Texture Reconstruction guided by a High-Resolution Patch

Mireille El Gheche, Jean-François Aujol, Yannick Berthoumieu, Charles-Alban Deledalle

► To cite this version:

Mireille El Gheche, Jean-François Aujol, Yannick Berthoumieu, Charles-Alban Deledalle. Texture Reconstruction guided by a High-Resolution Patch. IEEE Transactions on Image Processing, 2017, 26 (2), pp.549-560. 10.1109/TIP.2016.2627812 . hal-01276582v3

HAL Id: hal-01276582

<https://hal.science/hal-01276582v3>

Submitted on 27 Jul 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Texture Reconstruction guided by a High-Resolution Patch

Mireille El Gheche, Jean-François Aujol, *Member, IEEE*,
Yannick Berthoumieu, *Member, IEEE*, Charles-Alban Deledalle

Abstract—In this paper, we aim at super-resolving a low-resolution texture under the assumption that a high-resolution patch of the texture is available. To do so, we propose a variational method that combines two approaches, that are texture synthesis and image reconstruction. The resulting objective function holds a nonconvex energy that involves a quadratic distance to the low-resolution image, a histogram-based distance to the high-resolution patch, and a nonlocal regularization that links the missing pixels with the patch pixels. As for the histogram-based measure, we use a sum of Wasserstein distances between the histograms of some linear transformations of the textures. The resulting optimization problem is efficiently solved with a primal-dual proximal method. Experiments show that our method leads to a significant improvement, both visually and numerically, with respect to state-of-the-art algorithms for solving similar problems.

Index Terms—Super-Resolution, Texture Synthesis, Texture Reconstruction, Wasserstein distance, Histograms, Nonlocal regularization, Proximal algorithms, Nonconvex optimization.

I. INTRODUCTION

TEXTURES refer to the visual appearance of an object surface resulting from the size, shape, density, arrangement, proportion of its elementary parts. Human observers perceive textures as images displaying local spatial variations of features (like color, orientation and intensity) organized in some regular or repeated pattern. These visual patterns provide useful cues about the physical properties of the underlying surface. Consequently, the analysis of texture images plays a central role in multiple disciplines, such as neuroscience and psychophysics (for texture detection), computer vision (for defect detection, classification, and segmentation), and petrology (for physical property evaluation).

Among the several applications of texture analysis, the one that motivates our work is the physical parameter estimation of petroleum reservoirs from their textures. Indeed, one can

Mireille El Gheche is with Université de Bordeaux, IMS UMR 5218 and IMB UMR 5251, F-33400 Talence, France. E-mail: mireille.el-gheche@u-bordeaux.fr

Jean-François Aujol is with the CNRS, Institut Universitaire de France, Université de Bordeaux, IMB UMR 5251, F-33400 Talence, France. E-mail: jean-francois.aujol@math.u-bordeaux1.fr.

Yannick Berthoumieu is with the Institut Polytechnique de Bordeaux, Université de Bordeaux, IMS UMR 5218, F-33400 Talence, France. E-mail: yannick.berthoumieu@ims-bordeaux.fr.

Charles-Alban Deledalle is with the CNRS, Université de Bordeaux, IMB UMR 5251, F-33400 Talence, France. E-mail: charles-alban.deledalle@math.u-bordeaux1.fr.

This work was supported by the French National Research Agency through the Investments for the Future Programme IdEx BordeauxCPU under Grant ANR-10-IDEX-03-02. The work of J.-F. Aujol was supported by the Institut Universitaire de France.

Fig. 1. The left figure shows a complete low-resolution acquisition, which was up-sampled by a factor 2 for visualization purposes. The middle picture displays a high-resolution patch of the target texture. Our goal is to combine both observations in order to get a high-resolution texture (right figure).

extract a cylindrical sample of rock from a reservoir, scan it through computed tomography (CT), and obtain a stack of texture images, which can be then analyzed to find visual cues related to some petro-physical parameters. A typical example concerns the segmentation and classification of textures, which in turn relate to the porosity levels of the underlying material.

Due to the large size of a rock sample, the CT scanner produces a low-resolution image of the reservoir texture, leading to a loss of important details for the estimation of petro-physical parameters. Fortunately, the resolution can be improved by zooming into a specific part of the sample through a micro-CT scanner. Therefore, in order to make a high-resolution image of the reservoir texture, one can perform two CT acquisitions of the same rock sample. The first one supplies a low-resolution image of the entire texture, while the second one provides a high-resolution image of a small part of the texture. The same approach can be used for the analysis of micro-structure materials, whose scanned volume resolution should go to one nanometer to measure its porosity [1], or when using a digital camera for texture defect detection [2], which plays an important role in automatic inspection of most production processes (food, textile, steel industries, ...).

The main focus of this paper is on the super-resolution of a texture image, under the assumption that a small patch of the sought texture is available in high-resolution (the considered setting is illustrated in Fig. 1). Note that this problem can be related to:

- texture synthesis from a high-resolution patch, guided by a low-resolution texture;
- image reconstruction from a low-resolution texture, guided by a high-resolution patch;

However, neither of them exactly describes the considered

problem. Indeed, by texture reconstruction, we refer to the problem of generating a high-resolution image not necessary pixel-wise identical to the unknown underlying image, but only identical in terms of textural statistics. Obviously, the sought image should satisfy the constraint that its low-resolution version is faithful to the observed low-resolution one (at least up to the noise).

A. Related work

a) Reconstruction: The simplest way to upscale a low-resolution image is by mean of linear scaling (such as bi-cubic interpolation [3]) or image sharpening methods [4]. However, these methods introduce high frequency concealing, which blur the image and erase small texture details.

A number of single-image reconstruction methods are based on the assumption that a dataset of high-resolution textures are available, from which they learn a dictionary of low-resolution patches, a dictionary of high-resolution patches, and a correspondence map between the two dictionaries [5], [6], [7]. However, the results depend on the estimated dictionaries, and thus the reconstruction of true details is not guaranteed if the target texture does not appear in the training dataset. A similar reconstruction method was described in [8], where the k -nearest neighbors are searched into an external patch database. Instead of using external data, the authors in [9] extract correspondences at different scales directly from the low-resolution image. Some recent approaches learn mapping functions from external low- and high-resolution exemplar pairs [10], [11], [12]. Nevertheless, these techniques are not very robust to noise.

A different approach consists of introducing a suitable regularization into the reconstruction problem formulation, with the aim of conveying some prior knowledge about the signal to be recovered. In this context, Total Variation (TV) [13], [14], [15], [16] has emerged as a simple regularization, consisting in penalizing the gradient coefficients. However, TV fails to preserve textures, details, and structures, because they are hardly distinguishable from noise. To improve this behavior, the TV model has been extended by using higher-order spatial differences [17], [18], higher-degree directional derivatives [19], [20], or the nonlocality principle [21], [22], [23], [24], [25], [26], [27], [28]. The latter approach leads to the so-called nonlocal total variation (NLTV) regularization.

b) Texture synthesis: Texture synthesis techniques [29] can be broadly categorized into region-growing local methods and optimization-based global methods. Local approaches grow the texture one pixel (or patch) at a time, while maintaining the spatial coherence with nearby pixels by modeling the neighborhoods with Markov fields and fractal models [30]. A weakness of these methods is that the spatial coherence between pixels is enforced at a local scale. A possible approach to circumvent this limitation consists of resorting to a small patch from which to grow the texture, following a procedure that sequentially processes the pixels [31], [32], [33], [34], [35]. In such methods, however, small errors can accumulate over large distances, leading to inconsistencies in the synthesized texture. Another possible scenario is to use several exemplars

describing the texture at various scales and hence generate textures with arbitrary amount of details [36], [37].

