGV², and our proposed STV₁ regularizer. The quality of the reconstructions is measured in terms of increase in SNR w.r.t. the reconstructed image obtained by projecting the Fourier measurements back to the image domain. From these results we clearly observe that our proposed regularizer consistently outperforms TV and TGV² for all sampling strategies and noise levels. Further, we note that in the cases where the noise level is relatively low and the Fourier measurements are not very sparse, then the difference in SNR improvement is substantial. Indeed, on several occasions the SNR difference between STV₁ and the other two competing regularizers reaches




Figure 10. Reconstruction of the brain image (Figure 9(b)) from Fourier data sampled with 32 radial lines and 10-dB SNR. (a) Back-projected image ($PSNR = 23.30$), (b) TV reconstruction ($\tau = 3.84 \cdot 10^{-2}$, $PSNR = 25.88$), (c) TGV^2 reconstruction ($\tau = 4.01 \cdot 10^{-2}$, $PSNR = 26.36$), and (d) STV_1 reconstruction ($\tau = 2.62 \cdot 10^{-2}$, $PSNR = 26.90$).

more than 3 dB. In Figure 10 we show a representative reconstruction example for a radial sampling with 32 lines and a SNR of 10 dB. From this example, we can verify the superiority of our reconstruction over the ones of TV and TGV^2 both quantitatively and visually. Our STV_1 regularizer manages to better preserve the image edges than TV and at the same time avoid intensity smearing artifacts such as those appearing in the case of TGV^2 regularization.

5.5. Implementation details and computational runtimes. All methods were implemented in unoptimized MATLAB code that only makes use of the CPU.⁴ Using the average runtime of TV as baseline, the relative average runtimes of the rest of the tested methods in the conducted grayscale experiments were as follows: 0.96 for TGV^2 , 3.42 for STV_2 , and 3.54 for STV_1 . Regarding the color experiments and using the average runtime of VTV as baseline, the relative average runtimes of the tested methods were 1.11 for TV_J , 1.14 for TGV^2 , 1.09 for $STV_1(\text{loc})$, and 4.66 for STV_1 . Note that all runtimes were computed on a computer equipped with Intel Xeon Processor E5-1620 (Quad Core, 3.60-GHz Turbo, 10 MB).

We observe that TV, VTV, TGV^2 , TV_J , and the proposed $STV_1(\text{loc})$ regularizers demonstrate a similar runtime performance. The proposed STV_2 and STV_1 functionals are around 3.5–4 times slower in the current implementation, which is due to the presence of a convolution kernel. We believe that this additional computational effort is definitely worth spending in several applications, given the significant improvement in the reconstructions that we observed in the reported experiments. Furthermore, it is worth mentioning that the runtimes of the proposed algorithm could be significantly improved by exploiting the fact that it is highly parallelizable and admits fast GPU implementations.

6. Conclusions. In this work we have introduced a novel family of convex energy functionals that extend and generalize the total variation seminorm and vectorial extensions of it. The key feature of our functionals is that they incorporate information from the vicinity of every point in the image domain by penalizing the eigenvalues of the structure tensor at this point. Therefore, they provide a richer and more robust measure of image variation which translates to improved reconstruction performance. This consistent behavior has been verified

⁴The source code for STV regularization is publicly available from the authors' websites.

by extensive comparisons we performed against competing methods on several inverse imaging problems.

Since our regularizers can be used instead of TV or its vectorial extensions in any energy functional, we believe that there is a broad spectrum of imaging applications where they can be proven very useful. Another interesting research direction is related to the open theoretical issue regarding the continuous version of the STV regularizers. In this article, we have defined STV by assuming that the underlying image lies in the Sobolev space $W^{1,2}(\mathbb{R}^2, \mathbb{R}^M)$. Since TV, which can be viewed as a special case of STV, is defined for images that belong to the bounded variation (BV) space [15], an open theoretical question is whether STV can also be defined for functions of a BV space or a space containing BV.

Appendix A. Mathematical Proofs.

A.1. Compact Operators. In this section we review basic definitions and results from the theory of compact operators. The concept of a compact operator will be essential for the proof of Theorem 2.1.

Definition A.1 (compact operators). *An operator $\mathbf{A} : \mathcal{H}_1 \rightarrow \mathcal{H}_2$, where \mathcal{H}_1 and \mathcal{H}_2 are two separable Hilbert spaces is compact if and only if it can be specified as*

$$(A.1) \quad f \mapsto \mathbf{A}f = \sum_{n=1}^{\infty} \sigma_n u_n \langle \tilde{u}_n, f \rangle_{\mathcal{H}_1}$$

for any $f \in \mathcal{H}_1$. In (A.1) $\tilde{u}_1, \tilde{u}_2, \dots$ and u_1, u_2, \dots are orthonormal elements (functions) of \mathcal{H}_1 and \mathcal{H}_2 , respectively, and $\sigma_1, \sigma_2, \dots$ is a sequence of positive numbers with limit zero, called the singular values of the operator. The same property applies to the adjoint operator $\mathbf{A}^* : \mathcal{H}_2 \rightarrow \mathcal{H}_1$ with the role of u_n and \tilde{u}_n being interchanged.

While all finite-dimensional linear operators (i.e., matrices) are compact, the property is more restrictive in infinite dimensions. However, the property is robust in the sense that it is conserved through subsequent linear transformation; i.e., if $\mathbf{A} : \mathcal{H}_1 \rightarrow \mathcal{H}_2$ is compact, then $\mathbf{B}\mathbf{A}$ is compact as well for any bounded operator $\mathbf{B} : \mathcal{H}_2 \rightarrow \mathcal{H}_1$ (not necessarily compact).

In order to stay closer to the standard (finite-dimensional) matrix-vector notation, we shall not distinguish between the operator \mathbf{A} and its kernel, which is an element of the Hilbert space $\mathcal{H}_2 \otimes \mathcal{H}_1$, and simply write

$$(A.2) \quad \mathbf{A} = \sum_{n=1}^{\infty} \sigma_n(\mathbf{A}) u_n \otimes \tilde{u}_n$$

with the understanding that $u_n \in \mathcal{H}_2$ is in the ‘‘column’’ space of the operator, while \tilde{u}_n is in the ‘‘row’’ space \mathcal{H}_1 and is acting as a linear functional (inner product) on the input.

