LOCALLY PARALLEL TEXTURE MODELING*

PIERRE MAUREL*, JEAN-FRANÇOIS AUJOL†, AND GABRIEL PEYRÉ‡

Abstract. This article presents a new adaptive framework for locally parallel texture modeling. Oscillating patterns are modeled with functionals that constrain the local Fourier decomposition of the texture. We first introduce a texture functional which is a weighted Hilbert norm. The weights on the local Fourier atoms are optimized to match the local orientation and frequency of the texture. This adaptive model is used to solve image processing inverse problems, such as image decomposition and inpainting. The local orientation and frequency of the texture component are adaptively estimated during the minimization process. To improve inpainting performances over large missing regions, we introduce a highly non-convex generalization of our texture model. This new model constrains the amplitude of the texture and it allows one to impose an arbitrary oscillation profile. Numerical results illustrate the effectiveness of the method.

Key words. Locally parallel textures, image decomposition, inpainting, total variation, curvelets, cartoon image.

AMS subject classifications. 68U10, 65K10, 65F22

1. Introduction. Texture modeling is fundamental for a large number of image processing problems, such as image segmentation, object recognition and image restoration. Image restoration methods take advantage of a texture model which imposes constraints on the geometry of textures present in the image. This paper proposes a framework for modeling locally parallel oscillating patterns based on local Fourier decompositions. Figure 1.1 shows examples of typical locally parallel textures.

Our framework parametrizes the geometry of the texture using a frequency field $\xi(x)$ which gives the orientation and frequency of the texture around a point x . Two different models based on this frequency field are introduced. A first one is used for image separation and for inpainting small holes. To improve the performance of inpainting large missing regions, we extend our framework using a highly non-convex texture model.

FIG. 1.1. *Examples of locally parallel natural textures.*

*This work has been done with the support of the French “Agence Nationale de la Recherche” (ANR), under grant NatImages (ANR-08-EMER-009), “Adaptivity for natural images and texture representations”, while the two first authors of the paper were with CMLA, ENS Cachan, CNRS, UniverSud, 61 avenue du Président Wilson, 94235 Cachan Cedex, France.

*VisAGeS, Université de Rennes 1, IRISA, UMR CNRS-6074 Campus de Beaulieu, F-35042 Rennes, France (pierre.maurel@irisa.fr)

†LATP, CMI, Université de Provence, 39 rue F. Joliot-Curie, 13453 Marseille cedex 13, France (aujol@cmi.univ-mrs.fr).

‡Ceremade, Université Paris-Dauphine, Place du Maréchal De Lattre De Tassigny, 75775 Paris Cedex 16, France (gabriel.peyre@ceremade.dauphine.fr).

1.1. Previous Works. Decomposing an image into meaningful components is an important and challenging inverse problem in image processing. A variational decomposition algorithm seeks a decomposition of an image f into various components representing different information in the image. In this paper we focus on the cartoon-texture decomposition problem and we seek for a decomposition of $f = u + v$ where u should capture the sketch of the image and v the texture content. The definition of texture is vague and depends on the local image scale. As a matter of fact, a structure at one scale can be regarded as a texture at another scale. This article is focused on locally parallel textures, that correspond to oscillating patterns with a local geometric orientation.

Inverse problem regularization is an active area of research in image processing. It aims at restoring a high resolution image f_0 from possibly incomplete observations $f = \Phi f_0 + w$ where Φ is a non-invertible operator and w is an additive noise. Some of these problems, such as image denoising, image deconvolution or image inpainting, can be solved efficiently by seeking the solution as the sum of two components, $f_0 \approx u + v$, where u and v capture two different meaningful components, e.g. the structure and the texture.

The decomposition and regularization are done simultaneously by solving the following minimization problem

$$(u, v) = \underset{\tilde{u}, \tilde{v}}{\operatorname{argmin}} \frac{1}{2} \|f - \Phi(\tilde{u} + \tilde{v})\|_{L^2}^2 + \lambda J(\tilde{u}) + \mu T(\tilde{v}) \quad (1.1)$$

where the functionals J and T model respectively the cartoon and the texture contents, λ and μ are two positive real parameters balancing the importance of each term, and $\|f - \Phi(\tilde{u} + \tilde{v})\|_{L^2}$ is a fidelity term which takes into account the presence of noise in the image. The solution of the inverse problem is then given by $u + v$.

When $\Phi = \operatorname{Id}$ is the identity operator, the solution of (1.1) provides a decomposition of f between two components u and v such that $J(u)$ and $T(v)$ are small. If the functionals J and T do not capture the noise, then $u + v$ is a denoised version of f . When Φ is a blurring operator, then problem (1.1) is a deconvolution problem. Image decomposition into a cartoon part and a texture part has already been proposed in [40, 24, 25, 33] for image deconvolution.

Cartoon Model. In the early 90's and for denoising purpose, Rudin, Osher and Fatemi [45] proposed to use the total variation of an image to model its cartoon part. The TV norm, $J(u) = \|u\|_{\operatorname{TV}}$, is $\int |\nabla u|$ for a continuously differentiable function u and is extended to discontinuous functions. It allows one to recover piecewise smooth functions without smoothing sharp discontinuities. This is related to the wavelet thresholding algorithm of Donoho and Johnstone [26] where the wavelet decomposition of a cartoon image is assumed to be sparse. The cartoon functional $J(u)$ is therefore the ℓ^1 norm of the coefficients of this decomposition. To exploit the geometric image regularity along edge curves, Candès and Donoho proposed to decompose an image over curvelet atoms having both an elongated support and vanishing moments [14, 15]. It brings a mathematical and algorithmic solution to the problem of approximating geometric $\mathcal{C}^2$ images whose contours are $\mathcal{C}^2$. The bandlet transform, [36, 43], uses an adaptive scheme to approximate an image. A dictionary of bandlet bases is parametrized by the local orientation of edges and an algorithm finds a best basis adapted to the function to approximate.

Texture Model. Following [45], Yves Meyer [38] pushed forward the idea of using more sophisticated norms to capture oscillating patterns. In particular he proposed a weak norm dual of the TV norm, $T(v) = \|v\|_G$ where the Banach space G contains signals with large oscillations, and thus in particular textures and noise. This model has been successfully implemented in [3, 51]. Meyer’s idea inspired several works. In [40] Osher, Solé, and Vese use the H^{-1} norm to extract high frequency patterns. In [5], Aujol and Gilboa propose a general framework based on a Hilbert norm defined by some symmetric positive kernel K , $T(v) = \langle Kv, v \rangle_{L^2}$. They showed an example where the Hilbert norm promotes a single frequency in the extraction of the texture. Similarly to the cartoon case, several approaches are based on the sparsity of the texture in a well chosen dictionary. The morphological component analysis of Starck et al. [47] uses a local cosine dictionary to model oscillating patterns. Peyré improves this fixed sparse regularization by using an adaptive grouplet frame for geometric textures [42] that makes use of a local orientation field.

Inpainting. Inpainting aims at restoring an image f_0 from which a set $\Omega \subset \{0, \dots, n-1\}^2$ of pixels is missing. It corresponds to the inversion of the ill-posed problem $f = \Phi f_0 + w$ where Φ is defined as

$$(\Phi f_0)(x) = \begin{cases} 0 & \text{if } x \in \Omega, \\ f_0(x) & \text{if } x \notin \Omega \end{cases} \quad (1.2)$$

and w is some additive noise. The solution of this inverse problem can be obtained as $u + v$ by solving (1.1) with the operator Φ defined above.

The J and T functionals mentioned above for structure and texture separation can also be used to solve the inpainting problem. These regularization approaches are successful for missing regions Ω of small size, because inpainting corresponds to an interpolation problem. Let us for example cite the total variation inpainting of Chan and Shen [19] inspired from [37], or the morphological component analysis [31, 32] which relies on sparse regularization in several transform domains such as wavelet for structure and local cosine for texture. When the size of the missing parts is larger than the characteristic length of the structure or texture, more constrained models are required. Several variational methods have been developed [37, 10, 46], often adding a curvature term to the total variation to guide the diffusion along the level lines. These higher-order methods improve the inpainting results but are also more unstable and slower.

For images where large areas are missing, inpainting corresponds to a problem of image synthesis with prescribed boundary conditions, and non-convex methods are required. Indeed, convex regularization approaches suffer from contrast attenuation in the center of large missing areas. Exemplar-based methods [30, 52] enables such a constrained synthesis of missing data. They perform the inpainting by using a consistent recopy of patches. These methods reconstruct well non-geometric textures. However in images where edges and directional textures are present, a geometry-oriented process seems necessary. Different approaches have been proposed in combination with an exemplar-based inpainting, by using an optimized ordering of pixels to copy [22, 41, 27, 16], a manual intervention by the user [48], or combining texture and geometric interpolation [11]. Exemplar-based methods can be casted as non-convex variational minimizations, see for instance [8, 1, 34, 35]. This non-convexity is crucial to cope with the attenuation effect that plagues convex regularization approaches.

1.2. Contributions. The main contribution of this work is a new adaptive framework for modeling locally parallel oscillating patterns. A first adaptive tex-

ture norm which promotes locally parallel oscillating patterns is constructed in Section 3. This texture model is based on a local Fourier decomposition and an adaptive frequency and orientation field ξ which is estimated during the minimization, $T(v) = T_\xi(v)$. $T_\xi(v)$ is a convex functional (with respect to v) which is small for a texture v if its main frequency around a point x is close to $\xi(x)$. Our model has therefore two important properties: it is spatially variant since different frequencies are locally extracted and it is adaptive since the frequency field is optimized to fit the features of the texture to extract.

This field $\xi : \{1, \dots, n\}^2 \rightarrow \mathbb{R}^2$ associates to a point x in the image $v \in \mathbb{R}^{n \times n}$ the instantaneous frequency of the oscillating texture around x . $\|\xi(x)\|$ gives the local frequency and $\xi(x)/\|\xi(x)\|$ the local orientation of the texture around x . Figure 1.2 shows an example of the field ξ for several points of a locally parallel texture.

