

Convex Generalizations of Total Variation Based on the Structure Tensor with Applications to Inverse Problems^{*}

Stamatios Lefkimmiatis¹, Anastasios Roussos², Michael Unser¹, and Petros Maragos³

¹ Biomedical Imaging Group, EPFL, Lausanne, Switzerland

² School of EECS, Queen Mary University of London, United Kingdom

³ School of ECE, National Technical University of Athens, Greece

Abstract. We introduce a generic convex energy functional that is suitable for both grayscale and vector-valued images. Our functional is based on the eigenvalues of the structure tensor, therefore it penalizes image variation at every point by taking into account the information from its neighborhood. It generalizes several existing variational penalties, such as the Total Variation and vectorial extensions of it. By introducing the concept of patch-based Jacobian operator, we derive an equivalent formulation of the proposed regularizer that is based on the Schatten norm of this operator. Using this new formulation, we prove convexity and develop a dual definition for the proposed energy, which gives rise to an efficient and parallelizable minimization algorithm. Moreover, we establish a connection between the minimization of the proposed convex regularizer and a generic type of nonlinear anisotropic diffusion that is driven by a spatially regularized and adaptive diffusion tensor. Finally, we perform extensive experiments with image denoising and deblurring for grayscale and color images. The results show the effectiveness of the proposed approach as well as its improved performance compared to Total Variation and existing vectorial extensions of it.

1 Introduction

This work deals with image reconstruction problems such as denoising and deblurring. We adopt their classical formulation as linear inverse problems: Let $\mathbf{u}(\mathbf{x}) = [u_1(\mathbf{x}) \dots u_M(\mathbf{x})] : \Omega \rightarrow \mathbb{R}^M$ be a generic vector-valued image with M channels that we seek to estimate. We consider that the observed image \mathbf{v} is a degraded version of \mathbf{u} according to the model: $\mathbf{z} = \mathbf{A}\mathbf{u} + \varepsilon$, where \mathbf{A} is a linear operator and ε is the measurement noise. Following the common variational approach, we estimate \mathbf{u} by minimizing a cost functional. This functional is typically the sum of a *data term* and a *regularization term*. The former measures the consistency between the estimate and the measurements, while the latter promotes certain solutions. A *regularization parameter* $\tau \geq 0$ balances the contributions of the two terms.

^{*} S.L. and A.R. contributed equally and have joint first authorship. S.L. and M.U. were supported (in part) by the Hasler Foundation and the Indo-Swiss Joint Research Program. A.R. was supported by the ERC Starting Grant 204871-HUMANIS. P.M. was partially supported by the Greek research grant COGNIMUSE under the ARISTEIA action.

A widely used choice for the regularizer is the *Total Variation* (TV) [1], which is applied on grayscale images u ($M=1$) and is defined as:

$$\text{TV}(u) = \int_{\Omega} \|\nabla u\|_2 \, d\mathbf{x}. \quad (1)$$

TV owes its popularity to its ability to reconstruct images with well-preserved and sharp edges. This is due to the fact that it involves the gradient magnitude $\|\nabla u\|_2$ and it thus undergoes an L^1 -type of behavior that does not over-penalize high variations of u . Its downside, however, is that it oversmooths homogeneous regions and creates strong staircase artifacts [2]. This behavior stems from its tendency to favor piecewise-constant solutions. Another drawback of TV 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 information from its neighborhood.

TV has been extended to general vector-valued image data in several ways, see e.g. [3–7]. Another related regularizer is the Beltrami functional [8], which has been recently generalized and unified with the Mumford-Shah functional [9]. In [10], TV is extended in an anisotropic way by incorporating the structure tensor of the image. But as in the image-driven anisotropic regularization of [5], this tensor is considered fixed and computed by the observed image. In all the above cases, the regularizers integrate a penalty of image variation that, as in TV, is completely local. On the contrary, in [11] a non-local version of TV is proposed, while in [12] an extension of the Beltrami framework that uses image patches is introduced. In [13] the authors propose a generic regularizer for vector-valued images that is based on the eigenvalues of the structure tensor, therefore it also takes into account the vicinity of each point. They show that its minimization is connected to tensor-based anisotropic diffusions.

In this work, to overcome the limitations of TV we adopt more sophisticated descriptors of image variations that generalize the gradient magnitude. We build upon the work of [13] and propose a generic convex energy functional that is based on the eigenvalues of the structure tensor. However, the current work departs from [13] in several ways. First, we provide more intuition about why the usage of the structure tensor's eigenvalues leads to effective generalizations of Total Variation. Also, the focus of [13] was in gradient descent flows of solely the regularizers, whereas in this work we combine the regularizers with data terms and we focus on the minimum rather than the flow towards the minimization. Further, we prove convexity of the proposed regularizers and we design an efficient algorithm for their minimization, which copes with their non-differentiability. Finally, in [13] the regularizers were applied only on image denoising, whereas our regularization framework is applied on more general linear inverse problems.

To the best of our knowledge, this is the first work that establishes a connection between **1)** generic anisotropic diffusion that is based on a spatially regularized and adaptive diffusion tensor (in the sense that this tensor contains convolutions with a kernel and is steered by the structure tensor field of the evolving image, as e.g. in [6, 14]) and **2)** minimization of **convex** regularizers that can be incorporated in an optimization framework and implemented efficiently using convex optimization algorithms.

