



HAL
open science

Automated removal of quasiperiodic noise using frequency domain statistics

Frédéric Sur, Michel Grediac

► **To cite this version:**

Frédéric Sur, Michel Grediac. Automated removal of quasiperiodic noise using frequency domain statistics. *Journal of Electronic Imaging*, 2015, 24 (1), pp.013003/1-19. 10.1117/1.JEI.24.1.013003 . hal-01116309

HAL Id: hal-01116309

<https://inria.hal.science/hal-01116309v1>

Submitted on 13 Feb 2015

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

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

Journal of Electronic Imaging

JElectronicImaging.org

Automated removal of quasiperiodic noise using frequency domain statistics

Frédéric Sur
Michel Grédiac

SPIE



Automated removal of quasiperiodic noise using frequency domain statistics

Frédéric Sur^{a,*} and Michel Grédiac^b

^aUniversité de Lorraine, LORIA UMR CNRS 7503; CNRS, INRIA project-team Magrit; Campus Scientifique, BP 239, Vandœuvre-lès-Nancy Cedex 54506, France

^bUniversité Blaise Pascal, Institut Pascal UMR CNRS 6602, CNRS, BP 10448, Clermont-Ferrand 63000, France

Abstract. Digital images may be impaired by periodic or quasiperiodic noise, which manifests itself by spurious long-range repetitive patterns. Most of the time, quasiperiodic noise is well localized in the Fourier domain; thus it can be attenuated by smoothing out the image spectrum with a well-designed notch filter. While existing algorithms require hand-tuned filter design or parameter setting, this paper presents an automated approach based on the expected power spectrum of a natural image. The resulting algorithm enables not only the elimination of simple periodic noise whose influence on the image spectrum is limited to a few Fourier coefficients, but also of quasiperiodic structured noise with a much more complex contribution to the spectrum. Various examples illustrate the efficiency of the proposed algorithm. A comparison with morphological component analysis, a blind source separation algorithm, is also provided. A MATLAB® implementation is available. © 2015 SPIE and IS&T [DOI: 10.1117/1.JEI.24.1.013003]

Keywords: quasiperiodic noise; power spectrum distribution; notch filter design.

Paper 14340 received Jun. 11, 2014; accepted for publication Jan. 7, 2015; published online Feb. 11, 2015.

1 Introduction

Digital images may be affected by spurious repetitive (or at least structured) patterns covering the entire image. This unwanted feature is called quasiperiodic noise in the literature, since it is quasiperiodically distributed. The phenomenon may arise from electrical interferences during image acquisition, miscalibrated sensors, or missing data, causing, for instance, line dropout, striping, banding, or more complex background noise. Remote sensing applications are especially affected.¹ Particular reprographic techniques such as halftone printing or cathode ray tube monitors (which are affected by visible scanlines) are also likely to cause this phenomenon.

While white noise estimation and removal is the subject of a huge body of the literature (see, e.g., recent review papers^{2,3}), it turns out that quasiperiodic noise has been addressed by relatively few studies. The basic method has not changed since the dawn of image processing for robotic probe imaging.⁴ Periodic noise gives more or less sharp spikes in the Fourier domain, which have to be filtered out through notch filters to remove the corresponding noisy patterns. A reminder of the available approaches can be found in Ref. 5. The difficulty is to automate spike detection or, equivalently, notch filter design. Each Fourier coefficient, indeed, depends on the whole image content. Distinguishing between spikes caused by a localized texture or a repetitive structure (common in man-made environments) and spurious ones caused by periodic noise often requires prior knowledge. Some authors^{6,7} suggest detecting spikes in the Fourier domain as large deviations with respect to a localized median value. However, some parameters are still to be fixed by the user. It is important to mention that the efficient

stripe removal approach described in a recent paper⁸ is adapted to a subclass of stationary processes, namely convolved white noises. Even if periodic noise can be seen as a stationary process, it cannot be modeled as the convolution of a white noise with a simple kernel; therefore, it is not within the scope of the paper in question.

Distinguishing between the underlying image and pseudoperiodic noise can also be seen as a blind source separation problem.⁹ For instance, these two components can reasonably be assumed to be morphologically distinct in the sense of Ref. 10; that is, they have a sparse representation in two incoherent bases, one for the image component, the other for the periodic component. Reference 11 provides some image decomposition examples using curvelet representation for the image component and discrete cosine transform (DCT) for the periodic component. However, a localized, high-frequency texture is likely to be represented in the periodic component, and consequently to be discarded with the periodic noise. The sparsity assumption is also debatable. This is discussed in the experimental assessment.

The proposed contribution is the automated design of a notch filter for pseudoperiodic noise removal, which is especially interesting for blindly processing a series of images. It is based on the observation that, considering an image impaired by periodic noise, the sole periodic pattern present in every small patch extracted from this image is the periodic noise itself. Thus, an average power spectrum is calculated by averaging the power spectra of a set of patches from the noisy image. This basic idea is, in fact, not new and dates back to the 1980s. It was actually suggested in a short paper¹² and later applied to satellite image destriping.¹³ We propose to systematize this approach and to consider spikes as deviations in a model of the expected average

*Address all correspondence to: Frédéric Sur, E-mail: frederic.sur@loria.fr




Fig. 1 Mariner 4 experiment (1). (a) A 600×581 image from the Mariner 4 probe, impaired by pseudo-periodic noise. (b) Its power spectrum (log scale). (c) Power spectrum distribution against frequency. (d) Average power spectrum calculated for a set of 870 128×128 patches from the original image. (e) Average power spectrum distribution against frequency (logarithmic scales).

spectrum. It is known¹⁴⁻¹⁸ that the expected power spectrum of natural images can be modeled by a function which decreases with the inverse of some positive power of the frequency. Such a function should, thus, fit the average power spectrum, except at certain spikes expected to be due to the

noise pattern, which can subsequently be detected as statistical outliers.