Proposition A.2 (Hölder’s inequality for the trace of operators). *Let $\mathbf{A}, \mathbf{B} : \mathcal{H}_1 \rightarrow \mathcal{H}_2$ be two compact operators. Then, for any $1 \leq p, q \leq +\infty$ with $\frac{1}{p} + \frac{1}{q} = 1$,*

$$(A.3) \quad |\text{trace}(\mathbf{A}\mathbf{B}^*)| \leq \|\boldsymbol{\sigma}(\mathbf{A})\|_{\ell_p} \|\boldsymbol{\sigma}(\mathbf{B})\|_{\ell_q},$$

where $\boldsymbol{\sigma}(\mathbf{A})$ and $\boldsymbol{\sigma}(\mathbf{B})$ are the vectors formed from the singular values of \mathbf{A} and \mathbf{B} .

Proof. By direct calculation, we can express the trace($\mathbf{A}\mathbf{B}^*$) in terms of the singular values of \mathbf{A} and \mathbf{B} as

$$(A.4) \quad \text{trace}(\mathbf{A}\mathbf{B}^*) = \sum_{m=1}^{\infty} \sum_{n=1}^{\infty} \sigma_n(\mathbf{A})\sigma_m(\mathbf{B}) \langle u_n, v_m \rangle_{\mathcal{H}_2} \langle \tilde{u}_n, \tilde{v}_m \rangle_{\mathcal{H}_1},$$

where $u_n, v_n \in \mathcal{H}_2$ (resp., $\tilde{u}_n, \tilde{v}_n \in \mathcal{H}_1$) are the left-sided (resp., right-sided) eigenfunctions associated with the singular decomposition of \mathbf{A} and \mathbf{B} , respectively.

Using the Cauchy–Schwarz inequality, we first show that

$$(A.5) \quad \begin{aligned} \sum_{n=1}^{\infty} |\langle u_n, v_m \rangle_{\mathcal{H}_2} \langle \tilde{u}_n, \tilde{v}_m \rangle_{\mathcal{H}_1}| &\leq \left(\sum_{n=1}^{\infty} |\langle u_n, v_m \rangle_{\mathcal{H}_2}|^2 \right)^{1/2} \left(\sum_{n=1}^{\infty} |\langle \tilde{u}_n, \tilde{v}_m \rangle_{\mathcal{H}_1}|^2 \right)^{1/2} \\ &\leq \|v_m\|_{\mathcal{H}_2} \|\tilde{v}_m\|_{\mathcal{H}_1} = 1, \end{aligned}$$

where the latter norm estimate follows from Bessel’s inequality based on the property that $\{u_n\}$ and $\{\tilde{u}_n\}$ are orthonormal sets of \mathcal{H}_2 and \mathcal{H}_1 , respectively. There is also the dual counterpart of (A.5) where the summation is performed over m rather than n . Next, starting from the explicit trace formula (A.4), we apply the weighted Hölder inequality [20] for double sums, which gives

$$\begin{aligned} &|\text{trace}(\mathbf{A}\mathbf{B}^*)| \\ &\leq \sum_{m=1}^{\infty} \sum_{n=1}^{\infty} \sigma_n(\mathbf{A})\sigma_m(\mathbf{B}) |\langle u_n, v_m \rangle_{\mathcal{H}_2} \langle \tilde{u}_n, \tilde{v}_m \rangle_{\mathcal{H}_1}| \\ &\leq \left(\sum_{n=1}^{\infty} \sum_{m=1}^{\infty} \sigma_n^p(\mathbf{A}) |\langle u_n, v_m \rangle_{\mathcal{H}_2} \langle \tilde{u}_n, \tilde{v}_m \rangle_{\mathcal{H}_1}| \right)^{1/p} \left(\sum_{n=1}^{\infty} \sum_{m=1}^{\infty} \sigma_m^q(\mathbf{B}) |\langle u_n, v_m \rangle_{\mathcal{H}_2} \langle \tilde{u}_n, \tilde{v}_m \rangle_{\mathcal{H}_1}| \right)^{1/q} \\ &= \left(\sum_{n=1}^{\infty} \sigma_n^p(\mathbf{A}) \sum_{m=1}^{\infty} |\langle u_n, v_m \rangle_{\mathcal{H}_2} \langle \tilde{u}_n, \tilde{v}_m \rangle_{\mathcal{H}_1}| \right)^{1/p} \left(\sum_{m=1}^{\infty} \sigma_m^q(\mathbf{B}) \sum_{n=1}^{\infty} |\langle u_n, v_m \rangle_{\mathcal{H}_2} \langle \tilde{u}_n, \tilde{v}_m \rangle_{\mathcal{H}_1}| \right)^{1/q} \\ &\leq \left(\sum_{n=1}^{\infty} \sigma_n^p(\mathbf{A}) \|u_n\|_{\mathcal{H}_2} \|\tilde{u}_n\|_{\mathcal{H}_1} \right)^{1/p} \left(\sum_{m=1}^{\infty} \sigma_m^q(\mathbf{B}) \|v_m\|_{\mathcal{H}_2} \|\tilde{v}_m\|_{\mathcal{H}_1} \right)^{1/q}, \end{aligned}$$

where the availability of (A.5) is essential for reducing the double sum to a single one. ■