2 Structure Tensor-Based Regularization

2.1 Directional Variation and Structure Tensor Revisited

In this section, we revisit and reformulate the well-established theory behind the structure tensor [14, 15], in a way that better motivates the regularizers that we will propose. The *vectorial directional derivative* of the vector-valued image \mathbf{u} in an arbitrary 2D direction \mathbf{n} ($\|\mathbf{n}\|_2=1$) is: $\partial\mathbf{u}/\partial\mathbf{n} = (J\mathbf{u})\mathbf{n}$, where $J\mathbf{u}$ is the *Jacobian matrix* of \mathbf{u} :

$$J\mathbf{u} = [\nabla u_1 \ \dots \ \nabla u_M]^T . \quad (2)$$

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 specific point \mathbf{x} . This measure is typically unreliable since it is computed by concentrating completely at the point \mathbf{x} . In order to be more robust and capture also the behavior of the image \mathbf{u} in the neighborhood of \mathbf{x} , we consider the weighted *root mean square* (RMS) of $\|\partial\mathbf{u}/\partial\mathbf{n}\|_2$, which we call (local) *directional variation*:

$$\text{RMS}_K \{ \|\partial\mathbf{u}/\partial\mathbf{n}\|_2 \} = \sqrt{K * \|\partial\mathbf{u}/\partial\mathbf{n}\|_2^2} = \sqrt{\mathbf{n}^T (S_K \mathbf{u}) \mathbf{n}} . \quad (3)$$

In the above equation $K(\mathbf{x})$ is a non-negative, rotationally symmetric convolution kernel (e.g., a 2D Gaussian) that performs the weighted averaging and $S_K \mathbf{u}$ is the so-called *structure tensor* of the image \mathbf{u} defined as:

$$S_K \mathbf{u} = K * (J\mathbf{u}^T J\mathbf{u}) . \quad (4)$$

Similarly to [6, 13], in the above definition we do not consider any pre-smoothing of the image before computing its Jacobian, since the single convolution with K seems sufficient for the needs of image regularization. Let $\lambda_+ \geq \lambda_-$ be the **eigenvalues** of $S_K(\mathbf{u})$ and θ_+, θ_- be the corresponding unit **eigenvectors**. Also, let $\omega \in (-\pi, \pi]$ be the angle between the direction vector \mathbf{n} and the eigenvector θ_+ . Using the eigendecomposition of $S_K(\mathbf{u})$, we can express the directional variation (3) as a function of the angle ω :

$$V(\omega) \triangleq \text{RMS}_K \{ \|\partial\mathbf{u}/\partial\mathbf{n}\|_2 \} = \sqrt{\lambda_+ \cos^2 \omega + \lambda_- \sin^2 \omega} . \quad (5)$$

If we consider the parametric equation $\mathbf{X}(\omega) = (\sqrt{\lambda_+} \cos \omega, \sqrt{\lambda_-} \sin \omega)$ of an ellipse with semi-major axis $\sqrt{\lambda_+}$ and semi-minor axis $\sqrt{\lambda_-}$, $V(\omega)$ can be interpreted as the distance of any point $\mathbf{X}(\omega)$ from the center of the ellipse. Therefore, $\sqrt{\lambda_+}$ corresponds to the **maximum** of the directional variation $V(\omega)$ (which is achieved for $\omega=0, \pi$), whereas $\sqrt{\lambda_-}$ to the **minimum** of $V(\omega)$ (which is achieved for $\omega=\pm\pi/2$).

2.2 Proposed Class of Regularizers

Based on the above analysis, we conclude that the vector $\sqrt{\lambda} \triangleq (\sqrt{\lambda_+}, \sqrt{\lambda_-})$ is a synopsis of the function of local directional variation $V(\omega)$: it consists of the upper and lower bounds of this function. Therefore, we propose to generalize the Total Variation (1)

via replacing the gradient magnitude $\|\nabla u\|_2$ by ℓ_p norms of $\sqrt{\lambda}$. More precisely, we propose the following type of regularizers, with $p \geq 1$:

$$E_p(\mathbf{u}) = \int_{\Omega} \|\sqrt{\lambda}\|_p \, d\mathbf{x} = \int_{\Omega} \|(\sqrt{\lambda_+}, \sqrt{\lambda_-})\|_p \, d\mathbf{x}. \quad (6)$$

These norms measure the local variation of the image at each point more 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 not only on the maximum but also on the minimum of the directional variation. For instance, the response of these measures behaves differently at image edges than image corners.

We note that all the regularizers of the type (6) generalize TV. The reason is that for $M=1$ (grayscale images), if $K(\mathbf{x})$ is chosen to be the Dirac delta $\delta(\mathbf{x})$ (degenerated case where no convolution takes place at the computation of the structure tensor), then $\lambda_+ = \|\nabla u\|_2^2$ and λ_- is always 0. Therefore $\|\sqrt{\lambda}\|_p = \|\nabla u\|_2$ for any $p \geq 1$.

