

Image deblurring, Spectrum Interpolation and Application to Satellite Imaging

Sylvain Durand^{*†}, François Malgouyres^{‡§}, Bernard Rougé[¶]

February 2, 2000

Abstract

This paper deals with two complementary methods in noisy image deblurring: a nonlinear shrinkage of wavelet-packets coefficients called FCNR and Rudin-Osher-Fatemi's variational method. The FCNR has for objective to obtain a restored image with a white noise. It will prove to be very efficient to restore an image after an invertible blur but limited in the opposite situation. Whereas the Total Variation based method, with its ability to reconstruct the lost frequencies by interpolation, is very well adapted to non-invertible blur, but that it tends to erase low contrast textures. This complementarity is highlighted when the methods are applied to the restoration of satellite SPOT images.

Abstract

Cet article porte sur deux méthodes complémentaires de déconvolution d'images bruitées: une contraction non linéaire de coefficients de paquets d'ondelettes appelée FCNR et la méthode variationnelle introduite par Rudin, Osher et Fatemi. La FCNR a pour objectif de reconstruire une image avec un bruit blanc. Elle s'avère très efficace dans la restauration d'images convoluées par un opérateur inversible mais moins performante dans les autres cas. Par contre, bien qu'elle ait tendance à effacer les textures peu contrastées, la seconde méthode est parfaitement adaptée aux déconvolutions non inversibles car la variation totale permet de reconstruire, par interpolation, les fréquences perdues. Cette complémentarité est mise en évidence lorsque les méthodes sont appliquées à la restauration des images satellites SPOT.

1 Introduction

This paper deals with two complementary methods in image deblurring: a nonlinear shrinkage of wavelet-packets coefficients that we will call Fixed Chosen Noise Restoration (FCNR) and a variational method (known as Rudin-Osher-Fatemi's method). The aim of these methods is to recover an original image u_∞ describing a real scene from a blurred and noisy satellite image u_0 (the methods can however be adapted to other kinds of images and to signal deblurring). Generally, the degradation of the image occurs during the image acquisition and is modeled by a linear and translation invariant blur and an additive noise. In satellite imaging, the deterioration is due to the motion of the satellite and the imperfection of the optics for the blur and the electronics of the captors for the noise. We obtain the measured image from the "ideal" one from the following linear model:

$$u_0 = H(u_\infty) + n, \quad (1)$$

where H denotes the operator of convolution by an impulse response h : $H(u) = h * u$ and n is an additive white noise of standard deviation σ ("white" is used in the sense that pixels are independent, identically distributed, random variables). In practice, the noise can be considered as Gaussian.

^{*}CMLA, ENS Cachan, 61 avenue du président Wilson, 94235 Cachan Cedex, France; sdurand@cmla.ens-cachan.fr

[†]LAMFA, Université de Picardie, Jules Verne, 33, rue Saint Leu 80039 Amiens Cedex 1, France

[‡]CMLA, ENS Cachan, 61 avenue du président Wilson, 94235 Cachan Cedex, France; malgouy@cmla.ens-cachan.fr

[§]F. Malgouyres would like to thank the French DGA which has partially financed that work.

[¶]CNES, DGAT/SH/QTIS, 18 Avenue E. BELIN, 31401 Toulouse CEDEX4, FRANCE; rouge@cnes.fr

Even when the blurring operator is invertible, such a problem is known to be ill-posed. Assume that there exists an inverse operator H^{-1} defined by

$$H^{-1}(u) = \left(\frac{1}{\hat{h}}\hat{u}\right)^\vee,$$

where \hat{u} denotes the Fourier transform of u and u^\vee its inverse Fourier transform. Applying H^{-1} to (1), we obtain

$$H^{-1}(u_0) = u_\infty + \left(\frac{\hat{n}}{\hat{h}}\right)^\vee. \quad (2)$$

Looking at this equation, it appears clearly that the noise might blow up at the frequencies for which \hat{h} vanishes, or becomes very small.

Several approaches have been used to recover u_∞ . The reader is referred to [2] for most of the linear ones and to [12, 18] for overviews on the subject. In a few words, the first approach consists in enhancing images without regard to the convolution kernel [23]. The other methods are based on regularization approaches of the problem: using statistical properties (Wiener and Kalman filters) or regularity measurements of the images such as the entropy ([12] and its references), the Total Variation [31] or the characterization of Besov spaces by wavelets coefficients [13, 14].

We present here two approaches of the problem. Looking back to (2), note that the difficulty in image deblurring is directly related to the issue of denoising an image with highly colored noise. Our first objective will be therefore to control the noise in order to obtain a restored image with a white noise (the human eye is less sensitive to white noise than to colored noise) with a reasonable standard deviation. This will lead us to the two versions of the FCNR which are described in the next section.

Although it is very efficient to restore an image after an invertible blur, this technic will prove to be limited in the case when H is no longer invertible. If \hat{h} vanishes on some part of the Fourier domain, the informations supported by these frequencies will be lost during the degradation. Replacing them by zero creates Gibbs effects (see Figure 20).

The second method described in this paper has the ability to reconstruct the lost frequencies by interpolation and therefore avoids the Gibbs effects. It involves a regularization functional that makes the interpolation realisable. This functional is the Total Variation that is used in Rudin-Osher-Fatemi's method described in Section 3.1. The counterpart of this power to interpolate is the creation of homogeneous zones on the image and the erasure of low contrast textures that are well restored with the first method.

The complementarity of these two methods is highlighted in Section 4 where they are applied to the restoration of satellites SPOT 1 and SPOT 5 images.

Very recent papers deal with the equivalence between wavelet shrinkage methods and variational methods (see [6], [4]). The superiority of the wavelet based methods is therefore on a numerical point of view. Our purpose here is, among other things, to show that, in some situations, a variational method might remain more adapted.

2 Fixed Chosen Noise Restoration (FCNR)

The Fixed Chosen Noise Restoration which was introduced for the first time by one of the authors in [29, 30] is a version of Donoho and Johnstone's wavelet shrinkage method which consists in thresholding the coefficients of a noisy image in a wavelet basis [13, 14]. From the decomposition of the image in the orthonormal basis $(\psi_{j,k})_{j \in \mathcal{J}, k \in \mathcal{K}}$

$$u_0 = \sum_{j \in \mathcal{J}} \sum_{k \in \mathcal{K}} \langle u_0, \psi_{j,k} \rangle \psi_{j,k},$$

we obtain a denoised image by applying the following formula

$$u_1 = \sum_{j \in \mathcal{J}} \sum_{k \in \mathcal{K}} \tau_\lambda(\langle u_0, \psi_{j,k} \rangle) \psi_{j,k},$$

where the soft-thresholding function τ_λ is given by

$$\tau_\lambda(x) = \begin{cases} x - \text{sgn}(x)\lambda & \text{if } |x| \geq \lambda \\ 0 & \text{if } |x| < \lambda \end{cases} \quad (3)$$

and the threshold λ depends on the characteristics of the noise. In the case of a Gaussian noise, Donoho and Johnstone [13] chose a threshold which is proportional to the standard deviation σ of the noise,

$$\lambda = \sigma\sqrt{2\log N},$$

where N is the number of samples of the image. For N large, the L^∞ -norm of the noise is known to have a high probability of being close to λ [3].

Looking at (2), it seems opportune, in the case of a deblurring, to adapt the thresholding to frequencies. The deconvolution will indeed emphasize the noise at the frequencies for which h is small (in general, high frequencies). But fine scale wavelets (which are located in high frequencies) are badly concentrated in the Fourier domain. Thus it makes the adapted thresholding impossible when $1/h$ varies rapidly on high frequencies. Wavelets are not appropriate for the deblurring which needs a high resolution in frequency.

On the other hand, they are very well adapted to denoising since they concentrate the energy of the image over few coefficients (this is not the case for Fourier coefficients which offer an optimal resolution in frequency).

Wavelet-packets are a good alternative to wavelets since they offer a good resolution in frequency and a decomposition in “wavelets style” coefficients which ensures an efficient denoising.

2.1 The wavelet-packets

Let us start this section with a short introduction to wavelet-packets theory. For more details the reader is referred to [11]. We consider a pair of quadrature mirror filters (m_0, m_1) with

$$m_0(\xi) = \frac{1}{\sqrt{2}} \sum_{k \in \mathbb{Z}} c_k e^{-ik\xi},$$

and

$$m_1(\xi) = \frac{1}{\sqrt{2}} \sum_{k \in \mathbb{Z}} d_k e^{-ik\xi},$$

(for instance, $d_k = (-1)^k c_{1-k}$) which is connected with a multi-resolution analysis [26].