FIG. 1.2. *On the right, a graphic representation of the instantaneous frequency field ξ for the texture on the left. The length of the lines is proportional to the local frequency of the oscillating patterns and the orientation of the lines gives their local orientation.*

Section 4 shows how our adaptive texture model can be used to regularize inverse problems such as image decomposition, denoising or inpainting. The minimization is done with respect to u and v , respectively the structure and texture part, but also with respect to ξ , the frequency field, and this problem is written as

$$(u, v, \xi) \in \underset{\tilde{u}, \tilde{v}, \tilde{\xi} \in \mathcal{C}}{\operatorname{argmin}} \frac{1}{2} \|f - \Phi(\tilde{u} + \tilde{v})\|_{L^2}^2 + \lambda J(\tilde{u}) + \mu T_{\tilde{\xi}}(\tilde{v}) \quad (1.3)$$

where $\mathcal{C}$ is a set of constraints on the field ξ and Φ is the linear operator to invert. For the decomposition and denoising problem Φ is the identity operator Id . For the inpainting case, Φ is the masking operator given by (1.2). This method computes a separation of the image into a structure component u and a texture component v ; the denoising or inpainting result is then obtained by adding these two components, i.e. by considering $u + v$. The energy to be minimized in (1.3) is non-convex, and Section 4 details a block coordinate descent that converges (up to a sub-sequence) to local minimum.

Section 5 shows numerical examples for image decomposition, denoising and inpainting. Let us remark that some of the existing decomposition frameworks (such as TV- G [38] or TV- H^{-1} [40]) are not suitable for denoising. As a matter of fact, the G and the H^{-1} norms are low for any high-frequency patterns, and they are also low for a large part of the noise. On the other hand, the TV norm penalizes strongly oscillating patterns and therefore these models are not able to separate efficiently the texture from the noise. On the contrary our norm is low for patterns which present a

certain frequency and orientation, but it is high for noise without any significant oriented patterns. Our texture model is therefore more appropriate for denoising locally parallel textures.

To tackle the inpainting of large missing regions, Section 6 introduces a highly non convex texture functional $T_{A,\xi}(v)$. It extends our texture modeling framework by integrating an amplitude field $A(x)$ that imposes the contrast of the texture patterns around a point x inside the area to inpaint. This non-convex model makes also use of a “rendering function” h which is a change of contrast that maps the inpainted texture v to a more general oscillating profile $h(v)$. This enables the inpainting of arbitrary locally parallel textures over large missing areas. An algorithm finds a local minimum of this non-convex energy and numerical examples show the efficiency of this approach.

2. Cartoon Model. This section presents the cartoon model $J(u)$ that we use to constrain the cartoon layer u in our regularization framework. Following the idea of [9] and [47], it mixes the total variation norm and the ℓ^1 -norm of the curvelet decomposition of u .

Total Variation Norm. In the following, $u \in \mathbb{R}^N$ is a discrete image of $N = n \times n$ pixels. A discretized gradient for such an image u is defined as

$$\nabla u[i, j] = (\partial_x u[i, j], \partial_y u[i, j]), \quad (2.1)$$

where

$$\partial_x u[i, j] = \begin{cases} u[i+1, j] - u[i, j] & \text{if } 0 \leq i < n-1, \\ 0 & \text{otherwise,} \end{cases} \quad (2.2)$$

$$\partial_y u[i, j] = \begin{cases} u[i, j+1] - u[i, j] & \text{if } 0 \leq j < n-1, \\ 0 & \text{otherwise.} \end{cases} \quad (2.3)$$

The gradient is thus a vector field $\nabla u \in \mathbb{R}^{N \times 2}$.

The discrete total variation of an image u is the ℓ^1 -norm of the gradient

$$J_{TV}(u) = \|\nabla u\|_1 \quad (2.4)$$

where the ℓ^1 norm of a vector field $z = (z^1, z^2) \in \mathbb{R}^{N \times 2}$ is

$$\|z\|_1 = \sum_{1 \leq i \leq N} \sqrt{(z_i^1)^2 + (z_i^2)^2}. \quad (2.5)$$

Curvelets. The curvelet transform [14, 15] is a decomposition on multiscale oriented atoms $c_m = c_{j,l,k}$ designed to represent images at different scales and angles. The curvelets atoms form a redundant tight frame of $\mathbb{R}^N$. A curvelet atom $c_{j,l,k} \in \mathbb{R}^N$ is parametrized by a scale j , an orientation l and a position k . $c_{j,l,k}(x)$ is of rapid decay away from a 2^{-j} by $2^{-j/2}$ rectangle (width = length²), with major axis pointing in the direction $\theta_l = 2\pi \cdot 2^{-\lfloor j/2 \rfloor} \cdot l$ and centered on a point x_k depending on k , l and j .

Cartoon images that are $\mathcal{C}^2$ outside a set of $\mathcal{C}^2$ edge curves have a sparse decomposition in the curvelet frame [14]. The ℓ^1 -norm of the curvelet coefficient

$$J_{Cuv}(u) = \sum_{j,k,l} |\langle u, c_{j,l,k} \rangle| \quad (2.6)$$

is thus an appealing functional to characterize such cartoon images, as advocated by Starck et al. [47].

We use the CurveLab 2.1.2 toolbox, which implements the discrete curvelet transforms described in [15]. It corresponds to a non-normalized Parseval tight frame. The ℓ^2 -norm of the atoms is then approximately $1/\sqrt{P/N}$ where P is the number of curvelet atoms.

TV-Curvelets. J_{TV} is well suited for piecewise constant images whereas J_{Curv} captures efficiently piecewise smooth images with $\mathcal{C}^2$ -smooth contours. However, contrary to J_{TV} , J_{Curv} also captures oscillating patterns of locally parallel textures. This is an issue to perform an efficient texture/structure decomposition. Following the idea of [9] and [47], we thus use a linear combination of these two energies in our cartoon model

$$J(u) = \rho J_{TV}(u) + (1 - \rho) J_{Curv}(u) \quad (2.7)$$

where $\rho \in [0, 1]$ balances the importance of the two functionals and allows one to find the best compromise between the effects of each energy. For the decomposition or denoising case, ρ is chosen to be equal to 1, $J(u) = J_{TV}(u)$. For the inpainting case ρ is chosen function of the size of the holes. If the size of the missing parts is large, the total variation inpainting fails indeed to reconstruct the geometric structure and the curvelet transform helps to improve the result of inpainting.

3. Texture Modeling Using an Adaptive Hilbert Norm. This section introduces a new texture model based on an adaptive vector field ξ which represents the instantaneous frequencies of the texture. Figure 1.2 shows a graphical representation of this field. This first model is a functional $T_\xi(v)$ (convex with respect to v), depending on ξ , which is small for a texture v if its main frequency around a point x is close to $\xi(x)$.

In [5], Aujol and Gilboa use a linear Hilbert norm defined by a symmetric positive kernel K ,

$$T(v) = \langle Kv, v \rangle_{L^2}.$$

This norm can be computed using a frame $\{\psi_\ell\}_\ell$ that is a possibly redundant family of $P \geq N$ atoms $\psi_\ell \in \mathbb{R}^N$. The decomposition of an image in this frame reads

$$\Psi v = \{\langle v, \psi_\ell \rangle\}_{\ell=0}^{P-1} \in \mathbb{R}^P \quad (3.1)$$

where $\Psi : \mathbb{R}^N \rightarrow \mathbb{R}^P$ is the frame operator.

Given a set of positive weights $\gamma_\ell \geq 0$, a norm is then defined as

$$T(v) = \sum_{\ell} \gamma_\ell^2 |\langle v, \psi_\ell \rangle|^2 = \|\Gamma \Psi v\|_{L^2}^2 \quad (3.2)$$

where $\Gamma = \text{diag}_\ell(\gamma_\ell)$. This corresponds to a Hilbert space associated to the kernel $K = \Psi^* \Gamma^2 \Psi$.

Defining a norm in this framework requires to define a transform Ψ well-suited to capture oscillating textures, and to compute a set of weights Γ adapted to the texture content v to extract from f . Aujol and Gilboa [6] proposed to use the Fourier basis so that Ψ corresponds to the discrete Fourier transform. This defines a translation-invariant kernel K . This paper proposes to replace the global Fourier basis by a redundant local Fourier basis, to capture the spatially and frequently varying structures

of locally parallel textures and to use the frequency field ξ in the definition of the weights $\Gamma(\xi)$ to promote locally the main frequency of the texture,

$$T_\xi(v) = \|\Gamma(\xi)\Psi v\|_{L^2}^2 \quad (3.3)$$

where Ψ is the decomposition on the local Fourier frame presented in Section 3.1 and $\Gamma(\xi)$ are the weights depending on ξ defined in Section 3.2.

3.1. Local Fourier Frame. A discrete short time Fourier atom, located around a position $x_p = p\Delta_x$ and with local frequency $\xi_k = k\Delta_\xi = k/q$ is defined as

$$\psi_{p,k}[y] = q^{-1}g[y - p\Delta_x]e^{\frac{2i\pi}{q}(y_1k_1 + y_2k_2)} \quad (3.4)$$

for $k \in \{-q/2, \dots, q/2 - 1\}^2$ and $p \in \{0, \dots, n/\Delta_x\}^2$, where g is a smooth window, centered around 0, and the size of its support is $q \times q$ pixels with $q > \Delta_x$.

The local Fourier frame $\{\psi_{p,k}\}_{p,k}$ is a redundant family of $P = (q/\Delta_x)^2 N$ vectors of $\mathbb{R}^N$. The decomposition operator Ψ defined by (3.1) which decomposes an image v in this frame, $\Psi v = \{\langle v, \psi_{p,k} \rangle\}_{p,k} \in \mathbb{R}^P$, is computed with the 2D Fast Fourier Transform of the $q \times q$ images $v[y]g[\Delta_x p - y]$. The computation of Ψv thus requires $O(Nq^2 \log_2(q^2)/\Delta_x^2)$ operations.

The dual operator Ψ^* reconstructs an image $\Psi^*c \in \mathbb{R}^N$ from a set of coefficients $c[p, k] \in \mathbb{R}^{q^2 \times N}$

$$\Psi^*c = \sum_{p,k} c[p, k]\psi_{p,k}. \quad (3.5)$$

This dual operator is implemented using N/Δ_x^2 inverse Fast Fourier Transforms.

