Variational Texture Synthesis with Sparsity and Spectrum Constraints

Guillaume Tartavel · Yann Gousseau · Gabriel Peyré

Received: date / Accepted: date

Abstract This paper introduces a new approach for texture synthesis. We propose a unified framework that both imposes first order statistical constraints on the use of atoms from an adaptive dictionary, as well as second order constraints on pixel values. This is achieved thanks to a variational approach, the minimization of which yields local extrema, each one being a possible texture synthesis. On the one hand, the adaptive dictionary is created using a sparse image representation rationale, and a global constraint is imposed on the maximal number of use of each atom from this dictionary. On the other hand, a constraint on second order pixel statistics is achieved through the power spectrum of images. In a sense, this method reconciles the first and second theory proposed by B. Julesz for texture discrimination. An advantage of the proposed method is its ability to truly synthesize textures, without verbatim copy of small pieces from the exemplar. In an extensive experimental section, we show that the resulting synthesis achieves state of the art results, both for structured and small scale textures.

Keywords texture synthesis · variational methods · statistical constraints · Fourier spectrum · sparse decomposition ·

dictionary learning · random phase textures · Gaussian textures

1 Introduction

1.1 Texture Synthesis

Texture synthesis is the process of generating new texture images from a given image sample. Being able to perform such a synthesis is of practical interest both for computer graphics, where it is needed to give a realistic aspect to 3D models, or for image restoration, where it may help the reconstruction of missing parts through inpainting.

One of the most difficult aspect of texture synthesis is the ability to deal with the wide variety of texture images. In particular, it is relatively difficult to synthesize textures at all observation scales using a given model. Before being more specific in the next section, we can roughly split texture synthesis methods in two groups. On the one hand, non-parametric patch-based methods, initiated by the seminal works from Efros and Leung [8] and Wei and Levoy [36], have proven their efficiency to synthesize highly structured textures, such as brick walls or pebbles seen at a close distance. On the other hand, methods relying on statistical constraints, initiated by the work of Heeger and Bergen [14], provide efficient methods for the synthesis of textures with small scales oscillations, or micro-textures, such as sand or stone.

In this work, we propose to rely on a variational approach to build a synthesis method that allows at the same time for the reconstruction of geometric textures, thanks to the use of sparse, patch-based dictionary, and the reconstruction of small scale oscillations, thanks to constraints on the spectrum of images. Before presenting in more details the contributions of this paper, and in order to explain both its

Guillaume TARTAVEL
Telecom ParisTech / LTCI CNRS
46 rue Barrault – F-75634 Paris Cedex 13 – France
Tel.: +33 1 45 81 73 91
E-mail: tartavel@telecom-paristech.fr

Yann GOUSSEAU
Telecom ParisTech / LTCI CNRS
46 rue Barrault – F-75634 Paris Cedex 13 – France
Tel.: +33 1 45 81 78 41
E-mail: gousseau@telecom-paristech.fr

Gabriel PEYRÉ
CNRS / Ceremade, Université Paris-Dauphine
Pl. Maréchal De Lattre De Tassigny – 75775 Paris Cedex 16 – France
Tel.: +33 1 44 05 48 71
E-mail: peyre@ceremade.dauphine.fr

inspiration and its originality, we start with a short presentation of previous works in this field.

1.2 Previous Works

1.2.1 Patch-Based Approaches

The first methods yielding realistic texture synthesis rely on a Markovian hypothesis: the distribution of a pixel given the rest of the texture only depends on the values in a (small) neighborhood. These methods are mostly parametric, following the work in [5]. The law of a pixel given its neighborhood is chosen *a priori* and its parameters can be learned from an exemplar using a maximum-likelihood estimation.

Non-parametric approaches [8, 36] were introduced later to handle different kind of textures without designing the law of each of them by hand. For a given pixel, its value is sampled directly from an exemplar image, by seeking pixels whose square neighborhood (called a *patch*) is similar to the one of the pixel to be synthesized. This approach was a clear breakthrough in the field of texture synthesis, producing results that are visually almost perfect on difficult and structured examples. It was followed by a large body of works and many extensions were proposed, to perform real-time [21], dynamic, or volumetric texture synthesis, for instance. In particular, it is common to sample a small square area (a patch) at a time, instead of a single pixel [7]. The reader may refer to the state-of-the-art [35] for an overview of these approaches applied in different contexts.

One common drawback of these approaches, however, is that the resulting texture is often a plain juxtaposition of small pieces from the exemplar, even when sampling pixels one at a time [1]. In this sense they have a limited innovation capacity. An efficient way to avoid this drawback is to perform the synthesis from a dictionary of patches that is learned from the exemplar. Sparse decomposition methods such as [9] are classically applied to image denoising or enhancement. They rely on the assumption that each patch of an image can be decomposed as a sum of a small number of elements from a dictionary. The dictionary is usually learned from the image to be restored. These approaches were introduced in [24] as an efficient coding of natural images, inspired by the human vision system. More efficient algorithms for dictionary learning were proposed later: for instance the Method of Optimal Direction (MOD) [10], and the K-SVD algorithm [2]. Recently, one author of the present paper has shown that this framework is well suited to texture modeling [26]. In particular, sparse dictionary learning has the remarkable ability to adaptively extract basic elements, or *textons*, from a texture. In [26], it is shown that new texture patches can be created from a learned dictionary, simply by imposing a sparsity constraint on their use of the dictionary atoms.

1.2.2 Statistical Approaches

Another set of approaches is based on statistical constraints. A set of statistics is identified and the synthesis process consists in generating a new image whose selected statistics match those of the original texture. The basic principle of the approach is in agreement with early works on texture discrimination [16], that spotted statistical constraints as a fruitful tool to investigate the human perception of textures.

In 1995, Heeger and Bergen [14] also proposed to rely on some statistical constraints to perform texture synthesis. Starting from an exemplar image, a new image is generated by imposing the marginal of wavelet coefficients, separately at each scale and orientation, as well as the image color distribution. The resulting method, although limited to relatively weakly structured textures, has the ability to produce realistic results in a computationally efficient manner. Later, it has been proposed [29] to impose some second order statistics on wavelet coefficients. More precisely, the constraints are based on the correlations between coefficients of atoms that are neighbor, either at a given scale or across scales. This allows for a better preservation of structured textures. To the best of our knowledge, this relatively old method is the one permitting to synthesize the largest class of textures, without simply copy-pasting pieces from an exemplar.

In a different direction, one may synthesize a texture by imposing constraints on its Fourier transform. In [12], it is shown that an efficient synthesis method is achieved by imposing the spectrum of synthesized images, through *random phase* textures. The resulting synthesis is fast and reliable, but limited to non-structured textures with small scales oscillations. Several works have extended this approach, either simplifying the reference from which the synthesis is performed [6, 13] or developing video-texture synthesis methods [38].

In the field of computer graphics, procedural noises are widely used to yield a realistic aspect for materials. Contrary to other methods, a procedural noise is generated on-the-fly in any point of $\mathbb{R}^n$, and thus has a very low memory requirement and can be evaluated at any scales with any resolution. The Perlin noise [25] is the most classical example of such a noise. It is obtained as a linear combination of colored noises at different scales, resulting in a multi-scale noise whose spectrum is controlled by the weights of the combination. Sparse convolutive noise is another class of procedural noise, defined as the convolution of a compact kernel with a sparse convolution map. The Gabor noise is a particular case, using Gabor kernels: the distribution of scales and frequencies of the kernels can be either imposed by hand [20] or learned from a sample texture [13]. The state-of-the-art [19] provides a complete review of procedural noise functions.