Next, we describe some interesting cases of the proposed regularizers (6). For the following three cases of $E_p(\mathbf{u})$, the corresponding norms describe specific measures of the directional variation $V(\omega)$:

- $p=1$: $\|\sqrt{\lambda}\|_1 = \sqrt{\lambda_+} + \sqrt{\lambda_-}$ corresponds (up to the scale factor 1/2) to the **mid-range** of $V(\omega)$, i.e. the average of minimum and maximum values of $V(\omega)$.
- $p=2$: $\|\sqrt{\lambda}\|_2 = \sqrt{\lambda_+ + \lambda_-}$ corresponds (up to the scale factor $1/\sqrt{2}$) to the **RMS value** of $V(\omega)$, as it can be easily verified using Eq. (5).
- $p=\infty$: $\|\sqrt{\lambda}\|_{\infty} = \sqrt{\lambda_+}$ corresponds to the **maximum** of $V(\omega)$.

Invariance Properties. Since our regularizers are generalizations of TV, one should expect that they also share the same invariance properties. Two of the most favorable ones are the rotation invariance and contrast covariance (1-homogeneity), which according to Proposition 1 are indeed preserved (see Supplementary Material for the proof).

Proposition 1. *The energy functional (6) is rotation invariant and contrast covariant.*

2.3 Connections to Tensor-Based Anisotropic Diffusion and Previous Work

The proposed class of regularizers $E_p(\mathbf{u})$ (6) is a special case of the more generic form proposed in [13]: $E(\mathbf{u}) = \int_{\Omega} \psi(\lambda_+, \lambda_-) \, d\mathbf{x}$. This special case corresponds to cost functions of the form $\psi(\lambda_+, \lambda_-) = \|(\sqrt{\lambda_+}, \sqrt{\lambda_-})\|_p$.

In order to make the cost function in the proposed regularizers differentiable, let us consider the relaxation $E_{p,\epsilon}(\mathbf{u})$ that arises by setting $\psi(\lambda_+, \lambda_-) = \varphi_{p,\epsilon}(\lambda_+, \lambda_-) \triangleq \|(\sqrt{\epsilon + \lambda_+}, \sqrt{\epsilon + \lambda_-})\|_p$, where $\epsilon > 0$ is a small constant. Note that we need this relaxation only to establish connections to anisotropic diffusion and not for the actual optimization, since our optimization algorithm, described in Section 4, can cope with the non-differentiability of the functionals. By applying [13, Theorem 1], we find the relation of minimizing the proposed regularizers with anisotropic diffusion:

Corollary 1. *The functional gradient of $E_{p,\epsilon}$ w.r.t. each image component u_i is:*

$$\frac{\delta E_{p,\epsilon}}{\delta u_i} = -\operatorname{div}(D\nabla u_i), \quad D = K * \left(2 \frac{\partial \varphi_{p,\epsilon}}{\partial \lambda_+} \boldsymbol{\theta}_+ \otimes \boldsymbol{\theta}_+ + 2 \frac{\partial \varphi_{p,\epsilon}}{\partial \lambda_-} \boldsymbol{\theta}_- \otimes \boldsymbol{\theta}_- \right). \quad (7)$$

This gradient is a nonlinear anisotropic diffusion term, where the diffusion tensor D contains convolutions with the kernel K and depends on the structure tensor of the image. For the following characteristic choices of p , the diffusion tensor is given as:

– $p=1$: $D = K * \left(\frac{1}{\sqrt{\epsilon+\lambda_+}} \boldsymbol{\theta}_+ \otimes \boldsymbol{\theta}_+ + \frac{1}{\sqrt{\epsilon+\lambda_-}} \boldsymbol{\theta}_- \otimes \boldsymbol{\theta}_- \right)$. This tensor is adapting on the image structures in a conceptually similar way to tensor-based anisotropic diffusion methods, such as [6, 14]: 1) in the homogeneous regions (small λ_+, λ_-) it is strong and isotropic, 2) near the edges (large λ_+ , small λ_-) it is weaker and mainly oriented by the edges, whereas 3) near the corners (large λ_+, λ_-) it is even weaker.

– $p=2$: $D = \left(K * \frac{1}{\sqrt{2\epsilon + K * \sum_i \|\nabla u_i\|_2^2}} \right) I_{2 \times 2}$. This tensor is always isotropic, it thus corresponds to a diffusion coefficient. Similarly to nonlinear diffusion methods, such as [16, 17], this coefficient is strong in the homogeneous regions, whereas weaker near edges.

– $p=\infty$: $D = K * \left(\frac{1}{\sqrt{\epsilon+\lambda_+}} \boldsymbol{\theta}_+ \otimes \boldsymbol{\theta}_+ \right)$. This tensor is always highly anisotropic and oriented perpendicular to image edges.

Further Relations to Previous Work. As already stated, the proposed regularizers are special cases of the more generic functional of [13]. Furthermore, the special subcase of $p = 1$ corresponds to the so-called **Tensor Total Variation** of [13]. In addition, several other variational methods emerge as special cases of the proposed regularizers. The kernel K that corresponds to all these cases is the Dirac delta $\delta(\mathbf{x})$, which means that the regularization does not exploit information from the neighborhood of each point and is thus less coherent. As already described in Section 2.2, if we set $K(\mathbf{x})=\delta(\mathbf{x})$ and $M=1$ (grayscale images) then for all choices of $p \geq 1$ we recover the **Total Variation** [1]. The case of $K(\mathbf{x})=\delta(\mathbf{x})$, $M > 1$ and $p=2$ corresponds to the usually called **Vectorial TV** (TV-F) [3, 4], which is the most common extension of TV to vector-valued images. Finally, the case $K(\mathbf{x})=\delta(\mathbf{x})$, $M > 1$ and $p=\infty$ corresponds to the method of [7], which the authors call **Natural Vectorial TV** (TVJ).