The statistical background is explained in Sec. 2. The resulting algorithm is presented in Sec. 3 together with an illustrative example. Experiments are discussed in Sec. 4,

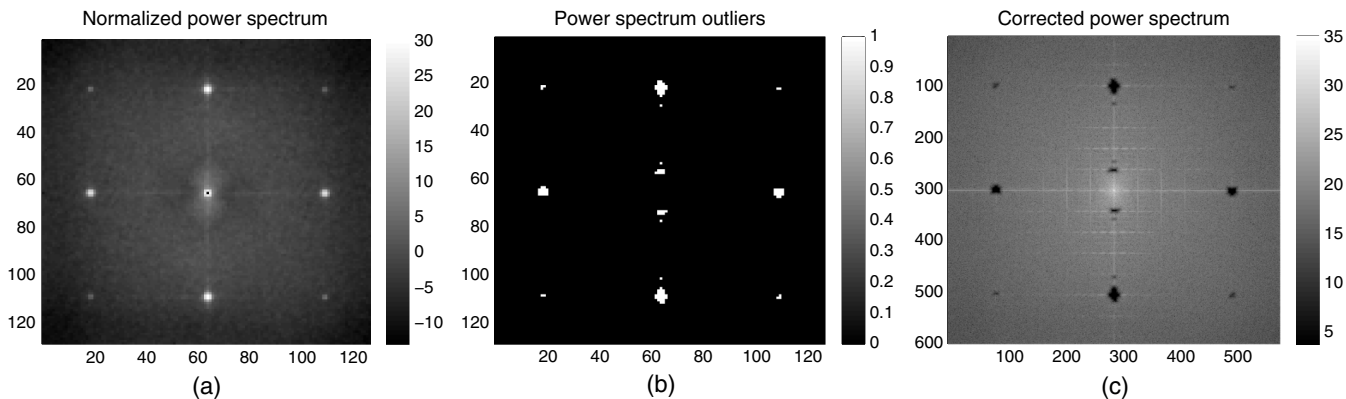


Fig. 2 Mariner 4 experiment (1). (a) Normalized average spectrum. (b) Outliers. (c) Corrected power spectrum.

and a comparison with morphological component analysis¹⁰ (MCA) is provided. We conclude with Sec. 5. A MATLAB® code is available.¹⁹

2 Automated Detection of Spurious Spikes Caused by Pseudoperiodic Noise

We propose to detect the spikes in the Fourier domain by averaging power spectra over a set of patches covering the original image: the spikes in this average spectrum are likely to be caused by the periodic noise covering the whole image. After providing some notations in Sec. 2.1, a stochastic model of the average power spectrum is proposed in Sec. 2.2.

2.1 Reminder and Notations

The discrete Fourier transform (DFT) of an image i of size $X \times Y$ is defined as the (X, Y) -periodic function

$$I(\xi, \eta) = \sum_{x=0}^{X-1} \sum_{y=0}^{Y-1} i(x, y) e^{-2\pi j(x\xi/X + y\eta/Y)}, \quad (1)$$

where $j^2 = -1$ and (ξ, η) have integer values. With $(\xi, \eta) \in [-X/2, X/2 - 1] \times [-Y/2, Y/2 - 1]$, $|I(\xi, \eta)|$ is the amplitude of a periodic component of frequencies ξ/X and η/Y (units are cycles per pixel) along each direction. $|I|$ is the amplitude spectrum of i and $|I|^2$ is the power spectrum of i .

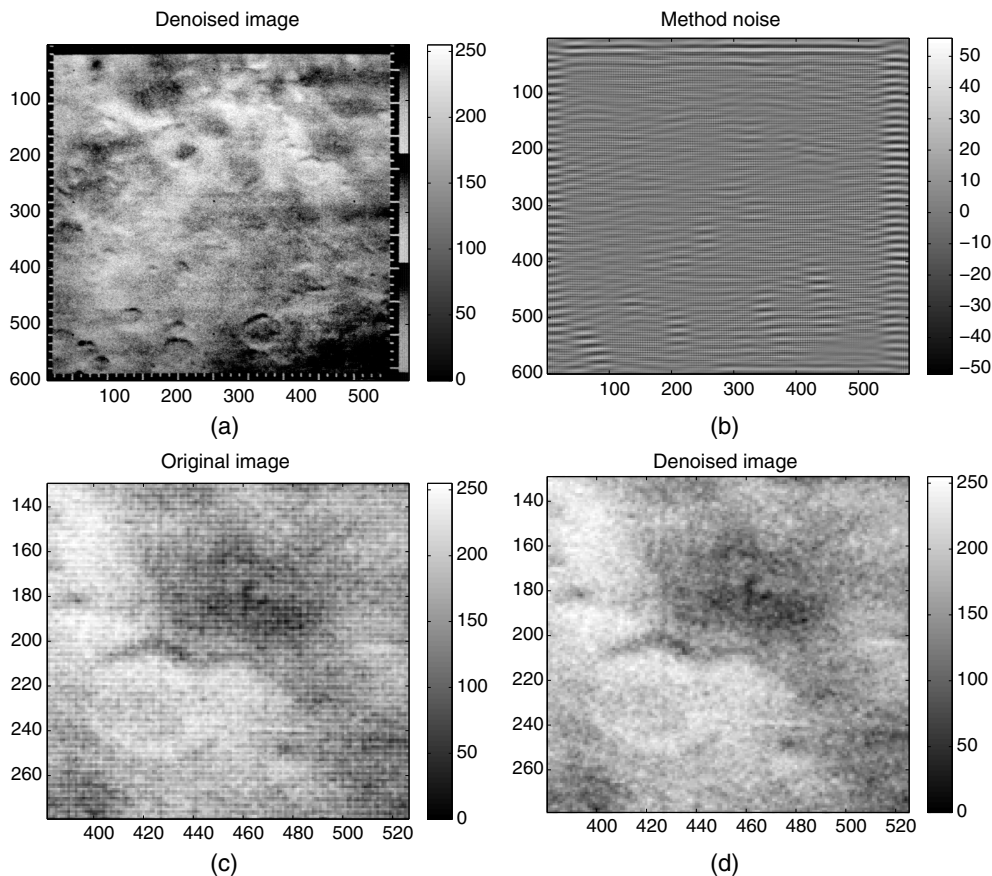


Fig. 3 Mariner 4 experiment (1). (a) Denoised image \hat{i} . (b) Noise component \hat{n} . (c) Close-up view of the noisy image. (d) Close-up view of the denoised image.

Since the pixel intensities $i(x, y)$ are real numbers, the (amplitude or power) spectrum is symmetric with respect to $(\xi, \eta) = (0, 0)$.

Both these spectra are affected by image edges and textures. For example, a part of image i with a periodic texture gives large $|I(\xi, \eta)|$ with (ξ, η) corresponding to the frequencies of the texture.

We also recall that a periodized translation by a vector (t_x, t_y) of any image i amounts to a multiplication of $I(\xi, \eta)$ by $e^{2\pi j(t_x \xi / X + t_y \eta / Y)}$. Thus, the amplitude and power spectra are not affected by translations.