Global methods process the entire texture as a whole, using some criteria for measuring its similarity with a small texture patch. For example, the latter can be modeled with a statistical descriptor based on histograms [38], [39], wavelet coefficients [40], or Fourier coefficients [41]. A similar approach was recently proposed in [42], [43], which introduces a preliminary step of dictionary learning for exploiting the given patch, and (not least) the Wasserstein distance for comparing the histograms of the entire texture with an extended version of the small patch. While the Wasserstein distance is well-known in image processing and computer vision under the name “earth mover distance” [44], it was only recently expanded to the context of texture synthesis [45], [43].

Regarding the problem considered in this paper (recall Fig. 1), an extension of NLTV regularization was recently proposed in [46]. This prior consists in building a pointwise estimate of the image, where each missing pixel is connected to a local neighborhood and to the most similar pixels in the given high resolution patch. The author in [47] proposed a statistical prior in addition to the standard NLTV regularization. This prior is based on a nonparametric spatial covariance structure from empirical estimation, and a parametric generalized Gaussian model learned from the high-resolution patch.

B. Contributions

In this paper, we follow a variational approach that aims at super-resolving a low-resolution texture (possibly degraded by blur and zero-mean additive Gaussian noise) by explicitly taking into account a high-resolution patch. To do so, we propose an algorithm for optimizing a nonconvex energy that involves two problem-specific terms in addition to the standard quadratic distance to the low-resolution image. The first term is a sum of Wasserstein distances to the histograms of the high-resolution patch, of its gradient, and of its Laplacian. The second term is a NLTV regularization that measures the distance of each pixel with a selected subset of pixels in the high-resolution patch. The variational approach adopted in this paper was inspired from [48], which uses the Wasserstein distance for comparing the histograms of the entire texture with a parametric model learned from the small patch, and introduces a preliminary step of dictionary learning for exploiting the given patch, in order to perform texture synthesis and image restoration.

There exist few methods in the literature that consider the texture synthesis problem guided by a low-resolution texture [46], [47]. These approaches have some similarities with the proposed one with notable important differences:

- In [46], the problem is solved by taking into account a nonlocal model. The approach consists in connecting each missing pixel to a set of neighboring pixels and some other pixels in the high-resolution patch. However, this approach neglects the statistical information that can be extracted from the high-resolution patch. The main novelty of the proposed approach w.r.t. [46] lies in the introduction of statistical prior that preserves the texture details, as well as the design of a new nonlocal graph that

provides better connections between the missing pixels and the high-resolution patch pixels.

- To interpolate the missing data, the work in [47] exploits a nonlocal regularization and a Generalized Gaussian distributions model of the texture gradients, whose parameters are learned on the high-resolution patch. Differently from [47], we enforce a histogram prior using the Wasserstein distance, linking our approach with the optimal transport [49]. The originality of our technique consists of the ability to consider multi-histogram priors (such as the intensity values, the gradients and the Laplacians histograms [48, Chapter 5]) without being constrained to a parametric model, regardless of the histogram shapes (unimodal, bimodal or multimodal mixture).

The paper is organized as follows. Section II describes the degradation model, the histogram-based distance, and the new graph for the NLTV regularization. Section III presents an algorithm for solving the proposed optimization problem via proximal tools. Section IV provides an experimental validation in the context of texture images. Finally, the conclusion is given in Section V.

C. notation

Let $\|\cdot\|$ be the standard Euclidean norm, and let Id be the identity matrix. The domain of a function $f : \mathbb{R}^N \rightarrow]-\infty, +\infty]$ is $\text{dom } f = \{x \in \mathbb{R}^N | f(x) < +\infty\}$. $\Gamma_0(\mathbb{R}^N)$ is the class of lower semi-continuous convex functions from $\mathbb{R}^N$ to $]-\infty, +\infty]$ such that $\text{dom } f \neq \emptyset$. Let $f \in \Gamma_0(\mathbb{R}^N)$. The conjugate of f is the function $f^* \in \Gamma_0(\mathbb{R}^N)$ defined by $f^* : \mathbb{R}^N \rightarrow]-\infty, +\infty]$: $u \mapsto \sup_{x \in \mathbb{R}^N} x^\top u - f(x)$. When f is Gâteaux-differentiable at $y \in \mathbb{R}^N$, $\partial f(y) = \{\nabla f(y)\}$ where $\nabla f(y)$ is the gradient of f at y . A differentiable convex function has β -Lipschitz continuous gradient ∇f if $(\forall (x, y) \in \mathbb{R}^N \times \mathbb{R}^N) \quad \|\nabla f(x) - \nabla f(y)\| \leq \beta \|x - y\|$, where $\beta \in]0, +\infty[$. Let C be a nonempty subset of $\mathbb{R}^N$, then ι_C is the indicator function of C , equal to 0 on C and $+\infty$ otherwise.

II. PROPOSED APPROACH

The high-resolution signal of interest is denoted by $\bar{x} \in \mathbb{R}^N$, which generally corresponds to an image of size $N = N_1 \times N_2$. The degradation model that we consider is the following:

$$\begin{cases} z^{(1)} = \text{DB}\bar{x} + \eta_1, \\ z^{(2)} = \text{M}\bar{x}. \end{cases} \quad (1)$$

Hereabove, $z^{(1)} \in \mathbb{R}^Q$ is the complete low-resolution image, $\text{B} \in \mathbb{R}^{N \times N}$ is a linear operator modeling some blur, $\text{D} \in \mathbb{R}^{Q \times N}$ stands for spatial down-sampling by a dyadic factor in each direction, yielding $Q = 2^{-r}N$, and $\eta_1 \in \mathbb{R}^Q$ is a realization of an additive zero-mean white Gaussian noise with standard deviation τ . Moreover, $z^{(2)} \in \mathbb{R}^M$ denotes a small patch of the high-resolution image $\bar{x}$, and $\text{M} \in \mathbb{R}^{M \times N}$ is a selection operator that extracts the high-resolution patch from $\bar{x}$.

We propose to recover $\bar{x}$ from the observations $z^{(1)}$ and $z^{(2)}$ through a variational approach that leads to solving the following optimization problem:

$$\begin{aligned} & \underset{x \in \mathbb{R}^N}{\text{minimize}} \quad \|\text{DB}x - z^{(1)}\|^2 + H(x, z^{(2)}) + \lambda R(x, z^{(2)}) \\ & \text{s. t.} \quad \text{M}x = z^{(2)}, \end{aligned} \quad (2)$$

where $\lambda > 0$ is a regularization parameter. Beside the data fidelity terms w.r.t. the observations $z^{(1)}$ and $z^{(2)}$, we use two additional pieces of information: a term $H(\cdot, z^{(2)})$ conveying some histogram-based statistics, and a term $R(\cdot, z^{(2)})$ enforcing a new type of nonlocal regularization. The histogram-based term is modeled through the Wasserstein distance discussed in Section II-A, while the regularization is grounded on the nonlocal approach presented in Section II-B.

The presented approach takes into account an image regularization and a global statistical information in one model. Note that, minimizing the term $R(\cdot, z^{(2)})$ entails spatial regularization, whereas $H(\cdot, z^{(2)})$ involves statistical information that is not spatially indexed. However, taking into account the statistical information through histogram priors, the proposed approach infers the correct structures and details of the high-resolution patch to the estimated structure, leading to improved results w.r.t. the classical regularizations considered alone.

A. Wasserstein distance to the high-resolution patch

Wasserstein distance is a well-known dissimilarity measure between probability distributions, as it is easier to optimize than other distances, such as the φ -divergences (Kullback-Leibler, Hellinger, Jeffreys-Kullback, $\dots$) [45], [43].