3 Patch-Based Jacobian and the Discrete Structure Tensor

In this section, we introduce a generalization of the Jacobian of an image, based on local weighted patches (see e.g. [12, 13]). 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 convolution kernel K . Then, we employ it to express the structure tensor in a novel way, which finally leads us to derive an equivalent definition of the proposed regularizers. This alternative definition provides more intuition, facilitates the proof of convexity and opens the way for an efficient optimization strategy.

Hereafter, we will focus on the discrete formulation of the image reconstruction problem. We consider that the discretized vector-valued image \mathbf{u} is defined on a rectangular grid with unary steps and that the corresponding intensities of each channel m of \mathbf{u} ($m = 1, \dots, M$) are rasterized in the vector \mathbf{u}_m of size N . 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 the coordinates of that pixel. Furthermore, we consider

that the convolution kernel K (see Eq. (4)) has been discretized and truncated in order to have compact support $\mathcal{S} = \{-L_K, \dots, L_K\}^2$, where L_K is a non-negative integer.

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 $L = (2L_K + 1)^2$. For each pixel n we denote by $[\mathbf{J}_K \mathbf{u}]_n$ the element of $\mathbf{J}_K \mathbf{u}$ that corresponds to that pixel and we construct it by: **1**) taking the discrete versions of the $M \times 2$ Jacobian matrices (2) of \mathbf{u} for all the pixels $\{\mathbf{x}_n - \mathbf{y} : \mathbf{y} \in \mathcal{S}\}$ in the \mathcal{S} -neighborhood of pixel \mathbf{x}_n , **2**) weighting these matrices with the window function $\mathbf{w}[\mathbf{y}] \triangleq \sqrt{K[\mathbf{y}]}$ and **3**) stacking all these matrices vertically in the matrix $[\mathbf{J}_K \mathbf{u}]_n$, whose dimension is $(LM) \times 2$. Formally, the patch-based Jacobian can be defined as:

$$[\mathbf{J}_K \mathbf{u}]_n^T = \begin{bmatrix} [P_{\mathbf{y}_1} \circ \mathbf{D}_h \mathbf{u}_1]_n \cdots [P_{\mathbf{y}_L} \circ \mathbf{D}_h \mathbf{u}_1]_n \cdots [P_{\mathbf{y}_1} \circ \mathbf{D}_h \mathbf{u}_M]_n \cdots [P_{\mathbf{y}_L} \circ \mathbf{D}_h \mathbf{u}_M]_n \\ [P_{\mathbf{y}_1} \circ \mathbf{D}_v \mathbf{u}_1]_n \cdots [P_{\mathbf{y}_L} \circ \mathbf{D}_v \mathbf{u}_1]_n \cdots [P_{\mathbf{y}_1} \circ \mathbf{D}_v \mathbf{u}_M]_n \cdots [P_{\mathbf{y}_L} \circ \mathbf{D}_v \mathbf{u}_M]_n \end{bmatrix}, \quad (8)$$

where $\mathbf{D}_h, \mathbf{D}_v$ are the two components of the discrete gradient, the shift vectors \mathbf{y}_l ($l = 1, \dots, L$) are the elements of the lattice \mathcal{S} , and $P_{\mathbf{y}_l}$ are weighted shift operators. The latter are designed to properly handle the image boundaries according to the assumed extension (*e.g.*, mirroring) and are defined as:

$$[P_{\mathbf{y}_l} \circ \mathbf{D}_h \mathbf{u}_m]_n = \mathbf{w}[\mathbf{y}_l] \mathbf{D}_h \{\mathbf{u}_m\}[\mathbf{x}_n - \mathbf{y}_l]. \quad (9)$$

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 $\mathbf{X}_n, \mathbf{Y}_n \in \mathbb{R}^{(LM) \times 2} \forall n = 1, \dots, N$. Then we have:

$$\langle \mathbf{X}, \mathbf{Y} \rangle_{\mathcal{X}} = \sum_{n=1}^N \text{tr}(\mathbf{Y}_n^T \mathbf{X}_n) \quad (10) \quad \text{and} \quad \|\mathbf{X}\|_{\mathcal{X}} = \sqrt{\langle \mathbf{X}, \mathbf{X} \rangle_{\mathcal{X}}}, \quad (11)$$

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

The adjoint of \mathbf{J}_K is the discrete linear operator $\mathbf{J}_K^* : \mathcal{X} \mapsto \mathbb{R}^{NM}$, defined by:

$$\langle \mathbf{Y}, \mathbf{J}_K \mathbf{u} \rangle_{\mathcal{X}} = \langle \mathbf{J}_K^* \mathbf{Y}, \mathbf{u} \rangle_2. \quad (12)$$

The following Proposition expresses \mathbf{J}_K^* in a form that facilitates its computation (see Supplementary Material for the proof).

Proposition 2. *The adjoint operator \mathbf{J}_K^* of the patch-based Jacobian is given by:*

$$[\mathbf{J}_K^* \mathbf{Y}]_{(n,m)} = \sum_{l=1}^L -\text{div} [P_{\mathbf{y}_l}^* \circ \mathbf{Y}^{((m-1)L+l,:)}]_n, \quad (13)$$

where div is the discrete divergence, P^* is the adjoint of the shift operator P , and $\mathbf{Y}_n^{(k,:)}$ corresponds to the k -th row of the n -th matrix component, $\mathbf{Y}_n \in \mathbb{R}^{(LM) \times 2}$, of \mathbf{Y} .