Let us finish with some notations used in the rest of the paper. We note (f, ϕ) the polar coordinates of a frequency pair (ξ, η) where $f (= \sqrt{(\xi/X)^2 + (\eta/Y)^2})$ if measured in cycles per pixel is the radial distance from $(0,0)$ and ϕ is the counterclockwise angle from the ξ -axis. For any complex number z , \bar{z} is the conjugate and $\text{Re}(z)$ is the real part. Fourier transforms are capitalized.

2.2 Stochastic Modeling

We assume that a set of patches spanning the whole image of interest is available, each patch being affected by the underlying periodic noise and by additional white noise. The aim of this section is to argue that the average power spectrum

enables the periodic noise component to be characterized. The patches are assumed to be realizations of the following observation model.

Definition 2.1 A random patch p is the sum of the noise-free component p_0 , the periodic noise is n , and Gaussian white noise is w , i.e.,

$$p(x, y) = p_0(x, y) + n(x, y) + w(x, y), \tag{2}$$

where

- p_0 is a stochastic image patch;
- $n = \tau_{t_x, t_y}(n_0)$ is the translation of a pattern n_0 [with possible mild variations, hence modeled itself as a random image of expectation $E(n_0)$] by a random displacement (t_x, t_y) uniformly distributed in $[1, X] \times [1, Y]$. Since n_0 is the periodic noise component, $E(n_0)$ is the inverse Fourier transform of a sparse map (i.e., made of a few spikes);
- w is a Gaussian white noise of variance $\text{Var}(w)$;
- The aforementioned random variables are independent.

The Fourier transform being linear, the same relation as Eq. (2) holds in the Fourier domain

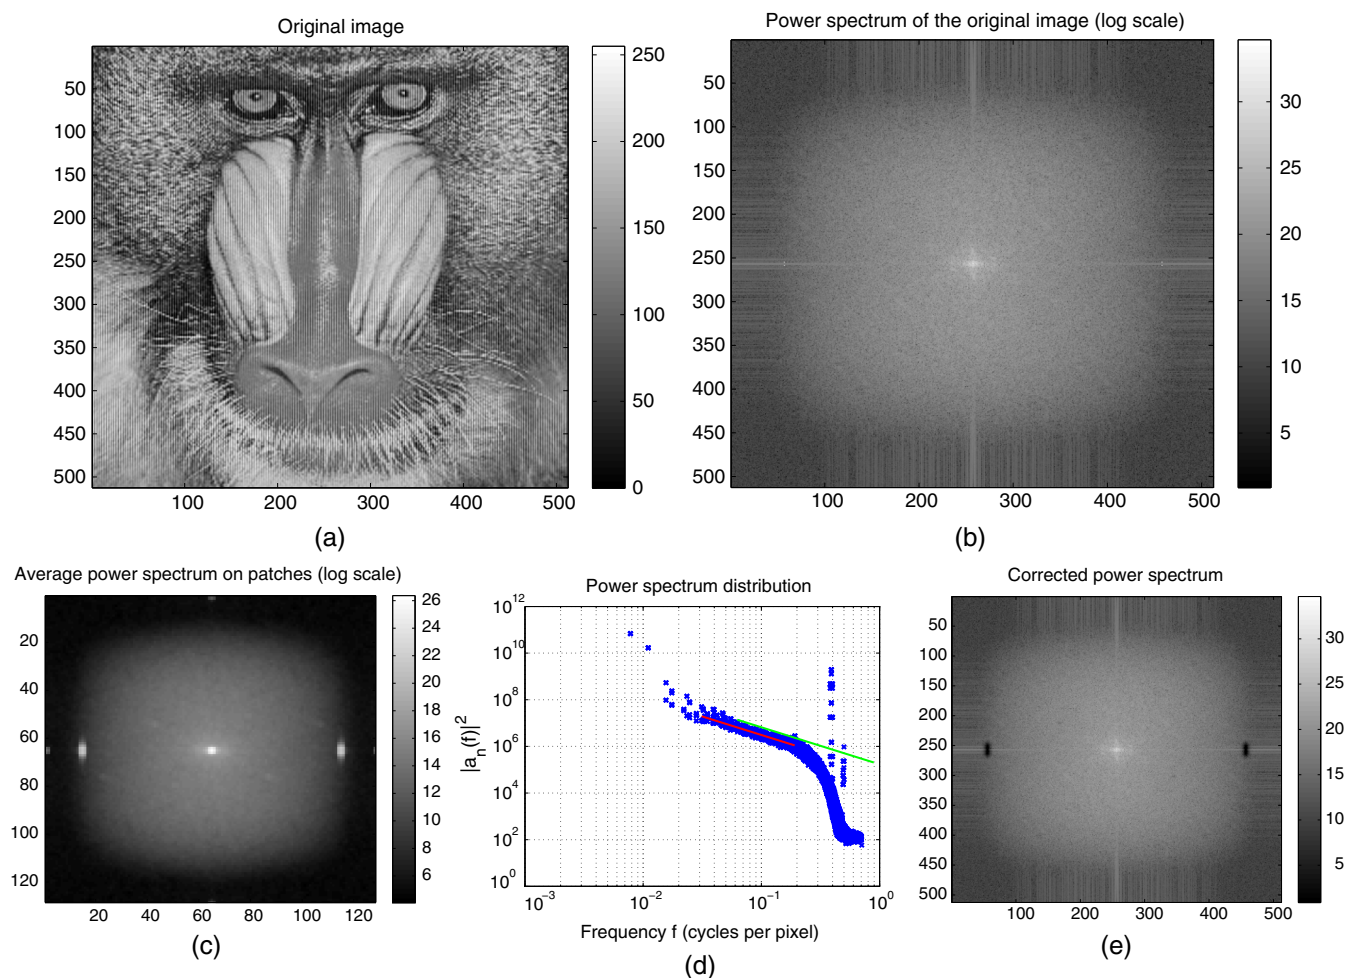


Fig. 4 Mandrill experiment. (a) Noisy image. (b) Power spectrum. (c) Average power spectrum. (d) Average power spectrum coefficient distribution. (e) Corrected power spectrum.