A.2. Proof of Theorem 2.1. While in (2.10) we define the STV functional for a 2D image; next we provide a proof of Theorem 2.1 for the more general case where \mathbf{u} is defined on \mathbb{R}^d . To show the translation and rotation invariance property of the STV functional, we consider the transformation of the image coordinates $T(\mathbf{x}) = \mathbf{R}\mathbf{x} + \boldsymbol{\alpha}$, where \mathbf{R} is a rotation matrix and $\boldsymbol{\alpha}$ represents a translation. Applying the chain rule to the Jacobian matrix, we get

$$(A.6) \quad J\{\mathbf{u} \circ T\}(\mathbf{x}) = J\mathbf{u}(T(\mathbf{x})) \mathbf{R},$$

where \circ indicates the composition operation. Now, if we denote $\mathbf{S} = S_K \{ \mathbf{u} \circ T \} (\mathbf{x})$ we have

$$\begin{aligned}
 \mathbf{S} &= \mathbf{R}^T \left(\int_{\mathbb{R}^d} K(\mathbf{y}) J\mathbf{u}^T(T(\mathbf{x} - \mathbf{y})) J\mathbf{u}(T(\mathbf{x} - \mathbf{y})) d\mathbf{y} \right) \mathbf{R} \\
 &= \mathbf{R}^T \left(\int_{\mathbb{R}^d} K(\mathbf{y}) J\mathbf{u}^T(T(\mathbf{x}) - \mathbf{R}\mathbf{y}) J\mathbf{u}(T(\mathbf{x}) - \mathbf{R}\mathbf{y}) d\mathbf{y} \right) \mathbf{R} \\
 &\stackrel{\mathbf{y}' = \mathbf{R}\mathbf{y}}{=} \mathbf{R}^T \left(\int_{\mathbb{R}^d} K(\mathbf{R}^{-1}\mathbf{y}') J\mathbf{u}^T(T(\mathbf{x}) - \mathbf{y}') J\mathbf{u}(T(\mathbf{x}) - \mathbf{y}') d\mathbf{y}' \right) \mathbf{R} \\
 &= \mathbf{R}^T \left(\int_{\mathbb{R}^d} K(\mathbf{y}') J\mathbf{u}^T(T(\mathbf{x}) - \mathbf{y}') J\mathbf{u}(T(\mathbf{x}) - \mathbf{y}') d\mathbf{y}' \right) \mathbf{R} \\
 \text{(A.7)} \quad &= \mathbf{R}^T S_K \{ \mathbf{u} \} (T(\mathbf{x})) \mathbf{R}.
 \end{aligned}$$

We note that in order to derive the above result we have considered the convolution kernel K to be isotropic, i.e., $K(\mathbf{x}) = K(|\mathbf{x}|)$. From (A.7) and since the eigenvalues of the structure tensor in the r.h.s. are not affected by the presence of the rotation matrix \mathbf{R} and its inverse \mathbf{R}^T , we have that $\lambda^\pm(\mathbf{x}) = \lambda^\pm(T(\mathbf{x}))$. Using this result and having in mind that in order to evaluate the STV_p functional in (2.10) we need to integrate over the whole image domain, it is now clear that the energy will be preserved despite the transformation of the coordinate system. In other words, it holds that $\text{STV}_p(\mathbf{u} \circ T) = \text{STV}_p(\mathbf{u})$.

Next, in order to prove the convexity and 1-homogeneity properties, let us first introduce the $d \times d$ matrix $\mathbf{G}_K(\mathbf{u})$ defined as

$$\text{(A.8)} \quad \mathbf{G}_K(\mathbf{u}) = \int_{\mathbb{R}^d} K(\mathbf{y}) J\mathbf{u}^T(\mathbf{y}) J\mathbf{u}(\mathbf{y}) d\mathbf{y},$$

which corresponds to evaluating the structure tensor (2.7) at the origin of \mathbf{u} . This matrix is well-defined over the functions that belong to the Sobolev space $W^{1,2}(\mathbb{R}^d, \mathbb{R}^M)$. Next, we show the following result.

Theorem A.3. *Let λ_n denote the n th eigenvalue of the matrix $\mathbf{G}_K(\mathbf{u})$. Then $\rho(\mathbf{u}) = (\sum_{n=1}^d \lambda_n^{\frac{2}{p}})^{1/p}$ is a 1-homogenous and convex function of \mathbf{u} in the sense that*

$$\text{(A.9)} \quad \rho(a\mathbf{u}) = |a| \rho(\mathbf{u}),$$

$$\text{(A.10)} \quad \rho(t\mathbf{u}_1 + (1-t)\mathbf{u}_2) \leq t\rho(\mathbf{u}_1) + (1-t)\rho(\mathbf{u}_2)$$

$\forall a \in \mathbb{R}, t \in [0, 1]$, and $\mathbf{u}, \mathbf{u}_1, \mathbf{u}_2 \in W^{1,2}(\mathbb{R}^d, \mathbb{R}^M)$.

Proof. To prove convexity of $\rho(\mathbf{u})$, let us first introduce the $(\infty \times d)$ -dimensional matrix defined as $\mathbf{A} = \mathbf{A}(\mathbf{u}, \mathbf{y}) = [\mathbf{f}_1(\mathbf{y}) \dots \mathbf{f}_d(\mathbf{y})]$ for arbitrary $\mathbf{y} \in \mathbb{R}^d$, where

$$\mathbf{f}_n(\mathbf{y}) = \left(\underbrace{k(\mathbf{y}) u_{x_n}^1(\mathbf{y})}_{\mathbf{f}_{n,1}(\mathbf{y})}, \dots, \underbrace{k(\mathbf{y}) u_{x_n}^M(\mathbf{y})}_{\mathbf{f}_{n,M}(\mathbf{y})} \right), 1 \leq n \leq d,$$

is a vector function $\mathbf{f}_n \in \mathcal{H} = L_2(\mathbb{R}^d, \mathbb{R}^M)$, $u_{x_n}^m \forall 1 \leq m \leq M$ denotes the derivative of the m th component of \mathbf{u} w.r.t. the n th dimension, and $k(\mathbf{y}) = \sqrt{K(\mathbf{y})}$. Note that the last

assignment is valid since the convolution kernel K is assumed to be nonnegative valued. Since the functions $\mathbf{f}_n = \mathbf{f}_n(\mathbf{y})$ belong to the Hilbert space \mathcal{H} , the matrix \mathbf{A} defines a linear map: $\mathbb{R}^d \rightarrow \mathcal{H}$. The Hilbert space \mathcal{H} is associated with the inner product

$$\langle \mathbf{f}_n, \mathbf{g}_n \rangle_{\mathcal{H}} = \sum_{m=1}^M \int_{\mathbb{R}^d} \mathbf{f}_{n,m}(\mathbf{y}) \mathbf{g}_{n,m}(\mathbf{y}) d\mathbf{y}.$$