1.2.3 Variational Formulation for Texture Synthesis

In this work, we advocate a generic variational approach to texture synthesis, as presented in detail in Sect. 2. This is not the first such approach in the context of texture manipulations. Previous variational approaches aim at transferring the texture of an image onto the content of another image. The functional to be minimized take into account both the fidelity to the texture image and the fidelity to the content image. The method presented in [18] optimizes an image so that all its patches look similar to a texture given as an example. The transfer method [31] generates an image A' following the description: “ A' must be to A what B' is to B ”.

1.3 Contributions of the Paper

In this paper, we introduce a generic texture synthesis method that is both suited to highly structured textures and textures that exhibit small scale oscillations.

First, we rely on adaptive sparse dictionary for the synthesis, following the ideas presented in [26]. As in this original approach, we impose a sparsity constraint on a decomposition of the patches into a dictionary. In addition, we also impose that atoms from the dictionary are used in the same proportions as in the original texture samples. This allows for a faithful reproduction of the structured parts (edges, corners, etc.) of the input texture. Second, we impose spectrum constraints on the synthesis, in the manner of [12]. This allows for the reproduction of high frequency components of the texture, but also, more surprisingly, for the low frequency regularity of the texture, as demonstrated in the experimental section. We also impose a constraint on the global color content of the image, that has a strong visual impact on the results. In order to combine the strength of the adaptive patch dictionary to reproduce geometry and the fidelity of frequency reproduction offered by spectrum constraints, we rely on a variational approach. Each constraint is enforced by controlling the distance of the synthesis to a set of compatible images. The minimization of the corresponding energy is performed using an averaged projections method. The approach is compared to the more classical alternating projection method proposed in [29], and the benefit of the proposed constraints is demonstrated. The resulting synthesis approach achieves, to the best of our knowledge, the best results obtained so far that truly generate a new texture without performing copy-pasting of the input. Another interesting asset of the proposed approach is that it only relies on first order statistical constraints between atoms from a dictionary. On the contrary, methods relying on a fixed dictionary necessitates second order statistics to synthesize structured textures [29]. This somehow accredits the second psychophysical theory proposed by B. Julesz [16], stating that texture discrimination is related to first order statis-

tics between *textons*. Observe nevertheless that the synthesis also relies on second order statistics between pixel values (through the power spectrum), therefore establishing a link between Julesz first [15] and second [16] theory for texture discrimination.

The plan of the paper is as follows. We define in Sect. 2 the energy function whose minimization leads our synthesis method and present in details the three constraints we impose for the synthesis: on the histogram, on the spectrum, and on the sparse decomposition of image patches. We then propose in Sect. 3 an algorithm to find a critical point of this energy function. Numerical results are shown, parameters are discussed and comparisons with the state of the art are provided in Sect. 4. Preliminary results of our work were presented in the conference paper [33]. More results can be found at:

<http://perso.enst.fr/~tartavel/research/jmiv14.html>,

and the Matlab code used for the experiments is publicly available at:

https://bitbucket.org/gtartavel/variational_synthesis.

1.4 Notations

A matrix $A = (A_j^i)_{i,j}$ is made of columns A_j and rows A^i . The ℓ^0 pseudo-norm of a vector b is $\|b\|_0 = \#\{i \mid b_i \neq 0\}$.

The indicator function $\iota_{\mathcal{C}}$ of a set $\mathcal{C}$ is equal to 0 on $\mathcal{C}$ and $+\infty$ outside. The distance of a point $x \in \mathbb{R}^m$ to a compact set $\mathcal{C}$ is $\mathcal{D}(x, \mathcal{C}) = \min_{y \in \mathcal{C}} \|x - y\|$ and a projection $\mathcal{P}_{\mathcal{C}}(x)$ of x on $\mathcal{C}$ is a minimizer of this distance.

The orthogonal discrete Fourier transform of an image u defined on a rectangular domain $\Omega \subset \mathbb{Z}^2$ of size $M_1 \times M_2$ is the image $\hat{u}$ made of the coefficients

$$\hat{u}(m) = \frac{1}{\sqrt{M_1 M_2}} \sum_{x \in \Omega} u(x) \exp \left\{ -2i\pi \left(\frac{x_1 m_1}{M_1} + \frac{x_2 m_2}{M_2} \right) \right\}. \quad (1)$$

The patches of size $\tau \times \tau$ in u are defined as

$$p_k = (u(x_k + t))_t \quad \text{for } t \in \{0, \dots, \tau - 1\}^2. \quad (2)$$

Their locations x_k lie on a regular grid $x_k \in \Delta \mathbb{Z}^2$ of step $\Delta > 0$. We denote by $\Pi(u) = (p_k)_k \in \mathbb{R}^{L \times K}$ the matrix, made of the K patches of u , where $L = d\tau^2$ is their dimension, with $d = 1$ for gray images and $d = 3$ for color. The adjoint operator of Π is denoted by Π^* .

2 Variational Formulation

As explained above, we define an energy functional E to assert the quality of a synthesis. The synthesis process then consists in finding local minima of this functional.

The functional E is designed to account for three constraints: one on the first order statistics on the use of atoms

from an adaptive, sparse dictionary, one on the spectrum of images, that is on second order pixel values statistics, and one on color histograms, that is on first order pixel values statistics. This is achieved through the distances to 3 sets: the set $\mathcal{C}_p$ of patches being sparse in a dictionary D_0 learned from u_0 and whose atoms frequency matches the ones of the decomposition of u_0 in this dictionary, the set $\mathcal{C}_s$ of images whose Fourier frequencies are the same as in u_0 , and the set $\mathcal{C}_h$ of images whose color histogram is the same as u_0 .

We therefore define $E(u)$ equals to

$$\frac{\alpha_p}{2} \mathcal{D}^2(\Pi(u), \mathcal{C}_p) + \frac{\alpha_s}{2} \mathcal{D}^2(u, \mathcal{C}_s) + \frac{\alpha_h}{2} \mathcal{D}^2(u, \mathcal{C}_h), \quad (3)$$

where $(\alpha_p, \alpha_s, \alpha_h)$ are weighting terms. The adjoint operator Π^* , involved in the optimization process, reverberates the constraints on the patches to the image. The sets $(\mathcal{C}_p, \mathcal{C}_s, \mathcal{C}_h)$ are detailed hereafter. The projection operators $\mathcal{P}_\mathcal{C}$ on these sets are also defined and presented in the following paragraphs since they are involved in the iterative optimization process detailed in Sect. 3. We start with the most straightforward histogram constraint, then proceed with the Fourier constraint, and eventually present the most involved, dictionary-based constraint.

2.1 Histogram Constraint

The constraint $\mathcal{C}_h$ imposes the histogram to match the one of the exemplar. In other words, we want to preserve the gray-level or color distribution of the texture. This requirement is common in texture synthesis: it is used in [14] or [26] to ensure that the synthesized image has a similar dynamic and color distribution as the original image.

2.1.1 Definition

We define the set of images whose color histogram is the same as the one of u_0 as:

$$\mathcal{C}_h = u_0(\Sigma_{|\Omega|}) = \{u_0 \circ \sigma \mid \sigma \in \Sigma_{|\Omega|}\}, \quad (4)$$

where $\Sigma_{|\Omega|}$ is the set of permutations of Ω . Two images have the same histogram if they only differ by a permutation of their pixels.

2.1.2 Projection

The projection on $\mathcal{C}_h$ is called ‘‘histogram transfer’’ [28]. In the case of grey-level images ($d = 1$), the projection $u_0 \circ \sigma^*$ of an image u on the set $\mathcal{C}_h$ is given by

$$\sigma^* = \sigma_{u_0} \circ \sigma_u^{-1} \quad (5)$$

where we denote by σ_v the permutation sorting the pixel values of an image v :

$$v_{\sigma_v(0)} \leq \dots \leq v_{\sigma_v(i)} \leq v_{\sigma_v(i+1)} \leq \dots$$

When u and u_0 have different numbers of pixels, the sorted sequence of values $u_0 \circ \sigma_{u_0}$ is oversampled using a nearest-neighbor approximation. Note that a linear or cubic interpolation is not adapted since it creates new gray level values. This may create artifacts if the histogram to be interpolated has a gap (as for example with a mixture of Gaussians).