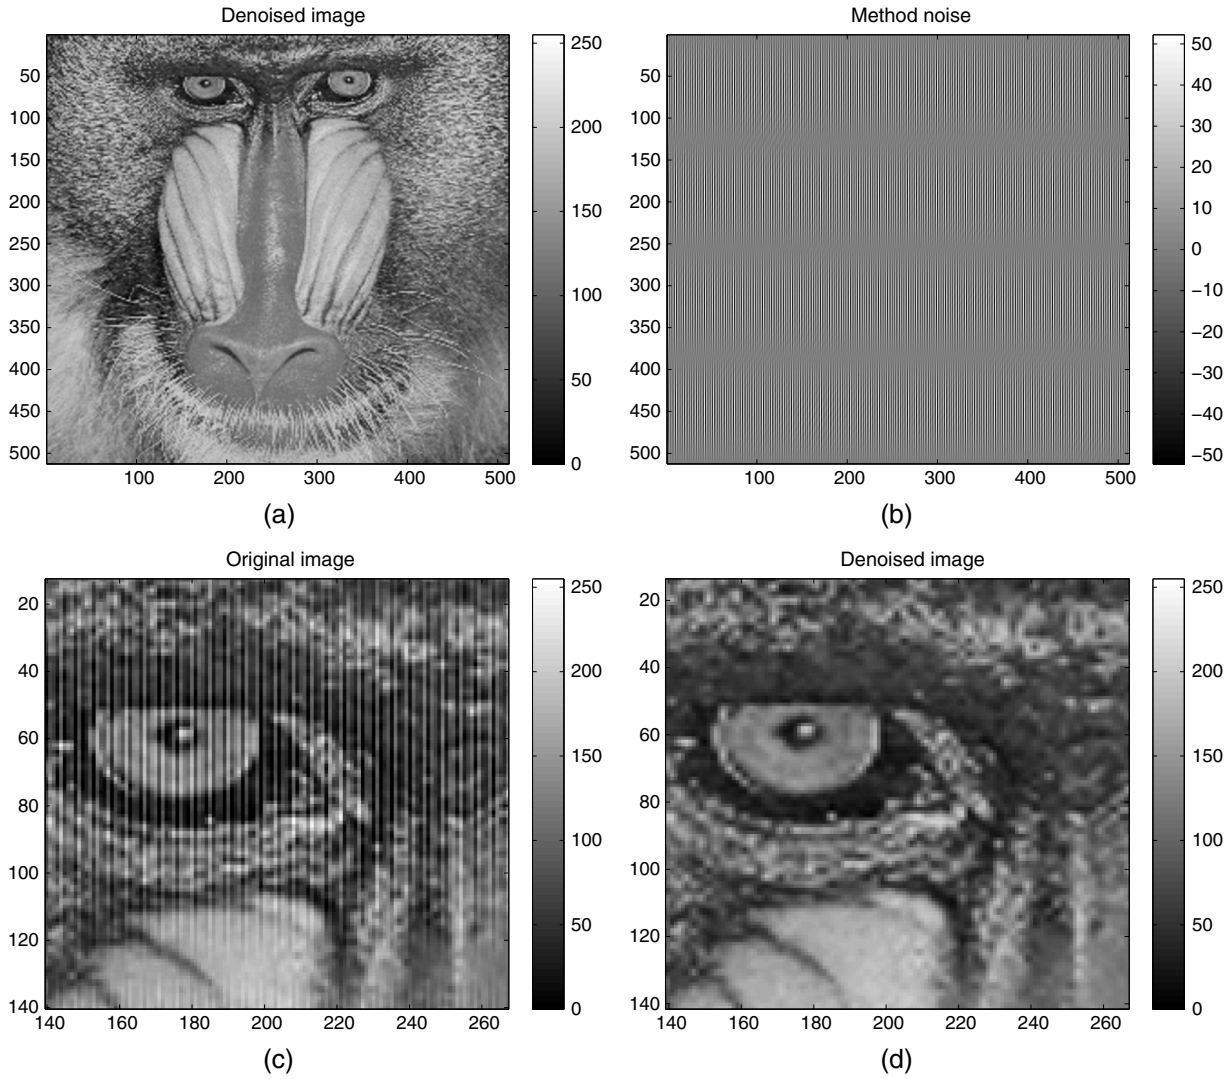


Fig. 5 Mandrill experiment. (a) Denoised image. (b) Estimation of the noise. (c) Close-up view of the noisy image. (d) Close-up view of the denoised image.




Fig. 6 Mandrill experiment. Map of the root-mean-square error (RMSE) as a function of periodic noise frequencies.

$$P(\xi, \eta) = P_0(\xi, \eta) + N(\xi, \eta) + W(\xi, \eta) \quad (3)$$

for any frequency pair (ξ, η) . Therefore, the power spectra satisfy

$$\begin{aligned} |P(\xi, \eta)|^2 &= |P_0(\xi, \eta)|^2 + |N(\xi, \eta)|^2 + |W(\xi, \eta)|^2 \\ &\quad + 2 \operatorname{Re}[P_0(\xi, \eta)\overline{N(\xi, \eta)} + W(\xi, \eta)\overline{N(\xi, \eta)} \\ &\quad + P_0(\xi, \eta)\overline{W(\xi, \eta)}]. \end{aligned} \quad (4)$$

Taking the expectation in Eq. (4), since the random variables are assumed independent, we obtain [mentions of (ξ, η) are skipped for brevity]

$$\begin{aligned} E(|P|^2) &= E(|P_0|^2) + E(|N|^2) + E(|W|^2) + 2 \operatorname{Re}[E(P_0)E(\overline{N}) \\ &\quad + E(W)E(\overline{N}) + E(P_0)E(\overline{W})]. \end{aligned} \quad (5)$$

From the definition in Sec. 2.1, $E(|W|^2) = XY \operatorname{Var}(w)$, $E(\overline{W}) = E(W) = E(w) = 0$, and, with the translation

property, $E(\bar{N}) = E(\bar{N}_0)E(e^{-2\pi j(t_x \xi/X + t_y \eta/Y)})$ and $|\bar{N}| = |\bar{N}_0|$. Consequently,

$$E(|P|^2) = E(|P_0|^2) + E(|N_0|^2) + XY \text{Var}(w) + 2 \text{Re}[E(P_0)E(\bar{N}_0)E(e^{-2\pi j(t_x \xi/X + t_y \eta/Y)})]. \quad (6)$$

Since t_x (respectively t_y) is assumed to be uniformly distributed in $[0, X]$ (respectively $[0, Y]$), $E(e^{-2\pi j(t_x \xi/X + t_y \eta/Y)}) = 0$. We conclude that

$$E(|P|^2) = E(|P_0|^2) + E(|N_0|^2) + XY \text{Var}(w). \quad (7)$$

Thus, the expected power spectrum of patch p comprises three components, namely:

1. The expected power spectrum of the noise-free patch p_0 . A large body of the literature has been dedicated to the modeling of the power spectrum of natural images (see, e.g., Refs. 14 and 17). This spectrum is known to be well modeled by a $1/f$ power law; that is, there exist positive real values α and A such that

$$E(|P_0(f, \phi)|^2) = \frac{A}{f^\alpha}, \quad (8)$$

where f and ϕ are defined in Sec. 2.1. Most natural images have $\alpha \simeq 2$ (or $\alpha \simeq 1$ when considering the amplitude spectrum). More sophisticated models are available: in Refs. 15 and 16; for example, it was observed that both A and α actually depend on the angle ϕ . We do not use such an advanced model in this paper, and simply use the $1/f$ power law given by Eq. (8).