The adjoint of \mathbf{A} is the unique operator $\mathbf{A}^* : \mathcal{H} \rightarrow \mathbb{R}^d$ that satisfies

$$(A.11) \quad \langle \mathbf{f}, \mathbf{A}\mathbf{v} \rangle_{\mathcal{H}} = \langle \mathbf{A}^* \mathbf{f}, \mathbf{v} \rangle$$

for every $\mathbf{v} \in \mathbb{R}^d$ and every $f \in \mathcal{H}$. Having introduced the adjoint operator of \mathbf{A} , we are in the position of specifying the Gram matrix of size $d \times d$:

$$\begin{aligned} \mathbf{A}^* \mathbf{A} &= \begin{pmatrix} \langle \mathbf{f}_1, \mathbf{f}_1 \rangle_{\mathcal{H}} & \cdots & \langle \mathbf{f}_1, \mathbf{f}_d \rangle_{\mathcal{H}} \\ \vdots & \ddots & \vdots \\ \langle \mathbf{f}_d, \mathbf{f}_1 \rangle_{\mathcal{H}} & \cdots & \langle \mathbf{f}_d, \mathbf{f}_d \rangle_{\mathcal{H}} \end{pmatrix} \\ &= \begin{pmatrix} \sum_{m=1}^M \int_{\mathbb{R}^d} k^2(\mathbf{y}) (u_{x_1}^m(\mathbf{y}))^2 d\mathbf{y} & \cdots & \sum_{m=1}^M \int_{\mathbb{R}^d} k^2(\mathbf{y}) u_{x_1}^m(\mathbf{y}) u_{x_d}^m(\mathbf{y}) d\mathbf{y} \\ \vdots & \ddots & \vdots \\ \sum_{m=1}^M \int_{\mathbb{R}^d} k^2(\mathbf{y}) u_{x_d}^m(\mathbf{y}) u_{x_1}^m(\mathbf{y}) d\mathbf{y} & \cdots & \sum_{m=1}^M \int_{\mathbb{R}^d} k^2(\mathbf{y}) (u_{x_d}^m(\mathbf{y}))^2 d\mathbf{y} \end{pmatrix} \\ &= \begin{pmatrix} \int_{\mathbb{R}^d} K(\mathbf{y}) \sum_{m=1}^M (u_{x_1}^m(\mathbf{y}))^2 d\mathbf{y} & \cdots & \int_{\mathbb{R}^d} K(\mathbf{y}) \sum_{m=1}^M u_{x_1}^m(\mathbf{y}) u_{x_d}^m(\mathbf{y}) d\mathbf{y} \\ \vdots & \ddots & \vdots \\ \int_{\mathbb{R}^d} K(\mathbf{y}) \sum_{m=1}^M u_{x_d}^m(\mathbf{y}) u_{x_1}^m(\mathbf{y}) d\mathbf{y} & \cdots & \int_{\mathbb{R}^d} K(\mathbf{y}) \sum_{m=1}^M (u_{x_d}^m(\mathbf{y}))^2 d\mathbf{y} \end{pmatrix} \\ (A.12) \quad &= \mathbf{G}_K(\mathbf{u}). \end{aligned}$$

From (A.12) it turns out that the $d \times d$ Gram matrix, which provides a means to compute the spectrum of \mathbf{A} , is equal to the $\mathbf{G}_K(\mathbf{u})$ matrix defined in (A.8). This relation further implies that the vector consisting of the singular values of the infinite-dimensional but finite rank matrix (or compact operator) \mathbf{A} , which we will denote as $\boldsymbol{\sigma}(\mathbf{A})$, is directly related to the eigenvalues of $\mathbf{G}_K(\mathbf{u})$. Indeed, we have

$$(A.13) \quad \boldsymbol{\sigma}(\mathbf{A}) = \sqrt{\boldsymbol{\lambda}(\mathbf{A}^* \mathbf{A})} = \sqrt{\boldsymbol{\lambda}(\mathbf{G}_K(\mathbf{u}))},$$

where $\boldsymbol{\lambda}(\mathbf{A}^* \mathbf{A})$ is the vector with the eigenvalues of the matrix $\mathbf{A}^* \mathbf{A}$. Thus, we conclude that the n th singular value of \mathbf{A} can be derived as $\sigma_n = \sqrt{\lambda_n}$ with $1 \leq n \leq d$. This means that we can equivalently write the function $\rho(\mathbf{u})$ as

$$(A.14) \quad \rho(\mathbf{u}) = \left(\sum_{n=1}^d \sigma_n^p \right)^{1/p} = \|\boldsymbol{\sigma}(\mathbf{A})\|_{\ell_p}.$$

Next, we shall show that for two compact operators $\mathbf{A}, \mathbf{B} : \mathcal{H}_1 \mapsto \mathcal{H}_2$ and for any $1 \leq p, q \leq +\infty$ with $\frac{1}{p} + \frac{1}{q} = 1$, it holds

$$(A.15) \quad \|\boldsymbol{\sigma}(\mathbf{A})\|_{\ell_p} = \sup_{\|\boldsymbol{\sigma}(\mathbf{B})\|_{\ell_q}=1} \text{trace}(\mathbf{A}\mathbf{B}^*).$$

We start by invoking Hölder's inequality for the trace of operators, whose proof is provided in Proposition A.2. According to this inequality we have that

$$(A.16) \quad |\text{trace}(\mathbf{A}\mathbf{B}^*)| \leq \|\boldsymbol{\sigma}(\mathbf{A})\|_{\ell_p} \|\boldsymbol{\sigma}(\mathbf{B})\|_{\ell_q}.$$

This, in turn, implies that

$$(A.17) \quad \sup_{\|\boldsymbol{\sigma}(\mathbf{B})\|_{\ell_q}=1} \text{trace}(\mathbf{A}\mathbf{B}^*) \leq \|\boldsymbol{\sigma}(\mathbf{A})\|_{\ell_p}.$$