The basic wavelet-packets $w_n(x)$ are defined by recursion. We start with $w_0(x)$ which is the unique $L^1 \cap L^2$ -solution of the equation

$$w_0(x) = \sqrt{2} \sum_{k \in \mathbb{Z}} c_k w_0(2x - k), \quad (4)$$

under the normalization constraint

$$\int_{\mathbb{R}} w_0(x) dx = 1.$$

This function w_0 is actually the scaling function or father wavelet, usually denoted φ , associated with the connected multi-resolution analysis. The other w_n 's are then defined as follows

$$w_{2m}(x) = \sqrt{2} \sum_{k \in \mathbb{Z}} c_k w_m(2x - k), \quad (5)$$

$$w_{2m+1}(x) = \sqrt{2} \sum_{k \in \mathbb{Z}} d_k w_m(2x - k). \quad (6)$$

The properties of the quadrature mirror filters ensure that, for all $j \in \mathbb{N}$, the family $\{2^{-j/2}w_m(2^{-j}x - k), k \in \mathbb{Z}\}$ is an orthonormal basis of

$$E_m^j := \overline{\text{span}\{2^{-j/2}w_m(2^{-j}x - k), k \in \mathbb{Z}\}}^{L^2(\mathbb{R})}$$

with the scalar product $\langle u, v \rangle = \int_{\mathbb{R}} u(x) \bar{v}(x) dx$. Moreover, we have

$$\forall j \in \mathbb{N}, \forall m \in \mathbb{N}, \quad E_m^j = E_{2m}^{j+1} \oplus E_{2m+1}^{j+1}, \quad (7)$$

and thus

$$\forall j \in \mathbb{N}, \quad E_0^0 = \bigoplus_{m=0}^{2^j-1} E_m^j$$

which means that, for all $j \in \mathbb{N}$, $\{2^{-j/2}w_m(2^{-j}x - k), k \in \mathbb{Z}, m = 0, 1, \dots, 2^j - 1\}$ is an orthonormal basis of E_0^0 . Hence, the wavelet-packets theory provides us with a collection of orthonormal bases of a Hilbert space E_0^0 that we can identify with $l^2(\mathbb{Z})$. Looking at (7), let us observe that there are many other bases than the ones just mentioned, which involve several values of j . But we will here restrict ourselves to the bases associated with a unique level of decomposition j .

All this can also be expressed in the Fourier domain. Equations (5) and (6) become

$$\hat{w}_{2m}(2\xi) = m_0(\xi) \hat{w}_m(\xi),$$

$$\hat{w}_{2m+1}(2\xi) = m_1(\xi) \hat{w}_m(\xi),$$

which implies

$$\hat{w}_m(2^j \xi) = M_m^j(\xi) \hat{w}_0(\xi),$$

with $m = \sum_{l=0}^{2^j-1} \epsilon_l 2^l$ ($\epsilon_l \in \{0, 1\}$) and $M_m^j(\xi) = \prod_{l=0}^{2^j-1} m_{\epsilon_l}(2^l \xi)$. The identification of E_0^0 with $l^2(\mathbb{Z})$ is translated by the identification of \hat{u} the Fourier transform of a function $u \in E_0^0$, with $m \in L^2(\mathbb{T})$ such that

$$\hat{u}(\xi) = m(\xi) \hat{w}_0(\xi).$$

The same way, the Fourier transform of the functions $u \in E_m^j$ will be of the form $m(\xi) 2^{j/2} \hat{w}_m(2^j \xi)$ where $m \in L^2(\mathbb{T})$. Therefore, using the notable property of the quadrature mirror filters, $|m_0(\xi)|^2 + |m_1(\xi)|^2 = 1$, we obtain, for any $u \in E_0^0$

$$\begin{aligned} \hat{u}(\xi) &= \sum_{m=0}^{2^j-1} |M_m^j(\xi)|^2 m(\xi) \hat{w}_0(\xi) \\ &= \sum_{m=0}^{2^j-1} m(\xi) M_m^j(\xi) \hat{w}_m(2^j \xi) \end{aligned} \quad (8)$$

which expresses u as the linear combination of its projections onto the spaces E_m^j . Let us observe that when $|m_0|^2$ is well localized around $[-\pi/2, \pi/2] \pmod{2\pi}$ (and therefore $|m_1|^2$ around $[\pi/2, 3\pi/2] \pmod{2\pi}$), $|M_m^j|^2$ will be localized around an interval of length $2^{-j}\pi$. Therefore, equation (8) provides us with a horizontal slicing of the time-frequency plane by the spaces E_m^j , imitating in some sense the Fourier transform.

The optimal slicing would be obtained with the filter

$$m_0(\xi) = \sum_{l \in \mathbb{Z}} \chi_{[-\pi/2, \pi/2]}(\xi + 2l\pi)$$

(χ_A denotes the characteristic function of a set A) that we will call, in the sequel, Shannon filter. But it does not generate wavelet-packets in the sense they have been defined above since there is no L^1 -solution of the associated scaling equation (4). A Shannon basis is however defined by starting with $\hat{w}_0 = \chi_{[-\pi, \pi]}$.

Going back to the time domain, note that we get from the decomposition of a signal into the wavelet basis of E_m^j to its projection onto E_{2m}^{j+1} and E_{2m+1}^{j+1} , with the benefit of the operators T_0 and T_1 defined as follows

$$\forall u \in l^2(\mathbb{Z}), \quad \begin{aligned} T_0(u)(k) &= \sum_{l \in \mathbb{Z}} c_{2k-l} u(l) \\ T_1(u)(k) &= \sum_{l \in \mathbb{Z}} d_{2k-l} u(l) \end{aligned} \quad (9)$$

and which satisfy the following property

$$\sum_{k \in \mathbb{Z}} u(k) w_{m,k}^j = \sum_{k \in \mathbb{Z}} T_0(u)(k) w_{2m,k}^{j+1} + \sum_{k \in \mathbb{Z}} T_1(u)(k) w_{2m+1,k}^{j+1},$$

with $w_{m,k}^j(x) = w_m(2^j x - k)$. The reconstruction is done with the dual operators T_0^* and T_1^* ,

$$T_0^*(T_0(u)) + T_1^*(T_1(u)) = u.$$

The 2D-wavelet-packets are obtained by tensor product but that, for convenience, we will think out in 1D in the sequel.

2.2 Study framework and approximation

We choose here another approach than Donoho and Johnstone's [13] and consequently our restoration is not optimal in Donoho's sense. There does not exist any method that is able to remove all the noise from an image without loosing a lot of data, because it is not possible to completely separate it from the structures of the image. Therefore, our goal should not be to remove the noise but to control it.

Using a "traditional" wavelet shrinkage method will lead to a restored image with a light but colored noise which is very bad visually. When the shrinkage is done before deblurring, we obtain a deconvolved-noise $H^{-1}(n_0) = (\hat{n}_0/\hat{h})^\vee$ where the thresholded noise n_0 is white in the homogeneous zones. When the shrinkage is done afterward, the noise is even more colored.

Our aim should be here to obtain a restored image with a white noise, and to make this possible, it is necessary to find a method of denoising that leaves a noise of the form $H(n_0)$ where n_0 is white so that further deconvolution leaves a white noise. The other possibility is to transform a colored noise $H^{-1}(n)$ into a white noise when the deblurring is done before the denoising.

To apply a wavelet-packet shrinkage method as it will be introduced below, we need to make the following assumption. We assume that the deconvolved-noise $H^{-1}(n)$ is "almost" white on the support of the packets E_m^j which means that the wavelet-packets we use are very well localized in frequency variable. But as the wavelet-packets do not give rise to a partition of the Fourier domain with arbitrary precision, even for a large j , the restoration would be all the more successful that $1/\hat{h}$ does not oscillate too much.

We must however use wavelet-packets that are very well localized in the Fourier domain. The optimal choice is the Shannon basis associated with Shannon filter (see Section 2.1). But it is not applicable for numerical reasons (the coefficients of the filter are not decreasing fast enough) and because the Shannon basis is not denoising very efficiently due to its bad localization in the space domain. Consequently, we will use cubic spline wavelets which have exponential decay and are well localized in frequency. They will therefore offer a good compromise.

Remark now that when the localization assumption is satisfied, we can as well approximate H by an operator which is diagonal in the wavelet-packets basis. And since wavelets associated with the same packet are deduced one from another by translation, the eigen values of H are approximately constant on this packet. Note therefore that deblurring by multiplying the projection of the image over every packet by a constant, ensures that the deconvolved-noise is white on the packets. In other terms, the coefficients of the noise on the wavelets of the packet are i.i.d variables.

Let us focus on the error made in this approximation. We start from a signal $U = \sum_{k \in \mathbb{Z}} u(k) w_0(\cdot - k) \in E_0^0$. For any $j \in \mathbb{N}$ and $m \in \mathbb{N}$, the projection of U on E_m^j is given by

$$P_m^j U = \sum_{k \in \mathbb{Z}} T_m^j(u)(k) w_m^j(\cdot - k),$$

where $T_m^j = T_{\epsilon_{j-1}} \circ \dots \circ T_{\epsilon_0}$ with $(\epsilon_0, \dots, \epsilon_{j-1}) \in \{0, 1\}^j$ and $m = \sum_{l=0}^{j-1} \epsilon_l 2^l$. The error made in the direction of each wavelet-packet is therefore given by

$$T_m^j(H^{-1}(u)) - C_m^j T_m^j(u).$$

And its L^2 -norm is minimized by

$$C_m^j = \frac{\langle T_m^j(H^{-1}(u)), T_m^j(u) \rangle}{\|T_m^j(u)\|_2^2}.$$

This formula is usable for users who have a statistical estimation of the images they deal with. But as the deblurring is relatively stable with regard to the variations of C_ϵ , it is more convenient to approximate them by values which are independent of u (see section 2.6).