2. The expected power spectrum of the quasiperiodic noise component is n_0 . It mainly consists of a few spikes, or relatively small regions spread along these spikes.
3. The contribution of the white noise w is constant over the whole spectrum. This contribution is dominated by the first two components, especially for low and medium frequencies. However, it is possible that it dominates the high frequencies because of the power law (first component) and the limited number of frequencies characterizing the quasiperiodic noise (second component).

Neglecting the white noise, we can see that the expected spectrum of the patches can essentially be modeled by the spikes caused by the pseudoperiodic noise, superimposed

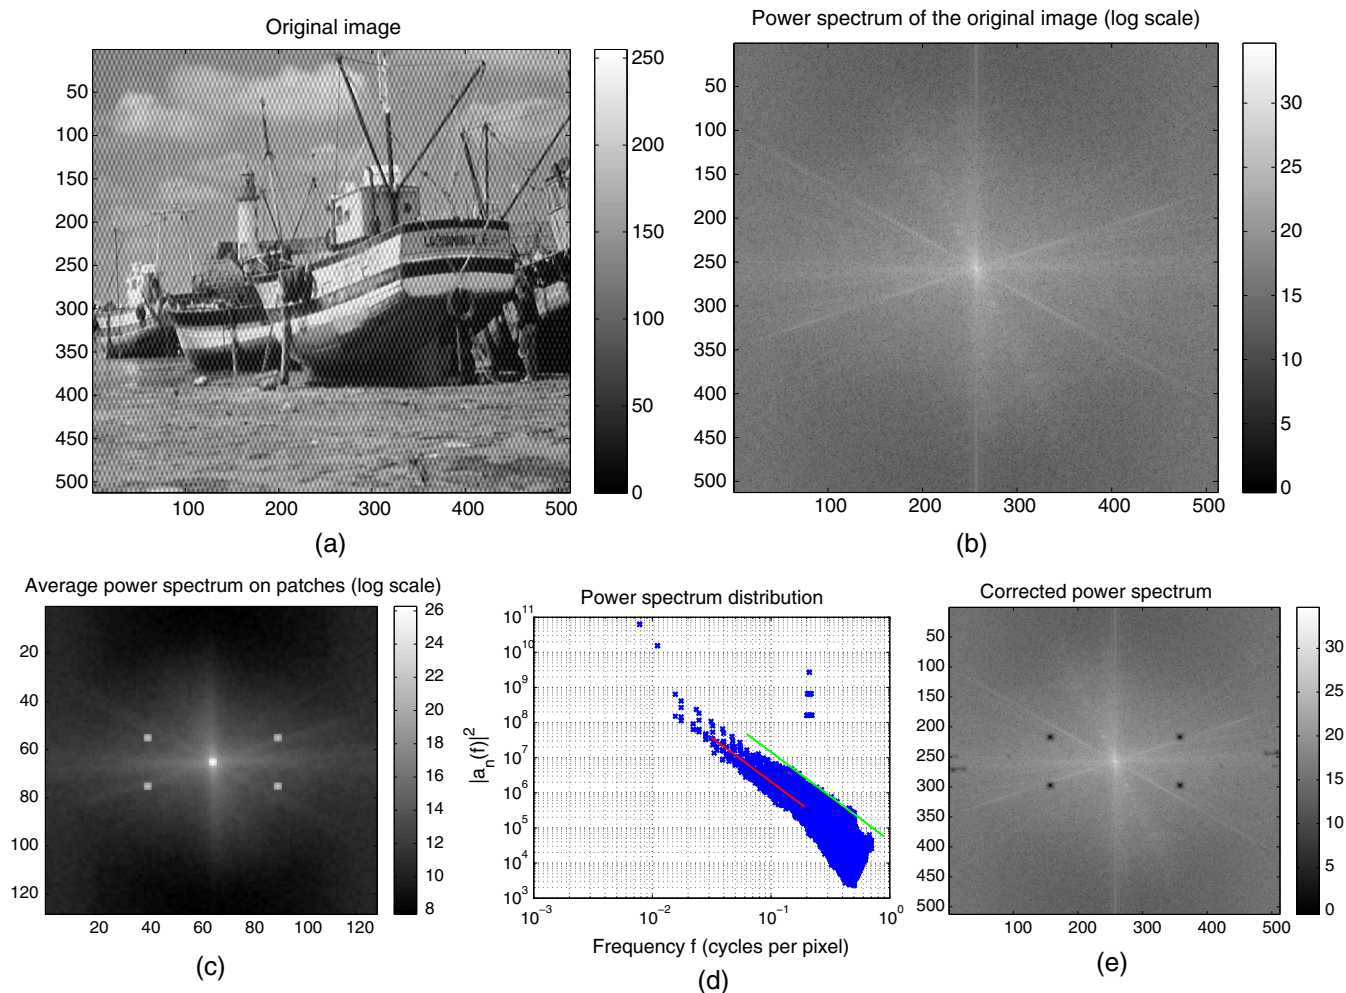


Fig. 7 Boat experiment. (a) Noisy image. (b) Power spectrum. (c) Average power spectrum. (d) Average power spectrum coefficient distribution. (e) Corrected power spectrum.

on the power law distribution. It turns out that this simplified model is realistic in all the images affected by pseudoperiodic noise that we have encountered.

Thus, fitting the power law to the average power spectrum gives an estimation of the parameters A and α , with the periodic noise components being statistical outliers. Furthermore, a notch filter n_f for an image i affected by quasiperiodic noise can be built from the localization of these outliers in the Fourier domain. An estimation of the denoised image is subsequently given by the inverse Fourier transform of $N_f I$ (or equivalently by the convolution of n_f and i). It should be noted that what we call “notch filter” does not remove a single (more accurately two because of the symmetry property) (ξ, η) frequency pair, but potentially removes the contributions to the image spectrum of a whole set of frequency pairs, all of them being outliers with respect to the expected spectrum. The following section gives an algorithm based on this idea to remove pseudoperiodic noise.




Fig. 9 Boat experiment. Map of the RMSE as a function of periodic noise frequencies.