Having introduced the necessary tools, we can now express the structure tensor in a novel way. This is done in Proposition 3 (see Supplementary Material for a proof).

Proposition 3. *Let $[\mathbf{S}_K \mathbf{u}]_n$ be the discretized structure tensor at pixel n , which is defined by adopting discrete derivatives and discrete convolution in (2) and (4), respectively. Then, it can be written in terms of the patch-based Jacobian as:*

$$[\mathbf{S}_K \mathbf{u}]_n = [\mathbf{J}_K \mathbf{u}]_n^T [\mathbf{J}_K \mathbf{u}]_n. \quad (14)$$

Since λ_+, λ_- are the eigenvalues of $[\mathbf{S}_K \mathbf{u}]_n$, the singular values of $[\mathbf{J}_K \mathbf{u}]_n$ are $\sqrt{\lambda_+}, \sqrt{\lambda_-}$. This connection permits us to use Schatten norms [18] and the patch-based Jacobian so as to write the proposed regularizers (6) (after discretization) as:

$$E_p(\mathbf{u}) = \sum_{n=1}^N \|[\mathbf{J}_K \mathbf{u}]_n\|_{\mathcal{S}_p}, \quad \text{with } p \geq 1. \quad (15)$$

Note that for a matrix \mathbf{Z} , its **Schatten norm** of order p (\mathcal{S}_p norm) denoted by $\|\mathbf{Z}\|_{\mathcal{S}_p}$, is defined as $\|\boldsymbol{\sigma}(\mathbf{Z})\|_p$, with $\boldsymbol{\sigma}(\mathbf{Z})$ the vector of the singular values of \mathbf{Z} . This equivalent formulation of $E_p(\mathbf{u})$ provides more intuition about the fact that the proposed regularizers are effective generalizations of TV. More precisely, $[\mathbf{J}_K \mathbf{u}]_n$ 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, an important contribution of the above result is that the expression (6), which involves the eigenvalues of the non-linear structure tensor, has been transformed to the expression (15) that is much easier to handle, since it depends on a norm of a linear operator acting on \mathbf{u} . We refer to the proposed regularizers as STV-[k] (Structure tensor Total Variation) where the character [k] denotes the order of the Schatten norm. For example, for the cases of $p=1,2$ and ∞ we use the notations STV-N (Nuclear norm), STV-F (Frobenius norm) and STV-S (Spectral norm) respectively.

It has now become straight-forward to show the following important result:

Proposition 4. *The regularizer $E_p(\mathbf{u})$ is convex w.r.t $\mathbf{u} \forall p \geq 1$.*

Proof. The regularizer of Eq. (15) is clearly convex since it results as the composition of a norm (mixed ℓ_1 - \mathcal{S}_p norm; see (17) for its definition) and the linear operator \mathbf{J}_K .

4 Energy Minimization Strategy

4.1 Proximal Map Evaluation

In this section we propose an efficient algorithm that provides a numerical solution to the following problem, for any $p \geq 1$:

$$\arg \min_{\mathbf{u} \in \mathbb{R}^{NM}} \frac{1}{2} \|\mathbf{u} - \mathbf{z}\|_2^2 + \psi(\mathbf{u}), \quad \text{with } \psi(\mathbf{u}) \triangleq \tau E_p(\mathbf{u}) + \iota_{\mathcal{C}}(\mathbf{u}), \quad (16)$$

where \mathcal{C} is a convex set that represents additional constraints on the solution 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. Note that the case of no constraints is simply the special case $\mathcal{C} = \mathbb{R}^{NM}$. The solution of (16) corresponds to evaluating the proximal map [19] of the function ψ at \mathbf{z} and arises in most linear inverse imaging problems, including the ones considered by this work.

To proceed with our minimization approach, we write the energy E_p in the compact form $E_p(\mathbf{u}) = \|\mathbf{J}_K \mathbf{u}\|_{1,p}$, where $\|\cdot\|_{1,p}$ 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

$$\|\mathbf{X}\|_{1,p} = \sum_{n=1}^N \|\mathbf{X}_n\|_{\mathcal{S}_p}. \quad (17)$$

Next, we rely on the following lemma to derive a dual formulation of our problem.

Lemma 1 ([20]). *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 norm $\|\cdot\|_{\infty,q}$ is dual to the mixed norm $\|\cdot\|_{1,p}$.*

Using Lemma 1 and the fact that the dual of the dual norm is the original norm [21], we write (17) in the equivalent form:

$$\|\mathbf{X}\|_{1,p} = \max_{\Omega \in \mathcal{B}_{\infty,q}} \langle \Omega, \mathbf{X} \rangle_{\mathcal{X}}, \quad (18)$$

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

$$\mathcal{B}_{\infty,q} \triangleq \{ \Omega = [\Omega_1^T, \dots, \Omega_N^T]^T \in \mathcal{X} : \|\Omega_n\|_{\mathcal{S}_q} \leq 1, \forall n = 1, \dots, N \}. \quad (19)$$

Note that from (19), it is clear that the orthogonal projection onto $\mathcal{B}_{\infty,q}$ can be obtained by projecting separately each submatrix Ω_n onto a unit-norm \mathcal{S}_q ball ($\mathcal{B}_{\mathcal{S}_q}$).