The Wasserstein distance between the histograms¹ of two images $u \in \mathbb{R}^N$ and $v \in \mathbb{R}^N$ is defined as follows:

$$\mathcal{W}_2^2(\nu_u, \nu_v) = \min_{\sigma \in \Sigma_N} \|u - v \circ \sigma\|^2, \quad (3)$$

where ν_u and ν_v are the normalized histograms of u and v , the symbol $v \circ \sigma$ denotes a permutation of the vector v , and Σ_N is the set of all the permutations of N -length vectors. For grayscale images, the optimal permutation σ^* is computed as

$$\sigma^* = \sigma_v \circ \sigma_u^{-1}, \quad (4)$$

where σ_v (resp. σ_u) denotes the permutation operator that arranges the pixels of v (resp. u) in ascending order (see Fig. 2). In our case, however, the two images have a different number of pixels ($u \in \mathbb{R}^N$ and $v \in \mathbb{R}^M$, with $M < N$). Hence, we replicate the patch v so as to obtain a larger image $\tilde{v} \in \mathbb{R}^N$ such that the normalized histogram $\nu_{\tilde{v}}$ is equal to ν_v .² Although the Wasserstein distance is nonconvex (due to the histogram transformation), its gradient is Lipschitz-continuous and takes the following form [42],[48, Chapter 5]

$$\nabla_u \mathcal{W}_2^2(\nu_u, \nu_{\tilde{v}}) = 2(u - \tilde{v} \circ \sigma_{\tilde{v}} \circ \sigma_u^{-1}). \quad (5)$$

This property allows us to employ the Wasserstein distance into the optimization algorithm presented in Section III.

¹The histogram of a signal refers to a histogram of the pixel intensity values.

²Another solution consists in oversampling the sorted sequence $v \circ \sigma_v$ using a nearest-neighbor approximation.

Fig. 2. A simple example showing histogram transfer from v to u .

Our first contribution is to use the Wasserstein distance for comparing the histograms of the intensity values, their gradient, and their Laplacian, in order to ensure that the synthesized image has similar details as the high-resolution patch. In other words, we want to restore the gray levels, the gradients and the Laplacian values of the textures. To do so, for every $s \in \{1, \dots, 4\}$, let L_s be a matrix in $\mathbb{R}^{N_s \times N}$ defined as follows.

- L_1 : the identity matrix ($N_1 = N$).
- L_2 : the concatenation of the horizontal and vertical difference operators ($N_2 = 2N$).
- L_3 : the concatenation of the diagonal difference operators ($N_3 = 2N$).
- L_4 : the isotropic Laplacian operator ($N_4 = N$).

We define the term H in (2) as

$$H(x, z^{(2)}) = \sum_{s=1}^4 \alpha_s \mathcal{W}_2^2(\nu_{L_s x}, \nu_{z_s^{(2)}}) \quad (6)$$

where, for every $\forall s \in \{1, \dots, 4\}$, $\alpha_s > 0$, $z_s^{(2)} = L_s z^{(2)}$ and $\tilde{z}_s^{(2)}$ is the extension of $z_s^{(2)}$ (after the linear transformation).

B. Nonlocal regularization

The quality of the images reconstructed with the nonlocal regularization highly depends on the graph used to model the similarity between pixels. Remarkable results can be obtained when a good estimate of the true similarity graph is available. However, in the context of super-resolution, this task presents some challenges, due to the missing information in the low-resolution image. One way to circumvent this issue consists of estimating the standard nonlocal graph over an interpolated version of the low-resolution image. Unfortunately, this approach performs poorly on texture images, as the graph connections are limited to local neighborhoods. A better solution was proposed in [46], where the missing pixels in the low-resolution image are linked (by patch similarity) to a set of pixels in the high-resolution patch. Since the original pixels are blurry, the pixels to which they are connected are also blurry, because the quadratic distance used to evaluate the similarity is not robust to this kind of degradation. Hence, the image reconstructed with such a nonlocal graph is necessarily blurred. The proposed approach is different, since each pixel is connected to the similar ones in the high-resolution patch, which are not blurry. Hence, the similarity scores are more likely to be reliable, resulting in a denoised image of better quality.

In the following, we present the way to build the new nonlocal graph (see Fig. 3), and we review two state of the art approaches: the conventional one and the one proposed in [46].

Fig. 3. Brodatz Bark image: bicubic interpolation of the low resolution image with (a) the high-resolution patch, (b) the location of the high-resolution patch (black square). Dashed Lines are the transitional links to build the final graph (straight lines).

1) *The proposed graph*: we define the function R in (2) as follows

$$R(x, z^{(2)}) = \sum_{i=1}^N \sqrt{\sum_{j \in \mathcal{N}_i} w_{i,j} (x_i - z_j^{(2)})^2}, \quad (7)$$

where $\mathcal{N}_i \subset \{1, \dots, M\}$ is a subset of positions w.r.t. the support of $z^{(2)}$, and $w_{i,j} > 0$. For every $i \in \{1, \dots, N\}$, the support $\mathcal{N}_i$ and the weights $(w_{i,j})_{j \in \mathcal{N}_i}$ are computed as explained in the following.

- Firstly, we select the pixel $z_{\bar{j}^{(i)}}^{(2)}$ that is the most similar to x_i . However, since the high-resolution image x is not available at this stage, we interpolate the low-resolution image $z^{(1)} \in \mathbb{R}^Q$ with a bicubic approach, so as to obtain a new image $\hat{z}^{(1)} \in \mathbb{R}^N$ that approximates x without loosing its spectral structure. Once we have $\hat{z}^{(1)}$, we find the index $\bar{j}^{(i)}$ according to the criterion

$$\bar{j}^{(i)} = \underset{m \in \{1, \dots, M\}}{\operatorname{argmin}} \|p_i(\hat{z}^{(1)}) - p_m(z^{(2)})\|^2, \quad (8)$$

where $p_i(\hat{z}^{(1)})$ is a block extracted from $\hat{z}^{(1)}$ and centered around the position i , and $p_m(z^{(2)})$ is defined similarly.

- Secondly, we define $\mathcal{N}_i$ as the set of positions $j_1, \dots, j_K$ such that $(z_{j_k}^{(2)})_{1 \leq k \leq K}$ are the K most similar pixels to $z_{\bar{j}^{(i)}}^{(2)}$. To do so, for every $m \in \{1, \dots, M\}$, we compute the Euclidean distance between the blocks $p_{\bar{j}^{(i)}}(z^{(2)})$ and $p_m(z^{(2)})$, and we select the positions $(j_k)_{1 \leq k \leq K}$ associated to the K lowest distances, which of course includes $\bar{j}^{(i)}$ itself (i.e., $j_1 = \bar{j}^{(i)}$).
- Finally, we compute the weights $w_{i,j}$ as follows

$$(\forall j \in \mathcal{N}_i) \quad w_{i,j} = e^{-\delta \|p_{\bar{j}^{(i)}}(z^{(2)}) - p_j(z^{(2)})\|^2}, \quad (9)$$

where δ is a positive constant.

The nonlocal energy proposed in (7) differs from the standard nonlocal regularization, as it takes into account the information carried by the high-resolution patch. To clarify this concept, we now review two nonlocal approaches related to our work.