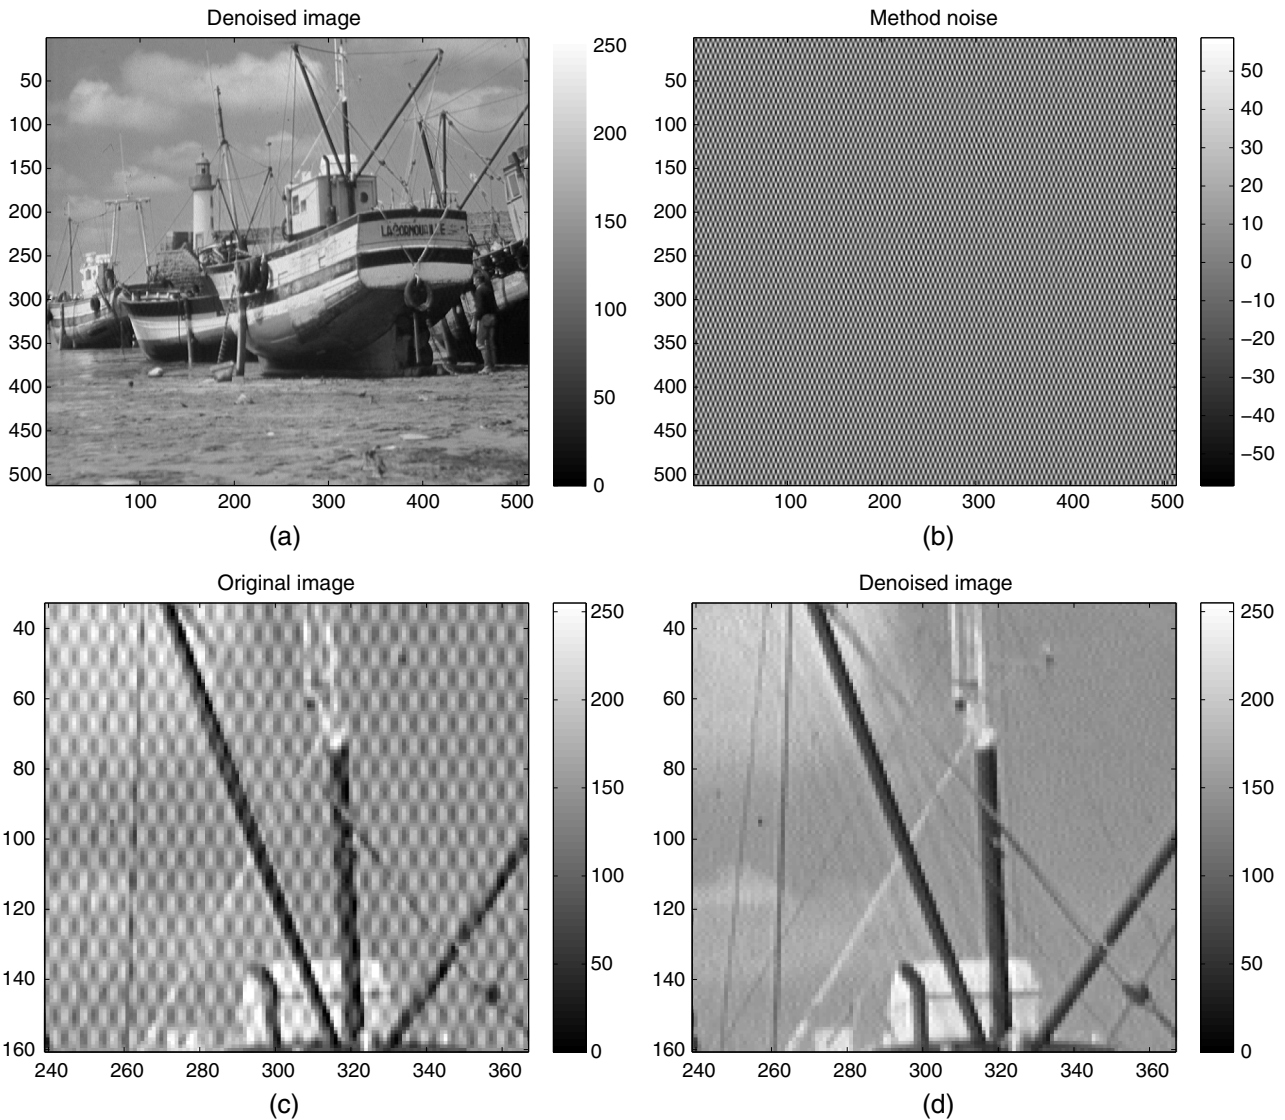


Fig. 8 Boat experiment. (a) Denoised image. (b) Estimation of the noise. (c) Close-up view of the noisy image. (d) Close-up view of the denoised image.

Table 1 Real images. From left to right: experiment number; image name and corresponding figures; size of the image; value of A ; value of α ; length L of the side of the patches; number N of patches covering the image to calculate the average power spectrum; computation time (in seconds, nonoptimized MATLAB code on Intel Xeon E3-1240 processor).

Number	Image	Size	A	α	L	N	Time
1	Mariner 4 (Figs. 1 to 3)	600 × 581	11.27	1.47	128	870	0.4
2	Apollo (Figs. 10 to 11)	474 × 630	10.86	1.38	128	704	0.4
3	Florida (Figs. 12 to 13)	808 × 754	11.46	1.91	256	288	0.5
4	Mariner 6 (Figs. 14 to 15)	461 × 471	9.60	2.31	128	462	0.2
5	Newspaper (Figs. 16 to 17)	831 × 730	12.41	1.25	200	572	0.6
6	Bending (Figs. 18 to 19)	1832 × 2382	7.36	2.88	300	2255	4.2

