Non-local Structure Tensor Functionals for Image Regularization

Stamatios Lefkimmiatis, Member, IEEE, and Stanley Osher

Authors' contact information:

University of California, Los Angeles (UCLA) Department of Mathematics Los Angeles, United States Tel: +1(310)8252601 Email: {stamatis,sjo}@math.ucla.edu

Abstract

We present a non-local regularization framework that we apply to inverse imaging problems. As opposed to existing non-local regularization methods that rely on the graph gradient as the regularization operator, we introduce a family of non-local energy functionals that involves the standard image gradient. Our motivation for designing these functionals is to exploit at the same time two important properties inherent in natural images, namely the local structural image regularity and the non-local image selfsimilarity. To this end, our regularizers employ as their regularization operator a novel non-local version of the structure tensor. This operator performs a non-local weighted average of the image gradients computed at every image location and, thus, is able to provide a robust measure of image variation. Further, we show a connection of the proposed regularizers to the Total Variation semi-norm and prove convexity. The convexity property allows us to employ powerful tools from convex optimization in order to design an efficient minimization algorithm. Our algorithm is based on a splitting variable strategy which leads to an augmented Lagrangian formulation. To solve the corresponding optimization problem we employ the alternating-direction methods of multipliers. Finally, we present extensive experiments on several inverse imaging problems, where we compare our regularizers with other competing local and non-local regularization approaches. Our results are shown to be systematically superior, both quantitatively and visually.

The authors are with the Department of Mathematics, University of California, Los Angeles, United States (email: stamatis@ math.ucla.edu;sjo@math.ucla.edu).

SL was supported by the Swiss National Science Foundation (SNF) under grant P300P2_151325 and SO was partially supported by the Keck Foundation.

Non-local Structure Tensor Functionals for Image Regularization

I. INTRODUCTION

Inverse problems typically arise in several image processing and computer vision applications, including image restoration, image inpainting, image segmentation, optical flow estimation, stereo and 3D reconstruction, etc. In these cases, the task is to estimate underlying features of interest from partial or indirect measurements. In practice, the majority of inverse imaging problems are ill-posed [1]. This implies that in order to obtain a physically or statistically meaningful solution, some type of prior information about the underlying image must be taken into account.

Among the available strategies that one can follow to deal with inverse problems, the variational approach is one of the most widely used. Under this framework, image recovery is cast as the minimization of an energy functional whose minimizer corresponds to the desired solution. Here, of significant importance is the proper selection of the regularizer, which is responsible to favor certain reconstructions. This has triggered an increasing research interest in the design of regularization functionals that can accurately model important properties of natural images.

One of the most successful regularization criteria that have been proposed in the literature is the Total variation (TV) semi-norm [2]. TV exploits the *local structural regularity* property, that is natural images are typically smooth everywhere except to edges, and employs the image gradient to penalize the image variation with an L_1 -type of penalty. This leads to reconstructions that feature sharp and well-preserved image edges. Inevitably, TV has also some limitations. The main one is that in certain cases it can over-smooth homogeneous regions and create staircase artifacts [3]. To address this issue several modifications/extensions of TV have been proposed. These involve either first-order functionals whose goal is to provide an improved measure of image variation [4]–[7], or higher-order functionals that favor piecewise-smooth instead of piecewise-constant solutions (see [3], [8], [9] and references therein).

The above functionals are considered as local or semi-local, since they involve a regularization operator that acts on a restricted region of the image domain. Recently, a new regularization paradigm has been introduced in [10]–[12], where non-local operators are used instead for defining energy functionals. These operators are non-local in the sense that they allow interactions between image points that can possibly be located far apart. The motivation behind this approach is that natural images exhibit a *non-local self*-

similarity property. This means that images often consist of localized patterns that repeat themselves at distant locations in the image domain. Therefore, non-local regularizers can effectively model long-range dependencies and lead to improved reconstruction results.

A. Contributions

In this work we combine ideas both from the local and non-local regularization settings and introduce a novel non-local family of regularizers. Our regularizers differ from the existing non-local ones in the sense that they involve a non-local operator which depends on the image gradient rather than the graph gradient. The motivation is that this way we can exploit at the same time both the local structural image regularity and the non-local image self-similarity properties. In detail, our key contributions are the following:

- 1. We extend our previous work [7] and design penalties that employ a non-local version of the structure tensor as the regularization operator. This operator performs a non-local weighted average of the image gradients computed at every image point and, thus, provides a more robust measure of image variation.
- 2. We prove that our regularization criteria are convex and show a connection with the total variation semi-norm.
- 3. We introduce the concept of the non-local gradient as the non-local extension of the discrete image gradient. We use this operator to re-express our discrete energy functionals in a form that facilitates its efficient minimization.
- 4. We develop an efficient optimization algorithm that is based on an augmented Lagrangian formulation of the problem.
- 5. We provide extensive comparisons of our method with alternative regularizers on several inverse imaging problems.

II. REGULARIZATION OF INVERSE PROBLEMS

A. Image-Formation Model and Variational Recovery

We will be concerned with the recovery of signals whose measurements are described by a linear observation model of the form:

$$\boldsymbol{v}\left(\boldsymbol{x}\right) = A\boldsymbol{u}\left(\boldsymbol{x}\right) + \boldsymbol{n}\left(\boldsymbol{x}\right). \tag{1}$$

Here, $u(x) = [u^1(x) \ u^2(x) \ \dots \ u^C(x)] : \mathbb{R}^d \mapsto \mathbb{R}^C$ represents the generic underlying vector-valued image consisting of C channels, v are the measurements, A is a linear operator that corresponds to the impulse response of the imaging device, and n is a term that accounts for all possible errors during the acquisition. Hereafter, we will assume n to be an independent and identically distributed (i.i.d) Gaussian term.

Despite the linear nature of the forward model in (1), the recovery of u from v is non-trivial. This is due to the measurement noise and the operator A, which typically is either ill-conditioned or singular. Therefore, a unique and stable solution does not exist. One way to tackle this problem is to cast the estimation of u as the minimization of an objective function of the form:

$$\mathcal{E}\left(\boldsymbol{u}\right) = \frac{1}{2} \left\|\boldsymbol{v} - A\boldsymbol{u}\right\|_{2}^{2} + \tau \psi\left(\boldsymbol{u}\right).$$
⁽²⁾

This cost function consists of the quadratic *data fidelity* term, which measures the proximity of the estimate to the measurements, and the *regularizer* $\psi(u)$, which encodes the regularity assumptions about the underlying image. The role of the latter is to narrow down the set of plausible solutions by favoring those that exhibit the expected properties. The *regularization parameter* $\tau \ge 0$ balances the influence of both terms to the solution.

B. Regularization Functionals Revisited

From the discussion above, it is apparent that the regularizer plays a crucial role in image reconstruction. Most of the regularization approaches that have been introduced so far in the literature can be assigned to two main categories, namely the *synthesis-based* and the *analysis-based* regularization. In the synthesis-based framework the reconstruction takes place in a sparsifying-base, such as the wavelet domain, where a penalty is imposed on the coefficients of the image in this base [13]. Then the final image is obtained by mapping the reconstructed coefficients back to the image domain through an inverse transform. On the other hand, the analysis-based framework involves regularizers that are directly applied on the image one wishes to reconstruct. Such regularizers have been shown to lead to better recovery results (see for example [14]) and, thus, are mostly preferred.

The analysis-based regularizers can be expressed in the following generic form:

$$R(\boldsymbol{u}) = \int_{\Omega} \Phi(L\boldsymbol{u}(\boldsymbol{x})) \, \mathrm{d}\boldsymbol{x}, \tag{3}$$

where $\Omega \subseteq \mathbb{R}^d$, L is the regularization operator (scalar or multi-component) acting on the image, and $\Phi(\cdot)$ is the potential function. Typical choices for L are differential operators such as the Laplacian (scalar

operator), the gradient (vectorial operator), the Hessian (matrix-valued operator) or wavelet-like operators (wavelets, curvelets, ridgelets, etc.), while the potential function usually involves a norm distance.

C. Total Variation

A very popular regularizer of the form (3) is the Total Variation (TV) semi-norm [2], which for a smooth grayscale image u (C = 1) corresponds to the L_1 norm of the gradient magnitude and it is formally defined as

$$TV = \int_{\Omega} \|\nabla u(\boldsymbol{x})\|_2 \, \mathrm{d}\boldsymbol{x}.$$
(4)

For the past two decades, TV has been extensively applied to several imaging and computer vision problems. One of the main reasons for its success is its ability to allow sharp features (discontinuities) in the solutions. This stems from the L_1 -type behavior of TV that does not over-penalize high intensity variations. For image reconstruction this means that the result will have well-preserved and sharp edges and, thus, it will be visually appealing. Since TV only applies to scalar images, it has been further extended in several ways to cover the case of vector-valued images [15]–[17]. The main requirement for its vectorial variants is that their definition should coincide with the scalar one in (4) when C = 1.

D. Semi-Local Regularization Functionals

While TV has been proven a very powerful regularizer, in several cases its applicability can be limited due to the fact that by design it promotes piecewise-constant solutions. Furthermore, the gradient magnitude, which is employed to penalize the image variation at every point in the image domain, is completely localized and thus it is not very informative of the geometric image structures.

To deal with these limitations of TV, in our prior work we have introduced a family of regularization functionals that involve more general descriptors of image variation which take into account information that is available in a local neighborhood of every point in the image domain [7]. Therefore, the resulting functionals exhibit a semi-local behavior and, thus, can provide a more robust measure of image variation. This is accomplished by employing in their formulation the structure tensor operator [18].

The structure tensor of an image u evaluated at a spatial location x, denoted as $S_k u(x) \in \mathbb{S}^d_+$, is a $d \times d$ symmetric positive semi-definite (PSD) matrix. It summarizes the dominant directions of the gradient in a neighborhood centered at the point x and it is defined as

$$S_{k}u(\boldsymbol{x}) = k_{\sigma} * \left(\nabla u(\boldsymbol{x}) \nabla u(\boldsymbol{x})^{\mathrm{T}}\right), \qquad (5)$$

where k_{σ} is a Gaussian convolution kernel of standard deviation σ . When dealing with vector-valued images the definition of the structure tensor can be extended by replacing the gradient in (5) with the Jacobian operator

$$J\boldsymbol{u}\left(\boldsymbol{x}\right) = \left[\nabla u^{1}\left(\boldsymbol{x}\right)\nabla u^{2}\left(\boldsymbol{x}\right)\dots\nabla u^{C}\left(\boldsymbol{x}\right)\right],\tag{6}$$

which is the gradient's extension for vector-valued functions.

The importance of the structure tensor lies in its eigenvalues which provide a rich and discriminative description of the local geometry of the image by summarizing the distribution of the image gradients in the neighborhood specified by the support of the kernel k_{σ} . Let us consider a 2D image and let $\lambda_1(x)$ and $\lambda_2(x)$ represent the largest and smallest eigenvalues of the structure tensor at a spatial point x, respectively. When both eigenvalues are relatively small there are small intensity variations in the neighborhood around x, indicating that the region is homogeneous. When $\lambda_1(x)$ is large and $\lambda_2(x)$ is small there are strong variations but only on a dominant orientation. Therefore, the point x is located close to an image edge. When both eigenvalues are large there are high variations on both directions specified by the corresponding eigenvectors and, thus, the point x is close to an image corner.

Based on the above, it appears that the eigenvalues of the structure tensor are more informative of the geometric structure of the image than the gradient magnitude which is computed completely locally. This observation has led to the introduction of the Structure tensor Total Variation (STV) functionals [7] which are defined as:

$$\operatorname{STV}_{p}(\boldsymbol{u}) = \int_{\Omega} \left(\sum_{i=1}^{d} \lambda_{i}^{p/2}(\boldsymbol{x}) \right)^{1/p} \mathrm{d}\boldsymbol{x},$$
(7)

where $\lambda_{i}(x)$ is the *i*th eigenvalue of the structure tensor evaluated at the point x and $p \geq 1$.