The operator $\Psi^*\Psi$ is diagonal

$$\Psi^*\Psi = \text{diag}_x \left(\sum_y g[\Delta_x y - x]^2 \right). \quad (3.6)$$

The window g is normalized to satisfy

$$\forall x, \quad \sum_y g[\Delta_x y - x]^2 = 1. \quad (3.7)$$

This implies that $\Psi^*\Psi = \text{Id}_N$ so that $\Psi^* = \Psi^+$ is the pseudo inverse of the decomposition operator Ψ . The corresponding Gabor family $\{\psi_{p,k}\}_{k,p}$ is then a tight frame of $\mathbb{R}^N$ and one has both the analysis-synthesis formula and the conservation of the energy

$$v = \sum_{p,k} \langle v, \psi_{p,k} \rangle \psi_{p,k} \quad \text{and} \quad \|v\|^2 = \sum_{p,k} |\langle v, \psi_{p,k} \rangle|^2, \quad (3.8)$$

which generalize the concept of orthogonal basis to a redundant family of $P > N$ vectors. Let us finally remark that $\Psi^*\Psi = \text{Id}_N$ but $\Psi\Psi^* \neq \text{Id}_N$, since the Gabor frame is a redundant family.

For the numerical results, we use a tensor product window $g[x] = g^o[x_1]g^o[x_2]$ where g^o is a normalized Hanning window

$$g^o[x_1] = \frac{\tilde{g}^o[x_1]}{\sqrt{\sum_y \tilde{g}^o[\Delta_x y - x_1]^2}} \quad \text{where} \quad \tilde{g}^o[x_1] = \sin(\pi x_1/q)^2. \quad (3.9)$$

The size q of the local Fourier windows should be set according to the smoothness of the geometry of the texture and to its frequency content. A locally parallel texture with irregular directions requires small windows, but this might be an issue to detect low frequency patterns. In practice, an estimation ξ_{min} of the lowest frequency present in the texture is given by the user, and we set $q = 3/\xi_{min}$.

The overlapping Δ_x of the windows should satisfy $\Delta_x < q$ and selecting a small value helps to reduce visually unpleasant artifacts. In practice we set $\Delta_x = q/4$.

3.2. Adaptive Weight Design. The Hilbert norm $T_\xi(\cdot)$ adapted to oscillating textures is a weighted norm over the local Fourier coefficients. It is parametrized by a vector field ξ

$$\xi : \{1, \dots, n/\Delta_x\}^2 \mapsto \mathbb{R}^2 \quad (3.10)$$

which represents the local frequency of the texture component v . For $p \in \{1, \dots, n/\Delta_x\}^2$, the local frequency around the point $x_p = p\Delta_x$ is given by $|\xi(p)|$ and the local orientation of the texture is given by $\xi(p)/|\xi(p)|$.

The general formulation (3.2) is instantiated using a local Fourier frame $\psi_\ell = \psi_{p,k}$ for $\ell = (p, k)$ as

$$T_\xi(v) = \sum_{p,k} \gamma_{p,k}(\xi)^2 |\langle v, \psi_{p,k} \rangle|^2 = \|\Gamma(\xi)\Psi v\|_{L^2}^2 \quad (3.11)$$

where

$$\Gamma(\xi) = \text{diag}_{\ell=(p,k)}(\gamma_{p,k}(\xi)). \quad (3.12)$$

Each $\gamma_{p,k}(\xi) \geq 0$ weights the influence of each local Fourier atom in the texture model.

The norm $T_\xi(\cdot)$ should be small for an oscillating pattern around the point x_p if its main frequency is close to $\xi(p)$. As a consequence the weight $\gamma_{p,k}(\xi)$ should be small if ξ_k is close to $\xi(p)$ or to $-\xi(p)$. Furthermore, a special case is considered when there is locally no significant parallel texture in a certain area of the image. By convention, $\xi(p)$ is set to $(0, 0)$ if there is no significant oriented patterns around x_p in the image.

The weights are therefore defined as

$$\gamma_{p,k}(\xi) = \begin{cases} 1 & \text{if } \xi(p) = (0, 0) \\ \gamma_0 + \left(1 - G_\sigma(\|\xi_k + \xi(p)\|)\right)\left(1 - G_\sigma(\|\xi_k - \xi(p)\|)\right) & \text{otherwise} \end{cases} \quad (3.13)$$

where $G_\sigma(x) = \exp(-(x/\sigma)^2/2)$, and $\gamma_0 > 0$ is a small constant that prevents the weights from vanishing, and is required to prove the convergence of our numerical scheme, see Theorem 1.

The scale parameter σ reflects the deviation we are expecting to find in the frequency spectrum of the texture compared to $\xi(p)$. In our numerical experiments we took $\sigma = 1/q$. When there is not a significant oriented texture around x_p , we choose $\gamma_{p,k} = 1$ for all k , in order not to promote an arbitrary orientation in the extraction. The decision between the presence or the absence of the oriented texture is made during the computation of ξ and will become clear in the next section, see (4.2).

3.3. Sparsity versus Unimodality for Texture Modeling. Our texture model (3.11) assumes an unimodality of the local Fourier expansion of the texture. This can be understood as a strong sparsity prior, since the local expansion of the texture is constrained to be highly localized near a single frequency. Our texture prior is thus less general than the MCA framework [47] that makes use of a less constrained ℓ^1 sparsity.

Imposing unimodality however helps to better recover strongly geometric textures, such as the natural textures displayed in Figure 1.1. For applications that require processing or extraction of this kind of patterns, our method is able to improve the results of more general sparsity priors.

Furthermore, our method introduces an explicit geometric flow parameter $\xi(x)$, that might be of interest for certain applications beyond texture restoration. Although this is out of the scope of this paper, this flow $\xi(x)$ might also be regularized by imposing additional constraints on its regularity.

3.4. Color Texture Processing. The adaptive texture energy $T_\xi(v)$ extends to the color setting, where $v = (v_1, v_2, v_3)$ and each $v_i(x) \in \mathbb{R}$ is a color component. Natural textures tend to have a single geometry across all the three color components, and we thus consider a single geometrical flow ξ . The energy (3.11) is extended as

$$T_\xi(v) = \sum_{p,k} \gamma_{p,k}(\xi)^2 \sum_{i=1}^3 |\langle v_i, \psi_{p,k} \rangle|^2.$$

Using a single geometry for all the components produces a coherent texture separation and inpainting.

For the TV part of the regularization term (J_{TV}), we can use any classical scheme for color TV minimization. It can be done with a regularized version of TV and the use of the CB or HSB spaces as in [20, 7]. Another possibility is to use the definition of $BV(\mathbb{R}^3)$ and then resort to a projection scheme as in [12, 29] (the mixing of the three channels is then done through the vectorial norm). For the curvelet part (J_{curv}) of the regularization term, the same idea can be used with a multichannel ℓ^1 -norm as in [49].

4. Adaptive Texture Regularization Algorithm. Taking advantage of the adaptivity of the texture model introduced in Section 3, an adaptive texture regularization algorithm is presented. It splits an image f into three components, $f = u + v + w$, where u captures the sketch of the image, v the texture content and w the noise. The minimization is done with respect to u and v , the structure and texture parts, but also with respect to ξ , the frequency field. This algorithm finds a stationary point of problem (1.3)

$$(u, v, \xi) = \underset{\tilde{u}, \tilde{v}, \tilde{\xi} \in \mathcal{C}}{\operatorname{argmin}} \mathcal{E}(\tilde{u}, \tilde{v}, \tilde{\xi}) = \underset{\tilde{u}, \tilde{v}, \tilde{\xi} \in \mathcal{C}}{\operatorname{argmin}} \frac{1}{2} \|f - \Phi(\tilde{u} + \tilde{v})\|_{L^2}^2 + \lambda J(\tilde{u}) + \mu T_{\tilde{\xi}}(\tilde{v}), \quad (4.1)$$

where J is the regularization term defined in (2.7) and T_ξ is our texture norm defined by (3.11). The residual noise layer is computed as $w = f - u - v$. Φ is a linear operator so that an inverse problem can be solved during the decomposition process. For a denoising and separation purpose, Φ is simply the identity operator, i.e. $\Phi = \operatorname{Id}$, and for the inpainting inverse problem, where one wants to reconstruct some missing parts, Φ is the masking operator given by (1.2).

$\mathcal{C}$ is a set of constraints on the orientation field ξ . Since the texture component v does not contain low frequency patterns, the frequency $|\xi|$ is forced to be large enough, i.e. $\forall p, |\xi(p)| > \tau$, for some real positive parameter $\tau > 0$. Furthermore, an oscillating pattern of frequency $\xi(p)$ is assumed to be present in the image f around the point x_p only if $|\langle f, \psi_{p,k} \rangle| > \eta_p$ where $k = \xi(p)/\Delta_\xi$ and $\eta_p > 0$ is a real positive parameter. This reflects the fact that one does not want to arbitrary select a frequency for an area of the image where there is no oscillating pattern. In short, we have

$$\mathcal{C} = \left\{ \xi : \{1, \dots, n/\Delta_x\}^2 \mapsto \mathbb{R}^2 \left| \begin{array}{l} \forall p, \quad \tau \leq |\xi(p)| \leq 1/2 \\ \forall p, \quad (\forall k, |\langle f, \psi_{p,k} \rangle| \leq \eta_p) \Rightarrow \xi(p) = (0, 0) \end{array} \right. \right\}. \quad (4.2)$$

In our numerical experiments we take $\tau = 2/q$, where q is the size of the local Fourier windows, and $\eta_p = 2|\overline{\Psi f_p}|$ where $|\overline{\Psi f_p}|$ is the average value of $|\langle f, \psi_{p,k'} \rangle|$ for $k' \in \{-q/2, \dots, q/2 - 1\}^2$.

4.1. Block coordinate descent optimization algorithm. The energy $\mathcal{E}$ defined in (1.3) is non-convex. We propose a block coordinate descent algorithm to minimize $\mathcal{E}$: we iteratively minimize $\mathcal{E}$ with respect to each of its variables ξ , u , v . Starting from $u^{(0)} = v^{(0)} = \xi^{(0)} = 0$, the following iterates are defined :

$$u^{(\ell+1)} \in \underset{u}{\operatorname{argmin}} \mathcal{E}(\tilde{u}, v^{(\ell)}, \xi^{(\ell)}), \quad (4.3)$$

$$v^{(\ell+1)} = \underset{v}{\operatorname{argmin}} \mathcal{E}(u^{(\ell+1)}, \tilde{v}, \xi^{(\ell)}), \quad (4.4)$$

$$\xi^{(\ell+1)} = \underset{\xi \in \mathcal{C}}{\operatorname{argmin}} \mathcal{E}(u^{(\ell+1)}, v^{(\ell+1)}, \tilde{\xi}). \quad (4.5)$$

Sections 4.2, 4.3 and 4.4 detail how to perform each one of these three minimizations.