2) *The conventional approach*: The classical NLTV consists of computing, for every $i \in \{1, \dots, N\}$, the most similar pixels lying in a local neighborhood, yielding the energy

$$\text{NLTV}(x) = \sum_{i=1}^N \sqrt{\sum_{j \in \mathcal{N}_i \subset \mathcal{W}_i} \hat{w}_{i,j} (x_i - x_j)^2}, \quad (10)$$

where $\mathcal{W}_i$ is the set of positions located into a $Q \times Q$ window $\mathcal{W}_i \subset \{1, \dots, N\} \setminus \{i\}$ centered at i . For every $i \in \{1, \dots, N\}$, the support $\mathcal{N}_i$ includes the positions of the K most similar pixels to x_i that are located within $\mathcal{W}_i$. Since the image x is not available, one can interpolate $z^{(1)} \in \mathbb{R}^Q$ and obtain $\hat{z}^{(1)} \in \mathbb{R}^N$, which is used for building $\mathcal{N}_i$ and for computing the associated weights

$$(\forall j \in \mathcal{N}_i) \quad \hat{w}_{i,j} = e^{-\delta \|p_i(\hat{z}^{(1)}) - p_j(\hat{z}^{(1)})\|^2}. \quad (11)$$

3) *The approach in [46]*: This approach consists of connecting each pixel $(x_i)_{i \in \{1, \dots, N\}}$ to its K -nearest observed pixels in the high-resolution patch $z^{(2)}$. The nonlocal regularization is expressed as

$$\text{NL-HR}(x) = \sum_{i=1}^N \sqrt{\sum_{j \in \mathcal{N}_i} \bar{w}_{i,j} (x_i - z_j^{(2)})^2}, \quad (12)$$

where $\mathcal{N}_i \subset \{1, \dots, M\}$ includes the positions of the pixels in $z^{(2)}$ that are the K most similar to x_i . Just like before, the interpolated image $\hat{z}^{(1)} \in \mathbb{R}^N$ is actually used for building $\mathcal{N}_i$ and for computing the associated weights

$$(\forall j \in \mathcal{N}_i) \quad \bar{w}_{i,j} = e^{-\delta \|p_i(\hat{z}^{(1)}) - p_j(\hat{z}^{(1)})\|^2}. \quad (13)$$

Note that the approach in [46] builds the similarity graph over an interpolated version of the low-resolution image. Conversely, we propose to build the graph by only using the high-resolution patch (see Fig. 3). This makes a very important difference in the quality of reconstructed images, as shown in Section IV.

III. OPTIMIZATION

The solution of Eq. (2) requires an efficient algorithm for dealing with problems involving nonsmooth functions and linear operators that are non-necessarily circulant. Recently, it has been shown experimentally that primal-dual proximal methods [50], [51], [52], [53], [54], [55], which were originally designed for convex optimization, can be also applied to nonconvex problems in some circumstances [56], [57], [58], [59]. In the convex setting, the key tool of these methods is the proximity operator [60] of a lower semicontinuous convex function $\varphi: \mathbb{R}^N \mapsto [-\infty, +\infty]$, defined as

$$(\forall y \in \mathbb{R}^N) \quad \text{prox}_\varphi(y) = \underset{z \in \mathbb{R}^N}{\text{argmin}} \quad \varphi(z) + \frac{1}{2} \|z - y\|^2. \quad (14)$$

The proximity operator can be interpreted as an implicit subgradient step for the function φ , since $p = \text{prox}_\varphi(y)$ is uniquely defined through the inclusion $y - p \in \partial\varphi(p)$. Proximity operators enjoy many properties [61]. In particular, they generalize the notion of projection onto a closed convex set C , in the sense that $\text{prox}_{\iota_C} = P_C$. Hence, proximal methods provide a unifying framework that allows one to address a wide class of convex optimization problems involving nonsmooth penalizations and hard constraints.

Among the wide array of existing proximal algorithms, we employ the Forward-Backward Primal Dual method (FBPD) [54] reported in Algorithm 1. The operators required by this algorithm are detailed below.

- The nonlocal regularization can be expressed as the $\ell_{1,2}$ -norm composed with a discrete difference operator, yielding

$$R(x, z^{(2)}) = \|\mathbf{T}x\|_{1,2}, \quad (15)$$

where

$$\mathbf{T}x = \begin{bmatrix} [w_{1,j}(x_1 - z_j^{(2)})]_{j \in \mathcal{N}_1} \\ \vdots \\ [w_{N,j}(x_N - z_j^{(2)})]_{j \in \mathcal{N}_N} \end{bmatrix} \begin{matrix} \in \mathbb{R}^K \\ \vdots \\ \in \mathbb{R}^K \end{matrix} \quad (16)$$

The proximity operator of the $\ell_{1,2}$ -norm, that is $\text{prox}_{\|\cdot\|_{1,2}}$ in Algorithm 1, can be found in [62].

- The projection onto the convex set associated to the constraint $\mathbf{M}x = z^{(2)}$ is expressed as

$$P_{\{\mathbf{M}x = z^{(2)}\}}(x) = x + \mathbf{M}^\top (z^{(2)} - \mathbf{M}x). \quad (17)$$

- The gradient of the sum of the remaining terms, that is $f(x) = \|\mathbf{D}\mathbf{B}x - z^{(1)}\|^2 + \sum_{s=1}^4 \alpha_s \mathcal{W}_2^2(\nu_{L_s x}, \nu_{\hat{z}_s^{(2)}})$, reads

$$\begin{aligned} \nabla f &= 2\mathbf{B}^\top \mathbf{D}^\top (\mathbf{D}\mathbf{B}x - z^{(1)}) \\ &+ 2 \sum_{s=1}^4 \alpha_s L_s^\top (L_s x - \hat{z}_s^{(2)} \circ \sigma_{\hat{z}_s^{(2)}} \circ \sigma_{L_s x}^{-1}), \end{aligned} \quad (18)$$

where ∇f is β -Lipschitz with $\beta = 2(1 + \sum_1^4 \alpha_s \|L_s\|^2)$.

Although there is no theoretical guarantee about the estimate produced by Algorithm 1, in our experiments we observed that it always converges to a stable solution.

Algorithm 1 FBPD [54]

INITIALIZATION

$$\begin{cases} \text{Choose } (x^{[0]}, y^{[0]}) \in \mathbb{R}^n \times \mathbb{R}^{K_n} \\ \text{set } \tau > 0 \text{ and } \omega > 0 \text{ such that} \\ \tau (\beta/2 + \omega \lambda \|\mathbf{T}\|^2) < 1 \end{cases}$$

FOR $l = 0, 1, \dots$

$$\begin{cases} \hat{x}^{[l]} = \nabla f(x^{[l]}) + \mathbf{T}^\top y^{[l]} \\ x^{[l+1]} = P_{\{\mathbf{M}x = z^{(2)}\}}(x^{[l]} - \tau \hat{x}^{[l]}) \\ \hat{y}^{[l]} = \mathbf{T}(2x^{[l+1]} - x^{[l]}) \\ y^{[l+1]} = \text{prox}_{\omega \lambda \|\cdot\|_{1,2}^*}(y^{[l]} + \omega \hat{y}^{[l]}) \end{cases}$$

IV. EXPERIMENTAL RESULTS

The numerical analysis is decomposed in three parts. Firstly, we assess the performance achieved by only using the nonlocal regularization proposed in Section II-B. Secondly, we extend our analysis to the histogram-based distance proposed in Section II-A, by evaluating its impact on the performance. Finally, we compare the full approach with several state-of-the-art methods.

As shown in Fig. 4, all the test images are grayscale textures, such as Nature (Grass, Leaves, Leaf, Flowers), Forest (Bark, Wall), Hair, Biology (Cell), and Petrology (CT scan).³

³The authors would like to thank ‘‘Sismage’’ team from TOTAL Group for providing us the CT data.

Fig. 4. Test images. From the left to the right, top line: Grass, Bark 1, Bark 2, Wall, Hair and Leaves. From the left to the right, bottom line: Leaf, Flowers, Cell 1, Cell 2, Cell 3, and CT image scan of a petroleum reservoir.

Fig. 5. **PSNR – SSIM – Err_{x̄}**. From the left to the right: noisy image (noise standard deviation $\tau = 50$), high-resolution patch (50% of the total image), true image and the results of: the proposed nonlocal regularization, the nonlocal regularization of [46], and the classical NLTV [28].