E. Non-local Regularization Functionals

Non-local regularization functionals have been recently introduced as a means of modeling complex image structures. These regularizers exploit the non-local self-similarity property which is inherent in natural images. The interest for designing such functionals was initiated by the introduction of the non-local means (NLM) filter, which has been used for the task of image denoising [19], [20].

The NLM filter is a patch-based extension of the bilateral filter [21] and it is defined as

NLM
$$(\boldsymbol{u})(\boldsymbol{x}) = \frac{1}{Z(\boldsymbol{x})} \int_{\Omega} e^{-\frac{\mathrm{d}_{\alpha}(\boldsymbol{u}(\boldsymbol{x}),\boldsymbol{u}(\boldsymbol{y}))}{\beta^2}} \boldsymbol{u}(\boldsymbol{y}) \mathrm{d}\boldsymbol{y},$$
 (8)

where

$$d_{\alpha}\left(\boldsymbol{u}\left(\boldsymbol{x}\right),\boldsymbol{u}\left(\boldsymbol{y}\right)\right) = \int_{\Omega} G_{\alpha}\left(t\right) \left|\boldsymbol{u}\left(\boldsymbol{x}+\boldsymbol{t}\right) - \boldsymbol{u}\left(\boldsymbol{y}+\boldsymbol{t}\right)\right|^{2} \mathrm{d}\boldsymbol{t},\tag{9}$$

 G_{α} is a Gaussian kernel of standard deviation α , β acts as a filtering parameter, and $Z(x) = \int_{\Omega} e^{-d_{\alpha}(u(x),u(y))/\beta^2} dy$ is the normalization factor. According to the definition of NLM, the filtered version of u at the spatial coordinates x corresponds to a weighted average of the values of all the image points whose Gaussian neighborhood (*image patch*) is similar to the neighborhood of x. The weights are determined according to the degree of similarity between two image patches, which is given by the distance metric defined in Eq. (9). This filtering strategy differs significantly from standard filtering approaches that involve convolutions and it exploits the simple but very effective idea that two points centered in similar patches are very likely to have the same intensity values. Since the search for similar points is not restricted in a local region of the image domain, the averaging in Eq. (8) can involve points that are located far from each other. This leads to a non-local filtering which has been shown to be very effective in eliminating noise.

Due to its success, NLM gathered great attention and it was first interpreted in [22] as a non-convex regularizer based on non-local functionals. These functionals are non-local in the sense that the involved operators allow a point to interact with any other point in the image domain. Then, in [23] NL-means was expressed as a non-local quadratic functional. Later, Gilboa and Osher following ideas from graph theory and specifically the gradient and divergence on graphs, which were first introduced in [11] in a discrete setting, they developed in [10] a non-local regularization framework defined in the continuous domain, while a discrete analog was considered in [12]. This framework made possible the extension of the non-local quadratic regularizers to non-smooth functionals.

The best representative of the existing non-local functionals is the non-local total variation (NLTV) [10]. NLTV is a convex regularizer that involves the non-local graph gradient operator. This operator is defined as

$$\nabla_{w} u\left(\boldsymbol{x}\right) = \left(u\left(\boldsymbol{y}\right) - u\left(\boldsymbol{x}\right)\right) \sqrt{w\left(\boldsymbol{x},\boldsymbol{y}\right)}, \quad \forall \boldsymbol{y} \in \Omega,$$
(10)

where $w(x, y) : \Omega \times \Omega \mapsto \mathbb{R}_+$ is a non-negative weighting function that assigns weights between a pair of spatial points (x, y) by taking into account their relative distance as well as the similarity of their corresponding values u(x), u(y). Based on the non-local graph gradient, NLTV can be expressed in the generic form of (3) as

NLTV
$$(u) = \int_{\Omega} \|\nabla_w u(\boldsymbol{x})\|_2 d\boldsymbol{x}$$

= $\int_{\Omega} \sqrt{\int_{\Omega} (u(\boldsymbol{y}) - u(\boldsymbol{x}))^2 w(\boldsymbol{x}, \boldsymbol{y}) d\boldsymbol{y}} d\boldsymbol{x}.$ (11)

From its definition it is clear that NLTV exhibits a non-local behavior, since it permits all the points in the image domain Ω to interact with each other. Moreover, the weights w(x, y) are chosen to be large for pairs of points whose neighborhoods are similar and smaller for the rest of the pairs. The non-local interactions and the L_1 -type penalty allow NLTV to model more efficiently complex geometric image structures than the local regularization functionals, such as TV. Therefore, its use can lead to improved image reconstructions. Discrete-domain extensions of NLTV for vector-valued images have been studied in [24], [25].

F. Proposed Non-Local Structure Tensor Total Variation

The non-local functionals that have been introduced so far in the literature are based on the notion of the non-local gradient which is defined on graphs. Therefore, under this approach an image is treated as a generic graph and, thus, its underlying structure is not fully exploited. In this work, we follow an alternative approach with our goal being to introduce non-local regularization functionals that employ the standard gradient operator and, thus, can be directly related to the classical Total Variation and provide a measure of non-local image variation. The motivation for designing such non-local functionals is that by combining ideas from local and non-local regularization strategies we can model at the same time two important image properties, namely the local structural regularity and the non-local self similarity.

The underlying idea of our approach is to define a non-local differential operator and use it as the regularization operator in a functional of the form (3). To do so, our starting point is the structure tensor operator which is defined in Eq. (5). As we mentioned earlier, the power of the structure tensor lies in its ability to encode richer information about the image variation than the gradient operator. This is because its computation involves a local neighborhood of a spatial point, as opposed to the gradient which is computed pointwise. Having in mind the non-local principle, a simple but useful observation is that the information encoded by the structure tensor could be further enriched by extending its scope to the entire image domain.

This idea can be formally presented as follows: Let us consider a 2D vector-valued image u and let r be an arbitrary 2D direction ($||r||_2 = 1$). The vectorial directional derivative of u in the direction

r is computed as $\partial_r u(x) = (Ju(x))^T r$ and its magnitude, $\|\partial_r u(x)\|_2$, yields a measure of intensity change in the location x of the image u at the direction r. For a more robust estimation of the intensity change at a specific direction one can instead use the following measure of *local directional variation* [7]

$$V_{\boldsymbol{r}}\left(\boldsymbol{u}\left(\boldsymbol{x}\right)\right) = \sqrt{\left(k_{\sigma} * \|\partial_{\boldsymbol{r}}\boldsymbol{u}\|_{2}^{2}\right)\left(\boldsymbol{x}\right)} = \sqrt{\boldsymbol{r}^{\mathrm{T}}S_{k}\boldsymbol{u}\left(\boldsymbol{x}\right)\boldsymbol{r}},\tag{12}$$

which involves the structure tensor. This measure is more informative, since it captures the behavior of u in a local neighborhood of x specified by the support of the convolution kernel k_{σ} . This measure of directional variation can be further improved by modifying it so that its range becomes non-local. Specifically, the reasoning here is that if two image patches are similar, then the gradients at the centers of these patches are also expected to be similar. Subsequently, the averaging of similar gradients can lead to a more robust and accurate estimation of image variations. Based on this observation, we consider a *non-local directional variation*, which is computed as the weighted average of the magnitude of the directional variation over points centered in similar neighborhoods. We define this non-local directional variation as

$$V_{\rm NL}\left(\boldsymbol{u}\left(\boldsymbol{x}\right)\right) = \left(\int_{\Omega} w\left(\boldsymbol{x},\boldsymbol{y}\right) \|\partial_{\boldsymbol{r}}\boldsymbol{u}\left(\boldsymbol{y}\right)\|_{2}^{2} \,\mathrm{d}\boldsymbol{y}\right)^{1/2} = \sqrt{\boldsymbol{r}^{\rm T}} \left(\int_{\Omega} w\left(\boldsymbol{x},\boldsymbol{y}\right) \,\mathrm{J}\boldsymbol{u}\left(\boldsymbol{y}\right) \left(\mathrm{J}\boldsymbol{u}\left(\boldsymbol{y}\right)\right)^{\rm T} \,\mathrm{d}\boldsymbol{y}\right) \boldsymbol{r}, \quad (13)$$

where w(x, y) is a non-negative weighting function that assigns appropriate weights between pairs of points (x, y). These weights, as opposed to the local directional variation, are computed not only based on the relative distance of the two points but also based on the photometric distance of their neighborhoods (similarity of their intensity values). A potential candidate for this weighting function is the one used in the NLM filter and reads as

$$w\left(\boldsymbol{x},\boldsymbol{y}\right) = e^{\frac{-d_{\alpha}\left(\boldsymbol{u}\left(\boldsymbol{x}\right),\boldsymbol{u}\left(\boldsymbol{y}\right)\right)}{\beta^{2}}},\tag{14}$$

where d_{α} is the "patch distance" defined in (9).

Motivated by the non-local measure of image variation, we introduce a non-local counterpart of the structure tensor which we define as:

$$S_{w}\boldsymbol{u}\left(\boldsymbol{x}\right) = \int_{\Omega} w\left(\boldsymbol{x},\boldsymbol{y}\right) \mathrm{J}\boldsymbol{u}\left(\boldsymbol{y}\right) \left(\mathrm{J}\boldsymbol{u}\left(\boldsymbol{y}\right)\right)^{\mathrm{T}} \mathrm{d}\boldsymbol{y}.$$
(15)

The non-local (NL) structure tensor S_w , similarly to S_k , when evaluated at a point x corresponds to a symmetric PSD matrix of dimensions $d \times d$ where its (i, j) entry, $1 \le i, j \le d$, is computed as

$$S_{w}^{(i,j)}\boldsymbol{u}\left(\boldsymbol{x}\right) = \int_{\Omega} w\left(\boldsymbol{x},\boldsymbol{y}\right) \sum_{c=1}^{C} \partial_{x_{i}} u_{c}\left(\boldsymbol{y}\right) \partial_{x_{j}} u_{c}\left(\boldsymbol{y}\right) \mathrm{d}\boldsymbol{y},\tag{16}$$



Fig. 1. Image edge information captured by the ℓ_1 -norm of the square rooted eigenvalues of (b) the structure tensor and (c) the proposed NL structure tensor when applied on image (a).

with $\partial_{x_i} u_m$ denoting the partial derivative of the *m*th channel of the vector-valued image *u* w.r.t the *i*-th dimension. To show the potential benefits of using such a non-local operator, in Fig. 1 we present the information captured by the standard structure tensor and its non-local counterpart when they are applied on a grayscale image. From this figure we observe that the edge content of the image is better encoded in the eigenvalues of the NL structure tensor. Indeed the image edges in Fig. 1(c) are sharper and better resolved than in Fig. 1(b).

Having introduced the NL structure tensor, we can now employ it to define our novel family of regularization functionals. Since most of the important information is encoded in the eigenvalues of the NL structure tensor, our non-local energy functionals will be expressed in the generic form

$$R_{\rm NL}\left(\boldsymbol{u}\right) = \int_{\Omega} \Phi\left(\varepsilon_{1}\left(\boldsymbol{x}\right), \dots, \varepsilon_{d}\left(\boldsymbol{x}\right)\right) \mathrm{d}\boldsymbol{x}, \tag{17}$$

where ε_i , i = 1, ..., d are the *d* eigenvalues of the NL structure tensor and $\Phi(\cdot)$ is a potential function. In this work we restrict our attention to potential functions that correspond to ℓ_p norms of the square rooted eigenvalues of the NL structure tensor and can provide a synopsis of the non-local image variation. This leads us to define the family of non-local structure tensor total variation (NLSTV) functionals for $p \ge 1$ as

$$\mathrm{NLSTV}_{p}\left(\boldsymbol{u}\right) = \int_{\Omega} \left(\sum_{i=1}^{d} \varepsilon_{i}^{p/2}\left(\boldsymbol{x}\right)\right)^{1/p} \mathrm{d}\boldsymbol{x}.$$
(18)

Next, we prove the following result.

Theorem 1. The NLSTV regularizers for a fixed weighting function $w(\cdot, \cdot)$ are 1-homogeneous and convex functionals of u for all $p \ge 1$.