A second step in our restoration issue comes down to say that if we are able to find a method that returns images with white noise, cannot we obtain a noise which has the same law as n ? An answer to this will be given in Section 2.4.

2.3 FCNR

As it was mentioned above, we are going to adapt the thresholding to every wavelet-packet. From a noisy image u_0 whose wavelet-packets decomposition is

$$u_0 = \sum_{\substack{m \in \mathbb{N} \\ k \in \mathbb{Z}}} \langle u_0, w_{m,k}^j \rangle w_{m,k}^j,$$

with $w_{m,k}^j(x) = w_m(2^j x - k)$, we obtain a denoised image

$$u_1 = \sum_{\substack{m \in \mathbb{N} \\ k \in \mathbb{Z}}} \tau_{\lambda(m,j)}(\langle u_0, w_{m,k}^j \rangle) w_{m,k}^j,$$

where the soft-thresholding operator τ_λ is defined in (3). The threshold $\lambda(m, j)$ depends naturally on the standard deviation σ of the noise but it must depend also on the wavelet-packet $w_{m,k}^j$ and the filter \hat{h} . Note that, since we get from $w_{m,k}^j$ to $w_{m,0}^j$ by translation and since the noise is white, there is no reason for λ to depend on the parameter k .

Let us choose to make the deblurring after denoising. When it is done by multiplying the projection of the image on each packet by a constant as it was mentioned in Section 2.2, the obtained result does actually not depend on that chosen order. Remark however that when the deconvolution is made properly, doing it afterwards might emphasize the artifacts generated by the denoising. In such a situation, it is therefore better to do the deblurring before denoising. Consequently, in the sequel, we will assume that H is approximated by a diagonal operation in the wavelet-packet basis but that the method can easily be generalized to the case when H is not approximated.

The standard deviation $\sigma_{\lambda_m^j}$ of the noise on each packet E_m^j , after thresholding and before deblurring, becomes, after restoration, $C_m^j \sigma_{\lambda_m^j}$. To obtain an image with a “white” noise of standard deviation σ_0 on homogeneous zones (“white” should be understood in the sense that on each packet, the standard deviation is σ_0), we must therefore resolve,

$$\sigma_0 = C_m^j \sigma_{\lambda_m^j},$$

or, in the case of a Gaussian noise,

$$\sigma_0 = C_m^j \frac{2}{\sqrt{2\pi}\sigma^2} \int_{\lambda_m^j}^{\infty} (x - \lambda_m^j)^2 e^{-\frac{x^2}{2\sigma^2}} dx.$$

We cannot give an explicit expression of $\lambda_m^j = \lambda(C_m^j)$ but a numerical estimation is represented on Figure 1.

Remark that when H is not invertible, C_m^j and therefore λ_m^j can be infinite. Thus, to generalize the method to non-invertible H 's, we must let $\tau_\infty(x) = 0$ for any real value of x . The packets associated with the infinite threshold vanish during the restoration.

2.4 Invertible FCNR

This version of the FCNR comes from the remark that it might be worthwhile to find a method of restoration that returns an image whose noise has the same law as the noise of the blurred image (which is




Figure 1: The function $\lambda(C)$ with $\sigma_0 = \sigma = 1$.

not the case with the preceding version - see Figures 2 and 3). For instance, in the case we are interested in here, we would like to obtain an image with a Gaussian noise.

Moreover, a low contrast texture drowned in a strong noise might still be detected by the eye of an observant, even when the standard deviation of the noise is much larger than the intensity of the contrast of the texture. Using a method which consists in thresholding the coordinates of the image in a given basis might erase this texture since its coordinates might be under the threshold. The thresholding is therefore not appropriate to recover an image with a lot of textures.

Hence, it seems advisable to replace the soft-thresholding by an invertible function that does not erase any data.

One way to match this constraint, together with the objective of recovering an image with a Gaussian noise, is to replace the soft-thresholding function τ by the following

$$\tilde{\tau}_{\lambda_1, \lambda_2}(x) = \begin{cases} (\lambda_2/\lambda_1)x & \text{if } |x| \leq \lambda_1 \\ x + \text{sgn}(x)(\lambda_2 - \lambda_1) & \text{if } |x| > \lambda_1 \end{cases} \quad (10)$$

This function applied on the coordinates of a noisy image, will replace, in homogeneous zones, the noise of standard deviation σ by a noise with the same law but of standard deviation $(\lambda_2/\lambda_1)\sigma$ (modulo a negligible part of the noise - see Figure 4), and it will not change the main structures of the image (if λ_1 is not too large). The choice of the parameter λ_2/λ_1 is therefore very clear. The role of λ_1 is more subjective. If the noise is bounded, it is natural to take λ_1 equal to this bound. In the case of a Gaussian noise, we can take $\lambda_1 = \sigma\sqrt{\log N}$ where N is the number of samples of the image (or the packet when it is applied on E_m^j).

But that (12) was not satisfied by the hard-thresholding function ($\tau_\lambda(x) = 0$ if $|x| \leq \lambda$, x otherwise), note that (13) is also satisfied by

$$\tilde{\tau}_{\lambda_1, \lambda_2}(x) = \begin{cases} (\lambda_2/\lambda_1)x & \text{if } |x| \leq \lambda_1 \\ x & \text{if } |x| > \lambda_1 \end{cases} \quad (11)$$

This second possible choice gives rise to better statistics when the restored image is compared to a reference but it might be a priori less satisfactory on a visual criterion since this function is not continuous. A transformation associated with a discontinuous function is unstable with regard to the perturbations on




Figure 2: Noise histogram before restoration.




Figure 3: Noise histogram in homogeneous zones after a soft-thresholding (with τ_λ).




Figure 4: Noise histogram in homogeneous zones after a shrinkage with the function $\tilde{\tau}_{\lambda_1, \lambda_2}$.




Figure 5: Two ways of shrinking the packet $E_m^{j-1} = E_{2m}^j \oplus E_{2m+1}^j$ (it is a case where $\lambda_2^{j,2m+1} \geq \lambda_2^{j,2m}$).

the initial image. The following would be a good compromise

$$\tilde{\tau}_{\lambda_1, \lambda_2}(x) = \begin{cases} (\lambda_2/\lambda_1) x & \text{if } |x| \leq \lambda_1 \\ ((C \lambda_1 - \lambda_2) x + \text{sgn}(x) C \lambda_1 (\lambda_2 - \lambda_1)) / (C \lambda_1 - \lambda_1) & \text{if } \lambda_1 < |x| \leq C \lambda_1 \\ x & \text{if } |x| > C \lambda_1 \end{cases}$$

Remark also that the method is very close to the one introduced in section 2.5 since, when $\tilde{\tau}$ is defined by (10), we have for $|x| \geq \lambda_1$,

$$\tilde{\tau}_{\lambda_1, \lambda_2}(x) = \tau_{\lambda_1 - \lambda_2}(x).$$

Remark finally that if we keep in memory the parameters λ_m^j (or C_m^j), this method is invertible in the sense that we can recover the blurred and noisy image from the restored image. There is therefore no lost of data.

2.5 Multi-level shrinkage

Yet, the model we are presenting here is not fully satisfactory because the basis we use does not de-correlate the noise from the significant signal very well. The support of the wavelets in a fine frequency decomposition is indeed too large. On the other hand, if the wavelet packet decomposition is stopped at a higher level, the noise will be more de-correlated, but the basis will not be adapted to the deconvolution since it will not have a fine frequency decomposition and thus $\mathcal{F}(H^{-1}(n))$ will not be “uniform” enough on the support of the \hat{w}_k^j ’s. We are meeting here Heisenberg uncertainty principle.

One way to get round this difficulty is to shrink wavelet coefficients at different levels of the decomposition. We get indeed from one decomposition level to the next one by a simple change of orthonormal basis. Thus, whatever the level at which the shrinkage is made, if the parameters are constant, the same amount of noise will have been removed in the homogeneous zones (a white noise remains fully de-correlated in any orthonormal basis). The sole difference is that we might have not removed the same amount of information. Moreover, using

$$\forall \lambda_0, \lambda_1 \in \mathbb{R}_+, \quad \tau_{\lambda_0} \circ \tau_{\lambda_1} = \tau_{\lambda_0 + \lambda_1}, \quad (12)$$

or

$$\forall \lambda_1, \lambda_2, \lambda_3 \in \mathbb{R}, \quad \tilde{\tau}_{\lambda_2, \lambda_3} \circ \tilde{\tau}_{\lambda_1, \lambda_2} = \tilde{\tau}_{\lambda_1, \lambda_3}. \quad (13)$$

we see that it is possible to make a shrinkage at every level of the decomposition in such a way that the composition of all the shrinkages is equal to the shrinkage computed in the preceding sections for the last level of decomposition.