The energy $\mathcal{E}$ is decreasing at each step. The following theorem ensures the convergence up to a sub-sequence of the algorithm towards a stationary point of $\mathcal{E}$.

THEOREM 1. *Suppose that $\Phi 1 \neq 0$. The sequence $(u^{(\ell)}, v^{(\ell)}, \xi^{(\ell)})$ defined in (4.3) is well defined and bounded. Every cluster point of this sequence is a stationary point of $\mathcal{E}$.*

Proof. The energy $\mathcal{E}$ is non-convex and non-smooth, and we use Theorem 4.1 of Tseng [50]. The energy $\mathcal{E}$ is re-written using the notations of [50] as

$$\begin{aligned} \mathcal{E}(u, v, \xi) &= \frac{1}{2} \|f - \Phi(u + v)\|^2 + \lambda T_\xi(v) + \mu J(u) + \chi_{\mathcal{C}}(\xi) \\ &= f_0(u, v, \xi) + f_1(u) + f_2(v) + f_3(\xi), \end{aligned}$$

where $f_1 = \mu J$, $f_2 = 0$ and $f_3 = \chi_{\mathcal{C}}$. The indicator function is defined as $\chi_{\mathcal{C}}(\xi) = +\infty$ if $\xi \notin \mathcal{C}$ and $\chi_{\mathcal{C}}(\xi) = 0$ otherwise.

First, the function $\mathcal{E}$ is continuous on the level set

$$X_0 = \left\{ (u, v, \xi) \mid \mathcal{E}(u, v, \xi) \leq \mathcal{E}(u^{(0)}, v^{(0)}, \xi^{(0)}) \right\}.$$

We then show that X_0 is bounded, which ensures the existence of convergence sub-sequences. First we note that the constraint set $\mathcal{C}$ defined in (4.2) is bounded. The texture norm with the weights defined in (3.13) satisfies

$$T_\xi(v) \geq \gamma_0 \|\Psi v\|^2 \geq C_1 \|v\|^2 \quad (4.6)$$

for some constant C_1 because Ψ is injective. If $\rho < 1$, the cartoon energy $J(u)$ defined in (2.7) is coercive because $J_{C_{uv}}$ is coercive. If $\rho = 1$, $J(u) = J_{TV}(u)$, and denoting $m(u) = (\sum_{i,j} u[i, j])/N$ the mean of u , one has

$$J(u) = \|\nabla(u - m(u))\|_1 \geq \|\nabla(u - m(u))\|_2 \geq C_2 \|u - m(u)\|, \quad (4.7)$$

where we have used the fact that $\|\cdot\|_1 \geq \|\cdot\|_2$, and where $C_2 > 0$ is the smallest non-zero singular value of the gradient ∇ . We decompose $f = \Phi \bar{f} + r$ where $r \in \text{Im}(\Phi)^\perp$, so that

$$\|f - \Phi(u + v)\|^2 \geq \|\Phi(\bar{f} - u - v)\|^2 \geq C_3^2 \|m(\bar{f}) - m(u) - m(v)\|^2 \quad (4.8)$$

where $C_3 > 0$ is the smallest non zero singular value of Φ , and where the last inequality makes use of the hypothesis that $\Phi 1 \neq 0$. Putting (4.6), (4.7) and (4.8) together proves that

$$\mathcal{E}(u, v, \xi) \geq \frac{C_3^2}{2} \|m(\bar{f}) - m(u) - m(v)\|^2 + \lambda C_2 \|u - m(u)\| + \mu C_1 \|v\|^2,$$

and hence the boundedness of X_0 .

The function $f_0(u, v, \xi)$ is of class C^1 , since the mapping

$$(u, \xi) \mapsto T_\xi(u) = \|\Gamma(\xi)\Psi u\|^2$$

is smooth because Γ is a smooth function. The function $\mathcal{E}$ satisfies hypothesis (A1) of [50], and hence to lemma 3.1 of [50] shows that $\mathcal{E}$ is regular at each point of X_0 .

This shows that $\mathcal{E}$ satisfies the hypotheses of Theorem 4.1 in [50] and it concludes the proof. $\square$

4.2. Minimization with respect to the Orientation Field ξ . If u and v are fixed, we search for the frequency field ξ satisfying

$$\xi = \underset{\xi \in \mathcal{C}}{\text{argmin}} T_\xi(v), \quad (4.9)$$

where T_ξ is given by (3.11). This requires, for each p , to compute

$$\xi(p) = \underset{|\theta| > \tau}{\text{argmin}} \sum_k \left(1 - G_\sigma(\|\xi_k + \theta\|)\right)^2 \left(1 - G_\sigma(\|\xi_k - \theta\|)\right)^2 |\langle v, \psi_{p,k} \rangle|^2. \quad (4.10)$$

If the width σ of the weights (3.13) is small enough, this minimization is approximately equivalent to compute $\max_k |\langle v, \psi_{p,k} \rangle|$, which allows us to speed up the computation by defining

$$\xi(p) = \Delta_\xi \underset{k > \frac{\tau}{|\Delta_\xi|}}{\text{argmax}} |\Psi v[p, k]|. \quad (4.11)$$

Figure 4.1 illustrates the underlying principle of this orientation estimation. For a given point x_p , a unique direction and frequency $\xi(p)$ is selected and the corresponding weights $\gamma_{p,k}(\xi)$ are constructed according to (3.13).

4.3. Minimization with respect to u . If v is fixed and if we set $y = f - \Phi(v^{(i)})$, one minimizes

$$u^{(i+1)} = \underset{\tilde{u}}{\text{argmin}} \frac{1}{2} \|y - \Phi(\tilde{u})\|_{L^2}^2 + \lambda J(\tilde{u}). \quad (4.12)$$

FIG. 4.1. Illustration of the orientation estimations: (a) the input image f , (b) the windowed image around some point x_p , (c) the corresponding local Fourier transform $(|\langle v, \psi_{p,k} \rangle|)_k$ and (d) the weights $(\gamma_{p,k}(\xi))_k$ corresponding to the ξ estimated from the local Fourier transform.

First case: $\Phi = \text{Id}$. The solution of (4.12) when $\Phi = \text{Id}$ is then given by

$$u^{(i+1)} = \text{prox}_{\lambda J}(y)$$

where the proximity operator of ωJ for $\omega > 0$ and $g \in \mathbb{R}^N$ is defined as

$$\text{prox}_{\omega J}(g) = \underset{\tilde{g}}{\text{argmin}} \frac{1}{2} \|g - \tilde{g}\|^2 + \omega J(\tilde{g}) \quad (4.13)$$

This corresponds to a denoising problem for which one can use for instance the algorithm proposed by Chambolle [17]. For sake of completeness, this algorithm is described in appendix A.

General case. The inverse problem corresponding to the general case of $\Phi \neq \text{Id}$ can be solved using gradient methods such as forward-backward splitting [21, 9] or Nesterov algorithms [39], see also [2] for an overview of these methods.

In the numerical experiments, we use the forward-backward splitting scheme. Starting from some initial $w^{(0)}$, one iterates between

- a gradient descent step of the minimization of the data term,

$$\bar{w}^{(n)} = w^{(n)} + \alpha \Phi^*(y - \Phi w^{(n)}) \quad (4.14)$$

- a denoising step over the current estimate $\bar{w}^{(n)}$,

$$w^{(n+1)} = \text{prox}_{\lambda \alpha J}(\bar{w}^{(n)}) \quad (4.15)$$

where $\text{prox}_{\lambda \alpha J}(\bar{w}^{(n)})$ is the proximity operator defined in (4.13) for $\omega = \lambda \alpha$, that is computed using the iterative algorithm detailed in Appendix A.

If the gradient step size α in (4.14) satisfies $0 < \alpha < \frac{2}{\|\Phi^* \Phi\|}$, then $w^{(n)}$ in (4.14) and (4.15) converges to $u^{(i+1)}$, a global minimizer of (4.12). In the inpainting case, when Φ is the masking operator given by (1.2), we have $\|\Phi^* \Phi\| = 1$.

4.4. Minimization with respect to v . If u is fixed, one minimizes

$$v^{(i+1)} = \underset{\tilde{v}}{\text{argmin}} \frac{1}{2} \|f - \Phi(u^{(i+1)}) - \Phi(\tilde{v})\|_{L^2}^2 + \mu \|\Gamma(\xi) \Psi \tilde{v}\|_{L^2}^2, \quad (4.16)$$

Computing the gradient of (4.16), we obtain that $v^{(i+1)}$ satisfies

$$(2\mu \Psi^* \Gamma^2 \Psi + \Phi^* \Phi) v^{(i+1)} = \Phi^*(f - \Phi(u^{(i+1)})) \quad (4.17)$$

and the solution $v^{(i+1)}$ is computed with a conjugate gradient descent, since $2\mu \Psi^* \Gamma^2 \Psi + \Phi^* \Phi$ is a positive symmetric operator.

4.5. Case of the Decomposition of a Noise Free Image. If the input image f does not contain any noise, one would like to decompose f into only two components, the sketch u and the texture $v = f - u$. This requires to solve

$$(u, \xi) = \operatorname{argmin}_{\tilde{u}, \xi \in \mathcal{C}} \frac{1}{2} T_{\xi}(f - \tilde{u}) + \lambda J(\tilde{u}). \quad (4.18)$$

The minimization step on ξ is the same as the one described in Section 4.2, but the second step on u should be modified. When ξ is fixed, using the definition (3.11) of T_{ξ} and setting $y = \Gamma(\xi)\Psi f$, the minimization on u can be rewritten

$$u = \operatorname{argmin}_{\tilde{u}} \frac{1}{2} \|y - \Gamma(\xi)\Psi(\tilde{u})\|_{L^2}^2 + \lambda J(\tilde{u}). \quad (4.19)$$

This minimization corresponds to a regularized inverse problem associated to the operator $\Gamma(\xi)\Psi$. It can therefore be solved using the forward-backward splitting algorithm described in Section 4.3.