Proof: Our proof follows closely the proof of [7][Theorem A.3]. Let us introduce the compact linear operator $T_x = T(u, y; x) = \sqrt{w(x, y)} [\nabla u^1(y), \dots, \nabla u^C(y)] \in \mathcal{H} = L_2(\Omega, \mathbb{R}^{d \times C})$ for arbitrary $x, y \in \Omega$ and a fixed weighting function $w(\cdot, \cdot)$. The adjoint of T_x is the unique operator $T_x^* : \mathcal{H} \mapsto \Omega$ that satisfies

$$\langle \boldsymbol{f}, \boldsymbol{T}_{\boldsymbol{x}} \boldsymbol{u} \rangle_{\mathcal{H}} = \langle \boldsymbol{T}_{\boldsymbol{x}}^* \boldsymbol{f}, \boldsymbol{u} \rangle.$$
 (19)

Now, we can specify the Gram matrix $G = T_x^*T_x$ of size $d \times d$, which characterizes the spectrum of T_x , as

$$G_{i,j} = \langle \sqrt{w (\boldsymbol{x}, \boldsymbol{y})} \partial_{x_i} \boldsymbol{u} (\boldsymbol{y}) , \sqrt{w (\boldsymbol{x}, \boldsymbol{y})} \partial_{x_j} \boldsymbol{u} (\boldsymbol{y}) \rangle_{\mathcal{H}}$$
$$= \int_{\Omega} w (\boldsymbol{x}, \boldsymbol{y}) \sum_{c=1}^{C} \partial_{x_i} u^c (\boldsymbol{y}) \partial_{x_j} u^c (\boldsymbol{y}) \, \mathrm{d}\boldsymbol{y}$$
$$= S_w^{(i,j)} \boldsymbol{u} (\boldsymbol{x}) .$$
(20)

From (20) it occurs that the singular values of T_x are related to the eigenvalues of $S_w u(x)$. Indeed, it holds that the *i*-th singular value of the finite rank operator T_x , denoted as $\sigma_i(T_x)$, $1 \le i \le d$, can be computed as $\sigma_i(T_x) = \sqrt{\varepsilon_i(x)}$. Consequently, we have that

$$\left(\sum_{i=1}^{d} \varepsilon_{i}^{p/2}(\boldsymbol{x})\right)^{1/p} = \left(\sum_{i=1}^{d} \sigma_{i}^{p}(\boldsymbol{T}_{\boldsymbol{x}})\right)^{1/p} = \|\boldsymbol{\sigma}(\boldsymbol{T}_{\boldsymbol{x}})\|_{p}.$$
(21)

The NLSTV functionals can now be equivalently written as

$$\mathrm{NLSTV}_{p}\left(\boldsymbol{u}\right) = \int_{\Omega} \left\|\boldsymbol{\sigma}\left(\boldsymbol{T}_{\boldsymbol{x}}\right)\right\|_{p} \mathrm{d}\boldsymbol{x}.$$
(22)

From this reformulation of NLSTV it is straight-forward to show the 1-homogeneity property. Next, to show the convexity of NLSTV we use that for two compact operators T_x , Λ_x and for any $1 \le p, q \le \infty$ with 1/p + 1/q = 1 it holds that [7]

$$\left\|\boldsymbol{\sigma}\left(\boldsymbol{T}_{\boldsymbol{x}}\right)\right\|_{p} = \sup_{\left\|\boldsymbol{\sigma}\left(\boldsymbol{\Lambda}_{\boldsymbol{x}}\right)\right\|_{q} \leq 1} \operatorname{trace}\left(\boldsymbol{\Lambda}_{\boldsymbol{x}}^{*}\boldsymbol{T}_{\boldsymbol{x}}\right).$$
(23)

Therefore, for $t \in [0, 1]$ we can show that:

$$\begin{aligned} \|\boldsymbol{\sigma} \left(t\boldsymbol{T}_{1,\boldsymbol{x}} + (1-t) \, \boldsymbol{T}_{2,\boldsymbol{x}}\right)\|_{p} \\ &\leq t \sup_{\|\boldsymbol{\sigma}(\boldsymbol{\Lambda}_{1,\boldsymbol{x}})\|_{q} \leq 1} \operatorname{trace} \left(\boldsymbol{\Lambda}_{1,\boldsymbol{x}}^{*} \boldsymbol{T}_{1,\boldsymbol{x}}\right) + (1-t) \sup_{\|\boldsymbol{\sigma}(\boldsymbol{\Lambda}_{2,\boldsymbol{x}})\|_{q} \leq 1} \operatorname{trace} \left(\boldsymbol{\Lambda}_{2,\boldsymbol{x}}^{*} \boldsymbol{T}_{2,\boldsymbol{x}}\right) \\ &= t \left\|\boldsymbol{\sigma} \left(\boldsymbol{T}_{1,\boldsymbol{x}}\right)\right\|_{p} + (1-t) \left\|\boldsymbol{\sigma} \left(\boldsymbol{T}_{2,\boldsymbol{x}}\right)\right\|_{p}. \end{aligned}$$

$$(24)$$

Since the operator T_x is linear in u, it turns out that $\|\sigma(T_x)\|_p$ is a convex functional of u. This further implies that the NLSTV functional in (22) is also a convex functional of u.

Finally, we can also show that our non-local functionals are directly related to the TV semi-norm. Specifically, if we consider a grayscale image u, then by choosing the weighting function so that it is completely localized, i.e,

$$w\left(\boldsymbol{x},\boldsymbol{y}\right) = \begin{cases} 0, & \text{if } \boldsymbol{x} \neq \boldsymbol{y} \\ \\ 1, & \text{if } \boldsymbol{x} = \boldsymbol{y}, \end{cases}$$
(25)

we get $\left\|\sqrt{\varepsilon(x)}\right\|_p = \left\|\nabla u(x)\right\|_2$ for any $p \ge 1$. It is then clear that in this degenerate case our regularizer in (18) reduces to the classical TV. In this sense, our regularizers can be considered as a non-local extension of TV.

III. DISCRETE NON-LOCAL STRUCTURE TENSOR TOTAL VARIATION

Typically, for most of the inverse imaging problems of interest we have to deal with discrete measurements. Therefore, in the rest of this paper we focus on the discrete problem formulation of (1) and its treatment. In this case v and u are discretized versions of the measurements and the underlying image, respectively, while the discrete analog of the operator A is the system matrix A that behaves as the impulse response of the imaging device.

A. Discrete Non-local Gradient

In this section we introduce a novel discrete non-local gradient operator. We will use this operator later to re-express the discrete version of our NLSTV functionals in an alternative form that will allow us to employ robust tools from convex optimization so that we can design an efficient minimization strategy. It is important to note that our non-local operator differs significantly from the non-local gradient that has been introduced in [10], [11]. The main difference is that in our case the non-local gradient operator is not based on the concept of derivatives on graphs but instead it involves the discrete counterpart of the standard gradient operator.

We assume that the discretized d-dimensional vector-valued image $\boldsymbol{u} = (\boldsymbol{u}^1, \dots, \boldsymbol{u}^C)$ is defined on a rectangular grid with unary steps and consists of $N = N_1 \cdot N_2 \cdot \dots \cdot N_d$ pixels. Each channel c of \boldsymbol{u} (c=1,..,C) is rasterized in a vector $\boldsymbol{u}^c \in \mathbb{R}^N$ and all the image channels are stacked together to form the single vector $\boldsymbol{u} \in \mathbb{R}^{NC}$. We use the notation \boldsymbol{u}_n^c to refer to the value of the *n*-th pixel of the grid for the *c*-th channel of \boldsymbol{u} and we use \boldsymbol{u}_n to refer to the tuple $(\boldsymbol{u}_n^1, \dots, \boldsymbol{u}_n^C)$. Now, let us define the discrete non-local gradient of a grayscale image u^c as the linear mapping $D_w : \mathbb{R}^N \mapsto \mathcal{X} \triangleq \mathbb{R}^{N \times d \times N}$. The non-local gradient D_w , when applied on the *n*-th pixel of u^c , corresponds to a matrix of the form¹:

$$\boldsymbol{D}_{w}\boldsymbol{u}_{n}^{c} = \left[\sqrt{w_{n,1}}\,\nabla\boldsymbol{u}_{1}^{c}\,\ldots\,\sqrt{w_{n,N}}\,\nabla\boldsymbol{u}_{N}^{c}\right] \in \mathbb{R}^{d \times N},\tag{26}$$

where ∇u_n^c is the discrete gradient of u^c evaluated at pixel n and $w_{n,k}$ is the discrete weight that has been assigned to the pair of pixels (n, k). Based on the above definition, the non-local gradient evaluated at pixel n is constructed by : (1) evaluating the discrete gradient of the image at all pixel locations in the image domain, (2) weighting these gradients with the square root of the weights $w_{n,k}$ with $1 \le k \le N$, and (3) stacking all these two-dimensional vectors horizontally to form the final matrix. To handle the image boundaries we assume symmetric boundary conditions and employ the discrete gradient as defined in [26]. Note however, that our framework is general enough to accommodate for different discretizations of the gradient operator with alternative assumptions on the image boundaries.

Since the non-local gradient operator is linear, we can further define its adjoint operator, which provides a reverse linear mapping from the space \mathcal{X} to \mathbb{R}^N . To do so, we first need to equip the space \mathcal{X} , which is the target space of D_w , with the inner product $\langle \cdot, \cdot \rangle_{\mathcal{X}}$ and the norm $\|\cdot\|_{\mathcal{X}}$. To define them, let X, $P \in \mathcal{X}$ with $X_n, P_n \in \mathbb{R}^{d \times N} \forall n = 1, 2, ..., N$. Then, we have:

$$\langle \boldsymbol{X}, \boldsymbol{P} \rangle_{\mathcal{X}} = \sum_{n=1}^{N} \operatorname{trace} \left(\boldsymbol{P}_{n}^{\mathrm{T}} \boldsymbol{X}_{n} \right)$$
 (27)

and

$$\|\boldsymbol{X}\|_{\mathcal{X}} = \langle \boldsymbol{X}, \, \boldsymbol{X} \rangle_{\mathcal{X}}^{1/2} = \left(\sum_{n=1}^{N} \|\boldsymbol{X}_{n}\|_{F}^{2}\right)^{\frac{1}{2}}, \qquad (28)$$

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

The definition of the adjoint operator $D_w^* : \mathcal{X} \mapsto \mathbb{R}^N$ is provided through the following relation of the inner products

$$\langle \boldsymbol{D}_w \boldsymbol{u}^c, \boldsymbol{P} \rangle_{\mathcal{X}} = \langle \boldsymbol{u}^c, \, \boldsymbol{D}_w^* \boldsymbol{P} \rangle_2.$$
 (29)

After carrying out some linear algebra calculations we can express D_w^* in a more suitable form that facilitates its numerical computation. This form is provided in the following proposition.

¹In practice, as we describe in Section III-B we use a sparse version of the weighting function $w_{n,k}$ and therefore the number of "active" neighbors for each pixel of the image will be equal to K which is significantly smaller than N. Consequently, the target space of the non-local gradient will be $\mathcal{X} \triangleq \mathbb{R}^{N \times d \times K}$ and $D_w u_n^c \in \mathbb{R}^{d \times K}$.

Proposition 1. The adjoint operator D_w^* of the discrete non-local gradient evaluated on the *n*-th matrix entry, $P_n \in \mathbb{R}^{d \times N}$, of the multidimensional matrix $P \in \mathcal{X}$ is given by:

$$\boldsymbol{D}_{w}^{*}\boldsymbol{P}_{n} = -\operatorname{div}\left(\sum_{k=1}^{N}\sqrt{w_{k,n}}\,\boldsymbol{P}_{k}\left(:,n\right)\right),\tag{30}$$

where div is the discrete divergence operator, defined in accordance to the discretization scheme of the gradient operator, and $P_k(:, n) \in \mathbb{R}^d$ refers to the *n*-th column of the *k*-th matrix entry of P.