A. Error measures

The simplest and the most popular quality measure is the Mean Squared Error (MSE) along with the related quantity of Peak Signal-to-Noise Ratio (PSNR). However, these indexes may not be good evaluation metrics for the considered problem, since we are dealing with texture reconstruction, and not with image reconstruction. In this paper, we consider the visual inspection to assess the image quality in terms of the down-sampling factor and the high-resolution patch size. In order to provide an objective measure, we will evaluate the results with the SSIM index (the closer to 1, the better the quality) [63]

and we propose to add a new statistical error measure $\text{Err}_{\bar{x}}$, defined as

$$\text{Err}_{\bar{x}} = \sum_{s=1}^4 \frac{\alpha_s}{N_s} \mathcal{W}_2^2(\nu_{L_s \check{x}}, \nu_{L_s \bar{x}}), \quad (19)$$

where $\check{x}$ denotes an estimate of $\bar{x}$. This error aims at showing how much the resulting image is statistically close to the true high-resolution image. Indeed, the closer $\text{Err}_{\bar{x}}$ to 0, the better the results.

To allow a reliable assessment, we will add the PSNR indices to the numerical evaluation. The aim of this step is twofold:

to be homogeneous to the state-of-the-art methods evaluation and to prove that the PSNR might not reflect the visual quality of reconstructed images.

B. Nonlocal regularization

In this section, we test in isolation the proposed regularization. To do so, we consider the texture denoising problem arising from (2) by setting $D = I$, $B = I$, and $(\alpha_s)_{\{1 \leq s \leq 4\}} = 0$, resulting in the following optimization problem

$$\underset{x \in \mathbb{R}^N}{\text{minimize}} \quad \|x - z^{(1)}\|^2 + \lambda R(x, z^{(2)}) \quad \text{s. t.} \quad Mx = z^{(2)}.$$

The 8-bit grayscale images are of size 256×256 , and $z^{(2)}$ is the half of the original image $\bar{x}$. The observation $z^{(1)}$ is generated by degrading $\bar{x}$ with an additive Gaussian noise whose standard deviation is equal to 50. The optimization is performed using Algorithm 1. The parameter λ is hand-tuned so as to obtain the best visual results. For the nonlocal graph detailed in Section II-B, K has been fixed to 14, Q to 25 and the blocks to 5×5 pixels.

Fig. 5 gives an example of natural textures, whose elements hold a rough, irregular and nonsmooth patterns. These delicate textures are difficult to recover when corrupted by noise, due to the concealing effect on the small details. Having a noiseless patch, we aim at recovering these details using a nonlocal regularization, whose graph is constructed as proposed in Section II-B. The visual inspection of the results shows the interest of considering the proposed nonlocal regularization $R(\cdot, z^{(2)})$ for texture denoising. Indeed, $R(\cdot, z^{(2)})$ proves to be more effective than NL-HR [46], which in turn outperforms the classical NLTV [28] that tends to average the pixel values. The better performance of $R(\cdot, z^{(2)})$, w.r.t. NL-HR, seems to be related to its ability to better choose connections to the noiseless patch, which in turn preserve the edges and the thin structures.

C. Discussion of the proposed approach

In this section, we compare the approach proposed in (2) w.r.t. the following methods

- when $\lambda = 0$ (without the nonlocal regularization)

$$\underset{x \in \mathbb{R}^N}{\text{minimize}} \quad \|DBx - z^{(1)}\|^2 + H(x, z^{(2)}) \quad \text{s. t.} \quad Mx = z^{(2)}, \quad (20)$$

- when, for every $s \in \{1, \dots, 4\}$, $\alpha_s = 0$ (without the statistical prior)

$$\underset{x \in \mathbb{R}^N}{\text{minimize}} \quad \|DBx - z^{(1)}\|^2 + \lambda R(x, z^{(2)}) \quad \text{s. t.} \quad Mx = z^{(2)}. \quad (21)$$

The parameters λ and $(\alpha_s)_{1 \leq s \leq 4}$ are set with the following strategy:

- The data fidelity term is normalized by the noise standard deviation τ ;
- λ is set to 100;
- For every $s \in \{1, \dots, 4\}$, $\alpha_s = \frac{\tau_s}{\tau_4}$, where τ_s is the variance of $L_s z^{(2)}$.

In these experiments, $\bar{x}$ corresponds to a 256×256 8-bit grayscale image. The observed image $z^{(1)}$ is generated by degrading $\bar{x}$ with a convolution operator B , which is equal to a truncated Gaussian function with standard deviation 1.2 and kernel size 3×3 , the blurry result is then downsampled by a factor $r = 2$ in each direction, and finally we add to this latter a white Gaussian noise of standard deviation $\tau = 10$. The observed high-resolution patch $z^{(2)}$ amounts to 25% of the total image $\bar{x}$. The quality of the results is evaluated in term of PSNR, SSIM and $\text{Err}_{\bar{x}}$.

Fig. 6 gives two examples of texture images: a leaf and a biological cell. One can observe that the visual results achieved with Problem (2) are better than the ones obtained with the standard formulations (20) and (21). Indeed, without the nonlocal regularization, we do not have a spatial prior. As the problem is nonconvex, we may get stuck in a local minimum that is near to the patch $z^{(2)}$, but far from the true image $\bar{x}$. However, without the statistical prior, even if the nonlocal regularization performs quite well to recover the smooth areas and to reconstruct the edges, it fails to infer the correct image structures. Using the proposed formulation (2), we gain the best visual results and the best $\text{SSIM} - \text{Err}_{\bar{x}}$, which indicates that the nonlocal regularization pushes the local minima of (20) closer to $\bar{x}$. While the best PSNR is given by the results of Eq. (21), this is not reflected on the visual inspection. Such a behaviour can be explained by the fact that the SSIM and $\text{Err}_{\bar{x}}$ measures take in consideration some features, while the PSNR fails to illustrate them (see Fig. 6).

In Table I, we evaluate the experiments on several images: Wall, Flowers, Leaves, Leaf, Hair, Cell 3 and Bark 2 (see Fig. 4). We report the PSNR, SSIM and the error measure to the true image. We present the numerical results obtained with (20) (without the nonlocal prior) and (2) (with the nonlocal prior). As one can see the latter approach achieves the best performance in term of PSNR and SSIM, as well as the smallest error in term of $\mathcal{W}_2^2(\nu_{L_s \bar{x}}, \nu_{L_s \bar{x}})$ for almost every $s \in \{1, \dots, 4\}$, and the better mean weighted value $\text{Err}_{\bar{x}}$ regardless of the degradation conditions. The best values for each case is put in boldface for the sake of clarity.

D. Comparison to the state-of-the-art methods

We now compare our results with the following approaches:

- Dong *et al.* [64]: the method is based on the assumption that a dataset of high-resolution textures is available, from which it is learned a dictionary of low-resolution patches, a dictionary of high-resolution patches, and a correspondence map between the two dictionaries.
- Hertzmann *et al.* [34]: given an image A and its transformation A' , the method provides a new image B' from an image B , such that the relationship between B and B' is the same as the one between A and A' .
- Portilla *et al.* [65]: this is a texture synthesis method based on a variational approach and the wavelet transform.

Fig. 7 and 8 present a comparison of the proposed solution with the state-of-the-art results for the following data:

- $\bar{x}$ corresponds to a 256×256 8-bit grayscale image.

Fig. 6. **PSNR – SSIM – Err_x**. From the left to the right: low-resolution image (down-sampling factor $r = 2$, noise standard deviation $\tau = 10$, blur 3×3), high-resolution patch (25% of the total image), true image, the proposed approach (with the statistical prior and the nonlocal regularization), the proposed approach without the nonlocal regularization, and the proposed approach without the statistical prior.

TABLE I

ERROR MEASURES ON THE PROPOSED APPROACH WITH AND WITHOUT THE NONLOCAL PRIOR $R(\cdot, z^{(2)})$ (DOWN-SAMPLING FACTOR $r = 2$, NOISE STANDARD DEVIATION $\tau = 10$, BLUR 3×3). THE IMAGES ARE EVALUATED USING THE PSNR, SSIM, AND THE WASSERSTEIN DISTANCE BETWEEN INTENSITY VALUE HISTOGRAMS ($s = 1$), HORIZONTAL-VERTICAL GRADIENT HISTOGRAMS ($s = 2$), DIAGONAL GRADIENT HISTOGRAMS ($s = 3$), OR LAPLACIAN HISTOGRAMS ($s = 4$), AND THE ERROR TO THE TRUE IMAGE $\text{Err}_{\bar{x}}$ (19).