5. Numerical Examples. For our numerical experiments, we use a non-optimized implementation in Matlab and, for an image of size 512×512 , the processing time is about 10 minutes. As mentioned in the section 4.3, we use a forward-backward splitting scheme for the minimizing step on u . The computation time should be reduced by using Nesterov algorithms [39] instead. However the computation of the decomposition in the local Fourier frame and its dual operator is the most time-consuming step and it should be optimized if better computational time is desired.

5.1. Image decomposition and denoising. For the decomposition and denoising problem, we take $\Phi = \text{Id}$ and solve (1.3) using the algorithm described above. We choose $J(u) = \|u\|_{\text{TV}}$ for the cartoon functional, which corresponds to taking $\rho = 1$ in (2.7).

Cartoon+texture decomposition.. Figure 5.1 shows an input image f , 256×256 generated by addition of a cartoon picture¹ and a synthetic texture whose orientation and frequency vary spatially ($\|f\|_{\infty} = 3$). Figure 5.2 presents the decomposition results obtained with different classical methods on this noise free image: the $TV - L^2$ model of [45], the $TV - G$ model of Y. Meyer [38] with the implementation of [3], the $TV - L^1$ model with the implementation of [18], the $TV - H^{-1}$ model [40] with the implementation of [4], and the TV -Gabor model [6]. For each of these methods we chose the first set of parameters which provides a total extraction of the texture. For our method we chose $\lambda = 0.1$, $q = 16$, $\Delta_x = 4$. As can be seen, the approach developed in this paper provides much better result. This is due to the fact that the texture in the image of Figure 5.1 lies exactly in the model of the paper. In particular, the $TV - L^1$ model performs poorly on this type of texture, since it splits an image into geometry and texture on a texton criterion [28].

Figure 5.3 presents an example of the decomposition of a natural image containing locally parallel textures. The input image f , 389×389 , is in the first column. The second column presents the result of the $TV - L^2$ method [45], we chose the smallest parameter λ (on the total variation norm) which provides a total extraction of the texture (here $\lambda = 0.4$). And finally our method is shown in the last column (we chose $\lambda = 0.1$, $q = 32$, $\Delta_x = 10$). As can be seen, whereas both texture components contain

¹This picture is a frame from the animated cartoon “Falling Hare” (1943). This cartoon is in the public domain.

FIG. 5.1. A synthetic example. (a) the input image ; (b) and (c) respectively the structure and texture components used to produce the image.

the locally parallel patterns of the fish, our method achieves a better decomposition in the sense of more details of the background are present in the geometric component.

Cartoon+texture+noise decomposition.. In Figure 5.4 an image f_0 composed of a cartoon picture and a fingerprint texture (the size of f_0 is 512×512 and $\|f_0\|_\infty = 3$) is degraded by a Gaussian noise of standard deviation $\sigma = 0.2$. The noisy image f is then decomposed into three components u , v , and w using our method with the following parameters $\lambda = 0.1$, $\mu = 0.3$, $q = 48$ and $\Delta_x = 16$. An orientation field ξ is therefore also computed. Since u captures the sketch of the image, v the locally parallel patterns and w the noise, we can reconstruct a restored version $u + v$ of the noisy image.

One should be careful that the comparison with TV-denoising is only intended to show the relevance of our texture model to extract oscillating components, since our cartoon model is the TV energy. The TV model should not be considered as a state of the art method for denoising, and we give below examples of state of the art denoising methods on a natural image.

Comparison with state of the art denoising methods.. Figure 5.5 shows the decomposition and denoising process applied to the “Barbara” image f_0 degraded by a Gaussian noise of standard deviation $\sigma = 0.15$. The size of f_0 is 512×512 and we have $\|f_0\|_\infty = 1$.

Figure 5.6 compares our result with Portilla et al. denoising method [44], based on scale mixtures of Gaussians over the wavelet domain, with the non-local means methods [13], and with the recent state of the art BM3D method [23].

We use the Signal-to-Noise Ratio (SNR) measure of distortion between the original noiseless image f_0 and the denoised image f^*

$$\text{SNR}(f_0, f^*) = 20 \log_{10} \frac{\|f_0\|_{L^2}}{\|f_0 - f^*\|_{L^2}}.$$

We used $\lambda = 0.2$, $\mu = 5$, $q = 32$ and $\Delta_x = 8$ as parameters for our method.

We note that our method improves the wavelet-based statistical modeling of Portilla et al. denoising [44] and the non-local means methods [13], but it fails to improve over a more recent state of the art denoising method, BM3D [23]. Zoom on the textured region of the image shows that our method gives satisfactory results on textured part, but does not performs as good as the state of the art in region containing either edges or complicated texture patterns. This is consistent with the fact that we use

FIG. 5.2. A synthetic example. (a) and (d): decomposition results with $TV - L^1$; (b) and (e): decomposition results with $TV - L^2$; (c) and (f): decomposition results with $TV - G$; (g) and (j): decomposition results with $TV - H^{-1}$; (h) and (k): decomposition results with $TV - \text{Gabor}$; (i) and (l) : decomposition results with our adapted $TV - \text{Hilbert}$ method. The obtained result is almost perfect.

FIG. 5.3. A natural image. (a): the input image ; (b) and (c): decomposition results with $TV - L^2$; (d): contrast enhanced version of (c) ; (e) and (f): decomposition results with our adapted TV -Hilbert method ; (g) contrast enhanced version of (f)

a TV regularization for the cartoon layer, and that our model is tailored for highly geometric textures.

5.2. Image Decomposition and Inpainting. Let us recall that inpainting aims at restoring an image f_0 from which a set $\Omega \subset \{0, \dots, n-1\}^2$ of pixels is missing. It corresponds to the inversion of the ill posed problem $f = \Phi f_0 + w$ where Φ is defined as

$$(\Phi f_0)(x) = \begin{cases} 0 & \text{if } x \in \Omega, \\ f_0(x) & \text{if } x \notin \Omega. \end{cases} \quad (5.1)$$

and w is some additive noise. We search for the image f_0 as a decomposition $f_0 \approx u + v$ where $J(u)$ and $T_\xi(v)$ are small, and ξ is optimized during the inpainting process. This

FIG. 5.4. (a) the original noise free image ; (b) the input noisy image f ; (c) the restored image $u + v$; (d) the estimated orientations of oscillating patterns ξ ; (e) the sketch u of the image ; (f) the texture content v ; (g) the noise w .

corresponds to the solution of (1.3) where the operator Φ is given by (1.2) and y is the image with missing parts one wants to inpaint. For the cartoon model, we take $J(u)$ given by (2.7), which mixes the total variation norm and the ℓ^1 -norm of the curvelet decomposition of u . An experimental exploration leads us to choose $\rho = 0.75$ which gives a good compromise between the effects of each energy.

Inpainting of a synthetic image.. Figure 5.7 shows an example of inpainting reconstruction of the image f_0 of size 512×512 from Figure 5.4 degraded by randomly placed holes, which consists in 350 squares of 15×15 pixels each. We use the same parameters as in Figure 5.4. The texture is well reconstructed thanks to the estimation of the orientations and to the overlapping of the local Fourier windows.

FIG. 5.5. (a) the original noise free image f_0 ; (b) the noisy input image f ; (c) the restored image $u + v$; (d) the sketch u of the image ; (e) the texture content v ; (f) the noise w .

Comparison with state of the art inpainting methods.. Figure 5.8 shows a second example of inpainting reconstruction for the “Barbara” image. We use the same parameters as in Figure 5.5. For comparison, we also show the result of inpainting using TV regularization, the “patchworks” method from Perez et al. [41] and the MCA method [32], using a curvelet dictionary for the cartoon component and a local discrete cosine transform for the texture part, that corresponds to take $T(v) = \|\Psi_{LC}v\|_1$ where Ψ_{LC} is the local cosine transform. The enlargements in the second and third columns illustrate the differences between the different methods and our framework.

As far as the MCA method is concerned, on the one hand, our reconstruction of the hatched chair (top-right corner of the original image) selects the main orientation since we are modeling locally parallel patterns. Our method is in fact designed to deal with strongly parallel textures which can be locally explained by a single frequency, whereas the MCA method based on sparsity seems to be able to handle more general cases such as the one presented in the second column. On the other hand, for parts of the image with locally parallel patterns, our method achieves a better reconstruction of the directions of the texture inside the missing parts thanks to its adaptivity. This is clearly visible in the third column.

Figure 5.8 also shows the result of a copy-paste based method [41], even though it is beyond the scope of this paper to make such comparisons, since the focus is on energy minimization based methods. Nevertheless, the enlargements in the second and third columns show that such a method reconstructs very well not-oriented texture parts of the image (first column), whereas it fails to reconstruct properly the

FIG. 5.6. Denoising of image f from Fig. 5.5. (a) TV-denoising ; (b) result of Portilla et al. method [44] ; (c) resultat of the Non-Local means method [13] (d) our method ; (e) result of BM3D method [23].

FIG. 5.7. (a) the image y to inpaint degraded by randomly chosen holes in black, the original image is f_0 in Fig. 5.4 ; (b) the inpainted geometric component u , (c) the inpainted texture component v ; (e) “TV Inpainting” using only a TV regularization ; (d) the reconstruction $u + v$ using our method.

orientation of the locally parallel texture present in the third column.

6. Highly non-Convex Generalization of our Model. Similarly to other existing convex regularization methods, our approach suffers from a contrast attenuation in the center of the inpainted area. Section 6.1 describes a highly non-convex functional that fixes this issue. Another issue with our locally parallel texture model is that it favors the apparition of patterns with a sinusoidal profile. Section 6.2 thus introduces a “rendering function” that enables the reconstruction of arbitrary locally parallel textures.