Now, the critical point is to show that for every operator \mathbf{A} we can always select another operator \mathbf{B} with $\|\boldsymbol{\sigma}(\mathbf{B})\|_{\ell_q} = 1$ such that $\text{trace}(\mathbf{A}\mathbf{B}^*) = \|\boldsymbol{\sigma}(\mathbf{A})\|_{\ell_p}$. To show this, we choose

$$(A.18) \quad \mathbf{B} = \sum_{n=1}^{\infty} \sigma_n(\mathbf{B}) u_n \otimes \tilde{u}_n$$

with

$$(A.19) \quad \sigma_n(\mathbf{B}) = \frac{\sigma_n^{p-1}(\mathbf{A})}{\|\boldsymbol{\sigma}(\mathbf{A})\|_{\ell_p}^{p-1}}.$$

Indeed, since \mathbf{B} has the same eigenvectors as \mathbf{A} , we have that

$$(A.20) \quad \text{trace}(\mathbf{A}\mathbf{B}^*) = \frac{\sum_{n=1}^{\infty} \sigma_n^p(\mathbf{A})}{\|\boldsymbol{\sigma}(\mathbf{A})\|_{\ell_p}^{p-1}} = \|\boldsymbol{\sigma}(\mathbf{A})\|_{\ell_p}.$$

Moreover, by noting that $p = q(p-1)$, we easily verify that

$$(A.21) \quad \|\boldsymbol{\sigma}(\mathbf{B})\|_{\ell_q}^q = \sum_{n=1}^{\infty} \sigma_n^q(\mathbf{B}) = \frac{\sum_{n=1}^{\infty} \sigma_n^{q(p-1)}(\mathbf{A})}{\|\boldsymbol{\sigma}(\mathbf{A})\|_{\ell_p}^{q(p-1)}} = 1.$$

This shows that the supremum in (A.17) achieves the equality, which proves (A.15).

Thanks to the dual definition in (A.15), we can now show that

$$(A.22) \quad \begin{aligned} \|\boldsymbol{\sigma}(t\mathbf{A}_1 + (1-t)\mathbf{A}_2)\|_{\ell_p} &= \sup_{\|\boldsymbol{\sigma}(\mathbf{B})\|_{\ell_q}=1} \text{trace}((t\mathbf{A}_1 + (1-t)\mathbf{A}_2)\mathbf{B}^*) \\ &\leq \sup_{\|\boldsymbol{\sigma}(\mathbf{B}_1)\|_{\ell_q}=1} \text{trace}(t\mathbf{A}_1\mathbf{B}_1^*) + \sup_{\|\boldsymbol{\sigma}(\mathbf{B}_2)\|_{\ell_q}=1} \text{trace}((1-t)\mathbf{A}_2\mathbf{B}_2^*) \\ &\leq t \sup_{\|\boldsymbol{\sigma}(\mathbf{B}_1)\|_{\ell_q}=1} \text{trace}(\mathbf{A}_1\mathbf{B}_1^*) + (1-t) \sup_{\|\boldsymbol{\sigma}(\mathbf{B}_2)\|_{\ell_q}=1} \text{trace}(\mathbf{A}_2\mathbf{B}_2^*) \\ &= t \|\boldsymbol{\sigma}(\mathbf{A}_1)\|_{\ell_p} + (1-t) \|\boldsymbol{\sigma}(\mathbf{A}_2)\|_{\ell_p}. \end{aligned}$$

Next, we recall that the kernel of \mathbf{A} is linear in \mathbf{u} , since it only involves derivatives which are linear operations and a multiplication with a constant kernel. This together with (A.22) proves that the function $\rho(\mathbf{u})$ is a convex function of \mathbf{u} .

Finally, the 1-homogeneity of $\rho(\mathbf{u})$ is easily verified by using the definition of ρ , given in (A.14). ■

Using the results of the above theorem, the convexity and 1-homogeneity properties of the STV_p functional can now be easily verified by first noting that we can write it as

$$(A.23) \quad \text{STV}_p(\mathbf{u}) = \int_{\mathbb{R}^d} \rho(\mathbf{u} \circ P_{\mathbf{x}}) d\mathbf{x},$$

where $P_{\mathbf{x}}(\mathbf{y}) = \mathbf{x} - \mathbf{y}$ and $\rho(\mathbf{u})$ is defined in Theorem A.3. If we apply the chain rule to the Jacobian matrix we get

$$(A.24) \quad J\{\mathbf{u} \circ P_{\mathbf{x}}\}(\mathbf{y}) = -J\mathbf{u}(P_{\mathbf{x}}(\mathbf{y})).$$

Now, for the matrix $\mathbf{G}_K(\mathbf{u} \circ P_{\mathbf{x}})$, which is the main component of the function $\rho(\mathbf{u} \circ P_{\mathbf{x}})$, we have

$$(A.25) \quad \begin{aligned} \mathbf{G}_K(\mathbf{u} \circ P_{\mathbf{x}}) &= \int_{\mathbb{R}^d} K(\mathbf{y}) \nabla \mathbf{u}^T(P_{\mathbf{x}}(\mathbf{y})) \nabla \mathbf{u}(P_{\mathbf{x}}(\mathbf{y})) d\mathbf{y} \\ &\stackrel{\mathbf{y}' = P_{\mathbf{x}}(\mathbf{y})}{=} \int_{\mathbb{R}^d} K(\mathbf{x} - \mathbf{y}') \nabla \mathbf{u}^T(\mathbf{y}') \nabla \mathbf{u}(\mathbf{y}') d\mathbf{y}' \\ &= \int_{\mathbb{R}^d} K_{\mathbf{x}}(\mathbf{y}) \nabla \mathbf{u}^T(\mathbf{y}) \nabla \mathbf{u}(\mathbf{y}) d\mathbf{y} = \mathbf{G}_{K_{\mathbf{x}}}(\mathbf{u}), \end{aligned}$$

where $K_{\mathbf{x}}(\mathbf{y}) = K(\mathbf{x} - \mathbf{y}) = K(\mathbf{y} - \mathbf{x})$. The above result implies that $\rho(\mathbf{u} \circ P_{\mathbf{x}}) = \rho_{\mathbf{x}}(\mathbf{u})$, where $\rho_{\mathbf{x}}(\mathbf{u})$ is defined similarly to $\rho(\mathbf{u})$, with the only difference being that its input eigenvalues correspond to those of the matrix $\mathbf{G}_{K_{\mathbf{x}}}(\mathbf{u})$ instead of $\mathbf{G}_K(\mathbf{u})$. Therefore, it holds that

$$(A.26) \quad \text{STV}_p(\mathbf{u}) = \int_{\mathbb{R}^d} \rho_{\mathbf{x}}(\mathbf{u}) d\mathbf{x}.$$