Having defined the non-local gradient and its adjoint operator, we can further introduce the discrete nonlocal Jacobian which applies on vector-valued images. The non-local Jacobian, $J_w : \mathbb{R}^{NC} \mapsto \mathbb{R}^{N \times d \times (NC)}$, is a matrix-valued operator which, similarly to the standard Jacobian operator, when evaluated on the *n*-th pixel of a vector-valued image u corresponds to a matrix of the form:

$$\boldsymbol{J}_{w}\boldsymbol{u}_{n} = \left[\boldsymbol{D}_{w}\boldsymbol{u}_{n}^{1}\ldots\boldsymbol{D}_{w}\boldsymbol{u}_{n}^{C}\right] \in \mathbb{R}^{d \times (NC)}.$$
(31)

By employing the adjoint of the non-local gradient, we can easily obtain the adjoint of the non-local Jacobian as

$$\boldsymbol{J}_{w}^{*}\boldsymbol{Y}_{n} = \begin{bmatrix} \boldsymbol{D}_{w}^{*}\boldsymbol{Y}_{n}^{1} \dots \boldsymbol{D}_{w}^{*}\boldsymbol{Y}_{n}^{C} \end{bmatrix}^{\mathrm{T}} \in \mathbb{R}^{C},$$
(32)

where $\boldsymbol{Y}_n = \begin{bmatrix} \boldsymbol{Y}_n^1 \dots \boldsymbol{Y}_n^C \end{bmatrix} \in \mathbb{R}^{d \times (NC)}$.

B. Non-local Weights Computation

In the definitions of the non-local operators that we provided above, a weight $w_{n,k}$ is assigned to every pair of pixels (n, k). In practice, mainly due to computational considerations, we use a sparse version of the discrete weighting function which is computed as follows: (1) For each pixel n in the image domain we extract a patch $\mathcal{P}(u_n)$ of size $\tau \times \tau$ centered around this pixel. (2) We compute the distance of this patch from all the patches whose centers lie inside a specified search window of size $r \times r$. To do so, we use $(d_{\alpha})_{n,k}$ which is the discrete version of (9) and it is defined as

$$(d_{\alpha})_{n,k} = \sum_{j=-\tau/2}^{\tau/2} (g_{\alpha})_j |\boldsymbol{u}_{n+j} - \boldsymbol{u}_{k+j}|^2, \qquad (33)$$

with g_{α} denoting a discrete weighting function of size $\tau \times \tau$ that defines the image neighborhood. Note that k is restricted to be at most r/2 pixels far from pixel n. (3) Out of all the computed distances we keep the K smallest ones while the rest are set to infinity. Then, we compute the corresponding weights as $w_{n,k} = e^{-(d_{\alpha})_{n,k}/\beta^2}$. We note that our strategy for computing the non-local weights is similar to the one used in NLTV [10]. An efficient computation of the patch distance $(d_{\alpha})_{n,k}$ was proposed in [27] and uses a similar idea with the integral image [28]. However, this method is only applicable when a uniform kernel g_{α} is used. Here, we consider an alternative fast implementation which is applicable to the more general case of a symmetric kernel, $(g_{\alpha})_j = (g_{\alpha})_{-j}$. This is based on the observation that given the symmetric nature of g_{α} we can re-write (33) as

$$(d_{\alpha})_{n,k} = \sum_{j=-\tau/2}^{\tau/2} (g_{\alpha})_j |\boldsymbol{u}_{n-j} - \boldsymbol{u}_{n+l-j}|^2, \qquad (34)$$

where $-r/2 \le l \le r/2$ is the relative distance of the pixel k from the pixel n. Clearly, (34) corresponds to a discrete convolution of g_{α} with $\mathbf{z}_{n}^{l} = |\mathbf{u}_{n} - \mathbf{u}_{n+l}|^{2}$. Therefore, it is now possible to compute with a single convolution the patch distances of all the pixel pairs (n, k) that have a relative distance of l pixels. We note that a similar observation was made in [29] in order to speed-up the computation time of the NLM method.

C. Discrete NL-STV

Let us indicate that a pixel k is linked to a pixel n with a non-zero weight $w_{k,n}$ by using the notation $k \in \mathcal{N}_n$ where $\mathcal{N}_n = \{k : w_{n,k} > 0\}$. Next, we define a "sparse" version of the discrete non-local structure tensor as

$$\boldsymbol{S}_{w}\boldsymbol{u}_{n} = \sum_{k \in \mathcal{N}_{n}} w_{n,k} \boldsymbol{J} \boldsymbol{u}_{k} \left(\boldsymbol{J} \boldsymbol{u}_{k} \right)^{\mathrm{T}}$$
(35)

where J is the discrete Jacobian operator. Based on the discrete NL structure tensor we can now formally define the discrete NLSTV regularizers as

$$\mathrm{NLSTV}_{p}\left(\boldsymbol{u}\right) = \sum_{n=1}^{N} \left(\sum_{i=1}^{d} \varepsilon_{n,i}^{p/2}\right)^{1/p},$$
(36)

where $\varepsilon_{n,i}$ denotes the *i*-th eigenvalue of the discrete non-local structure tensor S_w applied on u and evaluated at the pixel location n.

The current form of the proposed regularizers is difficult to work with, mainly because it involves the eigenvalues of a non-linear operator. Consequently, it is not clear how one could proceed with the minimization of such penalties in an inverse problem setting. To deal with this difficulty we derive an alternative definition of the proposed functionals that will prove very useful. This new definition arises by using the following result, whose proof is straight-forward and thus is omitted. **Proposition 2.** The discrete non-local structure tensor operator applied on u and evaluated at a pixel location n can be expressed in terms of the non-local Jacobian as:

$$\boldsymbol{S}_{w}\boldsymbol{u}_{n} = \boldsymbol{J}_{w}\boldsymbol{u}_{n}\left(\boldsymbol{J}_{w}\boldsymbol{u}_{n}\right)^{\mathrm{T}}.$$
(37)

Consequently, the eigenvalues of the NL structure tensor, $\varepsilon_{n,i}$ for i = 1, ..., d and n = 1, ..., N, can be computed from the singular values of the NL Jacobian, $\sigma_{n,i}$, as

$$\varepsilon_{n,i} = \sigma_{n,i}^2. \tag{38}$$

Proposition 2 directly implies that the discrete NLSTV functionals can be equivalently expressed in terms of the singular values of the non-local Jacobian. Indeed, from (36) and (38) we have

$$\operatorname{NLSTV}_{p}\left(\boldsymbol{u}\right) = \sum_{n=1}^{N} \left(\sum_{i=1}^{d} \sigma_{n,i}^{p}\right)^{1/p} = \sum_{n=1}^{N} \left\|\boldsymbol{J}_{w}\boldsymbol{u}_{n}\right\|_{\mathcal{S}_{p}},$$
(39)

where the r.h.s of (39) arises by using the definition of the Schatten norm of order p [30]. The Schatten matrix norms are directly related to the ℓ_p vector norms. In particular, for a generic complex matrix $X \in \mathbb{C}^{n_1 \times n_2}$ the S_p norm of X can be expressed as the ℓ_p norm of the vector $\sigma(X)$, whose entries are the singular values of X, i.e, $\|X\|_{S_p} = \|\sigma(X)\|_p$. Also note that the class of Schatten norms includes as its members the nuclear norm (p = 1), the Frobenius norm (p = 2), and the spectral/operator norm $(p = \infty)$.

This alternative formulation of the proposed functionals is much easier to handle in an optimization framework. The reason is that the non-linear NL structure tensor has been substituted by the linear operator J_w (D_w for the case of scalar images). Moreover, the expression in (39) better highlights the fact that the proposed regularizers are convex w.r.t $u \forall p \ge 1$. Indeed, it is straight-forward for one to show that these regularizers are convex, since they can be expressed as a composition of a norm, $\sum_n \|\cdot\|_{S_v}$, and a linear operator, J_w .

Based on (39) and Lemma 4.1 in [9] we further derive the following dual definition

$$\mathrm{NLSTV}_{p}\left(\boldsymbol{u}\right) = \max_{\boldsymbol{\Omega}\in\mathcal{B}_{\infty,q}}\langle\boldsymbol{u}\,,\,\boldsymbol{J}_{w}^{*}\boldsymbol{\Omega}\rangle_{2},\tag{40}$$

where $\Omega = [\Omega_1 \dots \Omega_N] \in \mathcal{X}$ and $\mathcal{B}_{\infty,q} = \{ \Omega \in \mathcal{X} : \|\Omega_n\|_{\mathcal{S}_q} \leq 1 \}$. This expression is extremely useful in case one needs to re-write the minimization of an objective function that includes NLSTV in a min-max formulation. Such reformulation of the minimization problem is necessary when a primal-dual approach, as those proposed in [31], [32], is employed to obtain the solution.

IV. NUMERICAL OPTIMIZATION

Based on the discrete version of the forward model in (1) and under the assumption that the noise perturbing the measurements is Gaussian, a solution of the regularized inverse problem is derived as

$$\boldsymbol{u}^{\star} = \operatorname*{arg\,min}_{\boldsymbol{u}} \frac{1}{2} \|\boldsymbol{v} - \boldsymbol{A}\boldsymbol{u}\|_{2}^{2} + \tau \|\boldsymbol{J}_{w}\boldsymbol{u}\|_{1,p} + \iota_{\mathcal{C}}(\boldsymbol{u}), \qquad (41)$$

for any $p \ge 1$. In Eq. (41) $\iota_{\mathcal{C}}$ is the indicator function of a convex set \mathcal{C} , while $\|\cdot\|_{1,p}$ is a shorthand notation for the mixed vector-matrix norm $\sum_n \|\cdot\|_{\mathcal{S}_p}$. The indicator function $\iota_{\mathcal{C}}$ takes the value 0 if $u \in \mathcal{C}$ and ∞ otherwise, and its role is to enforce the solution to lie in \mathcal{C} . The reason for including $\iota_{\mathcal{C}}$ in the overall objective function is that in many cases the need for such a constraint arises naturally. For example in many imaging applications it is common to require that the intensities of the reconstructed images should either be non-negative (non-negativity constraint) or lie in a specific range (box constraint). If one seeks for the unconstrained solution of (41), this can be obtained by simply setting $\mathcal{C} = \mathbb{R}^{NC}$.

The form of the problem in (41) is difficult to work with mainly for two reasons. The first one is that the objective function we want to minimize is non-smooth. This precludes the use of a gradient-based scheme. Given the non-smoothness of the objective function, the second reason is the coupling that exists among the different terms. This coupling makes the minimization task even more challenging. To circumvent these difficulties a common strategy that we will also follow here, is to decouple the different terms of the objective function by introducing a set of auxiliary variables. This will allow us to find the solution by solving a sequence of simpler problems.