6.1. Amplitude Boosting.

Contrast attenuation problem. Although convex methods provide good solutions to inpaint small or medium holes, this class of methods is not efficient to inpaint large holes. Notice that the functional presented in the previous section is convex with respect to the texture v . We therefore will refer to this previous model as *convex* energy, as opposed to the new one we introduce here. The minimization of convex energies causes some attenuation in the reconstruction, as shown on Figure 6.1. This texture inpainting is computed by solving

$$\operatorname{argmin}_{\tilde{v}, \tilde{\xi} \in \mathcal{C}} \frac{1}{2} \|y - \Phi \tilde{v}\|_{L^2}^2 + \mu T_{\tilde{\xi}}(\tilde{v})$$

which corresponds to taking $\lambda = \infty$ in (1.3).

FIG. 5.8. Inpainting of a degraded image, the original image is f_0 in Fig. 5.5. First column: (a) the image y to inpaint ; (b) reconstruction using a TV regularization ; (c) result of the “patchworks” method [41] ; (d) result of MCA [32] with curvelets and local discrete cosine dictionaries ; (e) our reconstruction. Second and third columns: enlargement of parts of the same images.

We thus propose a new non-convex energy to cope with large holes. The new texture functional $T_{A,\xi}$ depends both on the orientation field ξ and on some amplitude field

$$A : \{1, \dots, n/\Delta_x\}^2 \rightarrow \mathbb{R}^+.$$

For a point $x_p = p\Delta_x$ in the image plane, $A(p)$ is an estimation of the amplitude of the oscillating pattern around x_p . The new texture functional imposes a given amplitude field inside the missing parts which annihilates the attenuation problem.

FIG. 6.1. From left to right: the image y to inpaint, the result v_{convex} of our convex inpainting and the result $v_{non-convex}$ of our non convex inpainting.

Non-convex texture functional. The non-convex energy $T_{A,\xi}(v)$ sums over all points $x_p = p\Delta_x$ the distance between the magnitude A_p of the local Fourier coefficients of v and a pure pattern $A(p)\chi_{\xi(p)}$ of amplitude $A(p)$ and orientation $\xi(p)$. It is therefore defined as

$$T_{A,\xi}(v) = \sum_p \|A_p - A(p)\chi_{\xi(p)}\|^2 = \sum_{p,k} (A_p(k) - A(p)\chi_{\xi(p)}(k))^2 \quad (6.1)$$

Pure oriented pattern. The magnitude of the local Fourier transform of the texture v corresponds to

$$A_p = \{A_p(k)\}_k \quad \text{where} \quad A_p(k) = |\langle v, \psi_{p,k} \rangle|_\varepsilon.$$

We use a regularized absolute value

$$|a|_\varepsilon = \sqrt{|a|^2 + \varepsilon},$$

where $\varepsilon > 0$ is a small positive number chosen to avoid numerical instabilities.

The pure pattern χ_ξ is defined as the magnitude of the local Fourier transform of a pure sinusoidal function S_ξ

$$\chi_\xi(k) = |\langle S_\xi, \psi_{0,k} \rangle|_\varepsilon \quad \text{where} \quad S_\xi(x) = \sin(2\pi \langle x, \xi \rangle). \quad (6.2)$$

Figure 6.2 compares, for a given position x_p , the weights $\{\gamma_{p,k}(\xi)\}_k$ of the convex texture model defined in (3.13) with the pure pattern $\chi_{\xi(p)}$ defined in (6.2). Whereas the convex weights correspond to a double “potential well”, the non-convex pattern corresponds to a “double-bump” of fixed height. This difference is crucial since the amplitude A_p of the reconstructed texture is boosted in the non-convex setting due to the similarity term $\|A_p - A(p)\chi_{\xi(p)}\|$.

FIG. 6.2. *Left: graphical representation of the convex weights $\{\gamma_{p,k}(\xi)\}_k$; right: graphical representation of the pure pattern $\chi_{\xi(p)}$ used by the non-convex functional.*

Estimating the amplitude field. In our numerical experiments, this field $A(p)$ is estimated during the inpainting process from the texture $v = v^{(i)}$ obtained at iteration i of the algorithm. The automatic computation of a texture amplitude field with missing information is a difficult problem, and we assume here that this field varies smoothly over the image plane, which is a reasonable assumption for many natural textures.

For positions x_p located at a distance greater than $q/2$ from the mask Ω ,

$$p \in U_\Omega = \{p \mid d(x_p, \Omega) > q/2\} \quad \text{where} \quad d(x_p, \Omega) = \min_{y \in \Omega} \|x_p - y\|$$

the amplitude field $A(p)$ is computed by minimizing $\|A_p - A(p)\chi_{\xi(p)}\|^2$ since there is no missing information. If $v = v^{(i)}$ and $\xi(p)$ are fixed, the amplitude $A(p)$ is thus defined as

$$A(p) = \frac{\langle A_p, \chi_{\xi(p)} \rangle}{\|\chi_{\xi(p)}\|^2}. \quad (6.3)$$

For the other remaining positions $p \notin U_\Omega$, the amplitude $A(p)$ is computed by a smooth interpolation over the missing region, obtained by solving

$$\forall p \notin U_\Omega, \quad \Delta A(p) = 0, \quad \forall p \in U_\Omega, \quad A(p) = \frac{\langle A_p, \chi_{\xi(p)} \rangle}{\|\chi_{\xi(p)}\|^2} \quad (6.4)$$

where Δ is the Laplacian second derivative operator. The full field A is thus computed by conjugate gradient descent.

6.2. General Oscillation Profile.

Sinusoid texture profile problem. In the inpainting problem, the reconstruction inside the missing regions is only constrained by the two terms $J(u)$ and $T(v)$. There is no fidelity term inside these regions. For points x around $x_p = p\Delta_x$ close to the center of a missing region, the inpainted texture v is thus well approximated by a single frequency

$$v(x) \approx A(p) \sin(2\pi \langle x, \xi(p) \rangle + \phi) \quad (6.5)$$

where $\phi \in \mathbb{R}$ is a local phase.

Figure 6.3 (b) shows an example of this issue, where this texture inpainting result is obtained by solving

$$\operatorname{argmin}_{\tilde{v}, \tilde{\xi} \in \mathcal{C}} \frac{1}{2} \|y - \Phi \tilde{v}\|_{L^2}^2 + \mu T_{\tilde{\xi}}(\tilde{v}),$$

which corresponds to taking $\lambda = \infty$ in (1.3).

FIG. 6.3. (a) the image y to inpaint ; (b) the result v_{Id} of our inpainting with $h = \text{Id}$; the result v_h of our inpainting with $h = h_{0.1, -0.5}$ before the contrast change ; the result $v_h^{[h]}$ of our inpainting after the contrast change.

Rendering function. To cope with more general form of oscillating patterns, the reconstructed texture is switched from v to $v^{[h]}$ using a “rendering function”

$$h : [0, 1] \rightarrow \mathbb{R}.$$

This function allows one to change the contrast of a texture v having a local amplitude field $A(p)$ by computing

$$v^{[h]}(x) = A(p)h(v(x)/A(p)) \quad (6.6)$$

where p is the closest integer to x/Δ_x .

By analogy to (6.5), for points x around $x_p = p\Delta_x$ close to the center of a missing region, the inpainted texture v is approximately

$$v(x) \approx A(p)h(\sin(2\pi\langle x, \xi(p) \rangle + \phi)).$$

Figure 6.3 (c) and (d) shows an example of inpainting using a well chosen rendering function. The details of the inpainting method are given in Section 6.3.

Estimating the rendering function. Similarly to the computation of the amplitude field, computing an optimized rendering function to achieve the best visual result is difficult. We thus restrict ourself to a parametrized family $\{h_{a,b}\}_{a,b}$ depending on two shape parameters $a > 0$ and $-\frac{1}{2} < b < \frac{1}{2}$

$$\forall t, \quad h_{a,b}(t) = \text{sign}(t - b)|t - b|^a. \quad (6.7)$$

Figure 6.4 shows some examples of functions $h_{a,b}$ for various values of a and b and the effect of applying the change of contrast $v^{[h_{a,b}]}$ to a sinusoidal pattern v .

The parameter a controls the shape of the oscillating pattern: for $a < 1$ one obtains a “crenel” profile, for $a = 1$ one gets a sinusoidal profile, and for $a > 1$ the profile is flat around 0 and presents peaks at the minimum and maximum values. The

parameter b balances the importance between the negative and the positive values in the pattern profile.

In the numerical simulation, the parameters (a, b) are manually tuned by the user to achieve a good visual quality. More advanced adaptation strategies could be used, but we found it sufficient to explore manually the set of parameters.

FIG. 6.4. Examples of functions $h_{a,b}$ (top row) for various values of a and b and patterns $v^{[h_{a,b}]}$.

6.3. Non-Convex Inpainting. Taking into account the non convex energy $T_{A,\xi}$ defined in (6.1) and the change of contrast (6.6) gives rise to the following non-convex minimization problem:

$$(u, v, \xi) = \operatorname{argmin}_{\tilde{u}, \tilde{v}, \tilde{\xi} \in \mathcal{C}} \frac{1}{2} \|f - \Phi(\tilde{u} + \tilde{v}^{[h]})\|_{L^2}^2 + \lambda J(\tilde{u}) + \mu T_{A,\tilde{\xi}}(\tilde{v}). \quad (6.8)$$

The texture component extracted by this method is then given by $v^{[h]}$ and v is an approximation using pure sinusoidal oscillations of the texture present in the image. The inpainted image is computed as $u + v^{[h]}$.

Non-Convex Regularization Algorithm. This section describes an algorithm to minimize (6.8) for a given rendering function $h : \mathbb{R} \rightarrow \mathbb{R}^+$ and a given amplitude field $A : \{1, \dots, n/\Delta_x\}^2 \rightarrow \mathbb{R}^+$ with respect to ξ , u and v . Similarly to Section 4, one iterates between the minimization on ξ , on u and on v . In practice, the amplitude field A is updated during the minimization process to provide a more accurate estimation of the texture amplitude.

The energy is decreasing at each step but this algorithm is not guaranteed to converge to a local minimum of (6.8). However during our numerical experiments, we observed that the algorithm always converges.

Minimization with respect to ξ . Similarly to Section 4.2, if u and v are fixed, we compute the frequency field ξ that satisfies

$$\xi = \operatorname{argmin}_{\tilde{\xi} \in \mathcal{C}} T_{A,\tilde{\xi}}(v) \quad (6.9)$$

where $T_{A,\xi}$ is defined in (6.1).