Images	Regularization	PSNR	SSIM	$\mathcal{W}_2^2(\nu_{L_s \bar{x}}, \nu_{L_s \bar{x}})$				$\text{Err}_{\bar{x}}$
				$s = 1$	$s = 2$	$s = 3$	$s = 4$	
Wall	without $R(\cdot, z^{(2)})$	28.02	0.863	10.26	2.47	8.30	5.27	8.20
	with $R(\cdot, z^{(2)})$	28.84	0.876	5.68	0.99	3.17	5.19	4.56
Flowers	without $R(\cdot, z^{(2)})$	26.83	0.832	14.49	0.45	2.71	0.20	3.45
	with $R(\cdot, z^{(2)})$	27.53	0.835	12.43	0.48	2.24	0.25	2.97
Leaves	without $R(\cdot, z^{(2)})$	23.57	0.892	45.49	13.66	26.74	63.21	20.96
	with $R(\cdot, z^{(2)})$	24.06	0.895	21.03	12.56	22.77	42.19	12.15
Leaf	without $R(\cdot, z^{(2)})$	24.02	0.839	37.92	0.79	3.73	3.03	4.91
	with $R(\cdot, z^{(2)})$	25.29	0.845	8.64	0.48	2.35	5.36	1.70
Hair	without $R(\cdot, z^{(2)})$	19.43	0.679	45.12	10.87	26.51	39.84	5.49
	with $R(\cdot, z^{(2)})$	19.61	0.682	40.03	10.74	26.10	39.45	5.47
Cell 3	without $R(\cdot, z^{(2)})$	21.99	0.763	25.08	4.25	6.70	3.68	2.43
	with $R(\cdot, z^{(2)})$	22.36	0.769	3.81	2.92	8.04	2.17	0.82
Bark 2	without $R(\cdot, z^{(2)})$	24.77	0.858	47.99	1.32	5.09	0.84	20.14
	with $R(\cdot, z^{(2)})$	25.79	0.859	21.94	0.56	2.01	2.34	9.38

- The convolution operator is equal to a truncated Gaussian function with standard deviation 1.2 and kernel size 3×3 .
- The down-sampling factor $r = 4$ in Fig. 7 ($r = 2$ in Fig. 8).
- Noise standard deviation $\tau = 10$.
- $z^{(2)}$ matches 25% of $\bar{x}$ in Fig. 7 (6.25% in Fig. 8)

Fig. 7 gives two texture examples where the down-sampling factor is equal to 4, and the high-resolution patch amounts to 25% of the total image. As illustrated by these experiments, the proposed approach leads to a better texture reconstruction in the synthesized images, while the state-of-the-art super-resolution method [64] tends to smooth the textures, and the state-of-the-art image analogies technique [34] does not preserve textures and produces a final image that is highly noisy. With regard to texture synthesis approach [65], it produces a high-resolution image (by replicating the patch) without exploiting information provided by the low-resolution image. Hence, the results may be far from the ground truth. The improvement of the proposed method is confirmed both visually and numerically by the SSIM and $\text{Err}_{\bar{x}}$. However, the PSNR fails to assess the image quality.

Fig. 8 gives a similar example where the down-sampling factor is equal to 2 and the high-resolution patch amounts to 6.25% of the total image. The proposed approach gives the best visual results compared to the state-of-the-art methods. This is also confirmed by looking at the reconstructed textures and the SSIM – $\text{Err}_{\bar{x}}$ indexes. Again, the PSNR fails to assess the image quality.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we have proposed a variational approach to solve the problem of texture synthesis guided by a low-resolution image and a high-resolution patch. To do so, we have exploited two different sources of information: a nonlocal smoothness prior and a distance between histograms. The resulting model is well adapted to textures with sharp edges and small patterns, as shown in Fig 8. We have investigated a primal-dual proximal algorithm, which allows us to consider various convex and nonconvex functions. Interesting perspectives include a better modeling of the textures, possibly through constraints on the power spectrum of images as in [43], [47], as well as the use of a multi-scale strategy for a down-sampling factor greater than 4. Another perspective is to explore the convex relaxation of the histogram prior in order to have a convex minimization problem [66]. Finally, it could be interesting to explore the scenario in which the high-resolution patch is replaced by an intermediate image obtained with some existing state-of-art texture synthesis algorithm.

REFERENCES

- [1] D. Wildenschild and A. Sheppard, “X-ray imaging and analysis techniques for quantifying pore-scale structure and processes in subsurface porous medium systems,” *Advances in Water Resources*, vol. 51, pp. 217–246, 2013.
- [2] X. Xie, “A review of recent advances in surface defect detection using texture analysis techniques,” *ELCVIA Electronic Letters on Computer Vision and Image Analysis*, vol. 7, no. 3, 2008.
- [3] R. Keys, “Cubic convolution interpolation for digital image processing,” *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 29, no. 6, pp. 1153–1160, Dec. 1981.
- [4] R. Fattal, “Image upsampling via imposed edge statistics,” *ACM Trans. on Graphics*, vol. 26, no. 3, pp. 1–8, Jul. 1007.
- [5] J. Yang, Z. Wang, Z. Lin, S. Cohen, and T. Huang, “Coupled dictionary training for image super-resolution,” *IEEE Trans. Image Process.*, vol. 21, no. 8, pp. 3467–3478, Aug. 2012.
- [6] H. He and W.-C. Siu, “Single image super-resolution using gaussian process regression,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, Providence, USA, Jun. 2011, pp. 449–456.
- [7] K. Zhang, X. Gao, D. Tao, and X. Li, “Single image super-resolution with non-local means and steering kernel regression,” *IEEE Trans. Image Process.*, vol. 21, no. 11, pp. 4544–4556, Nov. 2012.
- [8] H. Chang, D.-Y. Yeung, and Y. Xiong, “Super-resolution through neighbor embedding,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, Washington, USA, Jun. 2004, pp. I–I.
- [9] D. Glasner, S. Bagon, and M. Irani, “Super-resolution from a single image,” in *Proc. IEEE Int. Conf. Comput. Vis.*, Kyoto, Japan, Sep. 2009, pp. 349–356.
- [10] C. Dong, C. Loy, K. He, and X. Tang, *Learning a Deep Convolutional Network for Image Super-Resolution*. Cham: Springer International Publishing, 2014, pp. 184–199.
- [11] S. Schuler, C. Leistner, and H. Bischof, “Fast and accurate image upscaling with super-resolution forests,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, Jun. 2015, pp. 3791–3799.
- [12] D. Dai, R. Timofte, and L. Van Gool, “Jointly optimized regressors for image super-resolution,” in *Eurographics*, 2015.
- [13] L. Rudin, S. Osher, and E. Fatemi, “Nonlinear total variation based noise removal algorithms,” *Physica D*, vol. 60, no. 1–4, pp. 259–268, Nov. 1992.
- [14] P. L. Combettes and J.-C. Pesquet, “Image restoration subject to a total variation constraint,” *IEEE Trans. Image Process.*, vol. 13, no. 9, pp. 1213–1222, Sep. 2004.
- [15] S. Farsiu, M. Robinson, M. Elad, and P. Milanfar, “Fast and robust multiframe super resolution,” *IEEE Trans. Image Process.*, vol. 13, no. 10, pp. 1327–1344, Oct. 2004.
- [16] D. Mitzel, T. Pock, T. Schoenemann, and D. Cremers, “Video super resolution using duality based tv-l1 optical flow,” in *Pattern Recognition*, ser. Lecture Notes in Computer Science, J. Denzler, G. Notni, and H. Se, Eds. Springer Berlin Heidelberg, 2009, vol. 5748, pp. 432–441.
- [17] K. Bredies, K. Kunisch, and T. Pock, “Total generalized variation,” *SIAM J. Imaging Sciences*, vol. 3, no. 3, pp. 492–526, Sep. 2010.
- [18] S. Ono and I. Yamada, “Second-order Total Generalized Variation constraint,” in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, Florence, Italy, May 2014, pp. 4938–4942.
- [19] Y. Hu and M. Jacob, “Higher degree total variation (HDTV) regularization for image recovery,” *IEEE Trans. Image Process.*, vol. 21, no. 5, pp. 2559–2571, May 2012.
- [20] —, “Generalized higher degree total variation (HDTV) regularization,” *IEEE Trans. Image Process.*, vol. 23, no. 6, pp. 2423–2435, Jun. 2014.
- [21] G. Peyré, S. Bougleux, and L. Cohen, “Non-local regularization of inverse problems,” in *Computer Vision ECCV*, ser. Lecture Notes in Computer Science, D. Forsyth, P. Torr, and A. Zisserman, Eds. Springer Berlin Heidelberg, 2008, vol. 5304, pp. 57–68.
- [22] G. Gilboa and S. Osher, “Nonlocal operators with applications to image processing,” *Multiscale Modeling and Simulation*, vol. 7, no. 3, pp. 1005–1028, Nov. 2009.
- [23] M. Protter, M. Elad, H. Takeda, and P. Milanfar, “Generalizing the nonlocal-means to super-resolution reconstruction,” *IEEE Trans. Image Process.*, vol. 18, no. 1, pp. 36–51, Jan. 2009.
- [24] G. Freedman and R. Fattal, “Image and video upscaling from local self-examples,” *ACM Trans. Graph.*, vol. 28, no. 3, pp. 1–10, 2010.
- [25] M. Maggioni, V. Katkovnik, K. Egiazarian, and A. Foi, “Nonlocal transform-domain filter for volumetric data denoising and reconstruction,” *IEEE Trans. Image Process.*, vol. 22, no. 1, pp. 119–133, Jan. 2013.
- [26] C. Couprie, L. Grady, L. Najman, J.-C. Pesquet, and H. Talbot, “Dual constrained TV-based regularization on graphs,” *SIAM J. Imaging Sciences*, vol. 6, no. 3, pp. 1246–1273, Oct. 2013.
- [27] G. Chierchia, N. Pustelnik, B. Pesquet-Popescu, and J.-C. Pesquet, “A nonlocal structure tensor based approach for multicomponent image recovery problems,” *IEEE Trans. Image Process.*, vol. 23, no. 12, pp. 5531–5544, Dec. 2014.
- [28] G. Chierchia, N. Pustelnik, J.-C. Pesquet, and B. Pesquet-Popescu, “Epigraphical projection and proximal tools for solving constrained convex optimization problems,” *Signal, Image and Video Processing*, vol. 9, no. 8, pp. 1737–1749, Nov. 2015.
- [29] D. Dai, H. Riemenschneider, and L. V. Gool, “The synthesizability of texture examples,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, Jun. 2014, pp. 3027–3034.