Since the kernel $K_{\mathbf{x}}$, which is used in $\rho_{\mathbf{x}}(\mathbf{u})$, is just a shifted version of the initial convolution kernel K and does not depend on \mathbf{u} , the convexity and homogeneity results of $\rho(\mathbf{u})$ are also inherited by $\rho_{\mathbf{x}}(\mathbf{u}) \forall \mathbf{x} \in \mathbb{R}^d$. Therefore, since the STV_p functional corresponds to the integral of convex and 1-homogeneous functionals of \mathbf{u} , it is also a convex and 1-homogeneous function of \mathbf{u} .

A.3. Proof of Proposition 3.2. To find the adjoint of the discrete patch-based Jacobian, we exploit the relation of the inner products of the spaces \mathbb{R}^{NM} and \mathcal{X} in (3.8). Using the definition of the inner product in \mathcal{X} (3.6), we can equivalently write (3.8) as

$$(A.27) \quad \sum_{n=1}^N \text{trace} \left([\mathbf{J}_K \mathbf{u}]_n^T \mathbf{Y}_n \right) = \sum_{n=1}^N \sum_{m=1}^M [\mathbf{u}_m]_n [\mathbf{J}_K^* \mathbf{Y}]_k,$$

where $k = (N - 1)m + n$. Then, we expand the l.h.s. of (A.27), to obtain

$$\begin{aligned}
& \sum_{n,m,l} [T_{s_l, \omega} \circ D_h \mathbf{u}_m]_n \mathbf{Y}_n^{(r,1)} + [T_{s_l, \omega} \circ D_v \mathbf{u}_m]_n \mathbf{Y}_n^{(r,2)} \\
&= \sum_{n,m,l} [\mathbf{u}_m]_n \left([D_h^* \circ T_{s_l, \omega}^* \circ \mathbf{Y}^{(r,1)}]_n + [D_v^* \circ T_{s_l, \omega}^* \circ \mathbf{Y}^{(r,2)}]_n \right) \\
\text{(A.28)} \quad &= \sum_{n,m} [\mathbf{u}_m]_n \left(\sum_{l=1}^L [-\mathbf{div} (T_{s_l, \omega}^* \circ \mathbf{Y}^{((m-1)L+l, :)})]_n \right),
\end{aligned}$$

where $r = (m - 1)L + l$ and D_h, D_v are the horizontal and vertical components of the discrete gradient, respectively. Also note that $\mathbf{Y}_n^{(i,j)}$ with $1 \leq n \leq N$, $1 \leq i \leq LM$, and $1 \leq j \leq 2$, corresponds to a single element of $\mathbf{Y} \in \mathcal{X}$, while $\mathbf{Y}_n^{(i,:)}$ is a vector extracted from the i th row of the n th matrix component, $\mathbf{Y}_n \in \mathbb{R}^{(LM) \times 2}$, of \mathbf{Y} . Now, by comparing the r.h.s. of (A.27) to the r.h.s. expansion of (A.28), we can verify that the adjoint of the patch-based Jacobian is indeed computed by the formula provided in Proposition 3.2.

A.4. Proof of Proposition 4.2. For any pair of variables $\Omega, \Psi \in \mathcal{X}$ and

$$d(\Omega) = \frac{1}{2} \|\mathbf{w} - \Pi_C(\mathbf{w})\|_2^2 + \frac{1}{2} (\|\mathbf{z}\|_2^2 - \|\mathbf{w}\|_2^2),$$

we can show that

$$\begin{aligned}
\|\nabla d(\Omega) - \nabla d(\Psi)\|_{\mathcal{X}} &= \|\tau \mathbf{J}_K (h(\Omega) - h(\Psi))\|_{\mathcal{X}} \\
&\leq \tau \|\mathbf{J}_K\| \|h(\Omega) - h(\Psi)\|_2 \leq \tau \|\mathbf{J}_K\| \|\tau \mathbf{J}_K^* (\Omega - \Psi)\|_2 \\
\text{(A.29)} \quad &\leq \tau^2 \|\mathbf{J}_K\|^2 \|\Omega - \Psi\|_{\mathcal{X}},
\end{aligned}$$

where $h(\Omega) = \Pi_C(\mathbf{z} - \tau \mathbf{J}_K^* \Omega)$. This result follows from the relation between the norms defined in the spaces \mathcal{X} and \mathbb{R}^{NM} and the induced operator norm, i.e., $\|\mathbf{J}_K \mathbf{u}\|_{\mathcal{X}} \leq \|\mathbf{J}_K\| \|\mathbf{u}\|_2$, and the fact that the projection operator onto a convex set $\mathcal{C} \subseteq \mathbb{R}^N$ is firmly nonexpansive [2, Proposition 4.8]. The latter translates to $\|\Pi_C(\mathbf{u}) - \Pi_C(\mathbf{v})\|_2 \leq \|\mathbf{u} - \mathbf{v}\|_2 \forall \mathbf{u}, \mathbf{v} \in \mathbb{R}^N$.

Now, in order to find an upper bound of the operator norm of the patch-based Jacobian we use that $\|\mathbf{J}_K\|^2 = \|\mathbf{J}_K^* \mathbf{J}_K\|$ [34] (general property of bounded linear operators) and denote $\mathbf{D} = -\mathbf{div} \circ \sum_l (T_{s_l, \omega}^* \circ T_{s_l, \omega}) \circ \nabla$. Then, we get

$$\text{(A.30)} \quad \|\mathbf{J}_K^* \mathbf{J}_K \mathbf{u}\|_2 = \|\mathbf{D} \mathbf{u}\|_2 \leq \|\nabla\|^2 \|T\|^2 \|\mathbf{u}\|_2,$$

where $T \triangleq \sum_{l=1}^L (T_{s_l, \omega}^* \circ T_{s_l, \omega})$. From [3] we have that $\|\nabla\|^2 \leq 8$ and we can further show that $\|T\|^2 \leq \sqrt{2}$, if we assume a unit-normalized kernel K and a shift operation with reflexive boundaries. These immediately imply that an upper bound of the Lipschitz constant of $\nabla d(\Omega)$ will be $L(d) \leq \tau^2 \|\mathbf{J}_K\|^2 \leq 8\sqrt{2}\tau^2$.