The color case is more difficult: exact algorithms like the Hungarian method [17] are polynomial and are not tractable. An iterative approximation has been proposed in [30] and can be applied to image histograms. We use a rougher but faster approximation to the optimal transport thanks to a change of color space. For $i = 1 \dots 3$, let $c^{(i)} \in \mathbb{R}^3$ be the colors obtained by a PCA on the RGB components of u_0 . Such a transform is adequate for our purpose. We denote respectively by $u_0^{(i)}$ and $u^{(i)}$ the components of u_0 and u in this color basis. We approximate the solution by solving 3 independent histogram transfers between each $u_0^{(i)}$ and $u^{(i)}$.

2.2 Fourier Spectrum Constraint

The correlation between the pixels of the image u_0 is given by its power spectrum $|\hat{u}_0|^2$, where $\hat{u}_0$ denotes the discrete Fourier transform of u_0 , defined in (1). As explained in the introduction, correlations between pixel values are an important characteristic of textures [16, 12], which motivates their inclusion in the functional (3).

2.2.1 Definition

The set of images whose spectrum is the same as u_0 is

$$\mathcal{C}_s = \left\{ u \mid \hat{u}(m) = e^{i\varphi(m)} \hat{u}_0(m) \quad \forall m \right\} \quad (6)$$

where the multiplication by $e^{i\varphi(m)}$ preserves the amplitude but changes the phase of the coefficient $\hat{u}_0(m)$ by $\varphi(m) \in [0, 2\pi[$. In the gray-level case ($d = 1$), the condition can be rewritten as $|\hat{u}(m)| = |\hat{u}_0(m)|$. In the color case ($d = 3$), it is important for color coherency that the multiplication of $\hat{u}_0(m) \in \mathbb{C}^3$ by $e^{i\varphi(m)}$ preserves the phase differences between the R, G, and B channels [12]. Therefore, texture synthesis is performed in [12] by preserving the spectrum of u_0 while generating phases $\varphi(m)$ randomly. It is equivalent to draw at random an element from $\mathcal{C}_s$.

To reduce the boundary artifacts caused by the inherent periodicity assumption of the discrete Fourier transform, we use the periodic decomposition of [23], so that the image u_0 is decomposed as the sum of a smooth image and a (circularly) periodic image. We drop the smooth component and only keep the periodic one, which is well suited to the discrete Fourier transform.

2.2.2 Projection

The projection of an image u on the set $\mathcal{C}_s$ consists in putting together the modulus of $\hat{u}_0$ and the phases of $\hat{u}$. The solution, as detailed in Appendix A.1, is the image u_s defined by

$$\hat{u}_s = \frac{\hat{u} \cdot \hat{u}_0}{|\hat{u} \cdot \hat{u}_0|} \hat{u}_0 \quad (7)$$

where the m variable has been dropped for the sake of readability and $x \cdot y = xy^*$ denotes the hermitian product of $\mathbb{C}^d$ ($d = 1$ for gray level images and $d = 3$ for color images).

Remark. As shown in the experimental section (Sect. 4), the spectrum constraint $\mathcal{C}_s$ handles a drawback of the patch sparsity constraint $\mathcal{C}_p$ described in the following paragraph: patches cannot handle the low frequencies of the image because they are too small to capture them. Patches also reduce high frequencies because of the sparse decomposition based on the ℓ^2 norm which promotes smooth approximations. The spectrum constraint is thus a good candidate to go along with the patch constraint.

2.3 Sparse Decomposition of the Patches

Decomposing the patches $\Pi(u)$ of an image u as a sparse linear combination $D_0 W$ of elements in an adaptive dictionary D_0 has first been proposed by [26] in the context of texture synthesis. As explained in the introduction, this framework enables us to efficiently synthesize highly structured aspects of textures, such as edges or regular patterns.

The set of patches $\Pi(u_0)$ of u_0 is first factorized into a dictionary D_0 and a sparse matrix W_0 of coefficients. In this first case, both the dictionary and the coefficients are unknown: the K-SVD algorithm [9] computes these two matrices D_0 and W_0 . Then, during the synthesis, any set of patches P is decomposed into this fixed dictionary D_0 and leads to a weight matrix W . The weights of this decomposition must satisfy two constraints: each patch can use only a few number of atoms from D_0 , and each atom of D_0 must be used as often in W as it is in W_0 .

2.3.1 Learning Stage

The learning consist in learning an over-complete dictionary D_0 from the exemplar u_0 . This stage is performed only once before the synthesis process.

A dictionary D_0 of N atoms is obtained by minimizing

$$(D_0, W_0) = \arg \min_{D, W} \|\Pi(u_0) - DW\|^2 \quad \text{s.t.} \quad \begin{cases} \|D_n\|_2 \leq 1 & \forall n = 1 \dots N, \\ \|W_k\|_0 \leq S & \forall k. \end{cases} \quad (8)$$

This non-convex combinatorial problem is NP-hard [34]. An approximated solution can be computed using the K-SVD [2] algorithm for instance, or the MOD [10] algorithm.

The number N of elements of the dictionary must be chosen so that D_0 has some redundancy. We choose, for instance, N to be approximately equal to $2\tau^2$ to have a redundancy factor of 2.

The sparsity constraint $\|W_k\|_0 \leq S$ imposes to each patch p_k to be approximated using at most S atoms from the dictionary. This constraint enforces the dictionary to represent the patterns and features of the texture, as shown in [26].

2.3.2 Definition of the Constraint

The dictionary D_0 being learned, we control both the sparsity of the decomposition and the number of occurrences of the atoms. As explained in the introduction, this is an algorithmic interpretation of the second Julesz principle, stating that first order statistics of *textons* (in the present context, of atoms) are important for the visual discrimination of textures. Precisely, the second constraint is that each atom d_n should be used at most F^n times. The bound $F^n = \frac{K}{K_0} \|W_0^n\|_0$ is learned from the decomposition W_0 of the exemplar u_0 and is normalized according to the numbers K and K_0 of patches in u and u_0 respectively (the exemplar and the synthesis may have different dimensions).

The resulting set of image patches satisfying these constraints is

$$\mathcal{C}_p = \{D_0 W \mid \|W_k\|_0 \leq S \text{ and } \|W^n\|_0 \leq F^n \quad \forall k, n\}. \quad (9)$$

Observe that this constraint results in a constraint on the number of non-zero coefficients both on the rows and the columns of the weight matrix W .

2.3.3 Projection

Computing the projection $D_0 W$ on $\mathcal{C}_p$ is a combinatorial problem quite similar to [9] which is known to be a NP-hard problem [34]. We approximate a solution of this problem using the greedy Algorithm 1. The result is then improved using a back-projection step, as detailed hereafter. This is inspired from the Matching Pursuit (MP) algorithm [22].

Algorithm 1: approximation of the projection on $\mathcal{C}_p$

Data: patches $P = \Pi(u)$, dictionary D_0 .

Input: sparsity S , # of occurrences F^n , iteration factor $\lambda \approx 1.5$.

Output: coefficients W .

Initialization: set $W = 0$ and $R = P$; compute $\Phi = D_0^T P$.

for $\ell = 1$ **to** λSK **do**

– find the best indices (k^*, n^*) defined by (13).

– compute the best weight w^* using (14).

– update W , R , and Φ using (12), (17), and (18).