Let us explain this more clearly in the case of a 1D-signal and for the invertible FCNR. Assume we want to apply a shrinkage at the level j of the decomposition. Every packet E_m^j , for $m = 0, \dots, 2^j - 1$, is associated with two parameters $\lambda_1^{j,m}$ and $\lambda_2^{j,m}$. As the noise is white (we have chosen to consider the case when the deblurring is done after the denoising) and as $\lambda_1^{j,m}$ corresponds, in some sense, to a bound of the noise, it does not depend, actually, on m . Let us denote it therefore λ_1 . Let us denote also, for $m = 0, \dots, 2^j - 1$,

$$\lambda_2^{j-1,m} = \min\{\lambda_2^{j,2m}, \lambda_2^{j,2m+1}\}.$$

Since the function $\tilde{\tau}$ satisfies (13), we have

$$\tilde{\tau}_{\lambda_2^{j-1,m}, \lambda_2^{j,2m}} \circ \tilde{\tau}_{\lambda_1, \lambda_2^{j-1,m}} = \tilde{\tau}_{\lambda_1, \lambda_2^{j,2m}}.$$

We can therefore shrink first the packet E_{2m}^j (respectively E_{2m+1}^j) from λ_1 to $\lambda_2^{j-1,m}$, and then from $\lambda_2^{j-1,m}$ to λ_{2m}^j (respectively λ_{2m+1}^j). Hence, in the first time, the packets E_{2m}^j and E_{2m+1}^j undergo the same shrinkage. Moreover, since the noise is white, this can, as well, be made in the wavelet basis $\{w_m(2^{j-1}x - k), k \in \mathbb{Z}\}$ associated with the packet $E_m^{j-1} = E_{2m}^j \oplus E_{2m+1}^j$ (see the remark above). So the problem has been moved to the issue of shrinking every packet E_m^{j-1} from λ_1 to λ_m^{j-1} (and afterwards, every packet E_{2m}^j and E_{2m+1}^j from λ_m^{j-1} to, respectively, λ_{2m}^j and λ_{2m+1}^j) (see Figure 5). We can now make the same reasoning to show that the shrinkage could, as well, have been started at the level $j-2$, every packet E_m^{j-2} being shrunk from λ_1 to $\min\{\lambda_{2m}^{j-1}, \lambda_{2m+1}^{j-1}\}$. Making again the reasoning for each $j' < j-2$, we finally show that we can apply the function $\tilde{\tau}$ at every step of the decomposition. Notice that we use here the fact that the noise remains “almost” Gaussian after shrinkage in such a way that its law does not change under orthonormal change of basis and that it remains white.

We can, the same way, apply a multi-level shrinkage in the non-invertible FCNR. But the thresholds must be computed more carefully as the noise is no longer Gaussian after thresholding. It remains however white (in the sense that the pixels are i.i.d random variables) after a change of level of decomposition.

Remark that, since the shrinkage is made at every level, we do not need to compute the best basis that is the wavelet-packets basis the most adapted to the operator H (the basis in which H is the most diagonal), but that it might have been interesting from a numerical point of view [19, 25].

In addition to the reason explained above, this technic of shrinking at different levels will also give a better preservation of the structures of the signal because the shrinkage is made in different bases.

In the same spirit, in order to fight against the lack of independence in translation of the wavelet-packets decomposition, one can apply the same treatment on several translated versions of the image and, afterward, make the mean of all the restored images [10].

Let us remark, finally, that in the case of a non-white noise which depends on the luminosity of the image, the FCNR can be generalized by computing, on every packet, shrinking parameters that depend on the mean luminosity of the image on the packet.

2.6 Numerical implementation

Let (m_0, m_1) denote the 1D-QMF associated with the cubic spline wavelets [25]. The 2D-QMF (respectively, wavelet-packets) are obtained by tensorial product $m_{\epsilon_1, \epsilon_2} = m_{\epsilon_1} \times m_{\epsilon_2}$ with $\epsilon_i \in \{0, 1\}$ (respectively, $w_{m,n}^j = w_m^j \times w_n^j$). From then on, we present the algorithm in the 2D point of view.

We get from \mathbb{Z}^2 to a finite grid (the support of the image) by periodization.

Let us consider the case when H is approximated by a diagonal operator in the wavelet-packets basis. For numerical reasons, we will not use the basis associated with the finest frequency level. Indeed, Wavelet-packets w_n for n large, are in general badly localized in the Fourier domain. Hence, there is no point of getting too far in the decomposition since we will not gain any precision in frequency. It will be stopped, at the fifth level. Therefore, the enhancement $C_{m,n}^5$ that is applied to the packet indexed by (m, n) on that level, has to be estimated.

To facilitate the measurement, the 1D QMF (m_0, m_1) is approximated by Shannon filters $(\sum_{k \in \mathbb{Z}} \chi_{[-\pi/2, \pi/2]}(\cdot + 2k\pi), \sum_{k \in \mathbb{Z}} \chi_{[\pi/2, 3\pi/2]}(\cdot + 2k\pi))$. The iterated filters $M_{m,n}^5$ are also characteristic functions, and we denote $S_{m,n}^5$ their supports. As it was said in section 2.2, the optimum $C_{m,n}^5$ depends on the image u . But for

convenience, the $C_{m,n}^5$'s are estimated independently of the image by ¹

$$C_{m,n}^5 = \text{card}(S_{m,n}^5) \left(\sum_{(k,l) \in S_{m,n}^5} |\hat{h}(k,l)| \right)^{-1},$$

where $\text{card}(S_{m,n}^5)$ denotes the cardinal of $S_{m,n}^5$. This enhancement is well defined unless \hat{h} is identically null on $S_{m,n}^5$. But when it is very large (or when it is not defined), we make the associated packet vanish since it should contain essentially noise without information.

For the non-invertible FCNR, the threshold $\lambda_{m,n}^5$ is obtained from $C_{m,n}^5$ with a numerical estimation such as the one yielding Figure 1.

$$\lambda_{m,n}^5 = \lambda(C_{m,n}^5).$$

But the packets $E_{0,0}^5$ does not need to be thresholded because it corresponds to very low frequencies. We let therefore

$$\lambda_{0,0}^5 = 0.$$

In the case of the invertible FCNR, we let

$$\lambda_2^{5,m,n} = \frac{\lambda_1 \sigma_0}{C_{m,n}^5 \sigma},$$

where λ_1 is a parameter which plays the role of a bound for the noise², σ is the standard deviation of the noise before restoration and σ_0 is the standard deviation we want to reach. As for the first method, we let

$$\lambda_2^{5,0,0} = 0.$$

In order to apply the shrinkage at the higher levels, we let also

$$\lambda_2^{j,m,n} = \min_{\epsilon_1, \epsilon_2 \in \{0,1\}} (\lambda_2^{j+1, 2m+\epsilon_1, 2n+\epsilon_2}),$$

for $j = 4, 3, 2, 1, 0$ (note that $\lambda_2^{j,0,0} = 0$). Then every packet $E_{2m+\epsilon_1, 2n+\epsilon_2}^{j+1}$ is associated with the parameters $(\lambda_2^{j,m,n}, \lambda_2^{j+1, 2m+\epsilon_1, 2n+\epsilon_2})$, for $j = 0, 1, 2, 3, 4$.

The wavelet-packets coefficients $(u_{m,n}^1)_{k,l}$ of the image at the 1st level are then computed and they are shrunk with respect to the rule exposed above. The operation is next repeated at the next levels until the fifth one.

$$u_{2m+\epsilon_1, 2n+\epsilon_2}^{j+1} = T_{\epsilon_1} \times T_{\epsilon_2}(\tilde{u}_{m,n}^j),$$

¹Remark that when the deblurring is done without approximation, the standard deviation of the noise on each packet is enhanced by

$$\|M_{m,n}^5 / \hat{h}\|_2.$$

Therefore the enhancement $C_{m,n}^5$ can be approximated by

$$C_{m,n}^5 = \text{card}(S_{m,n}^5)^{-1} \left(\sum_{(k,l) \in S_{m,n}^5} \frac{1}{|\hat{h}(k,l)|^2} \right)^{\frac{1}{2}},$$

Such $C_{m,n}^5$'s are less "stable" than the proposed one (they are not finite when \hat{h} vanishes on $S_{m,n}^5$). If \hat{h} vanishes on $S_{m,n}^5$, the coefficients on the packet are replaced by zero.

²In fact,

$$(C_{m,n}^5)^2 \int_{\lambda_1}^{\infty} (x - \lambda_1)^2 e^{-\frac{x^2}{2\sigma^2}} dx \ll (\sigma_0)^2$$

is sufficient.

with

$$(\tilde{u}_{m,n}^j)_{k,l} = \tilde{r}_{\lambda_2^{j-1, \lfloor m/2 \rfloor, \lfloor n/2 \rfloor}, \lambda_2^j, m, n}((u_{m,n}^j)_{k,l}),$$

where $\lfloor x \rfloor$ denotes the integer part of x .

Then, after deblurring, the reconstruction is done directly from level 5 to 0 with the help of the dual operators $T_m^{5*} \times T_n^{5*}$.

$$u = \sum_{m,n} T_m^{5*} \times T_n^{5*} (C_{m,n}^5 \tilde{u}_{m,n}^5).$$

3 Restoration by Rudin-Osher-Fatemi's method

3.1 Rudin-Osher-Fatemi's variational method

The second approach we are proposing was introduced for the first time by Rudin, Osher and Fatemi in [31]. This method consists in solving the following constrained minimization problem:

Minimize

$$J(u) = \int_{\Omega} |\nabla u|,$$

with

$$\int_{\Omega} |Hu - u_0|^2 = \sigma^2, \tag{14}$$

that is to say (if $\sigma \neq 0$), minimize

$$E(u) = J(u) + \lambda \int_{\Omega} |Hu - u_0|^2 \tag{15}$$

for some Lagrange multiplier λ .