Fig. 7. **PSNR – SSIM – Err_x**. From the left to the right: low-resolution image (down-sampling factor $r = 4$, noise standard deviation $\tau = 10$, blur 3×3), high-resolution patch (25% of the total image), true image, proposed approach, Dong *et al.* [64], Hertzmann *et al.* [34] and Portilla *et al.* [65].

- [30] G. Cross and A. Jain, “Markov random field texture models,” *IEEE Trans. Pattern Anal. Mach. Int.*, no. 1, pp. 25–39, Jan. 1983.
- [31] A. Efros and T. Leung, “Texture synthesis by non-parametric sampling,” in *Proc. IEEE Int. Conf. Comput. Vis.*, vol. 2, Kerkira, Greece, Sep. 1999, pp. 1033–1038.
- [32] L. Wei and M. Levoy, “Fast texture synthesis using tree-structured vector quantization,” in *Conf. on Comput. Graphics and Interactive Tech.*, New Orleans, USA, Jul. 2000, pp. 479–488.
- [33] S. Lefebvre and S. Hoppe, “Parallel controllable texture synthesis,” *Trans. on graphics*, vol. 24, no. 3, pp. 777–786, 2005.
- [34] A. Hertzmann, C. E. Jacobs, N. Oliver, B. Curless, and D. H. Salesin, “Image analogies,” in *Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH ’01. New York, NY, USA: ACM, 2001, pp. 327–340.
- [35] S. Darabi, E. Shechtman, C. Barnes, D. Goldman, and P. Sen, “Image Melding: Combining inconsistent images using patch-based synthesis,” *Conference on Computer Graphics and Interactive Techniques*, vol. 31, no. 4, pp. 82:1–82:10, 2012.
- [36] C. Han, E. Risser, R. Ramamoorthi, and E. Grinspun, “Multiscale texture synthesis,” *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2008)*, vol. 27, no. 3, pp. 51:1–51:8, 2008.
- [37] M. Eisemann and M. Magnor, “ZIPMAPS: zoom-into-parts texture maps,” in *Proc. Vision, Modeling and Visualization (VMV) 2010*, Siegen, Germany, Nov. 2010, pp. 291–297.
- [38] D. Heeger and J. Bergen, “Pyramid-based texture analysis/synthesis,” in *Conf. on Comput. Graphics and Interactive Tech.*, Washington, USA, 1995, pp. 229–238.
- [39] A. Kaspar, B. Neubert, D. Lischinski, M. Pauly, and J. Kopf, “Self tuning texture optimization,” *Comput. Graph. Forum*, vol. 34, no. 2, pp. 349–359, 2015.
- [40] J. Portilla and E. Simoncelli, “Image denoising via adjustment of wavelet coefficient magnitude correlation,” in *Proc. IEEE Int. Conf. Image Process.*, vol. 3, Vancouver, Canada, Sep. 2000, pp. 277–280.
- [41] B. Galerne, Y. Gousseau, and J.-M. Morel, “Random phase textures: Theory and synthesis,” *Image Process On Line*, vol. 1, 2011.
- [42] G. Tartavel, Y. Gousseau, and G. Peyré, “Constrained sparse texture synthesis,” in *Scale Space and Variational Methods in Computer Vision*, ser. Lecture Notes in Computer Science, A. Kuijper, K. Bredies, T. Pock, and H. Bischof, Eds. Springer Berlin Heidelberg, 2013, vol. 7893, pp. 186–197.
- [43] G. Tartavel, Y. Gousseau, and G. Peyré, “Variational texture synthesis with sparsity and spectrum constraints,” *J. Math. Imaging Vision*, vol. 52, no. 1, pp. 124–144, May 2015.
- [44] Y. Rubner, C. Tomasi, and L. J. Guibas, “The earth mover’s distance as a metric for image retrieval,” *Int. J. Comput. Vision*, vol. 40, no. 2, pp. 99–121, Nov. 2000.
- [45] J. Rabin and G. Peyré, “Wasserstein regularization of imaging problem,” in *Proc. IEEE Int. Conf. Image Process.*, Brussels, Belgium, Sep. 2011, pp. 1541–1544.
- [46] M. Hidane, M. El Gheche, J.-F. Aujol, Y. Berthoumieu, and C. Deledalle, “Image zoom completion,” *IEEE Trans. Image Process.*, 2015, HAL-01253124.
- [47] R. Fablet and F. Rousseau, “Missing data super-resolution using non-local and statistical priors,” in *Proc. IEEE Int. Conf. Image Process.*, Quebec, Canada, Sep. 2015, pp. 676–680.
- [48] G. Tartavel, “Modèles variationnels pour les textures : applications à la synthèse et à la restauration,” Ph.D. dissertation, Télécom ParisTech, 2015.
- [49] C. Villani, *Topics in optimal transportation*, A. M. Society, Ed. Graduate studies in mathematics, 2003.
- [50] A. Chambolle and T. Pock, “A first-order primal-dual algorithm for convex problems with applications to imaging,” *J. Math. Imaging Vision*, vol. 40, no. 1, pp. 120–145, May 2011.
- [51] L. M. Briceño-Arias and P. L. Combettes, “A monotone + skew splitting model for composite monotone inclusions in duality,” *SIAM J. Optim.*, vol. 21, no. 4, pp. 1230–1250, Oct. 2011.
- [52] P. L. Combettes and J.-C. Pesquet, “Primal-dual splitting algorithm for solving inclusions with mixtures of composite, lipschitzian, and parallel-sum type monotone operators,” *Set-Valued Var. Anal.*, vol. 20, p. 2, Jun. 2012, 22 pp.
- [53] B. C. Vũ, “A splitting algorithm for dual monotone inclusions involving cocoercive operators,” *Adv. Comput. Math*, vol. 38, no. 3, pp. 667–681, Apr. 2013.
- [54] L. Condat, “A primal-dual splitting method for convex optimization involving lipschitzian, proximable and linear composite terms,” *Journal of Optimization Theory and Applications*, vol. 158, no. 2, pp. 460–479, Aug. 2013.
- [55] N. Komodakis and J.-C. Pesquet, “Playing with duality: An overview of recent primal-dual approaches for solving large-scale optimization problems,” *IEEE Signal Processing Magazine*, vol. 6, no. 32, pp. 31–54, Nov. 2015.
- [56] M. Storath, A. Weinmann, and L. Demaret, “Jump-sparse and sparse