Back-projection: update W by solving the linear systems (20).

2.3.4 Algorithm Details

The algorithm is iterative. The coefficients W_k^n are updated one by one until the constraints $\mathcal{C}_p$ detailed in (9) are saturated. At each step, the choice of the couple patch/atom (k, n) to be updated is optimal. The non-zero coefficients of the resulting weight matrix W are then refined (during the back-projection step).

Greedy algorithm. We denote by $E_{n,k}$ the elementary matrix whose only non-zero coefficient is 1 at position (n, k) . At step ℓ of Algorithm 1, the current estimation of W is denoted by $W^{(\ell)}$. Both a patch index k^* and an atom index n^* are chosen according to

$$(k^*, n^*, w^*) = \arg \min_{k, n, w} \left\| P - D_0(W^{(\ell)} + wE_{n,k}) \right\|^2 \quad (10)$$

under the constraint

$$W^{(\ell)} + wE_{n,k} \in \mathcal{C}_p. \quad (11)$$

The coefficient (k^*, n^*) of $W^{(\ell)}$ is updated while the others are left unchanged:

$$W^{(\ell+1)} = W^{(\ell)} + w^* E_{n^*, k^*}. \quad (12)$$

As shown in Appendix A.2, the solution of (10) is

$$(k^*, n^*) = \arg \max_{(k, n) \in \mathcal{J}_{W^{(\ell)}}} \left| \langle R_k^{(\ell)}, D_n \rangle \right| \quad (13)$$

$$w^* = \langle R_k^{(\ell)}, D_n \rangle \quad (14)$$

where $R_k^{(\ell)}$ is the k^{th} column of the residual $R^{(\ell)}$ defined at step ℓ by

$$R^{(\ell)} = P - D_0 W^{(\ell)}, \quad (15)$$

and the set $\mathcal{J}_{W^{(\ell)}}$ of available indices is

$$\mathcal{J}_W = \{(k, n) \mid \|W_k^{\sim n}\|_0 < S \text{ and } \|W_{\sim k}^n\|_0 < F^n\}, \quad (16)$$

where we denote $W_k^{\sim n} = (W_k^{n'})_{n' \neq n}$ and $W_{\sim k}^n = (W_{k'}^n)_{k' \neq k}$.

In the very particular case where D_0 is orthogonal, this algorithm converges in at most KS iterations because the resulting (k^*, n^*) are different at each iteration: the K constraints $\|W_k\|_0 \leq S$ are saturated after KS iterations. In general, the algorithm does not converge in a finite number of iterations: we decide to stop after λKS iterations anyway with $\lambda = 1.5$.

Efficient computation. To save computation time, residuals $R^{(\ell)}$ and inner products $\Phi^{(\ell)} = (\langle R_k^{(\ell)}, D_n \rangle)_{n,k}$ can be pre-computed and updated by

$$R^{(\ell+1)} = R^{(\ell)} - w^* D_0 E_{n^*, k^*} \quad (17)$$

$$\Phi^{(\ell+1)} = \Phi^{(\ell)} - w^* D_0^T D_0 E_{n^*, k^*}. \quad (18)$$

Using a max-heap search for (13) and assuming that $S \ll L \leq N \ll K$, the time complexity of this algorithm is in $\mathcal{O}(KN(L + \lambda S \log K))$. Note as a reference that the computation of all the inner products $\langle P_k, D_n \rangle$ already requires $\Theta(KNL)$ operations.

Back-projection. The iterative greedy algorithm described above provides weights $\tilde{W}$ to decompose P as $D_0 \tilde{W} \in \mathcal{C}_p$. A drawback of the greedy approach is that the weights are estimated one by one: the approximation error may be quite large. The back-projection step —as introduced in [22]— improves this solution by refining all the non-zeros coefficients of $\tilde{W}$ while the constraints from $\mathcal{C}_p$ are still satisfied.

The support of $\tilde{W}_k$ is denoted by

$$I_k = \{n \mid \tilde{W}_k^n \neq 0\}. \quad (19)$$

Its cardinal satisfies $\#(I_k) \leq S$ since $\tilde{W} \in \mathcal{C}_p$. The back-projection step consists in computing the projection of P_k on $\text{Span}((D_n)_{n \in I_k})$. This is done by solving a linear system in dimension at most S : the back-projected weights W_k are

$$W_k^{I_k} = \arg \min_w \|P_k - D_{I_k} w\|^2 \quad (20)$$

and 0 elsewhere, where $W_k^I = (W_k^n)_{n \in I}$ and $D_I = (D_n)_{n \in I}$.

Patch extraction and image reconstruction. The patch extracting operator Π involved in the function (3) extracts the patches $P = \Pi(u)$ from any image u .

Its adjoint operator Π^* is involved in the optimization process described in Sect. 3. The effect of Π^* on a set of patches P is to merge all of them into an image $\tilde{u} = \Pi^*(P)$ by summing the overlapping parts of the patches.

In our previous work [33], we replaced this operator by a non-linear operator Π_{NL} , because the linear operator Π^* introduced some blur when averaging the overlapping parts of the patches. In the current paper, we introduced the spectrum constraint $\mathcal{C}_s$ which ensures the synthesized image to be sharp enough. As a result, we do not need the non-linear operator any longer. During the optimization process in Sect. 3, we use the true adjoint operator Π^* involved in the gradient of the function (3).

3 Optimization

In the previous section, we have defined and justified the variational framework we propose for texture synthesis. The corresponding energy is expressed as a weighted sum of distances to constraint sets. For each of these sets, a projection operator has been defined. In this section, we present a general optimization procedure that is suited to such an energy, resulting in a texture synthesis algorithm. Before proceeding, let us observe that the presented algorithm only yields local and non unique solutions to the proposed non-convex variational approach. From an optimization point of view, this is a pain. But from a synthesis point of view, this is good news: several (local) minimizers correspond to as many possible texture synthesis.

3.1 General Formulation

The penalty function E defined by (3) can be rewritten as

$$E(u) = \frac{1}{2} \sum_i \alpha_i \mathcal{D}^2(A_i u, \mathcal{C}_i) \quad (21)$$

where the indices i of the sum are (p,s,h), the linear operators A_i are $(\Pi, \text{Id}, \text{Id})$, and $\mathcal{C}_i$ and α_i are respectively the constraints and the weighting terms defined in Sect. 2.

This function measures how close some linear measurements $A_i u$ are to the subsets $\mathcal{C}_i$. We want these distances to be as small as possible and therefore look for a local minimizer of the energy E .

3.2 Minimization

We use a gradient descent scheme to minimize the energy E . As explained in the following paragraph, there is no theoretical guaranty about the convergence of this algorithm. However, local properties of the gradient are favorable for the cluster points of the algorithm to be critical points of the functional E , as desired.

Let $\mathcal{K}_i$ be the cut locus of $\mathcal{P}_{\mathcal{C}_i}$, that is, the points on which the projections are not uniquely defined, and $\mathcal{K} = \bigcup_i A_i^{-1}(\mathcal{K}_i)$ the union of their reciprocal images. In any $u \notin \mathcal{K}$, the functional E is $\mathcal{C}^1$. Its gradient is

$$\nabla E(u) = \sum_i \alpha_i A_i^* (A_i u - \mathcal{P}_{\mathcal{C}_i}(A_i u)), \quad (22)$$

where the projectors $\mathcal{P}_{\mathcal{C}_i}$ are given in Sect. 2. On the contrary, for $u \in \mathcal{K}$, the projections and thus the gradient are not uniquely defined; however, any projection provides a descent direction for such u . Observe also that $\mathcal{K}$ is a measure-zero set.