The existence of the minimum has been proved, under very general conditions, in [8]. When H is invertible, the uniqueness is guaranteed since E is strictly convex.

The functional J is the "regularization functional" that measures the irregularity of u . Among all the candidates in the set defined by (14), we choose the most regular one in the sense of J . In practice, we will not evaluate the Lagrangian λ ; it will be left as a parameter. It is clear that a large λ will lead to a noisy image with sharp edges. Conversely, for a weak λ , the image will be strongly denoised but blurred.

The term $\int_{\Omega} |\nabla u|$ denotes the Total Variation of the function u supported on an open subset Ω of \mathbb{R}^n (in practice, $n = 2$),

$$\int_{\Omega} |\nabla u| = \sup_{\varphi \in \mathcal{B}} \int_{\Omega} u(x) \operatorname{div} \varphi(x) dx,$$

where $\mathcal{B} = \{\varphi \in \mathcal{C}^1(\Omega, \mathbb{R}^2) / \varphi \text{ is compactly supported and } \|\varphi\|_{\infty} \leq 1\}$. In many papers, the notation $|Du|(\Omega)$ is also used. It reminds us that the derivative of u is not a function but a measure.

One of the most notable property of the method, when it is applied to a 1D-signal, is that it does not smooth, in general, the discontinuities or strong variations of the signal. We can easily see that for any monotone function u increasing from a to b , $\int_{\Omega} |\nabla u| = |b - a|$. In other terms

Proposition 1 *If $\Omega \subset \mathbb{R}$, and for any signal $u \in BV(\Omega) = \{u \in L^1(\Omega) / \int_{\Omega} |\nabla u| < \infty\}$, and any homeomorphism f from Ω to Ω ,*

$$\int_{\Omega} |\nabla(u \circ f)| = \int_{\Omega} |\nabla u|.$$




Figure 6: The continuous and dashed signals are deduced one from another by a change of variable, while the dot one corresponds to a limit case. The three signals have the same Total Variation.

Proof. We have

$$\begin{aligned}
 \int_{\Omega} |\nabla(u \circ f)| dx &= \sup_{\varphi \in \mathcal{B}} \int_{\Omega} u \circ f(x) \varphi'(x) dx \\
 &= \sup_{\varphi \in \mathcal{B}} \int_{\Omega} u(y) \varphi' \circ f^{-1}(y) (f^{-1})'(y) dy \\
 &= \sup_{\varphi \in \mathcal{B}} \int_{\Omega} u(y) (\varphi \circ f^{-1})'(y) dy \\
 &= \sup_{\phi \in \mathcal{B}} \int_{\Omega} u(y) \phi'(y) dy
 \end{aligned}$$

□

This means that the minimization of the Total Variation does no discrimination between two signals deduced one from another by a change of variable (it is done by the data fidelity term (14)). Note that the smoothing of discontinuities leaves the framework of this proposition since a continuous and a discontinuous functions are not deduced one from another by a change of variable. However, by making the homeomorphism f tends to a function which is constant on some intervals, we observe that a discontinuous signal might have the same Total Variation as a smooth version (see Figure 6).

This property is unfortunately false in 2-D. However, in spite of the global behavior of the Total Variation, we can consider that on the neighborhood of edges with small curvature, the 1-D behavior will dominate.

In 2-D, the Total Variation cannot only be seen as the measurement of the oscillations of an image, but also as the sum of its level lines length.

Theorem 1 (Coarea formula) For any $u \in BV(\Omega)$,

$$\int_{\Omega} |\nabla u| = \int_0^{\infty} P(\mathcal{X}_l u) dl$$

where $\mathcal{X}_l u$ denotes the level set associated to the gray level l , and $P(\mathcal{X}_l u) = \int_{\Omega} |\nabla \mathcal{X}_l u|$ its perimeter.

The proof of the theorem is in [15]. For almost every level line, $P(\mathcal{X}_l u)$ corresponds to the one dimensional Hausdorff measure \mathcal{H}^1 of $\partial \mathcal{X}_l u$. By minimizing the Total Variation, we tend to cancel the oscillations of the gray level (which is represented by the domain of integration), and to smooth the level lines (by minimizing their length).

With regard to the fidelity term, remark that, as it was also mentioned for the first method, we can adapt Rudin-Osher-Fatemi's method to the case when the noise depends on the luminosity of the image by incorporating in (14) the noise model.

3.2 Non-invertible blur and spectrum interpolation

This section is devoted to the case when H is not invertible. Let us express (14) in terms of frequency,

$$\int_{\text{supp}(\hat{h})} |\hat{h}(\xi)|^2 |\hat{u}(\xi) - \hat{u}_0(\xi) / \hat{h}(\xi)|^2 \xi = \sigma^2 - \int_{\Omega \setminus \text{supp}(\hat{h})} |n(\xi)|^2, \quad (16)$$




Figure 7: The functions $x \mapsto |\cos(x)|$ (solid line) and $x \mapsto |\cos(x)| + \alpha \cos(x)$ (with $\alpha \in [-1, 1]$) have the same Total Variation.

where $\hat{\Omega}$ denote the support of the Fourier domain. We observe in (16) that the smaller $\hat{h}(\xi)$ is, the less $\hat{u}(\xi)$ is constrained to be close to $\hat{u}_0(\xi)/\hat{h}(\xi)$. In the limit case, when $\hat{h}(\xi) = 0$, we do not have any constraint on $\hat{u}(\xi)$. Thus its value is interpolated by making $\int_{\Omega} |\nabla u|$ decrease. Such an idea has been explored by Guichard and one of the authors in [16] to get a better resolution of an image by reconstructing its high frequencies. The same way, we can study the reconstruction of intermediate frequencies.

For simplicity, let us restrict ourselves to the case when $\sigma = 0$ (or $n \equiv 0$) and $\hat{h} = \chi_{\hat{\Omega}} - \chi_K$ where K is a compact³ set of $\hat{\Omega}$. In other words,

$$\text{Minimize } J(u) = \int_{\Omega} |\nabla u|, \text{ with } \hat{u}(\xi) = \hat{u}_{\infty}(\xi) \text{ for all } \xi \in \hat{\Omega} \setminus K, \quad (17)$$

where u_{∞} is defined in (1).

A drawback of the model, is the non-uniqueness of the minimum, due to the fact that the Total Variation and the constraint are not strictly convex. For instance, let $u_{\infty} := x \mapsto |\cos kx| \in BV(\mathbb{T})$ for an integer k , and minimize $\int_{\mathbb{T}} |\nabla u|$ under the constraint that $\hat{u}(l) = \hat{u}_{\infty}(l)$ for all $l \notin \{-k, k\}$. We can easily check that $x \mapsto |\cos kx| + \alpha \cos kx$ is a minimizer for every $\alpha \in [-1, 1]$ (see Figure 7).

Therefore, we have to make a choice among all the minimizers. In the more general framework of Rudin-Osher-Fatemi's model, this choice can actually be arbitrary since the next proposition shows that minimizers are very "close" one from each other.

Proposition 2 *Let u_1 and u_2 be two \mathcal{C}^1 minimizers of J under a convex constraint, and $x \in \tilde{\Omega} := \{x \in \Omega / \nabla u_1(x) \neq 0 \text{ and } \nabla u_2(x) \neq 0\}$. Then there exists an open neighborhood V_x of x such that $\partial \mathcal{X}_{u_1(x)}(u_1) \cap V_x = \partial \mathcal{X}_{u_2(x)}(u_2) \cap V_x$.*

In other words, the sets of connected components in $\tilde{\Omega}$ of level lines of u_1 and u_2 are identical. Heuristically, this proposition means that we can get locally from one minimum to another by a "local change of contrast" (we can refer to [5] for an example of definition of "local change of contrast" related to level sets). Since the eye is not very sensitive to changes of contrast, we will not see much difference between every minimizer of E .

Proof. Let us consider two \mathcal{C}^1 minimizers u_1 and u_2 . Since J and the constraint are both convex, the image $(u_1 + u_2)/2$ is also a minimizer and we have

$$\int_{\Omega} |\nabla(u_1 + u_2)| = \int_{\Omega} |\nabla u_1| + \int_{\Omega} |\nabla u_2|.$$

And since for all $x \in \Omega$,

$$|\nabla(u_1 + u_2)(x)| \leq |\nabla u_1(x)| + |\nabla u_2(x)|,$$

³The fact that K is compact may appear unrealistic, but provides us with interesting examples and a good intuition of the way intermediate frequencies are reconstructed. Indeed, it is generally reasonable to assume that very large frequencies are null.

we must have, for all $x \in \Omega$ (the minimizers are \mathcal{C}^1),

$$|\nabla(u_1 + u_2)(x)| = |\nabla u_1(x)| + |\nabla u_2(x)|,$$

which implies that $\nabla u_1(x)$ and $\nabla u_2(x)$ are parallel. The conclusion of the proof comes from the fact that $\nabla u_i(x)$ is orthogonal to the level line of u_i passing by x . \square

Note finally that $\nabla u_1(x)$ and $\nabla u_2(x)$ have the same direction so that the change of contrast is monotone.