Acknowledgment. The authors would like to thank the anonymous reviewers and the associate editor for their useful comments and suggestions.

REFERENCES

- [1] M. V. AFONSO, J. M. BIOCAS-DIAS, AND M. A. T FIGUEIREDO, *Fast image recovery using variable splitting and constrained optimization*, IEEE Trans. Image Process., 19 (2010), pp. 2345–2356.
- [2] H. H. BAUSCHKE AND P. L. COMBETTES, *Convexity and nonexpansiveness*, in Convex Analysis and Monotone Operator Theory in Hilbert Spaces, J. M. Borwein and P. Borwein, eds., Springer, New York, 2011, pp. 59–74.
- [3] A. BECK AND M. TEBoulLE, *Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems*, IEEE Trans. Image Process., 18 (2009), pp. 2419–2434.
- [4] A. BECK AND M. TEBoulLE, *A fast iterative shrinkage-thresholding algorithm for linear inverse problems*, SIAM J. Imaging Sci., 2 (2009), pp. 183–202.
- [5] B. BENJAMIN, M. BURGER, M. DROSKE, O. NEMITZ, AND M. RUMPF, *Cartoon extraction based on anisotropic image classification vision, modeling, and visualization*, in Vision, Modeling, and Visualization, Akademische Verlagsgesellschaft, Berlin, 2006, pp. 293–300.
- [6] M. BERTERO AND P. BOCCACCI, *Introduction to Inverse Problems in Imaging*, IOP, Bristol, UK, 1998.
- [7] P. BLOMGREN AND T. F. CHAN, *Color TV: Total variation methods for restoration of vector-valued images*, IEEE Trans. Image Process., 7 (1998), pp. 304–309.
- [8] S. BOYD, N. PARIKH, E. CHU, B. PELEATO, AND J. ECKSTEIN, *Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers*, Now Publishers, Hanover, MA, 2011.
- [9] K. BREDIES, *Recovering Piecewise Smooth Multichannel Images by Minimization of Convex Functionals with Total Generalized Variation Penalty*, A. Bruhn et al. eds., in Global Optimization Methods, Lecture Notes in Comput. Sci. 8293, 2014, pp. 44–77.
- [10] K. BREDIES, K. KUNISCH, AND T. POCK, *Total generalized variation*, SIAM J. Imaging Sci., 3 (2010), pp. 492–526.
- [11] X. BRESSON AND T. F. CHAN, *Fast dual minimization of the vectorial total variation norm and applications to color image processing*, Inverse Probl. Imaging, 2 (2008), pp. 455–484.
- [12] A. CHAMBOLLE, *An algorithm for total variation minimization and applications*, J. Math. Imaging Vision, 20 (2004), pp. 89–97.
- [13] A. CHAMBOLLE AND P.-L. LIONS, *Image recovery via total variation minimization and related problems*, Numer. Math., 76 (1997), pp. 167–188.
- [14] A. CHAMBOLLE AND T. POCK, *A first-order primal-dual algorithm for convex problems with applications to imaging*, J. Math. Imaging Vision, 40 (2011), pp. 120–145.
- [15] T. CHAN, S. ESEDOGLU, F. PARK, AND A. YIP, *Total variation image restoration: Overview and recent developments*, Handbook of Mathematical Models in Computer Vision, Springer, New York, 2005, pp. 17–31.
- [16] T. CHAN, A. MARQUINA, AND P. MULET, *High-order total variation-based image restoration*, SIAM J. Sci. Comput., 22 (2000), pp. 503–516.
- [17] Q. CHENG, H. SHEN, L. ZHANG, AND P. LI, *Inpainting for remotely sensed images with a multichannel nonlocal total variation model*, IEEE Trans. Geosci. Remote Sensing, 52 (2014), pp. 175–187.
- [18] G. CHERCHIA, N. PUSTELNIK, B. PESQUET-POPESCU, AND J.-C PESQUET, *A non-local structure tensor based approach for multicomponent image recovery problems*, IEEE Trans. Image Process., 23 (2014), pp. 5531–5544.
- [19] P. L. COMBETTES AND V. R. WAJS, *Signal recovery by proximal forward-backward splitting*, Multiscale Model. Simul., 4 (2005), pp. 1168–1200.
- [20] Z. CVETKOVSKI, *Inequalities: Theorems, Techniques and Selected Problems*, Springer, Heidelberg, 2012.
- [21] S. DI ZENZO, *A note on the gradient of a multi-image*, Comp. Vision, Graphics, Image Process., 33 (1986), pp. 116–125.
- [22] Y. DONG AND M. HINTERMÜLLER, *Multi-scale total variation with automated regularization parameter selection for color image restoration*, in Scale Space and Variational Methods in Computer Vision, Springer, Berlin, 2009, pp. 271–281.
- [23] D. L. DONOHO, *Denosing by soft-thresholding*, IEEE Trans. Inform. Theory, 41 (1995), pp. 613–627.
- [24] A. ELMOATAZ, O. LEZORAY, AND S. BOUGLEUX, *Nonlocal discrete regularization on weighted graphs: A framework for image and manifold processing*, IEEE Trans. Image Process., 17 (2008), pp. 1047–1060.