In order to find a local minimum of (21), we perform a gradient descent from a random point $u^{(0)}$. The resulting sequence is defined by

$$u^{(\ell+1)} = u^{(\ell)} - \tau \nabla E(u^{(\ell)}) \quad (23)$$

where ∇E is the gradient (22) of E . This gradient is not globally nor locally Lipschitz, mainly because of the spectrum constraint $\mathcal{C}_s$. The same problem is encountered in phase retrieval problems [3]. In order for the gradient descent to converge, we should decrease the step size to 0 or choose it adaptively by line-search. In practice, we choose a constant step size $\tau = 1/c$ with $c = \sum_i \alpha_i \|A_i^* A_i\|$ and motivate this choice in the next paragraph. Even if there is no theoretical convergence guaranty, numerical experiments show that the gradient descent converges in practice to a stationary point of E (see Sect. 4).

3.3 Averaged Projections

This paragraph points out that the averaged projection algorithm is a particular case of the proposed gradient descent scheme.

We consider the case of a perfect tiling, i.e. when all the pixels belong to Z patches exactly where $Z \triangleq \lceil \tau/\Delta \rceil^2 = \|\Pi^* \Pi\|_{2,2}$, which is the case for integer ratio $\tau/\Delta \in \mathbb{N}^*$ of patch size over spacing. In this case, $\Pi^* \Pi = Z \cdot \text{Id}$. Gradient descent step (23) becomes in this case

$$u^{(\ell+1)} = c^{-1} \sum_i \alpha_i A_i^* \mathcal{P}_{\mathcal{C}_i}(A_i u^{(\ell)}). \quad (24)$$

where the constant c previously defined is simply $c = \tilde{\alpha}_p + \alpha_s + \alpha_h$ with $\tilde{\alpha}_p = \alpha_p Z$. Relation (24) is a step of averaged projections: the transforms $A_i u^{(\ell)}$ are projected on the respective sets $\mathcal{C}_i$ and are then averaged together.

A cluster point $\tilde{u}$ of the iterates (24) is a stationary point of the energy (21) and satisfies

$$\tilde{u} = c^{-1} \sum_i \alpha_i A_i^* \mathcal{P}_{\mathcal{C}_i}(A_i \tilde{u}). \quad (25)$$

This provides a geometrical interpretation of the solutions: each solution is the barycenter of the projections of its transforms on the sets $\mathcal{C}_i$ weighted by α_i .

Alternated projections. Instead of using averaged projections (24), it is possible to use alternated projections, which gives the following iterations:

$$u^{(3\ell+1)} = Z^{-1} \Pi^* \mathcal{P}_{\mathcal{C}_p}(\Pi u^{(3\ell)}),$$

$$u^{(3\ell+2)} = \mathcal{P}_{\mathcal{C}_s}(u^{(3\ell+1)}),$$

$$u^{(3\ell+3)} = \mathcal{P}_{\mathcal{C}_h}(u^{(3\ell+2)}),$$

in the case of $\tilde{\alpha}_p = \alpha_s = \alpha_h = 1$. The convergence properties of both averaged and alternated non-convex projections on smooth manifolds are analyzed in [19].

Observe however that this scheme has no chance to converge since the sets $\mathcal{C}_i$ are distinct in general. It leads to 3 cluster points, namely 1 point in each set $\mathcal{C}_i$. Recall that averaged projections lead to a cluster point $\tilde{u}$ (25), which is a compromise between the 3 constraints $\mathcal{C}_i$.

The experimental results in the next section (Fig. 8 in particular) show that the alternated projections algorithm is more likely to produce artifacts than the averaged projections algorithm. Note that several texture synthesis methods, like [14] and [32], are implicitly based on an alternated projections scheme. As explained in [14], it is better not to iterate alternated projections too many times because the results may suffer from artifacts.

3.4 Multi-scale Procedure

Since the energy (3) is non-convex, the choice of the initialization $u^{(0)}$ has a strong influence on the result. In order to avoid visually unsatisfying local minima, we perform the synthesis through J scales using a classical multi-scale scheme.

At scale $j \in \{J-1, \dots, 0\}$, the original image u_0 is decimated by a factor 2^j to give an image u_j . The synthesis is performed using u_j as an exemplar and a dictionary D_j learned on it. The result $u_j^{(\infty)}$ is oversampled by a factor 2 and is used as initialization $u_{j-1}^{(0)}$ for the next scale. The roughest scale is initialized with a white noise $u_{J-1}^{(0)}$. The resulting image is the synthesis $u_0^{(\infty)}$ obtained at the finest scale.

Decimation and oversampling operations are performed as follows. The image to be decimated is first filtered by a 2×2 box kernel and then decimated by a factor 2. The oversampling step is performed with a bi-cubic interpolation. The shift of half a pixel is obtained using the cubic convolution kernel $h = (-1, 9, 9, -1)/16$.

3.5 Implementation Details

Periodicity of the synthesis. During the synthesis process, because of the spectrum constraint, the image is assumed to be periodic. The histogram constraint and the white noise initialization are consistent with this assumption since all pixels are considered independent. The oversampling step is made periodic by using a circular convolution. The patch processing is made consistent by defining the patch extracting operator $\Pi : u \mapsto \Pi(u)$ on the periodic image. Some patches thus contain pixels both on the very left side and on the very right side of the image. As a result, the synthesized image is also periodic. It can be tiled or shift circularly.

Spacing between patches. As defined in the notation section (Sect. 1.4), the patches are located on a grid of step Δ . To avoid mosaic artifacts, the value of Δ must be as small as possible, 1 or 2 pixels for instance. But the lower the slower. A higher value, for instance 4 pixels, provides a trade-off between synthesis quality and computation time.

Note that the offset of the grid of patches can be freely chosen (recall that the image is assumed periodic). Choosing a random offset for each iteration reduces the mosaic artifact caused by a large Δ . In practice, a step size of $\Delta = 4$ pixels with patches of size $\tau = 12$ pixels does not produce mosaic artifacts with this technique and increases the speed of the algorithm (a few minutes for a 256×256 px image).

Post processing. To increase the grain of the resulting texture, we add an extra projection step at the very end of the synthesis process. The image $u_0^{(\infty)}$ synthesized at the finest step is projected on the spectrum constraint and then on the histogram constraint. The final image is $u^{(\infty)} = \mathcal{P}_{\mathcal{C}_h} \circ \mathcal{P}_{\mathcal{C}_s}(u_0^{(\infty)})$.

4 Numerical Results

This section provides numerical results of our synthesis algorithm. We illustrate the influence of each term of the energy function (3) and of the parameters. We compare our results to the most classical synthesis methods relying on statistical constraints. We also compare the innovation capacity of the approach to that of exemplar-based methods.

When not explicitly specified, the extracted patches are of size $\tau = 12$ and are centered on a grid with a step size of $\Delta = 4$ pixels. The synthesis is performed over $J = 3$ scales and the terms of the functional are given the same weight¹ $\alpha_h = \alpha_s = \tilde{\alpha}_p = 1$. The dictionary is made of $N = 384$ atoms and the sparsity is set to $S = 4$ coefficients per patch.

Figure 1 shows several results of our algorithm for textures from the VisTex [27] database. The sand and water textures are well reproduced thanks to the spectrum constraints. The regularity of quasi-periodic textures is ensured by both the spectrum constraint and the multi-scale scheme. The sharpness of the edges and the geometrical patterns are handled by the sparse decomposition of the patches into the adaptive dictionary.

The last row of Fig. 1 shows difficult examples to illustrate the limitations and the failure cases of our algorithm. Repetition of small sharp objects —like the pills texture— cannot be represented by our approach: the sparsity of the patches cannot handle several objects in a patch, and the spectrum constraint is not able to create sharp objects. A mixture of large scale patterns and small details —like the tiles or the pumpkins— are also difficult to generate because

¹ We recall the normalization $\tilde{\alpha}_p = \alpha_p Z$ with $Z = \lceil \tau/\Delta \rceil^2$

of the patch constraint: small patches cannot deal with large structures, whereas large patches cannot handle small details because of the sparsity constraint.