Fig. 8. **PSNR – SSIM – Err_x**. From the left to the right: low-resolution image (down-sampling factor $r = 2$, noise standard deviation $\tau = 10$, blur 3×3), high-resolution patch (6.25% of the total image), true image, proposed approach, Dong *et al.* [64], Hertzmann *et al.* [34] and Portilla *et al.* [65].

recovery using Potts functionals,” *IEEE Trans. on Signal Process.*, vol. 62, no. 14, pp. 3654–3666, Jul. 2014.

- [57] E. Strelakovsky and D. Cremers, *Computer Vision – ECCV 2014: 13th European Conference, Zürich, Switzerland, Sep. 6-12, 2014, Proceedings, Part II*. Cham: Springer International Publishing, 2014, ch. Real-Time Minimization of the Piecewise Smooth Mumford-Shah Functional, pp. 127–141.
- [58] T. Möllenhoff, E. Strelakovsky, M. Moeller, and D. Cremers, *Energy Minimization Methods in Computer Vision and Pattern Recognition: 10th International Conference, EMMCVPR 2015, Hong Kong, China, Jan. 13-16, 2015. Proceedings*. Cham: Springer International Publishing, 2015, ch. Low Rank Priors for Color Image Regularization, pp. 126–140.
- [59] T. Möllenhoff, E. Strelakovsky, M. Möller, and D. Cremers, “The primal-dual hybrid gradient method for semiconvex splittings,” *SIAM J. Imaging Sciences*, vol. 8, no. 2, pp. 827–857, 2015.
- [60] J. J. Moreau, “Fonctions convexes duales et points proximaux dans un espace hilbertien,” *C. R. Acad. Sci.*, vol. 255, pp. 2897–2899, 1962.
- [61] P. L. Combettes and J.-C. Pesquet, “Proximal splitting methods in signal processing,” in *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, H. H. Bauschke, R. Burachik, P. L. Combettes, V. Elser, D. R. Luke, and H. Wolkowicz, Eds. New York, USA: Springer-Verlag, 2011, pp. 185–212.
- [62] P. L. Combettes and V. R. Wajs, “Signal recovery by proximal forward-backward splitting,” *Multiscale Modeling and Simulation*, vol. 4, no. 4, pp. 1168–1200, Nov. 2005.
- [63] W. Malpica and C. A. Bovik, *Encyclopedia of Multimedia*. Boston: Springer US, 2008, ch. Range Image Quality Assessment by Structural Similarity, pp. 755–757.
- [64] W. Dong, L. Zhang, G. Shi, and X. Wu, “Image deblurring and super-resolution by adaptive sparse domain selection and adaptive regularization,” *IEEE Trans. Image Process.*, vol. 20, no. 7, pp. 1838–1857, Jul. 2011.
- [65] J. Portilla and E. P. Simoncelli, “A parametric texture model based on joint statistics of complex wavelet coefficients,” *Int. J. Comput. Vision*, vol. 40, no. 1, pp. 49–71, Oct. 2000.
- [66] P. Swoboda and C. Schnörr, “Convex variational image restoration with histogram priors,” *SIAM J. Imaging Sciences*, vol. 6, no. 3, pp. 1719–1735, 2013.

Mireille El Gheche received the engineering degree in Telecommunications from the Lebanese University (Tripoli, Lebanon) and the Master degree in Radio-communication from SupElec (Paris, France) in 2010. She received the Ph.D. degree in signal and image processing from the Université Paris Est (Marne-la-Vallée, France) in 2014. From 2015, she is a Postdoctoral Researcher at the IMS and IMB laboratories, Université de Bordeaux (Talence, France), where she is working on the topic of super-resolution of texture images and texture volume for Computed

Tomography applications. Her research interests are in convex and non-convex optimization applied to inverse problems, with a focus on image restoration, stereo matching, 3D reconstruction, and super-resolution.

Jean-François Aujol studied Mathematics at the Ecole Normale Supérieure (ENS) in Cachan. He received the Ph.D. degree in Mathematics from Nice-Sophia-Antipolis University in June 2004 and the Habilitation thesis in Applied Mathematics from ENS Cachan in May 2009. In 2004-2005, he was successively an assistant researcher in UCLA (Applied Mathematics Department), and then a research engineer at Telecom ParisTech. In 2005-2011, Jean-François Aujol was a scientist researcher with CNRS at ENS Cachan with CMLA (until 2009), and then in Marseille with LATP. Since March 2011, he has been a full professor of applied mathematics in Bordeaux with IMB. Since September 2013, he has been a member of IUF (Institut Universitaire de France). He is interested in mathematical problems related to image processing. Jean-François Aujol serves as an associate editor for IEEE Transactions on Image Processing, Computer Vision and Image Understanding, and Journal of Mathematical Imaging and Vision. He is an IEEE member.

Yannick Berthoumieu Yannick Berthoumieu received his Ph.D. degree in Electrical Engineering in 1996 and his Habilitation à Diriger des Recherches in signal and image processing in 2006 from the University of Bordeaux (UB). He joined the Bordeaux Institute of Technology in 1997 as an assistant professor. Since 2007 he is a full Professor of the Bordeaux Institute of Technology with the Department of Telecommunications. His research activities covers areas such as statistical image and video processing, information geometry, as well as probabilistic manifolds in machine learning, classification, segmentation, superresolution and inversion algorithms focused on textural content for optical, radar and hyperspectral imaging applications. Since 2003, he is with the Signal and Image processing group of IMS Lab, University of Bordeaux.

Charles-Alban Deledalle received the Engineer degree from EPITA, France, in 2008, the Master of Science and Technology degree from Université Paris VI, France, in 2008 and the Ph.D. degree in signal and image processing from Telecom ParisTech, France, in 2011. He made a postdoctoral fellowship in applied mathematics at Université Paris IX, France, in 2012. He is currently CNRS Researcher at Institut de Mathématiques de Bordeaux (IMB), Université de Bordeaux, France. His research interests include image denoising, reconstruction, variational and statistical modeling, and parameter estimation. He received the IEEE ICIP Best Student Paper Award in 2010, the ISIS/EEA/GRETSI Best PhD Award in 2012, and the IEEE GRSS Transactions Prize Paper Award in 2016.