- [25] E. ESSER, *Applications of Lagrangian-Based Alternating Direction Methods and Connections to Split Bregman*, Computational Applied Mathematics report, UCLA, Los Angeles, 2009.
- [26] M. A. T. FIGUEIREDO, J. M. BIOCAS-DIAS, AND R. D. NOWAK, *Majorization–minimization algorithms for wavelet-based image restoration*, IEEE Trans. Image Process., 16 (2007), pp. 2980–2991.
- [27] W. FÖRSTNER AND E. GÜLCH, *A fast operator for detection and precise location of distinct points, corners and centres of circular features*, in Proceedings ISPRS Intercommission Conference on Fast Processing of Photogrammetric data, Institut für Geodäsie und Photogrammetrie, Zurich, 1987, pp. 281–305.
- [28] G. GILBOA AND S. OSHER, *Nonlocal operators with applications to image processing*, Multiscale Model. Simul., 7 (2008), pp. 1005–1028.
- [29] G. GILBOA, N. SOCHEN, AND Y. ZEEVI, *Variational denoising of partly textured images by spatially varying constraints*, IEEE Trans. Image Process., 15 (2006), pp. 2281–2289.
- [30] B. GOLDLUECKE, E. STREKALOVSKIY, AND D. CREMERS, *The natural vectorial total variation which arises from geometric measure theory*, SIAM J. Imaging Sci., 5 (2012), pp. 537–563.
- [31] T. GOLDSTEIN AND S. OSHER, *The split Bregman method for L_1 -regularized problems*, SIAM J. Imaging Sci., 2 (2009), pp. 323–343.
- [32] M. GRASMAIR, *Locally adaptive total variation regularization*, in Scale Space and Variational Methods in Computer Vision, Springer, Berlin, 2009, pp. 331–342.
- [33] M. GRASMAIR AND F. LENZEN, *Anisotropic total variation filtering*, Appl. Math. Optim., 62 (2010), pp. 323–339.
- [34] V. HUTSON, J. S. PYM, AND M. J. CLOUD, *Applications of Functional Analysis and Operator Theory*, 2nd ed., Elsevier, Amsterdam, 2005.
- [35] B. JÄHNE, *Digital Image Processing*, Springer-Verlag, Berlin, 2002.
- [36] S. LEFKIMMIATIS, A. ROUSSOS, M. UNSER, AND P. MARAGOS, *Convex generalizations of total variation based on the structure tensor with applications to inverse problems*, in Scale Space and Variational Methods in Computer Vision, Springer, Berlin, 2013, pp. 48–60.
- [37] S. LEFKIMMIATIS, J. WARD, AND M. UNSER, *Hessian Schatten-norm regularization for linear inverse problems*, IEEE Trans. Image Process., 22 (2013), pp. 1873–1888.
- [38] F. LENZEN, F. BECKER, J. LELLMANN, S. PETRA, AND C. SCHNÖRR, *A class of quasi-variational inequalities for adaptive image denoising and decomposition*, Comput. Optim. Appl., 54 (2012), pp. 371–398.
- [39] J. LIU AND J. YE, *Efficient ℓ_1/ℓ_q Norm Regularization*, preprint, arXiv:1009.4766, 2010.
- [40] Y. NESTEROV, *A method for solving a convex programming problem with convergence rates $O(1/k^2)$* , Soviet Math. Dokl., 27 (1983), pp. 372–376.
- [41] K. PAPAFIGOROS AND C. SCHÖNLIEB, *A combined first and second order variational approach for image reconstruction*, J. Math. Imaging Vision, 48 (2014), pp. 308–338.
- [42] N. PARAGIOS, Y. CHEN, AND O. FAUGERAS, *Handbook of Mathematical Models in Computer Vision*, Springer, New York, 2006.
- [43] T. POCK AND A. CHAMBOLLE, *Diagonal preconditioning for first order primal-dual algorithms in convex optimization*, in IEEE International Conference on Computer Vision (ICCV), IEEE, Piscataway, NJ, 2011, pp. 1762–1769.
- [44] S. J. REEVES, *Fast image restoration without boundary artifacts*, IEEE Trans. Image Process., 14 (2005), pp. 1448–1453.
- [45] R. T. ROKCAFELLAR, *Convex Analysis*, Princeton University Press, Princeton, NJ, 1970.
- [46] A. ROUSSOS AND P. MARAGOS, *Tensor-based image diffusions derived from generalizations of the total variation and Beltrami functionals*, in Proceedings International Conference on Image Processing, IEEE, Piscataway, NJ, 2010, pp. 4141–4144.
- [47] L. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, Phys. D, 60 (1992), pp. 259–268.
- [48] G. SAPIRO, *Color snakes*, Comp. Vision Image Understanding, 68 (1997), pp. 247–253.
- [49] S. SETZER, G. STEIDL, AND T. TEUBER, *Restoration of images with rotated shapes*, Numer. Algorithms, 48 (2008), pp. 49–66.
- [50] S. SETZER, G. STEIDL, AND T. TEUBER, *Infimal convolution regularizations with discrete ℓ_1 -type functionals*, Commun. Math. Sci., 9 (2011), pp. 797–827.
- [51] N. SOCHEN AND L. BAR, *The Beltrami-Mumford-Shah functional*, Scale Space and Variational Methods in Computer Vision, Springer, Berlin, 2012, pp. 183–193.

- [52] N. SOCHEN, R. KIMMEL, AND R. MALLADI, *A general framework for low level vision*, IEEE Trans. Image Process., 7 (1998), pp. 310–338.
- [53] G. STEIDL AND T. TEUBER, *Anisotropic smoothing using double orientations*, in Scale Space and Variational Methods in Computer Vision, Springer, Berlin, 2009, pp. 477–489.
- [54] S. SUVRIT, *Fast projections onto $\ell_{1,q}$ -norm balls for grouped feature selection*, Machine Learning and Knowledge Discovery in Databases, Vol. 3, Springer, Berlin, 2011, pp. 305–317.
- [55] D. TSCHUMPERLÉ AND R. DERICHE, *Vector-valued image regularization with PDE's: A common framework for different applications*, IEEE Trans. Pattern Anal. Mach. Intelligence, 27 (2005), pp. 506–517.
- [56] C. R. VOGEL, *Computational Methods for Inverse Problems*, Front. Appl. Math. 23, SIAM, Philadelphia, 2002.
- [57] J. WEICKERT, *Anisotropic Diffusion in Image Processing*, Teubner, Stuttgart, 1998.
- [58] J. WEICKERT AND C. SCHNÖRR, *A theoretical framework for convex regularizers in PDE-based computation of image motion*, Int. J. Comput. Vision, 45 (2001), pp. 245–264.