In particular, we use the auxiliary variables $z_1 = J_w u \in \mathcal{X}$ and $z_2 = u \in \mathbb{R}^{NC}$ and we reformulate the problem in (41) in the constrained form

$$\underset{Ku=z}{\operatorname{arg\,min}} \underbrace{\frac{1}{2} \| \boldsymbol{v} - \boldsymbol{A} \boldsymbol{u} \|_{2}^{2}}_{f(\boldsymbol{u})} + \underbrace{\tau \| \boldsymbol{z}_{1} \|_{1,p}}_{g_{1}(\boldsymbol{z}_{1})} + \underbrace{\iota_{\mathcal{C}} (\boldsymbol{z}_{2})}_{g_{2}(\boldsymbol{z}_{2})}.$$
(42)

The augmented Lagrangian [33] associated to (42) is given by

$$\mathcal{L}_{\mu}\left(\boldsymbol{u},\boldsymbol{z},\boldsymbol{\eta}\right) = f\left(\boldsymbol{u}\right) + g\left(\boldsymbol{z}\right) + \langle\boldsymbol{\eta},\,\boldsymbol{K}\boldsymbol{u}-\boldsymbol{z}\rangle_{\mathcal{Y}} + \frac{\mu}{2} \|\boldsymbol{K}\boldsymbol{u}-\boldsymbol{z}\|_{\mathcal{Y}}^{2},\tag{43}$$

where $\mathcal{Y} \triangleq \mathcal{X} \times \mathbb{R}^N$ is a multilinear space, $\boldsymbol{\eta} = (\boldsymbol{\eta}_1, \boldsymbol{\eta}_2) \in \mathcal{Y}$ are the Lagrange multipliers, $\boldsymbol{z} = (\boldsymbol{z}_1, \boldsymbol{z}_2) \in \mathcal{Y}$, $g(\boldsymbol{z}) = g_1(\boldsymbol{z}_1) + g_2(\boldsymbol{z}_2)$, $\boldsymbol{K} = (\boldsymbol{J}_w, \boldsymbol{I})$ is a concatenated version of the non-local Jacobian \boldsymbol{J}_w and the identity operator \boldsymbol{I} , and $\mu \geq 0$ is a penalty parameter. Note that \mathcal{L}_0 corresponds to the standard Lagrangian of problem (42) while \mathcal{L}_μ corresponds to the Lagrangian of the problem

$$\underset{\substack{\boldsymbol{u},\boldsymbol{z}_{1},\boldsymbol{z}_{2}\\\boldsymbol{K}\boldsymbol{u}=\boldsymbol{z}}}{\operatorname{arg\,min}\,f\left(\boldsymbol{u}\right) + g\left(\boldsymbol{z}\right) + \left\|\boldsymbol{K}\boldsymbol{u} - \boldsymbol{z}\right\|_{\mathcal{Y}}^{2}.$$

$$\tag{44}$$

In any case, the two problems (42) and (44) are clearly equivalent, since for any feasible solution of (44) the additional quadratic term in (44) evaluates to zero.

To solve the problem in (42) we use the augmented Lagrangian and employ the *alternating-direction method of multipliers* (ADMM) [34]–[36]. This is an iterative method that attacks the constrained optimization by replacing it by a series of unconstrained problems. To this end, one ADMM iteration involves the minimization of \mathcal{L}_{μ} w.r.t z while keeping u fixed, the minimization of \mathcal{L}_{μ} w.r.t u while keeping z fixed, and an update of the Lagrange multipliers η . Formally, we can describe the steps involved in every ADMM iteration as:

$$\boldsymbol{z}^{t+1} = \operatorname*{arg\,min}_{\boldsymbol{z}\in\mathcal{Y}} g\left(\boldsymbol{z}\right) + \frac{\mu}{2} \left\|\boldsymbol{z} - \left(\boldsymbol{K}\boldsymbol{u}^{t} + \boldsymbol{s}^{t}\right)\right\|_{\mathcal{Y}}^{2}$$
(45a)

$$\boldsymbol{u}^{t+1} = \operatorname*{arg\,min}_{\boldsymbol{u} \in \mathbb{R}^{NC}} f\left(\boldsymbol{u}\right) + \frac{\mu}{2} \left\| \boldsymbol{K}\boldsymbol{u} - \left(\boldsymbol{z}^{t+1} - \boldsymbol{s}^{t}\right) \right\|_{\mathcal{Y}}^{2}$$
(45b)

$$s^t = s^t + Ku^{t+1} - z^{t+1}, \tag{45c}$$

where $s = \eta/\mu$ corresponds to a scaled version of the original Lagrange multipliers and in each one of the subproblems we have ignored constant terms that are irrelevant to the optimization task.

Next, we focus on the solution of the individual subproblems. First, based on the definition of g(z), we note that the *z*-update of ADMM (45a) can be decoupled in the following two independent problems

$$z_{1}^{t+1} = \underset{z_{1} \in \mathcal{X}}{\operatorname{arg\,min}} \frac{1}{2} \left\| z_{1} - \left(J_{w} u^{t} + s_{1}^{t} \right) \right\|_{\mathcal{X}}^{2} + \frac{\tau}{\mu} \left\| z_{1} \right\|_{1,p}$$
$$z_{2}^{t+1} = \underset{z_{2} \in \mathbb{R}^{NC}}{\operatorname{arg\,min}} \frac{1}{2} \left\| z_{2} - \left(u^{t} + s_{2}^{t} \right) \right\|_{2}^{2} + \iota_{\mathcal{C}} \left(z_{2} \right).$$
(46)

The first problem corresponds to the evaluation of the proximal map of the function $g_1(z_1)$ scaled by μ . Indeed, the proximal map of a function f with **dom** $f = \mathcal{X}$, evaluated at z is defined as [37]

$$\operatorname{prox}_{f}(\boldsymbol{z}) = \operatorname{arg\,min}_{\boldsymbol{u}\in\mathcal{X}} \frac{1}{2} \|\boldsymbol{u}-\boldsymbol{z}\|_{\mathcal{X}}^{2} + f(\boldsymbol{u}).$$
(47)

Therefore, it holds that

$$\boldsymbol{z}_{1}^{t+1} = \operatorname{prox}_{\frac{\tau}{\mu} \parallel \cdot \parallel_{1,p}} \left(\boldsymbol{J}_{w} \boldsymbol{u}^{t} + \boldsymbol{s}_{1}^{t} \right).$$

$$(48)$$

Due to the separability of $\|\cdot\|_{1,p}$ and of the quadratic term $\|\cdot\|_{\mathcal{X}}^2$, the above problem can be further decomposed in N independent subproblems of the form:

$$(\boldsymbol{z}_{1}^{t+1})_{n} = \operatorname{prox}_{\frac{\tau}{\mu} \parallel \cdot \parallel_{\mathcal{S}_{p}}} (\boldsymbol{\Omega}_{n})$$

=
$$\operatorname*{arg\,min}_{(\boldsymbol{z}_{1})_{n}} \frac{1}{2} \, \lVert (\boldsymbol{z}_{1})_{n} - \boldsymbol{\Omega}_{n} \rVert_{F}^{2} + \frac{\tau}{\mu} \, \lVert (\boldsymbol{z}_{1})_{n} \rVert_{\mathcal{S}_{p}},$$
(49)

where $\Omega = J_w u^t + s_1^t \in \mathcal{X}$ and $\Omega_n \in \mathbb{R}^{d \times KC}$ denotes the *n*-th matrix entry of Ω (*K* is the number of neighbors for a pixel *n* with non-zero weights $w_{n,k}$).

To compute the solution in (49) we use the following result, where $\mathbb{U}^n = \{ \mathbf{X} \in \mathbb{C}^{n \times n} : \mathbf{X}^{-1} = \mathbf{X}^H \}$ denotes the set of unitary matrices and $\mathbb{D}^{n_1 \times n_2} = \{ \mathbf{X} \in \mathbb{R}^{n_1 \times n_2}_+ : \mathbf{X}_{(i,j)} = 0 \ \forall i \neq j \}$ denotes the set of PSD diagonal matrices.

Proposition 3 ([38]). Let $Y \in \mathbb{C}^{n_1 \times n_2}$ be a generic complex matrix with an SVD decomposition $Y = U\Sigma V^{\mathrm{H}}$, where $U \in \mathbb{U}^{n_1}$, $V \in \mathbb{U}^{n_2}$, and $\Sigma \in \mathbb{D}^{n_1 \times n_2}$. The proximal map

$$\boldsymbol{X}^{\star} = prox_{\tau \parallel \cdot \parallel_{\mathcal{S}_p}} \left(\boldsymbol{Y} \right) \tag{50}$$

can be obtained as

$$\boldsymbol{X}^{\star} = \boldsymbol{U} diag\left(prox_{\tau \parallel \cdot \parallel_{p}} \left(\boldsymbol{\sigma} \right) \right) \boldsymbol{V}^{\mathrm{H}},$$
(51)

where diag (·) is the operator that transforms a vector to a diagonal matrix and σ denotes the vector whose entries are the singular values of Y.

In words, Proposition 3 states that the proximal map of an S_p matrix-norm can be computed by (1) decomposing the matrix Y in its singular values and singular vectors, (2) evaluating the proximal map of the ℓ_p norm at the vector σ consisting of the singular values of Y and, (3) deriving the final result by a singular value reconstruction that involves the singular vectors of Y and as singular values the result of step 2.

Efficient evaluation of the proximal map: Clearly the solution of (49) depends on our ability to evaluate the proximal map of an ℓ_p norm weighted by a constant τ . To do so for an arbitrary $p \ge 1$, we rely on the iterative proximal algorithm introduced in [39]. Furthermore, for the most interesting cases p = 1, 2 we can derive the solution in closed-form. Specifically, for p = 1 the proximal map corresponds to the soft-thresholding function $S_{\tau}^1(\sigma) = \max(\sigma - \tau, 0)$ [40] where the max is computed component-wise. For p = 2 the proximal map corresponds to another shrinkage function of the form $S_{\tau}^2(\sigma) = \max((\|\sigma\|_2 - \tau) / \|\sigma\|_2, 0) \cdot \sigma$ [35], [37]. Notably in this case, we can obtain the proximal map of the Frobenius norm (S_2 norm) without having to resort to an SVD decomposition of the input matrix. Specifically, we have

$$\operatorname{prox}_{\tau \parallel \cdot \parallel_{F}} \left(\boldsymbol{Y} \right) = \max \left(\frac{\|\boldsymbol{Y}\|_{F} - \tau}{\|\boldsymbol{Y}\|_{F}}, 0 \right) \cdot \boldsymbol{Y}.$$
(52)

Another consideration regarding the evaluation of the proximal map of (49) is the efficient computation of the singular values and vectors of the matrices $\Omega_n = U_n \Sigma_n V_n^{\mathrm{T}} \in \mathbb{R}^{d \times KC}$ with n = 1, 2, ..., N. Since even for images of moderate sizes the value of N can be of the order 10^5 , we need an efficient way to perform these SVDs. Here, we are interested in 2D images (d = 2) and we can follow an efficient strategy as described next. First, we observe that the matrix $\Omega_n \Omega_n^T$ is 2×2 symmetric with an eigenvalue decomposition $U_n \Sigma_n^2 U_n^T$. Therefore, both U_n and Σ_n are obtained in closed-form. Now, let Σ_n^+ be the pseudo-inverse of Σ_n . Then, based on Proposition 3 we can compute the proximal maps as

$$\operatorname{prox}_{\frac{\tau}{\mu} \parallel \cdot \parallel_{\mathcal{S}_p}} \left(\mathbf{\Omega}_n \right) = \left(U_n \mathbf{\Sigma}_n^* \mathbf{\Sigma}_n^+ U_n^{\mathrm{T}} \right) \mathbf{\Omega}_n$$
(53)

where $\Sigma_n^{\star} = \operatorname{diag}\left(\operatorname{prox}_{\tau \parallel \cdot \parallel_p}(\sigma_n)\right)$ and σ_n is the vector formed by the diagonal elements of Σ_n .

The second minimization problem in (46) corresponds to the projection of $u^t + s_2^t$ on the convex set C. For the imaging applications that we study later, we consider the convex set $C \triangleq \{u \in \mathbb{R}^N : 0 \le u_n \le \gamma, \forall n = 1, ..., N\}$. In this case the projection is simple to compute and is given by $\Pi_{\mathcal{C}}(u) = \min(\max(0, u), \gamma)$ where the min and max operations are computed component-wise.

The problem in (45b) is quadratic and thus the u-update is obtained as the solution of the set of linear equations

$$Bu^{t+1} = \frac{1}{\mu} A^{\mathrm{T}} v + J_w^* \left(z_1^{t+1} - s_1^t \right) + z_2^{t+1} - s_2^t,$$
(54)

where $B = \left(\frac{1}{\mu}A^{T}A + J_{w}^{*}J_{w} + I\right)$. In practice, the inversion of B is prohibitive due to its large size. Therefore, given that B is a symmetric positive definite matrix, we employ instead the conjugate gradient (CG) method [41]. We have experimentally observed that for several inverse imaging problems it suffices to run as few as two CG iterations in every ADMM iteration. This choice does not seem to compromise the convergence of the overall algorithm if we follow a "warm-start" strategy where in each ADMM iteration we initialize CG with the solution of u obtained in the previous iteration.