Combining (12) and (18) we re-write (16) as

$$\hat{\mathbf{u}} = \arg \min_{\mathbf{u} \in \mathcal{C}} \frac{1}{2} \|\mathbf{u} - \mathbf{z}\|_2^2 + \tau \max_{\Omega \in \mathcal{B}_{\infty,q}} \langle \mathbf{J}_K^* \Omega, \mathbf{u} \rangle_2. \quad (20)$$

This formulation naturally leads us to the following minimax problem:

$$\min_{\mathbf{u} \in \mathcal{C}} \max_{\Omega \in \mathcal{B}_{\infty,q}} \mathcal{L}(\mathbf{u}, \Omega), \quad (21)$$

where $\mathcal{L}(\mathbf{u}, \Omega) = \frac{1}{2} \|\mathbf{u} - \mathbf{z}\|_2^2 + \tau \langle \mathbf{J}_K^* \Omega, \mathbf{u} \rangle_2$. The function \mathcal{L} is strictly convex in \mathbf{u} and concave in Ω , and thus, we have the guarantee that a saddle-value of \mathcal{L} is attained [21]. Therefore, the order of the minimum and the maximum in (21) does not affect the solution and $\hat{\mathbf{u}}$ can be equivalently obtained by solving the problem:

$$\max_{\Omega \in \mathcal{B}_{\infty,q}} \min_{\mathbf{u} \in \mathcal{C}} \left(\frac{1}{2} \|\mathbf{u} - (\mathbf{z} - \tau \mathbf{J}_K^* \Omega)\|_2^2 + \frac{1}{2} \|\mathbf{z}\|_2^2 - \frac{1}{2} \|(\mathbf{z} - \tau \mathbf{J}_K^* \Omega)\|_2^2 \right). \quad (22)$$

The inner minimization in (22) has an exact solution:

$$\hat{\mathbf{u}} = \Pi_{\mathcal{C}} \left(\mathbf{z} - \tau \mathbf{J}_K^* \hat{\Omega} \right), \quad (23)$$

where $\Pi_{\mathcal{C}}$ denotes the orthogonal projection onto the convex set \mathcal{C} , while $\hat{\Omega}$ is the maximizer of the dual problem:

$$\max_{\Omega \in \mathcal{B}_{\infty,q}} \left(\phi(\Omega) \triangleq \frac{1}{2} \|\Pi_{\mathcal{C}}(\mathbf{c}) - \mathbf{c}\|_2^2 + \frac{1}{2} \|\mathbf{z}\|_2^2 - \frac{1}{2} \|\mathbf{c}\|_2^2 \right), \quad (24)$$

where $\mathbf{c} = \mathbf{z} - \tau \mathbf{J}_K^* \Omega$. Contrary to the primal problem (16), where the function to be minimized is not continuously differentiable, the dual one in (24) involves the function ϕ which is smooth and has a well defined gradient. To compute it, we use the result in [22, Lemma 4.1], according to which the gradient of a function $h(\mathbf{x}) = \|\mathbf{x} - \Pi_{\mathcal{C}}(\mathbf{x})\|_2^2$ is equal to: $\nabla h(\mathbf{x}) = 2(\mathbf{x} - \Pi_{\mathcal{C}}(\mathbf{x}))$. Based on that, we get:

$$\nabla \phi(\Omega) = \tau \mathbf{J}_K \Pi_{\mathcal{C}}(\mathbf{z} - \tau \mathbf{J}_K^* \Omega). \quad (25)$$

Algorithm 1: Evaluation of the proximal map of $\psi(\mathbf{u})$.

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

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

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

$$\begin{aligned} \mathbf{v} &= \Pi_C(\mathbf{z} - \tau \mathbf{J}_K^* \Psi_n); \\ \Omega_n &\leftarrow \Pi_{\mathcal{S}_q}\left(\Psi_n + \frac{1}{8\tau} \mathbf{J}_K \mathbf{v}\right); \\ t_{n+1} &\leftarrow \frac{1 + \sqrt{1 + 4t_n^2}}{2}; \\ \Psi_{n+1} &\leftarrow \Omega_n + \left(\frac{t_n - 1}{t_{n+1}}\right)(\Omega_n - \Omega_{n-1}); \\ n &\leftarrow n + 1; \end{aligned}$$

end

return $\hat{\mathbf{u}} = \Pi_C(\mathbf{z} - \tau \mathbf{J}_K^* \Omega_{n-1})$;

Then, we use (25) to design a gradient-based algorithm that solves (24). The solution of our primal problem (16) is obtained in two steps: **1**) we find the maximizer of the dual objective function (24), and **2**) we obtain the solution using (23).

Since (24) does not have a closed-form solution (\mathbf{J}_K has not a stable inverse), we employ Nesterov's iterative method [23] for smooth functions. This is a gradient-based scheme that exhibits state-of-the-art convergence rates of one order higher than the standard gradient-ascent method. A detailed description of the overall algorithm is provided in Algorithm 1. Note that for the implementation of Algorithm 1, we need to perform a projection of a matrix onto a Schatten norm ball. This is discussed in the next section.