Let us go back to the interpolation problem (17). The objective is, of course, to recover the whole spectrum of u_∞ . First of all, remark that the method cannot obviously restore textures whose Fourier spectrum is located on K . Secondly, the interpolation by the Total Variation does not consist in filling a hole in the spectrum of the image by interpolating the neighboring frequencies. If \hat{u}_∞ vanishes in the vicinity of K , the minimum might be reached by a function whose restriction of the Fourier transform to K is not trivial. For instance, if $u_\infty := x \mapsto \cos x \in BV(\mathbb{T})$ and $K = \{-k, k\}$ where k is an odd integer larger than 1, the minimizer is $x \mapsto \cos x - \alpha \cos kx$ where $\alpha \geq 1/k^2$ (compare the extrema). The reconstruction depends therefore on the whole spectrum of the image.

This remark shows however that the method is, in general, not very efficient in the restoration of textures. Whereas the following proposition proves that it is more powerful to restore edges.

Proposition 3 *If $u_\infty \in BV(\Omega)$ is piecewise constant and $0 \notin K$, then u_∞ is the unique solution of (17).*

Proof. This property comes from the fact that if $v \in BV(\Omega)$ has bounded frequencies (and is therefore \mathcal{C}^1), we have

$$\int_{\Omega} |\nabla(u_\infty + v)| = \int_{\Omega} |\nabla u_\infty| + \int_{\Omega} |\nabla v| > \int_{\Omega} |\nabla u_\infty|,$$

if $\int_{\Omega} |\nabla v| \neq 0$. Moreover, $0 \notin K$ ensures that if v is constant, $v \equiv 0$. \square

Unfortunately, this proposition cannot be generalized to piecewise affine functions. For instance, if $u_\infty := x \mapsto |x - x_0| \chi_{[-\pi, 0]}(x) + x \chi_{[0, \pi]}(x) \pmod{2\pi}$ and $K = \{-k, k\}$ where k is an odd integer, the minimizer is $x \mapsto u_\infty(x) + \alpha \cos kx$ with $\alpha > 1/k$ (see Figures 8 and 9).

We see therefore that if u_∞ can be decomposed in the sum of a piecewise constant function v_∞ and a continuous “texture” w_∞ , we will restore v_∞ perfectly (we will not touch the discontinuities) while the reconstruction of w_∞ is uncertain.

We must however moderate this remark because the textures are, in general, very different from one zone of the image to another one. And we can expect therefore that they will somehow make up for one another in order to reconstruct w_∞ . This opinion can be supported by the following example.

Let us go back to the example given above: $u_\infty = \cos \in BV(\mathbb{T})$ and $K = \{-k, k\}$. For $N \in \mathbb{N}^*$, we can inject u_∞ in $BV(N\mathbb{T})$, the compact K becoming $\{-Nk, Nk\}$. Let us now change u_∞ in $u_\infty := x \mapsto \sum_{l=0}^{N-1} \chi_{[2\pi l, 2\pi(l+1)]}(x) \cos(x + 2\pi l/N)$, in order to study the compensation of the different phases of the textures. In such a situation, we can show that the associated minimizer is u_∞ . Indeed, let us say it is $x \mapsto u_\infty(x) + \alpha \cos(kx + \varphi)$ with $\alpha, \varphi \in \mathbb{R}$. We can easily check that $x \mapsto u_\infty(x) + \alpha \cos(kx + \varphi + 2\pi kl/N)$ is also a minimizer for any $l \in \mathbb{Z}$. Since the set of minimizer is convex, we see that $u_\infty(x) = (1/N) \sum_{l=1}^N (u_\infty(x) + \alpha \cos(kx + \varphi + 2\pi kl/N))$ is a minimizer. Then, using the proof of Proposition 2 (the derivatives of the minimizers have the same sign at almost every point of \mathbb{T}), we show that u_∞ is the only one.

The case $k = 1$ is the most interesting. Remark that $u_\infty - \cos$ vanishes on $(0, 2\pi)$ but the “parasitic” cosine makes the Total Variation increase on the complementary interval in such a way that the minimizer is u_∞ .

More generally, we can expect that, if the phases of the textures are well distributed, the reconstruction will be successful. Similar arguments can also be applied to the case when the frequencies are well-distributed.

the reasoning above does not apply to the case of images supported on a finite grid. We can however see on Figure 10 that the spectrum interpolation by minimization of the Total Variation permits to cancel the ringing without touching edges.




Figure 8: The piecewise affine function $u_\infty(x)$.




Figure 9: The function $u_\infty(x) + \alpha\cos(x)$ with α minimizing the Total Variation.

3.3 Staircase effects, soft-thresholding and approximation of the functional

The Achilles' heel of Rudin-Osher-Fatemi's method is the well-known "stair-case" effect which, in dimension 2, is generalized by the creation of large homogeneous zones on the image. One of the consequences of this is the lost of low contrast textures.

Nikolova showed in [28] that this effect was due to a lack of regularity at the origin of the function $x \mapsto |x|$ used in $\int_\Omega |\nabla u|$. More precisely, this phenomenon can be explained, in the case of 1D discrete signal by the presence of a hidden soft-thresholding on the derivative of the signal. We minimize indeed the functional

$$E(u) = \sum_{k \in K} |u_{k+1} - u_k| + \lambda \sum_{k \in K} |(Hu)_k - (u_0)_k|^2,$$

or, letting $d = Du$ with $d_k = (Du)_k = u_{k+1} - u_k$,

$$\tilde{E}(d) = \sum_{k \in K} |d_k| + \lambda \sum_{k \in K} |(HD^{-1}d)_k - (u_0)_k|^2.$$

Deriving \tilde{E} , we can observe that the minimum d will satisfy

$$d_k = 0 \text{ when } |2(HD^{-1})_k^T \cdot (HD^{-1}d - u_0)| < \frac{1}{\lambda},$$

and we will have, otherwise

$$2(HD^{-1})_k^T \cdot (HD^{-1}d - u_0) = -\frac{\text{sign}(d_k)}{\lambda},$$

where $(HD^{-1})_k^T$ is the k^{th} column of HD^{-1} . In the particular case when $H = D$, the soft-thresholding appears very clearly

$$\begin{cases} d_k = 0 & \text{when } |(u_0)_k| < \frac{1}{2\lambda} \\ d_k = (u_0)_k - \frac{1}{2\lambda}(\text{sign}((u_0)_k)) & \text{otherwise} \end{cases}$$




Figure 10: Right-Left: an airplane image and its spectrum. Up-Down: u_∞ , u_0 and the restored image.

In the same spirit, Chambolle, De Vore, Lee and Lucier showed, in [6], that when $H = I$ (the identity), $\Omega = \mathbb{T}^2$ and the Total Variation is replaced by the Besov semi-norm $B_1^{1,1}(\Omega)$ (which is also a smoothness space of order 1 in $L^1(\Omega)$), the minimization of E comes down to a soft thresholding of the wavelets coefficients of the image. Moreover, when $\int_{\Omega} |\nabla u|$ is no longer approximated (but $H = I$), Cohen, De Vore, Petrushev and Xu proved, in [4], that one can obtain, by a wavelet shrinkage in the Harr basis, an image u_1 such that $E(u_1)$ is less than $C \min_{u \in BV(\Omega)} E(u)$, with C a constant not depending on u_0 and λ . All this makes, in some sense, the link between the two methods described in this paper.

Nikolova's remark suggests us to regularize the functional. Several variations have been made on the model. One of the most standard is the following.

Minimize

$$E_{\beta}(u) = \int \varphi_{\beta}(|\nabla u|) + \lambda \int |Hu - u_0|^2,$$

where φ_{β} is a regular approximation of the absolute value function as, for example, $\varphi_{\beta}(t) = \sqrt{\beta + t^2}$. This method will provide us with a better restitution of the low contrast textures and an image with less pixelization but it has the drawback to diffuse the gray levels on the surface of the image when the gradient is small. Nevertheless it is a good approximation of the original model for β small. Moreover, it is known to speed the convergence rate up, and therefore can be used to minimize the Rudin-Osher-Fatemi functional, by making β tend to 0[1].

Going back to the example given by Nikolova ($H = D$), when $\varphi_{\beta}(t) - \varphi_{\beta}(0)$ is equivalent to t^2 in the neighborhood of 0 (that is the case for $\varphi_{\beta}(t) = \sqrt{\beta + t^2}$), we can observe that d will satisfy

$$d_k \sim \frac{1}{1 + \lambda} (u_0)_k,$$

when $(u_0)_k$ is close to 0 (and $d_k \sim (u_0)_k$ when it is large enough)

At this point, we cannot avoid to draw a parallel with the method (and the function $\tilde{\tau}_{\lambda+1,1}$) introduced in section 2.4 in order to preserve the textures of the image by getting round the thresholding.

The reader can also refer to [7] for another technic to avoid the pixelization.