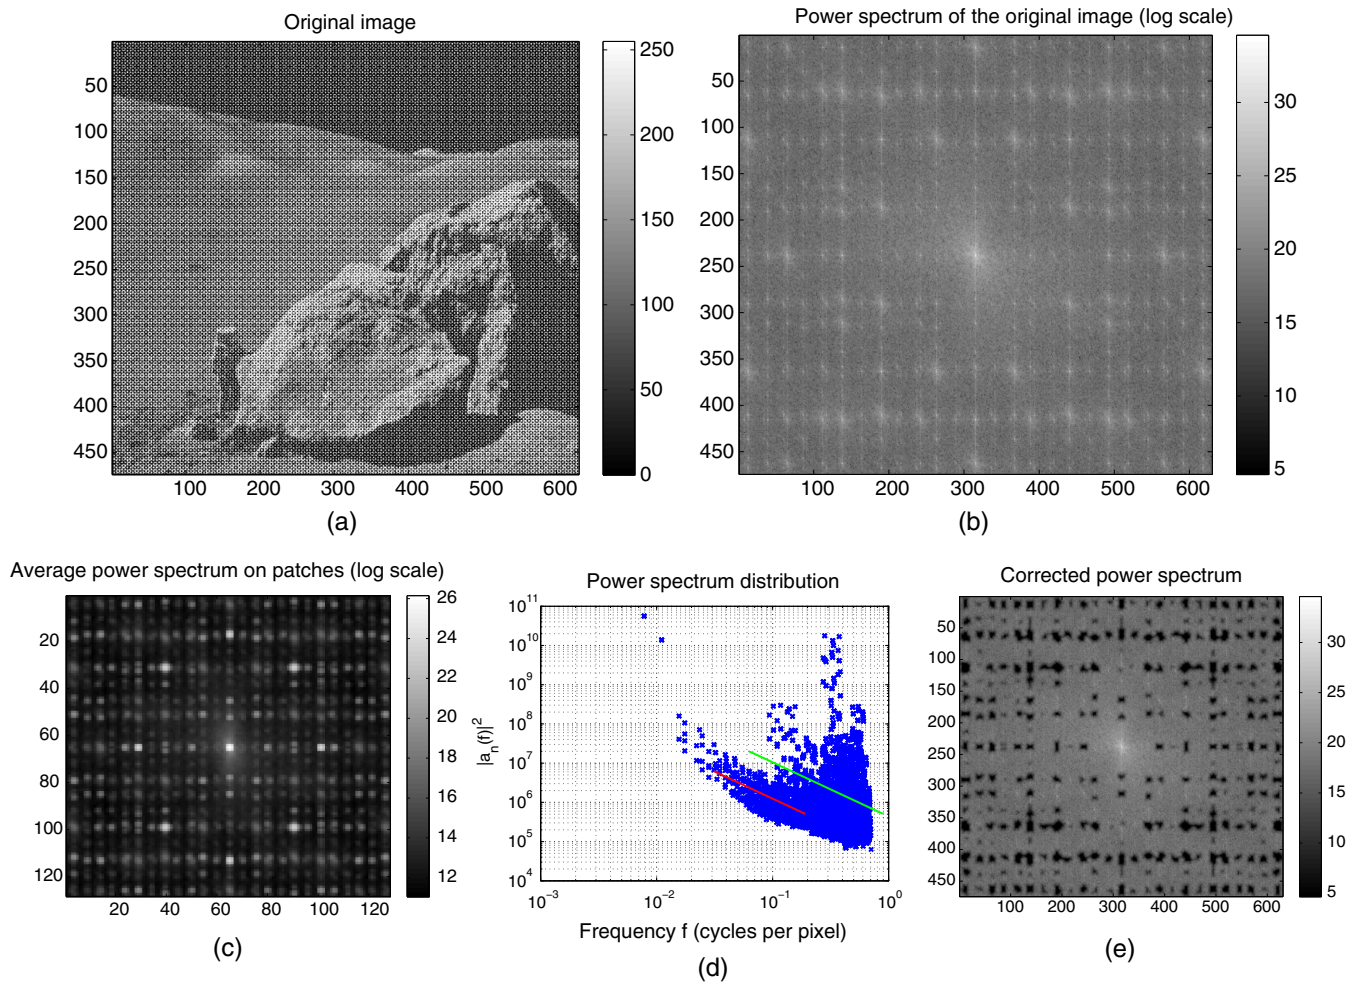


Fig. 10 Apollo experiment (2). (a) Noisy image. (b) Power spectrum. (c) Average power spectrum. (d) Power spectrum coefficient distribution. (e) Corrected power spectrum.

3 An Automated Algorithm to Remove Pseudoperiodic Noise from Natural Images

The algorithm is described in Sec. 3.1, practical considerations are discussed in Sec. 3.2, and an illustrative example is detailed in Sec. 3.3.

3.1 Algorithm

Given an image i of size $X \times Y$ impaired by quasiperiodic noise, the proposed algorithm consists of the following steps:

1. Calculate the average power spectrum (denoted $\widetilde{|P|^2}$) obtained by averaging the power spectrum of patches of size $L \times L$ regularly distributed in the image.
2. Fit a power law distribution to the power spectrum as a function of the frequency f , i.e., find A and α such that

$$\log[\widetilde{|P|^2}(f)] = A - \alpha \log(f). \quad (9)$$

The fitting is obtained by robust linear regression (we use iteratively reweighted least squares) on the $[\log(f), I(f, \phi)]$ scatter plot for f between f_0 and f_1 (to be specified) cycles per pixels. Robust regression gives consistent estimations which are not influenced by the spurious spikes due to pseudoperiodic noise. Least-squares estimation also gives the standard deviation σ of the residues.

3. Find the localization of upper outliers in the average power spectrum as frequency pairs (ξ, η) such that, under the common 3σ rule

$$\frac{\log(|\widetilde{P}|^2(\xi, \eta)) - [A - \alpha \log(f)]}{\sigma} > 3. \quad (10)$$

This results in an outlier map M_o^p such that $M_o^p(\xi, \eta) = 1$ if an outlier is present at (ξ, η) in the average spectrum of the patches, and $= 0$ otherwise. Note that a false-positive rate of 1% is expected under a Gaussian distribution. We restrict the outlier detection to frequencies $f > f_2$ (to be specified), since low frequencies do not correspond to repetitive patterns.

4. Resize the outlier map of size $L \times L$ to size $X \times Y$, giving a map M_o of the probable spurious spikes caused by quasiperiodic noise in the original image spectrum. Multiplying the initial image spectrum by $1 - M_o$ acts as a notch filter, eliminating the influence of the quasiperiodic noise.
5. Retrieve an estimation \hat{n} of the periodic noise component as the inverse Fourier transform of $M_o(\xi, \eta)I(\xi, \eta)$, and the estimated denoised image \hat{i} as $i - \hat{n}$ (i.e., the inverse transform of $[1 - M_o(\xi, \eta)]I(\xi, \eta)$).

3.2 Practical Considerations

The implementation details presented below do not play a crucial role in the good behavior of the algorithm, but are given in order to enable the algorithm to be recreated.