4.2 Efficient Projection of Rectangular Matrices

Let $\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}$ with an SVD decomposition $\mathbf{X} = \mathbf{U} \Sigma \mathbf{V}^T$ and $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n)$ with $n = \min(n_1, n_2)$. According to [20, Proposition 1], the projection of \mathbf{X} onto the unit-norm \mathcal{S}_q ball is computed as:

$$\Pi_{\mathcal{S}_q}(\mathbf{X}) = \mathbf{U} \Sigma_q \mathbf{V}^T, \quad (26)$$

where $\Sigma_q = \text{diag}(\sigma_q)$ and σ_q is the projection of the singular values of Σ onto the ℓ_q unit-norm ball $\mathcal{B}_q = \{\sigma \in \mathbb{R}_+^n : \|\sigma\|_q \leq 1\}$. The projection in (26) requires 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 note that the matrix $\mathbf{X}^T \mathbf{X}$ is $n_2 \times n_2$ symmetric with an eigenvalue decomposition $\mathbf{V} \Sigma^2 \mathbf{V}^T$. Therefore, for $n_2 = 2$ both \mathbf{V} and Σ can be computed in closed form. Now, if Σ^+ is the pseudoinverse matrix of Σ , defined as: $\Sigma^+ = \text{diag}(\sigma_1^{-1}, \dots, \sigma_k^{-1}, 0, \dots, 0)$, with σ_k the smallest nonzero singular value, then $\mathbf{U} = \mathbf{X} \mathbf{V} \Sigma^+$. Using this result we write (26) as:

$$\Pi_{\mathcal{S}_q}(\mathbf{X}) = \mathbf{X} \mathbf{V} \Sigma^+ \Sigma_q \mathbf{V}^T, \quad (27)$$

and we avoid the computation of \mathbf{U} . We note that the same idea was explored in [7] for efficiently computing the projection step that arises in the minimization of TVJ.

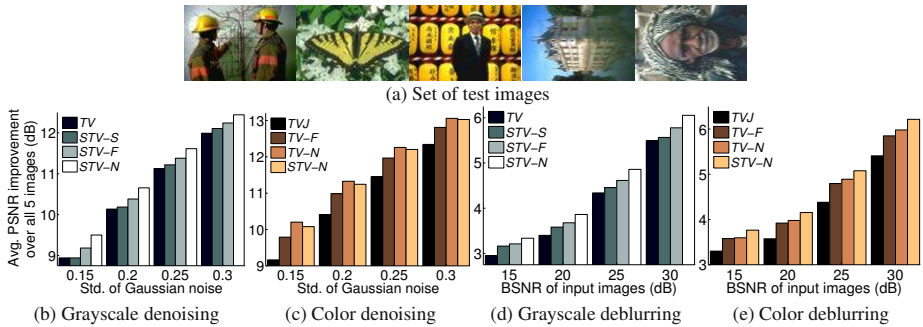


Fig. 1. Performance measures for different regularization methods

4.3 General Linear Inverse Problems

Algorithm 1 applies only in cases where no linear operator is involved in the data term. For general inverse problems, under the proposed regularization framework, one needs to solve a minimization problem of the form:

$$\arg \min_{\mathbf{u} \in \mathbb{R}^{NM}} \frac{1}{2} \|\mathbf{A}\mathbf{u} - \mathbf{z}\|_2^2 + \tau E_p(\mathbf{u}) + \iota_C(\mathbf{u}), \forall p \geq 1, \quad (28)$$

where \mathbf{A} is a linear degradation operator, which for most practical cases is ill-conditioned. To solve this type of problems we employ the MFISTA algorithm [22], which exhibits state-of-the-art convergence rates. Nevertheless, our algorithm is still a critical part, since the main step of MFISTA requires the evaluation of the proximal map that we investigated in Section 4.1. For a detailed description of the MFISTA approach we refer the readers to the Supplementary Material accompanying this paper.

5 Experimental Results

To evaluate the effectiveness of the proposed generic regularization framework, we report results for the problems of gray/color image denoising and deblurring. For both linear inverse problems we use the set of images shown in Fig. 1(a), taken from the Berkeley BSDS500 dataset. In the image denoising setting we consider four different standard deviations of Gaussian noise, $\sigma_w = \{0.15, 0.2, 0.25, 0.3\}$. In the image deblurring setting we consider a Gaussian blur kernel, which has a support of 9×9 pixel and a standard deviation $\sigma_b = 6$ pixels, and four noise levels corresponding to BSNR = $\{15, 20, 25, 30\}$ dB respectively, where BSNR is the *Blurred Signal to Noise Ratio*, defined as $\text{BSNR} = \text{var}(\mathbf{A}\mathbf{u}) / \sigma_w^2$.

In Figs. 1(b)-1(e) we report the average performance, in terms of *Peak Signal to Noise Ratio* (PSNR), over all tested images. For the grayscale experiments, we compare TV against three variants of our functional (STV-S, STV-F, STV-N). For the color case, we compare the results we obtained with our STV-N regularizer against those obtained using TVJ [7] and TV-F [3]. In these comparisons, we also include the variant of STV-N where no smoothing is involved in the computation of the structure tensor (TV-N), which is also a novel regularizer. For the sake of consistency among comparisons,




Fig. 2. Grayscale (*first row*) and Color (*second row*) image denoising examples. The PSNR and the Structural Similarity index (SSIM) are also reported.