3.4 Numerical implementation

One of the main difficulty in the computation of E_{β} on a discrete grid, is the lack of isotropy of the gradient. The total variation will be therefore approximated by the sum of four terms corresponding to four different directions of the gradient.

We use the usual notations:

1. $x_i = ih, y_j = jh$, pour $i, j = 1, 2, \dots, N$ et $h > 0$;
2. $u_{ij} = u(x_i, y_j)$;
3. $\Delta_+^x u_{ij} = u_{i+1,j} - u_{ij}$ et $\Delta_-^x u_{ij} = u_{ij} - u_{i-1,j}$;
4. $\Delta_+^y u_{ij} = u_{i,j+1} - u_{ij}$ et $\Delta_-^y u_{ij} = u_{ij} - u_{i,j-1}$.

The function E_{β} is calculated as follows.

$$\begin{aligned} E_{\beta}(u) &= \sum_{i,j=0}^N \varphi_{\beta} \left(\sqrt{(\Delta_+^x u_{ij})^2 + (\Delta_+^y u_{ij})^2} \right) + \sum_{i,j=0}^N \varphi_{\beta} \left(\sqrt{(\Delta_-^x u_{ij})^2 + (\Delta_+^y u_{ij})^2} \right) \\ &+ \sum_{i,j=0}^N \varphi_{\beta} \left(\sqrt{(\Delta_+^x u_{ij})^2 + (\Delta_-^y u_{ij})^2} \right) + \sum_{i,j=0}^N \varphi_{\beta} \left(\sqrt{(\Delta_-^x u_{ij})^2 + (\Delta_-^y u_{ij})^2} \right) \\ &+ 4 \lambda \sum_{i,j=0}^N |\hat{h}_{ij} \hat{u}_{ij} - (\hat{u}_0)_{ij}|^2, \end{aligned}$$

taking $\Delta_+^x u_{ij}$, $\Delta_-^x u_{ij}$, $\Delta_+^y u_{ij}$, $\Delta_-^y u_{ij}$ equal to 0 on the edge of the image, when this is necessary, and the term \hat{u} denotes the Discrete Fourier Transform of u .

This functional is minimized with a conjugate gradient descent algorithm. We start from the blurred image u_0 . To get u_{n+1} from u_n , the gradient of the function E_β at u_n , $\nabla E_\beta(u_n)$, is calculated at each step of the algorithm. Then the optimal amplitude of the variation of the image in the direction $-\nabla E_\beta(u_n)$ is estimated by the resolution of

$$\min_{s>0} E_\beta(u_n - s \nabla E_\beta(u_n))$$

using a trichotomy method. Once the optimal amplitude s_0 is calculated, we let

$$u_{n+1} = u_n - s_0 \nabla E_\beta(u_n),$$

and the process is iterated.

Note that in order to speed the convergence up, we can start from a deblurred version of u_0 instead of u_0 . Moreover, it is better to start with a large β and let it decrease to 0.

The reader can refer to T. Chan ([9] and other articles) for a general study on the minimization of Rudin-Osher-Fatemi's functional.

4 Numerical results and comparison

4.1 Description of the data (Physical model)

In order to compare the two methods described above they have been put in competition in a very practical situation that is the restoration of satellite SPOT images.

Earth observation satellites in the SPOT family are based on the PUSH-BROOM acquisition principle. A CCD linear array placed at the focal plane of a telescope image at a given moment, yields a scan-line perpendicular to the satellite track, while the satellite's motion on its orbit ensures acquisition of successive lines. The sampling time interval is adjusted so that the sampling grid thus generated is square. Its interval is then equal to the size of an elementary detector in the linear array as projected on the ground.

In order to analyze the problem from a "signal processing" point of view, the digital image is modeled as follows: The signal which is directly proportional to the luminance entering the instrument is convoluted by the impulse response h of the instrument, then sampled according to a given sampling grid which has a square mesh.

The Fourier transform \hat{h} of the impulse response h , is called the Modulation Transfer Function (MTF). It is the product of three MTF's corresponding to the imperfection of the optical, the imperfection and the size of the CCD detectors and the motion of the satellite.

At this point, let us consider two models of degradations which correspond to two different satellites: SPOT 5 and a specific mode of SPOT 1. In both cases, the Fourier transform of the impulse response h is supported on $[-1/2, 1/2] \times [-1/2, 1/2]$. Even if the real noise is the sum of three noises having different structures, it is approximated by a Gaussian noise. The standard deviation of this Gaussian noise is, in both cases, realistic and gives rise to the same difficulty as the real noise.

- The convolution kernel of the first model is given by

$$\hat{h}(\xi, \eta) = e^{-2\gamma_\xi|\xi| - 2\gamma_\eta|\eta|} \left(\frac{\sin(2\pi\xi)}{2\pi\xi} \right) \left(\frac{\sin(2\pi\eta)}{2\pi\eta} \right) \left(\frac{\sin(\pi\eta)}{\pi\eta} \right), \quad \xi, \eta \in [-1/2, 1/2],$$

where $\gamma_\xi = 1.505$, $\gamma_\eta = 1.412$ and the standard deviation of the noise $\sigma_1 = 2.4$ (see Figure 11).

- The convolution kernel of the second model is given by

$$\hat{h}(\xi, \eta) = e^{-2\gamma_\xi|\xi| - 2\gamma_\eta|\eta|} \left(\frac{\sin(4\pi\xi)}{4\pi\xi} \right) \left(\frac{\sin(4\pi\eta)}{4\pi\eta} \right), \quad \xi, \eta \in [-1/2, 1/2],$$

for the same values for γ_ξ and γ_η . The standard deviation of the noise is $\sigma_2 = 0.5$ (see Figure 12).




Figure 11: Profile of the “invertible” convolution kernel spectrum (for convenience, σ_1 is not at the real scale).




Figure 12: Profile of the non-invertible convolution kernel spectrum (σ_2 is not at the real scale).

These two models correspond to non-invertible operators H . However the first one can be considered as invertible since \hat{h} vanishes only at the frontier of the Fourier domain. The zones where the noise might blow up are therefore located in the neighborhood of the cut-off frequencies (see the hatching on Figure 11). Thus, The noise in these zones can be abruptly erased with a controlled risk of ringing.

Whereas the second FTM vanishes at intermediate frequencies. Putting at zero the bared zones on Figure 12 will generate ringing in the vicinity of edges. We can therefore expect to obtain better result with the variational method that will interpolate the spectrum of the image on these zones.

4.2 Method of comparison

In order to compare the restored images quantitatively and objectively, we have to measure them in front of a reference. For that reason, the blurred images have been re-created artificially from an airplane image. But in order to take into account the whole degradation suffered by the images (including aliasing), the airplane image must be over-sampled. Consequently, the restored images cannot be compared directly with the sharp image since they have not the same scale of sampling. The reference image has therefore been constructed from the over-sampled image by convolving it with a prolate function (see [22]) and afterwards making a sub-sampling. The obtained reference has an optimal resolution without aliasing. This is the best image we can hope to obtain from the degraded ones. The “artificial” images corresponding to the two models of degradation above are respectively called Blurred Image 1 and Blurred Image 2.

The difference between the restored images and the reference is measured with the l^1 and l^2 norms. Moreover, in order to evaluate the ability of each method to restore every kind of structure, we focus also the comparison on four zones which correspond to very different families of structures (see Figure 13).

- The first zone represents an oscillatory texture. This structure is well localized in high frequency. It gives information about the quality of the high frequency recovery. Moreover, this shows the ability of the algorithm to recover highly correlated textures.
- The second zone is nearly homogeneous. It permits to measure the ability of the method to remove noise.
- The third zone contains a strong edge which may produce Gibbs effect after the restoration.
- The fourth zone contains lots of sharp structures and thus does represent the quality of the sharpening.



Figure 13: Reference Image with zones on which statistics are computed.

In order to illustrate the difficulty to restore blurred and noisy images, a Wiener filter has also been applied on the images (see [2]). We have tried to fix the parameter of the filter (more properly, the assumed variance of the noise) such that the restored images have about the same quantity of noise in the homogeneous regions as the ones obtained by the other methods. Unfortunately, for the restoration of Blurred Image 1, we could not reach that goal without losing all the details of the image. So we have chosen a parameter yielding acceptable result (in spite of the remaining noise).

4.3 Results

The statistics are split into two tables. The first one sums-up the results obtained by the four methods, for Blurred Image 1 (see Table 1), and the second one, for Blurred Image 2 (see Table 2). However, we have to keep in mind that neither the FCNR nor the variational methods aim to minimize these errors (the FCNR controls the structure of the noise, and the variational method controls the regularity of the image). Let us indicate also that the parameters of these methods has not been fixed with regard to these statistics but on visual criterion.

In regard to Table 1, note that the invertible FCNR and the variational method yield, on the whole, comparable statistics while the results are a little worth for the non-invertible FCNR, and much worth, on any region, for the wiener filter. Looking at the different zones in detail, we observe that the FCNR's are more efficient than the variational method on Zones 1 and 2 and less on Zones 3 and 4, which confirms the fact that the Total Variation penalizes textures but preserves strong edges (it does not generate ringing in there vicinity).

If we turn now to Table 2, both wavelet-packets based algorithms yield poor results. The convolution operator of Blurred Image 2 is actually not invertible so that we left the framework in which these methods are valid.