Finally, the update of the scaled Lagrange multipliers in (45c) is decoupled and it is computed as

$$s_{1}^{t+1} = s_{1}^{t} + J_{w}u^{t+1} - z_{1}^{t+1},$$

$$s_{2}^{t+1} = s_{2}^{t} + u^{t+1} - z_{2}^{t+1}.$$
(55)

A summary of our overall minimization approach is provided in Algorithm 1.

V. APPLICATIONS AND EXPERIMENTS

To assess the potentials of the proposed non-local functionals, we compare their reconstruction performance with that of other related methods on several inverse imaging applications. In particular, we consider the problems of image denoising, image deblurring, and image reconstruction from undersampled

Algorithm 1 : Image Reconstruction Algorithm.

Input: $v, A, \tau > 0, \mu > 0, p \ge 1$. Initialization: $u^0 = v, s_1^0 = 0, s_2^0 = 0, t = 0$. while stopping criterion is not satisfied do $z_1^{t+1} \leftarrow \operatorname{prox}_{\overline{u} \parallel \cdot \parallel_{1,v}} (J_w u^t + s_1^t);$

$$\begin{split} \boldsymbol{z}_{2}^{t+1} &\leftarrow \Pi_{\mathcal{C}} \left(\boldsymbol{u}^{t} + \boldsymbol{s}_{2}^{t} \right); \\ \boldsymbol{B} &\leftarrow \left(\frac{1}{\mu} \boldsymbol{A}^{\mathrm{T}} \boldsymbol{A} + \boldsymbol{J}_{w}^{*} \boldsymbol{J}_{w} + \boldsymbol{I} \right); \\ \boldsymbol{\omega}_{1}^{t+1} &\leftarrow \boldsymbol{z}_{1}^{t+1} - \boldsymbol{s}_{1}^{t}; \\ \boldsymbol{\omega}_{2}^{t+1} &\leftarrow \boldsymbol{z}_{2}^{t+1} - \boldsymbol{s}_{2}^{t}; \\ \boldsymbol{u}^{t+1} &\leftarrow \boldsymbol{B}^{-1} \left(\frac{1}{\mu} \boldsymbol{A}^{\mathrm{T}} \boldsymbol{v} + \boldsymbol{J}_{w}^{*} \boldsymbol{\omega}_{1}^{t+1} + \boldsymbol{\omega}_{2}^{t+1} \right); \\ \boldsymbol{s}_{1}^{t+1} &\leftarrow \boldsymbol{s}_{1}^{t} + \boldsymbol{J}_{w} \boldsymbol{u}^{t+1} - \boldsymbol{z}_{1}^{t+1}; \\ \boldsymbol{s}_{2}^{t+1} &\leftarrow \boldsymbol{s}_{2}^{t} + \boldsymbol{u}^{t+1} - \boldsymbol{z}_{2}^{t+1}; \\ t &\leftarrow t+1; \end{split}$$

end

return u^t ;





Fourier measurements. For all the problems under study we report results on both grayscale and color versions of the images shown in Fig. 2, where their intensities have been normalized to lie in the range [0, 1]. The comparisons are performed among the first-order functionals TV [2], STV₁ [7], NLTV [10], and our proposed NLSTV₁. For the color case, since TV is applicable only to grayscale images we use instead for our comparisons the vectorial total variation (VTV) that was proposed in [16]. The same applies to NLTV where for the color case we use a vectorial version that has been studied in [24], [25]. Furthermore, for the problem of image denoising we also include the results obtained by the NLM filter [20].

The reconstruction performance of the regularization methods under comparison highly depends on

the choice of the regularization parameter τ . This also holds true for the NLM filter where the parameter β in (8) needs to be tuned. In order our comparisons to be consistent, we have experimentally fine-tuned these parameters and here we report only the results that led to the best peak signal-to-noise ratio (PSNR). For STV₁ we also had to choose the convolution kernel that is used for the computation of the structure tensor in (5). This was set to be a Gaussian of standard deviation $\sigma = 0.5$ and support of 3×3 pixel. For the non-local techniques we had to further specify the size of the neighborhood that is used for the computation of the non-local weights $w(\cdot, \cdot)$. He have considered a uniform neighborhood (patch) of size 5×5 and 7×7 and a search window of size 11×11 . For NLTV and NLM we experimentally observed that the 5×5 patch led to better results on average than the 7×7 patch, while for NLSTV₁ the 7×7 patch turned out to be the best choice. Finally, for NLTV and NLSTV₁ we considered a sparse version of the weight function, where only K neighbors are linked to each pixel of the image with a non-zero weight. In order to be able to compare directly the performance of STV₁ with our non-local extension NLSTV₁, we set K = 9. For NLTV we set K = 14 (the 4 closest neighbors plus the best 10 non-local neighbors) as suggested in [42].

For the minimization of the objective functions that are related to the regularizers we are comparing, we employ an optimization strategy similar to the one described in Algorithm 1. Our rationale for using a common optimization framework for all the studied regularizers is that this way we can make sure that the convergence behavior is comparable in all cases. Consequently, the image reconstruction quality depends on the choice of the regularizer rather than on the efficiency of the employed minimization scheme.

A. Image denoising

Image denoising is the simplest form of inverse imaging problem where the system matrix A reduces to the identity operator I. In our scenario we consider i.i.d Gaussian noise at three different noise levels (low, medium, and high) that correspond to a standard deviation of $\sigma_n = 0.05, 0.075, 0.1$, respectively. Regarding the stopping criterion of the minimization algorithms, this is set to either reaching a relative normed difference of $5 \cdot 10^{-5}$ between two successive image estimates, or a maximum of 150 iterations. In addition, for the non-local methods we need to compute the non-local weights. These weights according to the definitions we provided earlier depend on the underlying image. Here, however, we compute them from a smoothed version of the noisy image by a Gaussian filter. We have experimentally observed that the weights obtained in this way lead to better denoising results than the ones computed directly from the noisy image itself.

TABLE I

ISNR COMPARISONS ON GRAYSCALE/COLOR IMAGE DENOISING

$ \begin{array}{ c c c c c c c c c c c c c c c c c c c$	0.1 9.23/9.81 2.89/12.44
$ \begin{array}{ c c c c c c c c c c c c c c c c c c c$	0.1 9.23/9.81
$ \begin{array}{ c c c c c c c c c c c c c c c c c c c$	9.23/9.81
Img.2 9.21/9.95 10.85/11.64 12.06/11.72 9.87/10.40 11.42/12.00 12.57/13.18 10.01/10.29 11.61/12.03 12.74/13.28 10.04/10.48 11.62/12.03 12.72/13.04 10.19/10.60 11.76/12.25 12.57/13.18 10.01/10.29 11.61/12.03 12.74/13.28 10.04/10.48 11.62/12.03 12.72/13.04 10.19/10.60 11.76/12.55 12.57/13.18 10.01/10.29 11.61/12.03 12.74/13.28 10.04/10.48 11.62/12.03 12.72/13.04 10.19/10.60 11.76/12.55 12.57/13.18 10.01/10.29 11.61/12.03 12.74/13.28 10.04/10.48 11.62/12.03 12.72/13.04 10.19/10.60 11.76/12.55 12.57/13.18 10.01/10.29 11.61/12.03 12.74/13.28 10.04/10.48 11.62/12.03 12.72/13.04 10.19/10.60 11.76/12.55 12.57/13.18 10.01/10.29 11.61/12.03 12.74/13.28 10.04/10.48 11.62/12.03 12.72/13.04 10.19/10.60 11.76/12.55 12.57/13.18 10.01/10.29 11.61/12.03 12.74/13.28 10.04/10.48 11.62/12.03 12.72/13.04 10.19/10.60 11.76/12.55 12.57/13.18 10.01/10.29 11.61/12.03 12.74/13.28 10.04/10.48 11.62/12.03 12.72/13.04 10.19/10.60 11.76/12.55 12.57/13.18 10.01/10.29 11.61/12.03 12.74/13.28 10.04/10.48 11.62/12.03 12.72/13.04 10.57/13.57 12.57/13.58 12.57/13.	9 80/19 44
	.2.09/13.44
Img.3 5.22/5.86 6.41/7.08 7.34/7.96 5.90/6.34 7.07/7.54 7.99/8.46 5.08/5.09 6.66/6.75 7.89/8.06 5.93/6.05 7.18/7.12 8.13/7.93 6.17/6.59 7.30/7.81	8.17/8.77
Img.4 8.10/8.89 9.51/10.35 10.57/10.80 8.57/9.15 9.98/10.61 11.04/11.69 8.20/8.61 9.83/10.30 11.06/11.56 8.58/8.85 10.11/10.22 11.23/11.22 8.84/9.42 10.23/10.94 10.94 1	1.28/12.10
Img.5 6.01/6.65 7.13/7.93 8.01/8.72 7.00/7.41 8.10/8.64 8.92/9.55 6.79/6.84 8.18/8.42 9.20/9.57 7.06/7.09 8.28/8.24 9.18/8.97 7.41/7.67 8.53/8.98	9.34/9.94
Img.6 5.40/6.09 6.75/7.48 7.83/8.51 5.85/6.34 7.21/7.73 8.28/8.81 5.69/5.89 7.14/7.44 8.21/8.61 6.01/6.55 7.47/7.97 8.60/9.02 6.30/6.80 7.71/8.28	8.79/9.42
Img.7 4.12/4.76 5.43/6.13 6.51/7.24 4.39/4.85 5.72/6.22 6.80/7.32 3.99/4.23 5.44/5.72 6.56/6.91 4.42/4.82 5.88/6.26 7.03/7.35 4.67/5.09 6.05/6.57 6.05/6.57 6.67/6.91 6.67/6.9	7.15/7.75
Img.8 4.81/5.47 6.28/6.96 7.47/8.11 5.24/5.71 6.69/7.17 7.85/8.33 4.73/4.95 6.38/6.67 7.68/8.03 5.17/5.57 6.71/6.99 7.92/8.12 5.44/5.82 6.88/7.31	8.02/8.53
Img.9 3.32/3.92 4.65/5.31 5.73/6.42 3.66/4.12 5.02/5.50 6.10/6.60 3.10/3.35 4.69/5.03 5.98/6.36 3.56/3.91 5.02/5.37 6.18/6.49 3.80/4.14 5.18/5.59	6.26/6.75
Img.10 5.74/6.46 7.22/7.96 8.45/9.10 6.16/6.69 7.61/8.15 8.80/9.33 5.69/5.96 7.19/7.56 8.45/8.87 6.05/6.60 7.61/8.08 8.87/9.25 6.45/6.93 7.89/8.40	9.05/9.61
Img.11 3.83/4.51 5.07/5.80 6.11/6.86 4.19/4.70 5.45/5.98 6.48/7.02 3.83/4.10 5.15/5.50 6.27/6.68 4.16/4.67 5.46/5.95 6.53/6.95 4.47/4.99 5.73/6.33	6.73/7.41
Img.12 5.06/5.72 6.15/6.90 7.02/7.81 5.62/6.12 6.72/7.28 7.58/8.17 5.82/5.98 7.05/7.34 7.96/8.37 6.00/6.47 7.22/7.67 8.14/8.54 6.13/6.57 7.25/7.80	8.09/8.75
Avg. 5.55/6.23 6.89/7.63 7.96/8.52 6.06/6.55 7.39/7.92 8.44/8.98 5.73/5.94 7.22/7.53 8.38/8.77 6.11/6.49 7.53/7.84 8.63/8.85 6.37/6.81 7.72/8.24	8.75/9.36



Fig. 3. Image denoising examples. Close-ups of noisy inputs and their corresponding denoised versions. Top row: grayscale denoising of input image with noise level $\sigma_n = 0.075$. Bottom row: color denoising of input image with noise level $\sigma_n = 0.1$.