This requires, for each p , to compute

$$\xi(p) = \operatorname{argmin}_{|\theta| > \tau} \|A_p - A(p)\chi_\theta\| \quad \text{where} \quad A_p = \{|\langle v, \psi_{p,k} \rangle|_\varepsilon\}_k. \quad (6.10)$$

To speed up the computation, and similarly to Section 4.2, we compute an approximate solution

$$\xi(p) = \Delta_\xi \operatorname{argmax}_{k > \frac{\tau}{|\Delta_\xi|}} |\Psi v[p, k]|. \quad (6.11)$$

Minimization with respect to u . If ξ and v are fixed, the minimization on u is similar to (4.12) by setting $y = f - \Phi(v^{[h]})$, and it can therefore be solved using the algorithm described in Section 4.3.

Minimization with respect to v . When ξ and u are fixed, the minimization on v is a smooth non-convex problem. Defining $y = f - \Phi(u)$, one minimizes

$$E(v) = \frac{1}{2} \|y - \Phi(v^{[h]})\|_{L^2}^2 + \mu T_{A,\xi}(v) \quad (6.12)$$

For a given step size $\nu > 0$, we use a gradient descent scheme

$$v^{(l+1)} = v^{(l)} - \nu(\mathcal{G}_1 + \mu\mathcal{G}_2) \quad (6.13)$$

where $\mathcal{G}_1$ is the gradient of $\frac{1}{2} \|y - \Phi(v^{[h]})\|_{L^2}^2$ with respect to v and $\mathcal{G}_2$ is the gradient of $T_{A,\xi}(v)$ with respect to v .

These two gradients are computed as

$$\mathcal{G}_1 = -\left(y - \Phi(v^{[h]})\right) h'(v) \quad \text{and} \quad \mathcal{G}_2 = \Psi^* c \quad (6.14)$$

where c is a set of coefficients defined by

$$c[p, k] = 2 \left(1 - \frac{A(p) \chi_{\xi(p)}(k)}{|\langle v, \psi_{p,k} \rangle|_\varepsilon}\right) \langle v, \psi_{p,k} \rangle. \quad (6.15)$$

and where the dual operator Ψ^* is defined in (3.5).

If the gradient step size ν is small enough, the iterates $v^{(l)}$ converge to a local minimum of $E(v)$ defined in (6.12).

6.4. Numerical Experiments. Similarly to Section 5.2, we define the cartoon functional $J(u)$ using (2.7) with $\rho = 0.75$. As far as the computational cost is concerned, the minimizing steps on u and ξ are equivalent to the one studied in section 4. The minimizing step on v is yet different and the conjugate gradient descent is replaced by a standard gradient descent which is generally slower. For an image of size 512×512 , the processing time is about 30 minutes.

Figure 6.5 shows an example of inpainting reconstruction of the image f_0 from Figure 5.4 degraded by a large hole Ω . We used the same parameters as in Figure 5.4. To obtain a “crenel” profile similar to the profile of the fingerprint texture, we chose $h = h_{0.3,0}$.

Figure 6.6 compares our result with a copy-paste based method [41], even though it is beyond the scope of this paper to make such comparisons, since the focus is on

FIG. 6.5. (a) the image y to inpaint (the missing part is in black), the original image is f_0 in Fig. 5.4 ; (b) the reconstruction $u + v^{[h]}$ using our non convex method ; (c) the inpainted geometric component u , (d) the sinusoidal approximation v of the texture component, (e) the inpainted texture component $v^{[h]}$ after the change of contrast.

FIG. 6.6. (a) the image to inpaint (the missing part is in black), (b) the reconstruction using our non convex method, (c) result of the “patchworks” method [41].

energy minimization based methods. Nevertheless, this figure shows that, as expected, a copy-paste based method fails to reconstruct properly the orientation of the locally parallel texture.

Figure 6.7 shows the inpainting of a desert picture² of size 512×512 . We use our

²This picture is extracted from a photography of Cesar Fernandez and we thank him for allowing us to use it.

non-convex method with parameters $\lambda = 0.3$, $\mu = 0.01$, $q = 48$ and $\Delta_x = 16$. We choose as rendering function $h = h_{2,0}$ which gives a texture profile similar to the one present in the original picture (see Figure 6.4). Since the profile of the real texture is more sophisticated than the chosen rendering function, the difference between the inside and the outside of the mask is visually noticeable. However the reconstruction of the texture orientation inside the mask is coherent with the known part of the image.

FIG. 6.7. First row: (a) f_0 the original image. Second row: (b) y the image to inpaint (the missing part is in black), (c) the reconstruction $u + v^{[h]}$ using our non convex method, (d) the estimated amplitudes A . Third row: (e) the inpainted geometric component u , (f) the sinusoidal approximation v of the texture component, (g) the inpainted texture component $v^{[h]}$ after the change of contrast.

Figure 6.8 compares our result with a copy-paste based method [41], even though it is beyond the scope of this paper to make such comparisons, since the focus is on

energy minimization based methods. Nevertheless, this figure shows that, as expected, a copy-paste based method fails to reconstruct properly the orientation of the locally parallel texture.

FIG. 6.8. (a) the image to inpaint (the missing part is in black), (b) the reconstruction using our non convex method, (c) result of the “patchworks” method [41].

7. Conclusion. This paper presented a new adaptive framework for locally parallel texture modeling. Two new adaptive texture models well-suited for locally parallel oscillating patterns were studied. The use of these adaptive models provide satisfactory results both in decomposition and inpainting for images which contain oriented textures. Adaptivity is indeed crucial for this kind of images where the texture is anisotropic, since it allows one to take into account the texture geometry.

The adaptation to the texture geometry is obtained by optimizing an orientation field ξ which is computed iteratively by our algorithm. To inpaint large missing regions, we increase this adaptivity by considering also an amplitude field and a rendering function. Inpainting larges holes, which boils down to the synthesis of new information, is hence obtained by switching from a convex regularization to a non-convex boosting term controlled by an amplitude field.

Acknowledgment:. The authors would like to thank Yann Gousseau for providing them with the inpainting results of the patch-based inpainting algorithms.

Appendix A. Algorithm to Compute the Proximity Operator.

In the following we use the discrete gradient operator ∇ defined in (2.1) and denote the curvelet decomposition operator as

$$\mathcal{D}g = \{\langle g, c_m \rangle\}_m$$

where $c_m = c_{j,l,k}$ is a curvelet atom. As first noted by Bect et al. in [9], the cartoon model defined in Section 2 can be written as

$$J(g) = \|Qg\|_1$$

where the linear operator Q is defined as

$$Q : \begin{cases} \mathbb{R}^N & \longrightarrow \mathbb{R}^{N \times 2} \times \mathbb{R}^P \\ g & \longmapsto \begin{pmatrix} \rho \nabla g \\ (1 - \rho) \mathcal{D}g \end{pmatrix} \end{cases} \quad (\text{A.1})$$

where P is the number of curvelet atoms. The ℓ^1 norm of an element $p = (z, w) \in \mathbb{R}^{N \times 2} \times \mathbb{R}^P$ is defined as

$$\|(z, w)\|_1 = \|z\|_1 + \|w\|_1 = \sum_{1 \leq i \leq N} \sqrt{(z_i^1)^2 + (z_i^2)^2} + \sum_{1 \leq i \leq P} |w_i|. \quad (\text{A.2})$$

The dual operator Q^* is then given by

$$Q^* : \begin{cases} \mathbb{R}^{N \times 2} \times \mathbb{R}^P & \longrightarrow \mathbb{R}^N \\ \begin{pmatrix} z \\ w \end{pmatrix} & \longmapsto -\rho \operatorname{div} z + (1 - \rho) \mathcal{D}^* w \end{cases} \quad (\text{A.3})$$

where the divergence of a vector field $p \in \mathbb{R}^{N \times 2}$ is computed as described in [17]

$$\operatorname{div}(p) = -\nabla^* p = \partial_x^* p_1 + \partial_y^* p_2,$$

where

$$\begin{aligned} \partial_x^* f[i, j] &= \begin{cases} f[i, j] - f[i-1, j] & \text{if } 0 < i < n-1, \\ f[i, j] & \text{if } i = 0, \\ -f[i-1, j] & \text{if } i = n-1, \end{cases} \\ \partial_y^* f[i, j] &= \begin{cases} f[i, j] - f[i, j-1] & \text{if } 0 < j < n-1, \\ f[i, j] & \text{if } j = 0, \\ -f[i, j-1] & \text{if } j = n-1. \end{cases} \end{aligned}$$

Chambolle shows in [17] that for such a functional $J(g) = \|Qg\|_1$, the proximal operator satisfies

$$\operatorname{prox}_{\omega J}(g) = g - \omega Q^*(p^*) \quad (\text{A.4})$$

where $p^* = (z^*, w^*) \in \mathbb{R}^{N \times 2} \times \mathbb{R}^P$ is a solution of

$$\min_{\|p\|_\infty \leq 1} \|Q^* p - g/\omega\|^2 \quad (\text{A.5})$$

where the ℓ^∞ norm is defined as

$$\|(z, w)\|_\infty = \max \left(\max_{1 \leq i \leq N} \sqrt{(z_i^1)^2 + (z_i^2)^2}, \max_{1 \leq i \leq P} |w_i| \right). \quad (\text{A.6})$$

For the computation of p^* , Chambolle proposes a fixed point algorithm that can be replaced by a projected gradient descent. Starting from $p^{(0)} = (0, 0) \in \mathbb{R}^{N \times 2} \times \mathbb{R}^P$, $\ell = 0$, one then iterates between a gradient descent step to minimize $\|Q^* p - g/\omega\|^2$

$$\tilde{p}^{(\ell)} = p^{(\ell)} - \nu Q(Q^* p^{(\ell)} - g/\omega) \quad (\text{A.7})$$

and a projection onto the constraints $\{p \mid \|p\|_\infty \leq 1\}$

$$\forall i, \quad p_i^{(\ell+1)} = \frac{\tilde{p}_i^{(\ell)}}{\max(1, \|\tilde{p}_i^{(\ell)}\|)}. \quad (\text{A.8})$$

If the gradient step size satisfies $\nu < 2/\|QQ^*\|$, then $g - \omega Q^* p^{(\ell)}$ converges to $\operatorname{prox}_{\omega J}(g)$.