These comments are confirmed when we look at the images. The “ideal” one is represented on Figure 13. This is the image we would like to reach while starting from the blurred images of Figure 14. Looking at this, we can figure out the complexity of this issue since lots of details seems to be definitely lost, especially for Blurred Image 2. But, thanks to the low noise level, some of these details will be recovered,

		Wiener filter	Non-invertible FCNR	Invertible FCNR	Variational method
Image	l^1	4.44	4.21	4	3.66
	l^2	6.54	5.97	5.56	5.57
Zone 1	l^1	5.92	5.14	4.29	5.56
	l^2	7.3	6.38	5.39	6.86
Zone 2	l^1	2.82	1.79	1.82	1.87
	l^2	3.46	2.31	2.34	2.4
Zone 3	l^1	3.53	3.33	3.10	2.77
	l^2	4.51	4.25	3.89	3.81
Zone 4	l^1	5.15	5.30	4.99	4.47
	l^2	6.66	7.23	6.67	6.01

Table 1: Statistics for Blurred Image 1.

		Wiener filter	Non-invertible FCNR	Invertible FCNR	Variational method
Image	l^1	5.08	6.69	6.72	4.14
	l^2	7.38	9.27	9.31	5.8
Zone 1	l^1	5.79	6.05	6.08	4.64
	l^2	7.04	7.42	7.47	5.87
Zone 2	l^1	2	2.43	2.42	2
	l^2	2.56	3.17	3.16	2.62
Zone 3	l^1	4.27	6.98	7.01	2.92
	l^2	5.6	8.24	8.24	4.21
Zone 4	l^1	7.51	7.78	7.83	5.65
	l^2	10.21	10.42	10.48	7.59

Table 2: Statistics for Blurred Image 2.



Figure 14: Degraded images. Up-down: Blurred Image 1 and Blurred Image 2..

without creating artifact or noise enhancement.

The restored images are represented on Figures 15, 16, 17 and 18. The Wiener restoration yields noise that we cannot remove without loosing a lot of information, but the results are reasonably sharp. Wavelet-packets based methods are very successful with Blurred Image 1, the restored images looking like being free of any noise and with a good sharpness. The invertible FCNR is even better in the sense that it yields more informations (see the cars on the bridge and the texture of Zone 1). On the other hand, in the case of the Blurred Image 2, the results are still very blurred and present Gibbs effects. At last, the Total Variation based method yields good result in both cases. But it tends to remove the texture of Zone 1, and to create homogeneous zones. It is the sole of the four methods which does not create any ringing when deblurring the second image. All these remarks are confirmed on the blow-ups of Figures 19 and 20.

References

- [1] R. Acart and C. Vogel. “*Analysis of bounded variation methods for ill-posed problems*”. Inverse problems 10, pp. 1217-1229, 1994.
- [2] H.C. Andrews and B.R. Hunt. “*Digital signal processing*”. Tech. Englewood Cliffs, NJ: Prentice-Hall, 1977.
- [3] S.M. Berman. “*Sojournes and Extremes os Stochastic Processes*”. Wadsworth, Reading, MA, 1989.
- [4] A. Cohen, R. De Vore, P. Petrushev and H. Xu. “*Nonlinear Approximation and the Space $BV(\mathbb{R}^2)$* ”. Preprint.
- [5] V. Caselles, J.L. Lisani, J.M. Morel and G. Sapiro. “*Shape Preserving Local Histogram Modification*”. IEEE Transactions on Image Processing. Vol. 8, Number 2, Feb. 1999.
- [6] A. Chambolle, R.A. De Vore, N. Lee and B.J. Lucier. “*Nonlinear Wavelet Image Processing: Variational Problems, Compression and Noise Removal through Wavelet Shrinkage*”. Preprint CEREMADE No. 9728, September 1997, short version in: IEEE Trans. Image Processing, Vol. 7, No. 3, pp. 319-335, 1998.
- [7] A. Chambolle and P.L. Lions. “*Restauration de données par minimisation de la variation Total et variantes d’ordre supérieur*”. Proceedings of GRETSI, Juan-les-Pins, France, September 1995.
- [8] A. Chambolle and P.L. Lions. “*Image recovery via Total Variation minimisation and related problems*”. Numerische Mathematik vol. 76, pp. 167-188, 1997. IEEE Transactions on Image Processing.
- [9] R.H. Chan, T.F. Chan, C Wong. “*Cosine Transform Based Preconditioners for Total Variation Deblurring*”. UCLA Math Department CAM Report 95-23, 1995.
- [10] R.R. Coifman and D.L. Donoho. “*Translation-invariant de-noising*”. Technical Report 475, Standford University, May 1995.
- [11] R.R. Coifman, Y. Meyer and M.V. Wickerhauser. “*Wavelet analysis and signal processing. In Wavelets and their Applications*”. pp. 153-178, Boston, 1992. Jones and Barlett. B. Ruskai et al. eds.
- [12] G. Demoment. “*Image reconstruction and restoration: Overview of Common Estimation Structures and Problems*”. IEEE Trans. on acoustics, speech and signal processing, Vol. 37, No. 12, Dec. 1989.
- [13] D. Donoho and I.M. Johnstone. “*Minimax Estimation via wavelet shrinkage*”. Tech. Report, Dept. of Stat., Stanford Univ., 1992.
- [14] D.L. Donoho and I.M. Johnstone. “*Ideal spatial adaptation by wavelet shrinkage*”. Biometrika, 81 (3), pp. 425-455, 1994.
- [15] L.C. Evans and R.F. Gariepy. “*Measure theory and fine properties of functions*”. Studies in advanced mathematics, CRS Press Inc., 1992.



Figure 15: Restoration with a Wiener filter (Up-down: of Blurred Image 1 and Blurred Image 2).



Figure 16: Restoration with the non-invertible FCNR (Up-down: of Blurred Image 1 up Blurred Image 2).



Figure 17: Restoration with the invertible FCNR (Up-down: of Blurred Image 1 and Blurred Image 2).



Figure 18: Restoration with the variational method (Up-down: of Blurred Image 1 and Blurred Image 2).




Figure 19: Top: Reference Image. Left-Right: restoration of respectively Blurred Images 1 and 2. Up-down: with a Wiener filter, with the non-invertible FCNR, with the invertible FCNR, with the variational method.




Figure 20: Top: Reference Image. Left-Right: restoration of respectively Blurred Images 1 and 2. Up-down: with a Wiener filter, with the non-invertible FCNR, with the invertible FCNR, with the variational method.

- [16] F. Guichard and F. Malgouyres. “*Total Variation Based interpolation*”. Proceedings of European Signal Processing Conference 1998 (EUSIPCO-98). Vol. 3, pp. 1741-1744.
- [17] I.M. Johnstone and B.W. Silverman. “*Wavelet threshold estimators for data with correlated noise*”. Technical report, Standford University, Dec.1994.
- [18] T. Kailath. “*A View of Three Decades of Linear Filtering Theory*”. IEEE transaction on information theory, Vol. IT20, No. 2, March 1974.
- [19] J. Kalifa. “*Restauration minimax et déconvolution dans un base d’ondelettes miroirs*”. Thèse, Ecole Polytechnique, 1999.
- [20] J. Kalifa, S. Mallat and B. Rougé. “*Restauration d’Images par Paquets d’Ondelettes*”. 16eme Colloque GRETSI, 1997.
- [21] J. Kalifa, S. Mallat and B. Rougé. “*Image Deconvolution in Mirror Wavelet Bases*”. IEEE, ICIP’98.
- [22] H.J. Landau and H.O. Pollak. “*Prolate Spheroidal Wave Functions , Fourier Analysis and Uncertainty -III : The Dimension of the Space of Essentially Time and Band-Limited Signals*”. The Bell system technical Journal July 1962.
- [23] M. Lindenbaum, M. Fischer and A. Bruckstein “*On Gabor’s contribution to image enhancement*”. PR. Vol. 27, No. 1, pp. 1-8, Jan. 1994.
- [24] S. Mallat. “*A theory for multiresolution signal decomposition : the wavelet representation*”. IEEE Transaction on patern analysis and machine intelligence. Vol. II. No. 7, July 1989.
- [25] S. Mallat. “*A Wavelet Tour of Signal Processing*”. Academic Press, 1998.
- [26] Y. Meyer. “*Ondelettes et opérateurs*”. Hermann, 1990, tome 1.
- [27] Y. Meyer. “*les ondelettes, algorithmes et applications*”. 1992, Armand Colin.
- [28] M. Nikolova. “*Local strong homogeneity of a regularized estimator*”. To appear in SIAM.
- [29] B. Rougé. “*Remarks about space-frequency and space-scale to clean and restore noisy images in satellite frameworks. Progress in wavelets and applications*”. Ed. by Y.Meyer and S. Roques Frontières. Gif/Yvette 1993 (Proceedings Toulouse conference).
- [30] B. Rougé. “*Fixed Chosen Noise Restauration (FCNR)*”. IEEE 95 Philadelphia (USA).
- [31] L.Rudin, S. Osher and E. Fatemi. “*Nonlinear Total Variation based noise removal algorithms*”. Physica D 60, pp. 259-268, 1992.