4.1 Comparison

4.1.1 Statistical Methods

We first show a comparison with other statistical methods that we now recall.

The approach of Heeger and Bergen [14] imposes the histogram of the image and also the distributions of the coefficient of a multi-scale decomposition [32], scale by scale. An alternated projection method is used: the constraints are imposed turn by turn. As explained in [14], this should not be iterated more than 4 or 5 times because of convergence issues. The decomposition being based on gabor-like filters, this method is not adapted to synthesize edges or structures. This method may be tested online using the *IPOL* demo from [4].

The method from Portilla and Simoncelli [29] is the state of the art among statistical methods, to the best of our knowledge. Several constraints on the coefficients of the multi-scale decomposition [32] are imposed: mean, variance, and other moments. The main difference with [14] is that correlations are imposed between neighbor coefficients, whereas [14] only considers marginals at each scale and orientation. Some edges and sharp structures can be synthesized thanks to these dependencies. Let us emphasize that, in contrast, the method presented in this paper does not impose dependency between coefficients corresponding to different atoms in the dictionary. However, we used an *adaptive* dictionary learned on the exemplar, whereas [14, 29] use a non-adaptive pyramid transform.

In [12], synthesis is performed by preserving the amplitude of the Fourier spectrum and by randomly shuffling the phases of the Fourier Transform. This is strictly equivalent to generating a white noise and projecting it on our spectrum constraint $\mathcal{C}_s$. This approach is fast and is well adapted to smooth textures such as sand or a cloudy sky, but cannot produce edges. This method may be tested online using the *IPOL* demo from [11].

The algorithm of [26] is the first applying sparse decomposition of patches to texture synthesis. A redundant dictionary is learned beforehand on an exemplar of the texture. The synthesis is then performed from a white noise image by alternating histogram transfer and sparse coding into the dictionary. This is equivalent to alternated projection on our sets $\mathcal{C}_h$ and $\mathcal{C}_p'$ where $\mathcal{C}_p'$ has no constraint on the number of usage F^n of the elements in the dictionary.

In Fig. 2, we show several synthesis examples using our approach and the statistical methods recalled in this section.

Fig. 3 Results from the pixel-by-pixel copy-paste algorithm [8] (middle); the coordinate maps (bottom) show that the results are tilings from the exemplars (top), even if the synthesis is performed pixel after pixel

As we can see, the proposed approach has the best ability to reproduce both structures and fine details of textures.

4.1.2 Copy-Paste Methods

Figure 3 shows two synthesis results obtained with the Efros and Leung algorithm [8], using the *IPOL* accelerated implementation [1]. The synthesis is performed by choosing and copying pixels one by one from the exemplar. The concept of this algorithm may be encountered in a large number of more recent algorithms [35].

Figure 3 also displays on the bottom a map of coordinates which represents the location in the exemplar of the synthesized pixels. The structure of the synthesized image is clearly visible on these maps: the synthesis is made of a tiling from the exemplar even if pixels are copied one by one. The results of such copy-paste methods are often visually better than those obtained using statistical methods, but their innovation capacity is much lower. In contrast, results from our approach (see Fig. 1) include no repetition from the input image.

4.2 Detailed Analysis of the Synthesis Functional

The influence of each of the three terms of the function (3) is illustrated in Fig. 4. The effect of each constraint is demonstrated by suppressing one of them while preserving the two others. This experiment yields the following observations.

Fig. 1 Result of our algorithm (big images) for several exemplar (small images). The last examples are the most difficult because they simultaneously exhibit many small sharp details and large geometrical structures.

The *patch decomposition* term relies on the adaptive dictionary, which is good at representing geometrical features. The syntheses produced without this constraint have little geometrical content: edges and sharp details are completely lost.

On the other hand, a sparse decomposition cannot represent granularities of textures like sand or rocks. This is due to the noisy aspect of such textures which cannot be sparsely

represented. It is approximated by an almost-constant value and the texture generated is too smooth when using only histogram + patch sparsity terms.

The *spectrum* term acts as a complement of the patch sparsity term. It is powerful to represent granularities, as may be seen on the sand example. However, the spectrum cannot produce any edge, as illustrated by the synthesis with

Fig. 2 We present synthesis results using our approach and other statistical methods. From left to right: our approach penalizes a deviation of histogram, spectrum, and sparse approximation of patches; [29] imposes statistics including correlations in a wavelet frame; [14] imposes the image histogram and wavelet coefficient distributions (scale by scale); [12] preserves only the spectrum modulus; [26] imposes the image histogram and patch sparsity

only the spectrum and histogram terms. Preserving only the spectrum is proposed and illustrated in [12].

We also remark that the low frequencies imposed by the spectrum enhance the regularity of the texture. Without this constraint, the results have some brighter or darker regions, whereas the results using the spectrum are more stationary.

The *histogram* term deals with the contrast and the color faithfulness of the result. Note that the patch decomposition alone has no control on the patch dynamic.

4.3 Influence of the Parameters

This section illustrates the influence of the parameters of our synthesis algorithm: the size $\tau \times \tau$ of the patches, the number J of scales, the weights α of each term of the function, and the sparsity factor S .

4.3.1 Patch Size and Multi-scale

Figure 5 illustrates the influence of the multi-scale process and of the size of the patches.

The effect of the multi-scale is clearly visible on all the results: it increases the spatial coherence of the texture. In the case of quasi-periodic textures with small patterns, it increases the regularity of the texture: the patterns are more regularly distributed. In the case of textures with large patterns, it produces large structures such as long and continuous edges. This is not possible with the single scale synthesis.

The size of the patches must be roughly the size of the elements in the texture. A patch size of 12×12 pixels is a good compromise on the VisTex database.

If the patches are too small compared to the patterns of the texture, the pattern cannot be well synthesized and suffer from geometrical distortions. On the contrary, using too big patches makes the sparse approximation into D_0 rougher and the small details are lost. Moreover, bigger patches imply more coefficients, a bigger dictionary, and less sparsity: the sparse decomposition algorithm becomes far slower and less accurate (we recall that we use a heuristic since the problem is NP-hard).

4.3.2 Weights of the Function

Figure 6 shows results obtained when varying the weighting of the different terms in the functional (3) instead of choosing $\alpha_h = \alpha_s = \tilde{\alpha}_p = 1$.

This default choice provides reasonable results for most of the textures, but the weights can be tuned for each texture to be synthesized. We did not find an easy way to automatically compute a good set of weights for each texture, although this would be useful in practice.

The synthesis given in Fig. 6 are obtained, from left to right, with the following settings of $\alpha_s/\tilde{\alpha}_p$: 5/1, 3/1, 1/1, 1/3, and 1/5, and always with $\alpha_h = (\alpha_s + \tilde{\alpha}_p)/2$. Textures with large structures, long edges, or geometrical elements, are better synthesized with a higher weight for the patch sparsity term. On the contrary, a texture without sharp nor structured elements but with a granular aspect is better reproduced with more weight on the spectrum term. In the extreme case of a texture without any geometrical element, the patch term can be removed: see the examples of the sand texture in Fig. 2 with [12] (spectrum only) or in Fig. 4 without the patches sparsity term (spectrum and histogram).

The intermediate cases (a granular texture with some sharp patterns) are well synthesized with a balanced weighting $\alpha_s = \tilde{\alpha}_p$.

4.3.3 Sparsity

Figure 7 illustrates the effect of the sparsity parameter S .