REFERENCES

- [1] P. ARIAS, V. CASELLES, AND G. SAPIRO, *A variational framework for non-local image inpainting*, in Energy Minimization Methods in Computer Vision and Pattern Recognition, Springer, 2009, pp. 345–358.
- [2] J.-F. AUJOL, *Some first-order algorithms for total variation based image restoration*, J. Math. Imaging Vis., 34-3 (2009), pp. 307–327.
- [3] J.-F. AUJOL, G. AUBERT, L. BLANC-FÉRAUD, AND A. CHAMBOLLE, *Image decomposition into a bounded variation component and an oscillating component*, Journal of Mathematical Imaging and Vision, 22 (2005), pp. 71–88.
- [4] J.-F. AUJOL AND A. CHAMBOLLE, *Dual norms and image decomposition models*, International Journal of Computer Vision, 63 (2005), pp. 85–104.
- [5] J.-F. AUJOL AND G. GILBOA, *Constrained and SNR-based solutions for tv-hilbert space image denoising*, J. Math. Imaging Vis., 26 (2006), pp. 217–237.
- [6] J.-F. AUJOL, G. GILBOA, T. CHAN, AND S. OSHER, *Structure-texture image decomposition—modeling, algorithms, and parameter selection*, International Journal of Computer Vision, 67 (2006), pp. 111–136.
- [7] J.-F. AUJOL AND S. H. KANG, *Color image decomposition and restoration*, Journal of Visual Communication and Image Representation, 17 (2006), pp. 916–928.
- [8] J.-F. AUJOL, S. LADJAL, AND S. MASNOU, *Exemplar-based inpainting from a variational point of view*, SIAM Journal on Mathematical Analysis, 42 (2010), pp. 1246–1285.
- [9] J. BECT, L. BLANC FÉRAUD, G. AUBERT, AND A. CHAMBOLLE, *A ℓ_1 -unified variational framework for image restoration*, in Proc. of ECCV04, Springer-Verlag, 2004, pp. Vol IV: 1–13.
- [10] M. BERTALMIO, G. SAPIRO, V. CASELLES, AND C. BALLESTER, *Image inpainting*, in Siggraph 2000, 2000, pp. 417–424.
- [11] M. BERTALMIO, L. VESE, G. SAPIRO, AND S. OSHER, *Simultaneous structure and texture image inpainting*, IEEE Transactions on Image Processing, 12 (2003), pp. 882–889.
- [12] X. BRESSON AND T. CHAN, *Fast minimization of the vectorial total variation norm and applications to color image processing*, Inverse Problems and Imaging, 2 (2008), pp. 455–484.
- [13] A. BUADES, B. COLL, AND J. M. MOREL, *On image denoising methods*, SIAM Multiscale Modeling and Simulation, To Appear, (2005).
- [14] E. CANDÈS AND D. DONOHO, *New tight frames of curvelets and optimal representations of objects with piecewise C^2 singularities*, Comm. Pure Appl. Math., 57 (2004), pp. 219–266.
- [15] E. J. CANDÈS, L. DEMANET, D. L. DONOHO, AND L. YING, *Fast discrete curvelet transforms*, Multiscale Model. Simul., 5 (2005), pp. 861–899.
- [16] F. CAO, Y. GOUSSEAU, S. MASNOU, AND P. PÉREZ, *Geometrically guided exemplar-based inpainting*, Preprint hal-00392018, (2009).
- [17] A. CHAMBOLLE, *An algorithm for total variation minimization and applications*, J. Math. Imaging Vis., 20 (2004), pp. 89–97.
- [18] A. CHAMBOLLE AND T. POCK, *A first-order primal-dual algorithm for convex problems with applications to imaging*, Hal-00490826, <http://hal.archives-ouvertes.fr/hal-00490826.v1/>, (2010).
- [19] T. CHAN AND J. SHEN, *Mathematical models for local nontexture inpaintings*, SIAM J. Appl. Math., 62 (2002), pp. 1019–1043.
- [20] TONY F. CHAN, SUNG HA KANG, AND JIANHONG SHEN, *Total variation denoising and enhancement of color images based on the CB and HSV color models*, Journal of Visual Communication and Image Representation, 12 (2000), p. 2001.
- [21] P. L. COMBETTES AND V. R. WAJS, *Signal recovery by proximal forward-backward splitting*, Multiscale Modeling & Simulation, 4 (2005), pp. 1168–1200.
- [22] A. CRIMINISI, P. PEREZ, AND K. TOYAMA, *Region filling and object removal by exemplar-based image inpainting*, IEEE Trans. Image Processing, 13 (2004), pp. 1200–1212.
- [23] K. DABOV, A. FOI, V. KATKOVNIK, AND K. EGIAZARIAN, *Image denoising by sparse 3d transform-domain collaborative filtering*, IEEE Trans. Image Proc., 16 (2007).
- [24] I. DAUBECHIES AND G. TESCKE, *Variational image restoration by means of wavelets: simultaneous decomposition, deblurring and denoising*, Appl. Comp. Harm. Anal., 19 (2005), pp. 1–16.
- [25] I. DAUBECHIES, G. TESCKE, AND L. VESE, *On some iterative concepts for image restoration*, Advances in Imaging and Electronics, 150 (2008), pp. 2–52.
- [26] D. DONOHO AND I. JOHNSTONE, *Ideal spatial adaptation via wavelet shrinkage*, Biometrika, 81 (1994), pp. 425–455.
- [27] I. DRORI, D. COHEN-OR, AND H. YESHURUN, *Fragment-based image completion*, ACM Transactions on Graphics, 22 (2003), pp. 303–312.
- [28] V. DUVAL, J.-F. AUJOL, AND Y. GOUSSEAU, *The TVL1 model: a geometric point of view*, SIAM Journal on Multiscale Modeling and Simulation, 8 (2009), pp. 154–189.

- [29] V. DUVAL, J.-F. AUJOL, AND L. A. VESE, *Mathematical modeling of textures : application to color image decomposition with a projected gradient algorithm*, Journal of Mathematical Imaging and Vision, 37 (2010), pp. 232–248.
- [30] A. A. EFROS AND T. K. LEUNG, *Texture synthesis by non-parametric sampling*, in ICCV '99, 1999, p. 1033.
- [31] M. ELAD, J. STARCK, P. QUERRE, AND D. DONOHO, *Simultaneous cartoon and texture image inpainting using morphological component analysis (MCA)*, Applied and Computational Harmonic Analysis, 19 (2005), pp. 340–358.
- [32] M.J. FADILLI, J.-L. STARCK, AND F. MURTAGH, *Inpainting and zooming using sparse representations*, The Computer Journal, (2006). submitted.
- [33] Y. KIM AND L. VESE, *Image recovery using functions of bounded variation and sobolev spaces of negative differentiability*, Inverse problems and Imaging, 3 (2009), pp. 43–68.
- [34] N. KOMODAKIS, *Image completion using global optimization*, in CVPR, IEEE Computer Society, 2006, pp. 442–452.
- [35] V. KWATRA, I. ESSA, A. BOBICK, AND N. KWATRA, *Texture optimization for example-based synthesis*, ACM Transactions on Graphics, 24 (2005), pp. 795–802.
- [36] E. LE, PENNEC, AND S. MALLAT, *Bandelet image approximation and compression*, SIAM Journal of Multiscale Modeling and Simulation, 4 (2005), p. 2005.
- [37] S. MASNOU, *Disocclusion: a variational approach using level lines*, IEEE Trans. Image Processing, 11 (2002), pp. 68–76.
- [38] Y. MEYER, *Oscillating Patterns in Image Processing and Nonlinear Evolution Equations*, American Mathematical Society, Boston, MA, USA, 2001.
- [39] Y. NESTEROV, *Smooth minimization of non-smooth functions*, Math. Program., 103 (2005), pp. 127–152.
- [40] S. OSHER, A. SOLÉ, AND L. VESE, *Image decomposition and restoration using total variation minimization and the H^{-1} norm*, Multiscale Modeling & Simulation, 1 (2003), pp. 349–370.
- [41] P. PÉREZ, M. GANGNET, AND A. BLAKE, *Patchworks: Example-based region tiling for image editing*, tech. report, Microsoft Research, MSR-TR-2004-04, 2004.
- [42] G. PEYRÉ, *Texture Synthesis with Grouplets*, IEEE Transactions on Pattern Analysis and Machine Intelligence, (2009), p. to appear.
- [43] G. PEYRÉ AND S. MALLAT, *A Review of Bandelet Methods for Geometrical Image Representation*, Numerical Algorithms, 44 (2007), pp. 205–234.
- [44] J. PORTILLA, V. STRELA, M.J. WAINWRIGHT, AND SIMONCELLI E.P., *Image denoising using scale mixtures of Gaussians in the wavelet domain*, IEEE Trans Image Processing, 12 (2003), pp. 1338–1351.
- [45] L. I. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, Phys. D, 60 (1992), pp. 259–268.
- [46] J. SHEN, S. HA KANG, AND T. CHAN, *Euler's elastica and curvature-based inpainting*, SIAM Journal of Applied Mathematics, 63 (2003), pp. 564–592.
- [47] J.-L. STARCK, M. ELAD, AND D.L. DONOHO, *Redundant multiscale transforms and their application for morphological component analysis*, Advances in Imaging and Electron Physics, 132 (2004).
- [48] J. SUN, L. YUAN, J. JIA, AND H.-Y. SHUM, *Image completion with structure propagation*, in SIGGRAPH '05, 2005, pp. 861–868.
- [49] J. A. TROPP, *Algorithms for simultaneous sparse approximation. part II: Convex relaxation*, Signal Processing, special issue "Sparse approximations in signal and image processing", 86 (2006), pp. 589–602.
- [50] P. TSENG, *Convergence of a block coordinate descent method for nondifferentiable minimization*, Journal of Optimization Theory and Applications, 109 (2001), pp. 475–494.
- [51] L. VESE AND S. OSHER, *Modeling textures with total variation minimization and oscillating patterns in image processing*, J. Sci. Comput., 19 (2003), pp. 553–572.
- [52] L.-Y. WEI AND M. LEVOY, *Fast texture synthesis using tree-structured vector quantization*, in SIGGRAPH '00, 2000, pp. 479–488.