In Table I we report the grayscale and color denoising results for all test images and noise levels. The performance of all the methods is measured in terms of the PSNR improvement (ISNR) w.r.t the noisy image. By inspecting the grayscale results we observe that TV is the least performing method for all noise levels. NLM filtering leads to better results than TV by exploiting the non-local self-similarity property. However it performs worse than the semi-local STV_1 regularizer. Additionally, NLTV which is inspired by NLM performs better than all these methods. Finally, NLSTV₁ consistently outperforms all the tested methods and shows an improvement over its semi-local version of about 0.3 dBs on average. Similar observations are drawn from the color denoising results. However, in this case STV_1 outperforms

TABLE II

Gaussian PSF												
Method		TV/VT	V	STV_1			NLTV (5×5)			$NLSTV_{1(7\times7)}$		
BSNR	20 dB	25 dB	30 dB	20 dB	25 dB	30 dB	20 dB	25 dB	30 dB	20 dB	25 dB	30 dB
Img.1	4.19/4.81	5.28/5.84	4 6.49/6.97	4.55/4.99	5.67/6.04	6.88/7.20	4.63/4.9	8 5.70/6.00	6.88/7.11	4.75/5.34	5.91/6.43	3 7.13/7.52
Img.2	5.00/5.78	5.67/6.3	$1 \ 6.69/7.21$	5.29/5.81	6.01/6.39	7.04/7.32	5.52/6.1	5 6.21/6.67	7.16/7.54	5.74/6.41	6.47/6.96	5.7.46 / 7.81
Img.3	4.41/5.19	5.58/6.33	3 6.89/7.56	5.12/5.54	6.27/6.68	7.52/7.92	5.09/5.0	8 6.19/6.16	7.45/7.44	5.19/5.64	6.35/6.81	7.65/8.12
Img.4	4.65/4.81	5.75/5.82	2 6.83/6.86	5.02/5.00	6.07/6.00	7.06/6.97	4.90/4.8	$1 \ 6.12/6.00$	7.28 /7.17	5.17/5.26	6.26/6.32	2 7.28/ 7.31
Img.5	4.07/4.65	5.19/5.60	$0 \ 6.54/6.84$	4.78/5.09	5.95/6.09	7.35/7.40	4.78/4.7	9 5.96/5.80	7.44/7.15	4.94/5.14	6.16/6.19	7.65/7.58
Img.6	3.14/3.60	4.00/4.3	7 5.07/5.37	3.41/3.72	4.35/4.56	5.43/5.61	3.41/3.6	1 4.38/4.46	5.52/5.60	3.56/3.93	4.53/4.81	5.66/5.89
Img.7	2.58/2.92	3.08/3.40	0 3.81/4.13	2.74/2.99	3.28/3.50	4.02/4.24	2.75/2.8	9 3.31/3.44	4.10/4.26	2.81/3.09	3.38/3.67	4.18/4.44
Img.8	3.62/4.10	4.39/4.84	4 5.33/5.74	3.93/4.21	4.68/4.95	5.59/5.85	3.92/4.0	5 4.67/4.80	5.59/5.76	4.02/4.33	4.81/5.14	5.75/6.05
Img.9	3.72/4.01	4.70/4.92	2 5.73/5.89	3.96/4.10	4.98/5.07	5.98/6.03	4.20 /4.1	8 5.13/5.11	6.12/6.09	4.18/4.26	5.22/5.25	6.23/6.22
Img.10	3.16/3.64	3.49/3.90	0 4.17/4.56	3.42/3.77	3.73/4.01	4.42/4.67	3.50/3.7	8 3.77/3.93	4.40/4.53	3.52/3.90	3.82/4.14	4.51/4.79
Img.11	3.00/3.33	3.90/4.19	9 4.92/5.14	3.27/3.46	4.20/4.35	5.16/5.28	3.18/3.2	6 4.10/4.18	5.14/5.25	$\bf 3.34/3.51$	4.28/4.44	5.29/5.44
Img.12	3.08/3.62	4.11/4.60	0 5.32/5.74	3.48/3.88	4.60/4.94	5.82/6.10	3.63/3.9	6 4.79/5.07	6.12/6.38	3.68/4.11	4.84/5.21	6.13/6.42
Avg.	3.72/4.21	4.60/5.02	1 5.65/6.00	4.08/4.38	4.98/5.21	6.02/6.22	4.13/4.2	9 5.03/5.13	6.10/6.19	4.24/4.58	5.17/5.45	6.24/6.47
						Motion 1	PSF					
Method		TV/VTV			STV_1			NLTV (5×5))	Ν	$LSTV_{1 (7 \times$	7)
BSNR	20 dB	25 dB	30 dB	20 dB	25 dB	30 dB	20 dB	25 dB	30 dB	20 dB	25 dB	30 dB
Img.1	6.11/6.74	7.93/8.43	10.18/10.55	6.54/6.95	8.36/8.66	10.61/10.81	6.59/7.04	8.40/8.74	10.59/10.89	6.93/7.45	8.76/9.18	10.98/11.31
Img.2	6.62/7.34	8.02/8.60	9.94/10.42	7.12/7.57	8.62/8.94	10.66/10.86	7.19/7.73	8.76/9.14	10.78/11.04	7.56/8.02	9.04/9.37	11.00/11.21
Img.3	0.12/0.93 3 7 80/7 70 10	8.24/9.05	10.77/11.55	6.83/7.36 9.40/9.05.1	9.03/9.56	11.60/12.12	0.83/0.98 9.12/9.02	9.03/9.16	11.58/11.79	7.03/7.70	9.30/9.97	11.91/12.53
Img.4 Img.5	6 10/6 38	8 34/8 20	10.00/10.68	6.03/6.00	0.30/8.03	13.09/13.02 12.06/11.45	7.05/6.84	0.46/8.87	13.00/12.98 19.14/11.35	8.10/8.31 1 7 36/7 38	0.83/0.30	12.60/11.00
Img.0	5 14/5 47	6 85/7 03	9 10/9 11	5.56/5.69	7 33/7 30	9 60/9 43	5 61/5 78	7.40/0.01 7.44/7.54	9 70/9 78	5.88/6.10	7.73/7.83	10.03/10.03
Img.7	4.78/5.33	6.58/7.10	8.93/9.39	5.09/5.44	6.89/7.21	9.22/9.50	5.17/5.46	7.00/7.31	9.26/9.58	5.35/5.81	7.20/7.65	9.51/9.91
Img.8	5.68/6.32	7.37/7.97	9.47/10.02	6.06/6.50	7.79/8.20	9.92/10.29	5.89/6.16	7.59/7.89	9.71/10.03	6.19/6.61	7.95/8.33	10.10/10.43
Img.9	6.15/6.42	8.35/8.55	10.95/11.06	6.50/6.58	8.71/8.74	11.31/11.29	6.53/6.69	8.68/8.78	11.19/11.20	6.79/6.84	8.98/8.98	11.53/11.45
Img.10	4.81/5.38	6.03/6.53	7.92/8.35	5.18/5.55	6.47/6.75	8.40/8.61	5.13/5.32	6.32/6.51	8.18/8.39	5.32/5.68	6.62/6.91	8.56/8.79
Img.11	5.67/6.01	7.99/8.20	10.71/10.82	6.03/6.16	8.36/8.39	11.07/11.03	5.82/6.04	8.17/8.32	10.89/10.97	6.22/6.40	8.58/8.71	11.30/11.37
Img.12	5.47/6.02	7.72/8.19	10.51/10.83	6.08/6.40	8.40/8.63	11.17/11.31	6.38/6.96	8.89/9.35	11.64/12.01	6.59/7.00	9.03/9.37	11.84/12.08
Avg.	5.88/6.34	7.82/8.17	10.22/10.46	6.36/6.60	8.35/8.48	10.78/10.81	6.36/6.59	8.38/8.50	10.77/10.83	6.67/6.96	8.69/8.88	11.11/11.21

NLTV but not $NLSTV_1$.

Besides the quantitative comparisons, to allow for a visual assessment of the reconstruction performance, we provide in Fig. 3 representative grayscale and color image denoising results. From these results we observe that our non-local functional achieves a satisfactory denoising performance without introducing staircase artifacts, which are present in TV and NLTV reconstructions, or excessively smoothing important image structures, such as in the case of NLM. Moreover, due to its non-local nature NLSTV₁ proves more efficient in removing the noise than its semi-local counterpart, STV₁.

B. Image deblurring

In the image deblurring setting we consider a circulant system matrix A that models the point spread function (PSF) of the imaging device. We test all the methods for two blurring kernels, a Gaussian of support 9×9 pixel with a standard deviation $\sigma_d = 6$ and a motion kernel of support 19×19 pixel².

²This psf was obtained from http://www.wisdom.weizmann.ac.il/~levina/papers/LevinEtalCVPR09Data.rar

As an additional degradation we consider three different levels of Gaussian noise which correspond to a blurred SNR (BSNR) of the input image of 20, 25, and 30 dBs, respectively. The BSNR is defined as BSNR = var $(Au) / \sigma_n^2$, where var (Au) is the variance of the blurred input and σ_n is the standard deviation of the noise.

For this problem a typical strategy for computing the non-local weights is to minimize the objective function w.r.t to the underlying image and the weighting function [42]. This way the weights are updated in every iteration of the algorithm. In this case, however, we need to solve a non-convex problem and, thus, we lack any guarantees about reaching the global minimum. Here, we follow a different approach which leads to a convex optimization problem. Specifically, we estimate the non-local weights from a smooth version of the image that has been previously deblurred by a Wiener filter. Then, we use these fixed weights to minimize the objective function. Finally, the stopping criteria of the minimization algorithm remain the same with the ones used in the denoising problem.

In Table II we report the grayscale and color deblurring results for all test images, blurring kernels, and noise levels. Once again the performance of the methods under comparison is measured in terms of ISNR w.r.t to the blurred and noisy image. The conclusions that we can draw from these results are on par with those in the image denoising case. Indeed, TV is the least performing regularizer, NLTV and STV_1 perform comparably, while NLSTV₁ consistently outperforms all the other regularizers both on the grayscale and color images. Representative deblurring examples are shown in Fig. 4.

C. Image reconstruction from sparse Fourier samples

In this section we examine the problem of image reconstruction from a limited number of Fourier measurements. In this case the system matrix is expressed as A = MF, where F represents the Fourier transform while M is a masking operator that retains only a subset of the Fourier coefficients and discards the rest. For our comparisons we consider a mask consisting of 32 radial lines. This corresponds to retaining about 7% of the Fourier coefficients. The Fourier measurements are further corrupted by complex Gaussian noise at three different levels. These correspond to a SNR of the fully sampled image in the Fourier domain of 10, 20, and 30 dBs. The adopted forward model is closely related to the one encountered in magnetic resonance imaging (MRI). The main difference is that in our case the underlying image is real-valued rather than complex-valued. Similarly to the deblurring problem, we use fixed weights for the non-local regularizers. These weights are computed from the back-projected image $u_b = F^H M^T v$. Finally, regarding the stopping criterion of the minimization algorithm, the number of maximum iterations is set to 200.



Fig. 4. Image deblurring examples. Close-ups of blurred and noisy inputs and their corresponding deblurred versions. Top row: grayscale deblurring of input image degraded by motion blur and noise level of BSNR=20 dBs. Bottom row : color deblurring of input image degraded by Gaussian blur and noise level of BSNR=25 dBs.