Fig. 3. Grayscale (*first row*) and Color (*second row*) image deblurring examples. The PSNR and SSIM measures are also reported.

the reported results for each regularizer were obtained using the individualized regularization parameter that gives the best PSNR performance. Moreover, all reconstructions are performed under box constraints, meaning that the restored intensities must lie in the convex set $\mathcal{C} = \{\mathbf{u} \in \mathbb{R}^N | u_n \in [0, 1] \forall n = 1, \dots, N\}$. Finally, in all the STV regularizers, we choose the structure tensor's convolution kernel to be a Gaussian with a support of 3×3 pixels.

From the reported results, we observe that in the grayscale case the best performance for both image denoising and deblurring is achieved by STV-N. On the other hand, TV has the worst performance, especially in deblurring, since in denoising its performance is very close to STV-S. In the color denoising experiments, TV-N performs slightly better than STV-N, and both are superior than the competitive regularizers. However, when we consider the image deblurring problem, STV-N behaves better than TV-N and provides the best results. This can be attributed to the fact that deblurring is a more ill-conditioned problem and, thus, the use of a convolution kernel K is more critical.

Finally, apart from the quantitative comparisons, conclusions for the effectiveness of the proposed approach can be drawn by a visual inspection of the results provided in Figs. 2-3. From these examples we can verify that the proposed STV regularizers perform better in reducing the staircase effects of other Total Variation methods and better reconstruct the edges and the other image structures.

6 Conclusions

In this work we introduced a family of regularizers that is based on the eigenvalues of the structure tensor. In image denoising and deblurring problems these regularizers can more accurately restore image edges than TV and its vectorial extensions, and, thus, lead to improved results. Furthermore, based on a novel formulation of the structure tensor, we proved convexity for the regularizers and designed an efficient primal-dual algorithm for their minimization. Since TV-based reconstructions are used in a host of imaging applications, an interesting research direction is to investigate whether our regularizers can also lead to an improved performance in other inverse problems, as well. This will be the subject of future work.

References

1. Rudin, L., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Physica D* 60, 259–268 (1992)
2. Chan, T., Marquina, A., Mulet, P.: High-order total variation-based image restoration. *SIAM J. Sci. Comput.* 22, 503–516 (2000)
3. Sapiro, G.: Color snakes. *Comp. Vision and Image Understanding* 68(2) (1997)
4. Blomgren, P., Chan, T.: Color TV: Total Variation methods for restoration of vector-valued images. *IEEE Trans. on Image Processing* 7(3), 304–309 (1998)
5. Weickert, J., Schnörr, C.: A theoretical framework for convex regularizers in PDE-based computation of image motion. *Int. Journ. of Computer Vision* 45(3), 245–264 (2001)
6. Tschumperlé, D., Deriche, R.: Vector-valued image regularization with PDE's: A common framework for different applications. *IEEE T-PAMI* 27(4), 506–517 (2005)
7. Goldluecke, B., Strelakovski, E., Cremers, D.: The natural vectorial total variation which arises from geometric measure theory. *SIAM J. Imaging Sci.* 5, 537–563 (2012)
8. Sochen, N., Kimmel, R., Malladi, R.: A general framework for low level vision. *IEEE Trans. on Image Processing* 7, 310–338 (1998)
9. Sochen, N., Bar, L.: The Beltrami-Mumford-Shah functional. In: *Scale Space and Variational Methods in Computer Vision*, pp. 183–193 (2012)
10. Grasmair, M., Lenzen, F.: Anisotropic total variation filtering. *Applied Mathematics & Optimization* 62, 323–339 (2010)
11. Gilboa, G., Osher, S.: Nonlocal operators with applications to image processing. *Multiscale Modeling & Simulation* 7(3), 1005–1028 (2008)
12. Wetzler, A., Kimmel, R.: Efficient Beltrami flow in patch-space. In: *Scale Space and Variational Methods in Computer Vision*, pp. 134–143 (2012)
13. Roussos, A., Maragos, P.: Tensor-based image diffusions derived from generalizations of the total variation and Beltrami functionals. In: *Proc. Int. Conf. on Image Processing* (2010)
14. Weickert, J.: *Anisotropic Diffusion in Image Processing*. Teubner, Stuttgart (1998)
15. Jähne, B.: *Digital Image Processing*. Springer (2002)

16. Perona, P., Malik, J.: Scale space and edge detection using anisotropic diffusion. *IEEE T-PAMI* 12(7), 629–639 (1990)
17. Catté, F., Lions, P., Morel, J., Coll, T.: Image selective smoothing and edge detection by nonlinear diffusion. *SIAM Journ. Numer. Anal.* 29(1), 182–193 (1992)
18. Bhatia, R.: *Matrix Analysis*. Springer (1997)
19. Combettes, P.L., Wajs, V.R.: Signal recovery by proximal forward-backward splitting. *Multiscale Model. Simul.* 4(4), 1168–1200 (2005)
20. Lefkimiatis, S., Ward, J., Unser, M.: Hessian Schatten-norm regularization for linear inverse problems. *IEEE Trans. Image Processing* 22(5), 1873–1888 (2013)
21. Rockafellar, R.T.: *Convex Analysis*. Princeton Univ. Press, Princeton (1970)
22. Beck, A., Teboulle, M.: Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems. *IEEE Trans. Image Processing* 18, 2419–2434 (2009)
23. Nesterov, Y.: A method for solving a convex programming problem with convergence rates $O(1/k^2)$. *Soviet Math. Dokl* 27, 372–376 (1983)