A larger parameter S means that the patches are a linear combination of more elements from the dictionary D_0 . In the case of texture synthesis, a large value of S allows the superposition of several atoms and create artifacts, particularly visible on the example in Fig. 7.

On the contrary, the smallest value $S = 1$ imposes each patch to be proportional to 1 atom from D_0 . Imposing $S = 1$ and binary weights $W \in \{0, 1\}^{N \times K}$ is an interesting alternative. It forces each patch to be equal to an atom of the dictionary. The normalization constraint $\|d_n\|_2 \leq 1$ is no longer necessary and should be removed in this case.

Within this setting, the learning stage (8) becomes a K -means algorithm. The decomposition Algorithm 1 becomes a nearest-neighbor classification, with a constraint on the number F^n of use of each atom. The dictionary is in this case a resampled and weighted version of the set of all the patches in the original image. It is more compact because it contains far less patches than the image. The nearest-neighbor search would thus be far faster than an exhaustive search. Observe that in the case of an exhaustive patch search, the synthesis method is similar to the “Texture Optimization” algorithm [18] and to the video inpainting approach of Wexler et al. [37].

4.3.4 Other Settings

Figure 8 shows the effect of several options of our method.

The set of patches Πu can be subsampled to improve the speed of the algorithm as explained in Sect. 3.5.

We only consider the patches lying on a grid of step Δ with $\Delta > 1$. This leads to a mosaic artifact: it appears because the borders of the patches becomes visible. To avoid this artifact, we use random translations of the grid at each

Fig. 4 Synthesis results when dropping one of the three terms. The histogram term prevents from a loss of contrast. The spectrum term spatially regularizes the synthesis and generates granularity. The patch decomposition term handles sharp edges. Note that the post-process is not performed to allow a fair comparison of each term.

iteration: the border of the patches are not always at the same relative position, and the mosaic artifact disappears. Figure 8 shows that using $\Delta = 4$ with random offset gives a similar result than when using $\Delta = 1$, while being 16 times faster.

Alternated projections on the sets $\mathcal{C}_h$, $\mathcal{C}_s$, and $\mathcal{C}_p$ (Sect. 3.3) leads to the images in Fig. 8 (right). The result is quite similar to our result (2nd image from the right) but has some artifacts: some small edges appear, and some grain is missing. More generally, the results obtained using this alternated projection algorithm are sometimes similar to ours, and sometimes suffer from several kinds of artifacts, which point out the bad convergence of that method.

The post-processing step, that is the final projection on $\mathcal{C}_s$ and then $\mathcal{C}_h$, removes residual smoothness of the texture,

due to the patch sparsity constraint. This is visible in Fig. 8 (3rd and 4th images from the left).

5 Conclusion

In this paper, we proposed a texture synthesis framework which puts together both a Fourier spectrum constraint and a sparsity constraint on the patches. Thanks to these two constraints, our method is able to synthesize a wide range of textures without a patch-based copy-paste approach. The spectrum controls the amount of oscillations and the grain of the synthesized image, while the sparsity of the patches handles geometrical information such as edges and sharp patterns.

We propose a variational framework to take both constraints into account. This framework is based on a non-convex energy function defined as the sum of weighted dis-

Fig. 5 Synthesis results with different numbers of scales (top) and different sizes of patches (bottom). Multi-scale ensures spatial coherency of the texture. The patches must roughly have the size of the patterns of the texture.

Fig. 6 Synthesis results with different weighting: more spectrum on the left, more patch sparsity on the right. Textures with few geometrical content (top) are better reproduced with more spectrum; textures with large structures (bottom) need more patches sparsity.

Fig. 7 Synthesis results with different sparsity $S = 2, 4, 6, 8$. His example illustrate the artifacts caused by the superposition of several atoms for larger S .

Fig. 8 Synthesis results with different settings. From left to right: using all the patches ($\Delta = 1$) without post-processing (the final projection on $\mathcal{C}_s$ and then $\mathcal{C}_h$); using only the patches on a grid of step $\Delta = 4$; adding random translations of the grid at each iteration; adding the post-processing step; comparison with alternated projections instead of our gradient descent scheme.

tances between the image (or some linear transform of it) and these constraints. The synthesis consists in finding local minima of this function. The proposed framework is generic and can take into account other constraints: in addition to the sparsity and spectrum constraint, we use a histogram constraint to control the color distribution.

A Appendix: Proofs

A.1 Projection on $\mathcal{C}_s$

Here we demonstrate the expression (7) of the projection of u on $\mathcal{C}_s$. The projection is $u_s \in \mathcal{C}_s$ of the form $\hat{u}_s(m) = e^{i\varphi(m)} \hat{u}_0(m)$. Our goal is to minimize

$$\|u - u_s\|^2 = \sum_m \left\| \hat{u}(m) - e^{i\varphi(m)} \hat{u}_0(m) \right\|^2 \quad (26)$$

with respect to φ .

As the term of the sum are independent, the problem is to minimize

$$f(\psi) = \left\| x - y e^{i\psi} \right\|^2 \quad (27)$$

where $x = \hat{u}(m)$, $y = \hat{u}_0(m)$ and $\psi = \varphi(m)$ for any m . The hermitian product of $x, y \in \mathbb{C}^3$ is denoted by $x \cdot y = \sum_i x_i y_i^* \in \mathbb{C}$

The development of the expression of $f(\psi)$ gives

$$f(\psi) = \|x\|^2 - e^{i\psi} y \cdot x - e^{-i\psi} x \cdot y + \|y\|^2.$$

The function f being continuous and 2π -periodic on $\mathbb{R}$, it admits (at least) a minimum and a maximum which are critical points ψ_c satisfying $f'(\psi_c) = 0$. Let's write $x \cdot y = A e^{i\theta}$ with $A \geq 0$. The derivative

$$f'(\psi) = -i e^{i\psi} y \cdot x + i e^{-i\psi} x \cdot y.$$

gives $e^{2i\psi_c} = e^{2i\theta}$ and the critical points ψ_c are thus characterized by $e^{i\psi_c} = \pm e^{i\theta}$.

The second derivative

$$f''(\psi) = e^{i\psi} y \cdot x + e^{-i\psi} x \cdot y.$$

provides more information: we know $e^{i\theta}$ is a minimum since $f''(e^{i\theta}) = 2A \geq 0$, and $-e^{i\theta}$ is a maximum since $f''(-e^{i\theta}) = -2A \leq 0$.

The case $x \cdot y = 0$ leads to $A = 0$ and f being constant. In other cases, $A > 0$: the minimums $\psi_{\min}$ of the functions are strict and satisfy $e^{i\psi_{\min}} = e^{i\theta} = \frac{x \cdot y}{|x \cdot y|}$, hence the expression of $\hat{u}_s(m)$ given in (7).

A.2 Projection on $\mathcal{C}_p$

Here is the proof that (13) and (14) are the minimizers of (10).

Using (15), the expression (10) is $\|R^{(\ell)} - wD_0E_{n,k}\|^2$. Expressing this norm as the sum of the norms of the columns leads to

$$\min_{k,n,w} \|R_k^{(\ell)} - wD_n\|^2 + \sum_{k' \neq k} \|R_{k'}^{(\ell)}\|^2. \quad (28)$$

Let's fix k and n . The problem

$$\min_w \|R_k^{(\ell)} - wD_n\|^2 \quad (29)$$

is now an orthogonal projection. Since D_n is normalized, the solution is

$$w^* = \langle R_k^{(\ell)}, D_n \rangle \quad (30)$$

as stated in (14) and Pythagora's theorem gives the error

$$\|R_k^{(\ell)} - w^*D_n\|^2 = \|R_k^{(\ell)}\|^2 - \langle R_k^{(\ell)}, D_n \rangle^2. \quad (31)$$

Substituting w in (28) by its optimal value w^* function of (k, n) simplifies the problem to

$$\min_{k,n} -\langle R_k^{(\ell)}, D_n \rangle^2 + \sum_{k'} \|R_{k'}^{(\ell)}\|^2. \quad (32)$$

Hence the result (15) stating $(k^*, n^*) = \arg \max \langle R_k^{(\ell)}, D_n \rangle$ among the set $\mathcal{J}_{W^{(\ell)}}$ of admissible (k, n) .