TABLE III

ISNR COMPARISONS ON GRAYSCALE/COLOR SPARSE FOURIER RECONSTRUCTION

3.6 -1 - 1					CODE I						T OTTI	
Method	TV/VTV			STV_1			NLTV (5×5)			$NLSTV_{1}(7 \times 7)$		
PSNR	10 dB	20 dB	30 dB	10 dB	20 dB	30 dB	10 dB	20 dB	30 dB	10 dB	20 dB	30 dB
Img.1	2.39/2.79 2.	.98/3.35	3.78/3.98	2.53/2.86	3.29/3.56	4.16/4.31	2.32/2.77	2.88/3.52	3.55/4.24	2.60/2.93	3.48/3.87	4.39/4.66
Img.2	3.58/4.47 3.	.85/4.67	5.09/5.81	3.86/4.56	4.09/4.74	5.43/5.99	3.48/4.34	3.36/4.36	4.27/5.28	4.02/4.73	$\bf 4.37/5.11$	5.75/6.35
Img.3	1.82/2.19 2.	.58/2.91	3.07/3.32	2.14/2.39	2.92/3.15	3.42/3.61	1.68/2.18	2.43/2.99	2.83/3.44	2.14/2.34	3.09 / 3.44	$\bf 3.61/3.98$
Img.4	2.58/3.26 4.	.01/4.41	4.76/4.91	2.88/3.42	4.29/4.62	5.03/5.18	2.24/3.20	3.19/4.21	3.66/4.67	3.06/3.84	$\bf 4.53/5.20$	5.19/5.81
Img.5	2.67/3.37 3.	.69/4.45	4.23/5.07	3.08/3.70	4.21/4.89	4.88/5.64	2.60/3.35	3.70/4.68	4.29/5.42	3.10/3.64	$\bf 4.39/5.01$	5.12/5.79
Img.6	1.68/2.14 2.	.07/2.46	2.43/2.78	1.82/2.22	2.27/2.62	2.68/3.02	1.61/2.07	1.92/2.39	2.24/2.78	1.87/2.30	2.32 / 2.76	2.75/3.21
Img.7	1.48/1.82 1.	.95/2.21	2.34/2.46	1.64/1.89	2.16/2.34	2.59/2.66	1.33/1.65	1.86/2.09	2.29/2.44	1.63/ 1.94	2.20 / 2.40	2.63/2.72
Img.8	2.10/2.40 2.	.69/2.86	3.00/3.08	2.30/2.52	2.92/3.03	3.26/3.30	2.02/2.22	2.51/2.74	2.78/3.00	2.35/2.62	2.99 / 3.13	3.32/3.40
Img.9	1.56/1.72 1.	.91/1.92	2.04/2.02	1.69/1.78	2.06/2.03	2.22/2.16	1.46/1.60	1.80/1.85	1.96/1.99	1.75/1.84	2.10/2.01	2.26/2.10
Img.10	2.67/3.16 2.	.87/3.41	3.67/4.04	2.87/3.25	3.10/3.51	3.92/4.21	2.66/2.98	2.84/3.30	3.59/4.04	2.83/3.28	$\bf 3.14/3.57$	3.99/4.27
Img.11	1.05/1.34 1.	.53/1.69	1.80/1.85	1.21/1.43	1.74/1.83	2.02/2.05	0.90/1.19	1.37/1.60	1.64/1.82	1.16/1.47	1.73/1.87	2.01/2.09
Img.12	1.02/1.48 1.	.45/1.88	1.76/2.14	1.23/1.61	1.76/2.11	2.10/2.44	1.04/1.56	1.55/2.13	1.85/2.45	1.28/1.68	1.87/2.28	2.25/2.64
Avg.	2.05/2.51 2.	.63/3.02	3.16/3.46	2.27/2.64	2.90/3.20	3.48/3.71	1.94/2.43	2.45/2.99	2.91/3.46	2.32/2.72	3.02/3.39	$\bf 3.61/3.92$

In Table III we provide the ISNR scores of all the methods under comparison on both grayscale and color images. Similarly to the previous two inverse problems, we observe that on average our $NLSTV_1$ regularizer leads to the best reconstruction performance. On the contrary, NLTV shows a different behavior and does not perform as well as it did in the denoising and deblurring tasks. In fact, the results indicate that it is the worst performing method. This might be due to the computation of the non-local weights from the back-projected images whose quality are rather poor compared to the ground-truth data. If this



Fig. 5. Examples of image reconstruction from sparse Fourier measurements. Close-ups of compressed and noisy inputs and their corresponding reconstructed versions. Top row: grayscale input (back-projected) image sampled with a radial mask of 32 lines and at noise level of SNR=30 dBs. Bottom row: color input (back-projected) image sampled with a radial mask of 32 lines and at a noise level of SNR=20 dBs.

is the case, then $NLSTV_1$ is less sensitive in the choice of the non-local weights since it seems not to be affected as much as NLTV.

VI. CONCLUSIONS

In this work we combined ideas from local and non-local regularization strategies and proposed a novel family of non-local functionals to regularize inverse imaging problems. Our non-local regularizers differ from the existing ones in the sense that they employ a non-local version of the structure tensor as the regularization operator. Therefore, they depend on the standard image gradient rather than the graph gradient. This way we are able to exploit both the local structural regularity and the non-local self-similarity properties of natural images. Further, we proposed an efficient minimization algorithm that is based on a variable splitting strategy. Finally, we assessed the reconstruction performance of our regularizers on several inverse imaging problems. Our reconstruction results were shown to compare favorably to the ones obtained by other competing local and non-local regularization methods.

REFERENCES

- [1] M. Bertero and P. Boccacci, Introduction to Inverse Problems in Imaging. IOP Publishing, 1998.
- [2] L. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D*, vol. 60, pp. 259–268, 1992.

- [3] T. Chan, A. Marquina, and P. Mulet, "High-order total variation-based image restoration," SIAM J. Sci. Comput., vol. 22, pp. 503–516, 2000.
- [4] G. Gilboa, N. Sochen, and Y. Zeevi, "Variational denoising of partly textured images by spatially varying constraints," *IEEE Trans. Image Process.*, vol. 15, no. 8, pp. 2281–2289, 2006.
- [5] M. Grasmair, "Locally adaptive total variation regularization," in *Scale Space and Variational methods in Computer Vision*, 2009, pp. 331–342.
- [6] M. Grasmair and F. Lenzen, "Anisotropic total variation filtering," *Applied Mathematics & Optimization*, vol. 62, pp. 323–339, 2010.
- [7] S. Lefkimmiatis, A. Roussos, P. Maragos, and M. Unser, "Structure tensor total variation," SIAM J. Imaging Sci., 2015, in press.
- [8] K. Bredies, K. Kunisch, and T. Pock, "Total generalized variation," SIAM J. Imaging Sci., vol. 3, pp. 492–526, 2010.
- [9] S. Lefkimmiatis, J. Ward, and M. Unser, "Hessian Schatten-norm regularization for linear inverse problems," *IEEE Trans. Image Process.*, vol. 22, no. 5, pp. 1873–1888, 2013.
- [10] G. Gilboa and S. Osher, "Nonlocal operators with applications to image processing," *Multiscale Model. Simul.*, vol. 7, pp. 1005–1028, 2008.
- [11] D. Zhou and B. Schölkopf, "Regularization on discrete spaces," in Pattern Recognition. Springer, 2005, pp. 361–368.
- [12] A. Elmoataz, O. Lezoray, and S. Bougleux, "Nonlocal discrete regularization on weighted graphs: a framework for image and manifold processing," *IEEE Trans. Image Proces.*, vol. 17, pp. 1047–1060, 2008.
- [13] M. Figueiredo, J. Bioucas-Dias, and R. Nowak, "Majorization-minimization algorithms for wavelet-based image restoration," *IEEE Trans. Image Process.*, vol. 16, pp. 2980–2991, 2007.
- [14] I. Selesnick and M. Figueiredo, "Signal restoration with overcomplete wavelet transforms: Comparison of analysis and synthesis priors," in *SPIE (Wavelets XIII)*, 2009.
- [15] H. Attouch, G. Buttazzo, and G. Michaille, Variational Analysis in Sobolev and BV Spaces: Applications to PDEs and Optimization, ser. MPS-SIAM Series on Optimization. Society for Industrial and Applied Mathematics, 2006.
- [16] P. Blomgren and T. Chan, "Color TV: Total Variation methods for restoration of vector-valued images," *IEEE Trans. Image Process.*, vol. 7, pp. 304–309, 1998.
- [17] B. Goldluecke, E. Strekalovskiy, and D. Cremers, "The natural vectorial total variation which arises from geometric measure theory," SIAM J. Imaging Sci., vol. 5, pp. 537–563, 2012.
- [18] W. Förstner and E. Gülch, "A fast operator for detection and precise location of distinct points, corners and centres of circular features," in *Proc. ISPRS Intercommission Conference on Fast Processing of Photogrammetric data*, 1987, pp. 281–305.
- [19] A. Buades, B. Coll, and J.-M. Morel, "A review of image denoising algorithms, with a new one," *Multiscale Model. Simul.*, vol. 4, pp. 490–530, 2005.
- [20] —, "Image denoising methods. A new nonlocal principle," SIAM review, vol. 52, pp. 113-147, 2010.
- [21] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *IEEE Int. Conf. Comput. Vision*, 1998, pp. 839–846.
- [22] S. Kindermann, S. Osher, and P. W. Jones, "Deblurring and denoising of images by nonlocal functionals," *Multiscale Model. Simul.*, vol. 4, pp. 1091–1115, 2005.
- [23] G. Gilboa and S. Osher, "Nonlocal linear image regularization and supervised segmentation," *Multiscale Model. Simul.*, vol. 6, pp. 595–630, 2007.

- [24] Q. Cheng, H. Shen, L. Zhang, and P. Li, "Inpainting for remotely sensed images with a multichannel nonlocal total variation model," *IEEE Trans. Geoscience and Remote Sensing*, vol. 52, pp. 175–187, 2014.
- [25] G. Chierchia, N. Pustelnik, B. Pesquet-Popescu, and J.-C. Pesquet, *IEEE Trans. Image Processing*, vol. 23, pp. 5531–5544, 2014.
- [26] A. Chambolle, "An algorithm for total variation minimization and applications," J. Math. Imag. Vis., vol. 20, pp. 89–97, 2004.
- [27] J. Darbon, A. Cunha, T. Chan, S. Osher, and G. Jensen, "Fast nonlocal filtering applied to electron cryomicroscopy," in *IEEE International Symposium on Biomedical Imaging*, 2008, pp. 1331–1334.
- [28] P. Viola and M. Jones, "Rapid object detection using a boosted cascade of simple features," in *IEEE Int. Conf Comput. Vision Patt. Recogn. (CVPR)*, vol. 1, 2001, pp. 511–518.
- [29] L. Condat, "A simple trick to speed up and improve the non-local means," 2010. [Online]. Available: http://www.gipsa-lab.grenoble-inp.fr/~laurent.condat/publis/condat_resreport_NLmeansv3.pdf
- [30] R. Bhatia, Matrix Analysis. Springer, 1997.
- [31] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," J. Math. Imag. Vis., vol. 40, no. 1, pp. 120–145, 2011.
- [32] T. Goldstein, E. Esser, and R. Baraniuk, "Adaptive primal-dual hybrid gradient methods for saddle-point problems," arXiv preprint arXiv:1305.0546, 2013.
- [33] D. Bertsekas, Constrained optimization and Lagrange multiplier methods. Athena Scientific, 1996.
- [34] J. Eckstein and D. Bertsekas, "On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators," *Mathematical Programming*, vol. 55, no. 1, pp. 293–318, 1992.
- [35] E. Esser, "Applications of Lagrangian-based alternating direction methods and connections to split Bregman," CAM report, vol. 9, 2009.
- [36] 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, 2011.
- [37] P. L. Combettes and V. R. Wajs, "Signal recovery by proximal forward-backward splitting," *Multiscale Model. Simul.*, vol. 4, pp. 1168–1200, 2005.
- [38] S. Lefkimmiatis and M. Unser, "Poisson image reconstruction with Hessian Schatten-norm regularization," *IEEE Trans. Image Process.*, vol. 22, pp. 4314–4327, 2013.
- [39] J. Liu and J. Ye, "Efficient ℓ_1/ℓ_q norm regularization," Arxiv preprint arXiv:1009.4766, 2010.
- [40] D. L. Donoho, "Denoising by soft-thresholding," IEEE Trans. Inf. Theory, vol. 41, pp. 613–627, 1995.
- [41] J. R. Shewchuk, "An introduction to the conjugate gradient method without the agonizing pain," 1994. [Online]. Available: http://www.cs.cmu.edu/~jrs/jrspapers.html
- [42] X. Zhang, M. Burger, X. Bresson, and S. Osher, "Bregmanized nonlocal regularization for deconvolution and sparse reconstruction," *SIAM Journal on Imaging Sciences*, vol. 3, no. 3, pp. 253–276, 2010.