Acknowledgements We would like to thanks the authors of the Vis-Text database [27] for the set of textures they publicly provide. Gabriel Peyré also acknowledges support from the European Research Council (ERC project SIGMA-Vision). Yann Gousseau acknowledges support from the French National Research Agency (ANR project MATAIM).

References

1. Aguerrebere, C., Gousseau, Y., Tartavel, G.: Exemplar-based Texture Synthesis: the Efros-Leung Algorithm. *Image Processing On Line* **2013**, 213–231 (2013)
2. Aharon, M., Elad, M., Bruckstein, A.: K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *Signal Processing, IEEE Transactions on* **54**(11), 4311–4322 (2006)
3. Bauschke, H.H., Combettes, P.L., Luke, D.R.: Phase retrieval, error reduction algorithm, and fienuip variants: a view from convex optimization. *JOSA A* **19**(7), 1334–1345 (2002)
4. Briand, T., Vacher, J., Galerne, B., Rabin, J.: The heeger and bergen pyramid based texture synthesis algorithm. *Image Processing On Line*, preprint (2013)
5. Cross, G., Jain, A.: Markov random field texture models. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **5**(1), 25–39 (1983)
6. Desolneux, A., Moisan, L., Ronsin, S.: A compact representation of random phase and gaussian textures. In: *Acoustics, Speech and Signal Processing (ICASSP)*, 2012 IEEE International Conference on, pp. 1381–1384. IEEE (2012)
7. Efros, A., Freeman, W.: Image quilting for texture synthesis and transfer. In: *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pp. 341–346. ACM (2001)
8. Efros, A., Leung, T.: Texture synthesis by non-parametric sampling. In: *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on*, vol. 2, pp. 1033–1038. IEEE (1999)
9. Elad, M., Aharon, M.: Image denoising via learned dictionaries and sparse representation. In: *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, vol. 1, pp. 895–900. IEEE (2006)
10. Engan, K., Aase, S., Hakon Husoy, J.: Method of optimal directions for frame design. In: *Acoustics, Speech, and Signal Processing, 1999. Proceedings., 1999 IEEE International Conference on*, vol. 5, pp. 2443–2446. IEEE (1999)
11. Galerne, B., Gousseau, Y., Morel, J.M.: Micro-texture synthesis by phase randomization. *Image Processing On Line* (2011)
12. Galerne, B., Gousseau, Y., Morel, J.: Random phase textures: Theory and synthesis. *Image Processing, IEEE Transactions on* **20**(1), 257–267 (2011)
13. Galerne, B., Lagae, A., Lefebvre, S., Drettakis, G.: Gabor noise by example. *ACM Transactions on Graphics (TOG)* **31**(4), 73 (2012)
14. Heeger, D., Bergen, J.: Pyramid-based texture analysis/synthesis. In: *SIGGRAPH '95*, pp. 229–238 (1995)
15. Julesz, B.: Visual pattern discrimination. *Information Theory, IRE Transactions on* **8**(2), 84–92 (1962)
16. Julesz, B.: A theory of preattentive texture discrimination based on first-order statistics of textons. *Biological Cybernetics* **41**(2), 131–138 (1981)
17. Kuhn, H.: The hungarian method for the assignment problem. *Naval research logistics quarterly* **2**(1-2), 83–97 (1955)
18. Kwatra, V., Essa, I., Bobick, A., Kwatra, N.: Texture optimization for example-based synthesis. In: *ACM Transactions on Graphics (TOG)*, vol. 24, pp. 795–802. ACM (2005)
19. Lagae, A., Lefebvre, S., Cook, R., Derosé, T., Drettakis, G., Ebert, D., Lewis, J., Perlin, K., Zwicker, M.: State of the art in procedural noise functions. *EG 2010-State of the Art Reports* (2010)
20. Lagae, A., Lefebvre, S., Drettakis, G., Dutré, P.: Procedural noise using sparse gabor convolution. In: *ACM Transactions on Graphics (TOG)*, vol. 28, p. 54. ACM (2009)
21. Lefebvre, S., Hoppe, H.: Parallel controllable texture synthesis. *ACM Transactions on Graphics (TOG)* **24**(3), 777–786 (2005)
22. Mallat, S., Zhang, Z.: Matching pursuits with time-frequency dictionaries. *Signal Processing, IEEE Transactions on* **41**(12), 3397–3415 (1993)
23. Moisan, L.: Periodic plus smooth image decomposition. *Journal of Mathematical Imaging and Vision* **39**(2), 161–179 (2011)
24. Olshausen, B., Field, D.: Natural image statistics and efficient coding*. *Network: computation in neural systems* **7**(2), 333–339 (1996)
25. Perlin, K.: An image synthesizer. *SIGGRAPH Comput. Graph.* **19**(3), 287–296 (1985)
26. Peyré, G.: Sparse modeling of textures. *Journal of Mathematical Imaging and Vision* **34**(1), 17–31 (2009)
27. Pickard, R., Graszky, C., Mann, S., Wachman, J., Pickard, L., Campbell, L.: *Vistex database*. Media Lab., MIT, Cambridge, Massachusetts (1995)
28. Pitie, F., Kokaram, A., Dahyot, R.: N-dimensional probability density function transfer and its application to color transfer. In: *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on*, vol. 2, pp. 1434–1439. IEEE (2005)
29. Portilla, J., Simoncelli, E.: A parametric texture model based on joint statistics of complex wavelet coefficients. *International Journal of Computer Vision* **40**(1), 49–70 (2000)
30. Rabin, J., Peyré, G., Delon, J., Bernot, M.: Wasserstein barycenter and its application to texture mixing. *Scale Space and Variational Methods in Computer Vision* pp. 435–446 (2012)
31. Ramanarayanan, G., Bala, K.: Constrained texture synthesis via energy minimization. *Visualization and Computer Graphics, IEEE Transactions on* **13**(1), 167–178 (2007)
32. Simoncelli, E., Freeman, W., Adelson, E., Heeger, D.: Shiftable multiscale transforms. *Information Theory, IEEE Transactions on* **38**(2), 587–607 (1992)

33. Tartavel, G., Gousseau, Y., Peyré, G.: Constrained sparse texture synthesis. *Proceedings of SSVM'13* (2013)
34. Tropp, J.: Greed is good: Algorithmic results for sparse approximation. *Information Theory, IEEE Transactions on* **50**(10), 2231–2242 (2004)
35. Wei, L., Lefebvre, S., Kwatra, V., Turk, G.: State of the art in example-based texture synthesis. In: *Eurographics 2009, State of the Art Report, EG-STAR*. Eurographics Association (2009)
36. Wei, L., Levoy, M.: Fast texture synthesis using tree-structured vector quantization. In: *SIGGRAPH '00*, pp. 479–488. ACM Press/Addison-Wesley Publishing Co. (2000)
37. Wexler, Y., Shechtman, E., Irani, M.: Space-time video completion. In: *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, vol. 1, pp. I–120. IEEE (2004)
38. Xia, G.S., Ferradans, S., Peyré, G., Aujol, J.F.: Compact representations of stationary dynamic textures. In: *Image Processing (ICIP), 2012 19th IEEE International Conference on*, pp. 2993–2996. IEEE (2012)