First, since most images have discontinuities between their left/right (respectively top/bottom) borders, their spectrum shows dominant straight lines along the horizontal axis (respectively vertical axis). To reduce these boundary effects, we multiply the patches p by a two-dimensional

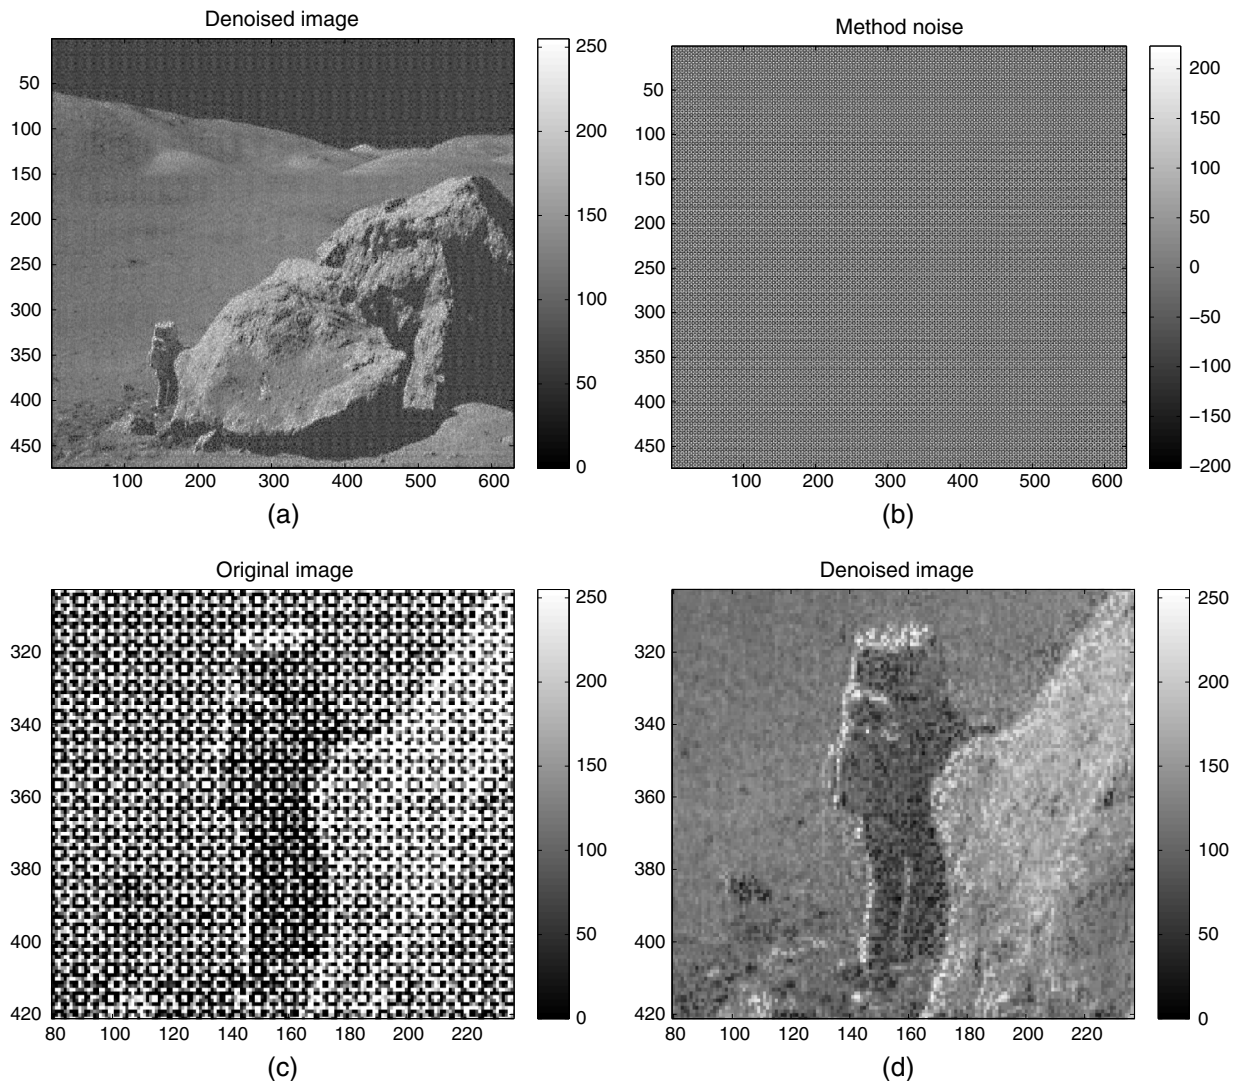


Fig. 11 Apollo experiment (2). (a) Denoised image. (b) Estimation of the noise. (c) Close-up view of the noisy image. (d) Close-up view of the denoised image.

Hann window with the same width L as the patches, as in earlier works.^{12,15,17}

In Step 1, the definition of the patches must be stipulated. Regarding the distribution of the sample of noncorrelated patches i , the absence of correlation is required to ensure the consistency of the sample mean estimator, hence patches with no or only limited overlapping. The size of the patches should be large enough to ensure both a good accuracy in the periodic noise spike detection (frequencies are distributed with $1/L$ steps in the power spectrum of a patch, cf. Sec. 2.1) and the detectability of low-frequency noise, but not too large, so as to make it possible to build enough independent patches from the noisy image of interest. Using $L \times L$ patches, we found that a good compromise is to take a sampling step of $L/8$ in both the horizontal and vertical directions, which gives a total number of patches equal to

$$\lceil 8(X - L)/L \rceil \times \lceil 8(Y - L)/L \rceil \quad (11)$$

with $\lceil \cdot \rceil$ rounding a number to the next larger integer. In addition, the average power spectrum is obtained by a geometric mean instead of an arithmetic mean, in order to limit the influence of large values, caused in particular by the periodic noise.

The values of f_0 , f_1 , and f_2 in Steps 2 and 3 are fixed as follows. We set $f_2 = 8/L$. This means that the sought noise frequency must be above $8/L$ cycles per pixel, i.e., eight observed cycles inside the patch shaped by the Hann window. The range in which the power law distribution of the spectrum coefficients should be satisfied is set to $f_0 = f_2/4$ (as only a few low frequencies are available) and $f_1 = 0.2$. Usual values for f_2 ensure that $f_1 > f_0$ is satisfied. In previous works,^{15,16} the upper limit is set to 0.35 cycle per pixel. The reason behind this is that high frequencies are mostly affected by white noise (cf. Sec. 2.2) and aliasing. In our patch-based average spectrum, we observed that the 0.35 value is rather optimistic in some experiments. This is probably explained by the particular nature of the considered images, and by the patch-based average spectrum, where high frequencies are penalized (they, indeed, are likely to correspond to localized phenomena), causing a fall-off with respect to the expected power law.

In Step 4, bilinear interpolation is used to expand the outlier map. The notch filter is obtained after convolving the outlier map M_o (after interpolation) by an isotropic Gaussian kernel of standard deviation 2 pixels, in order to limit ringing artifacts due to strong cut-offs in high frequencies.

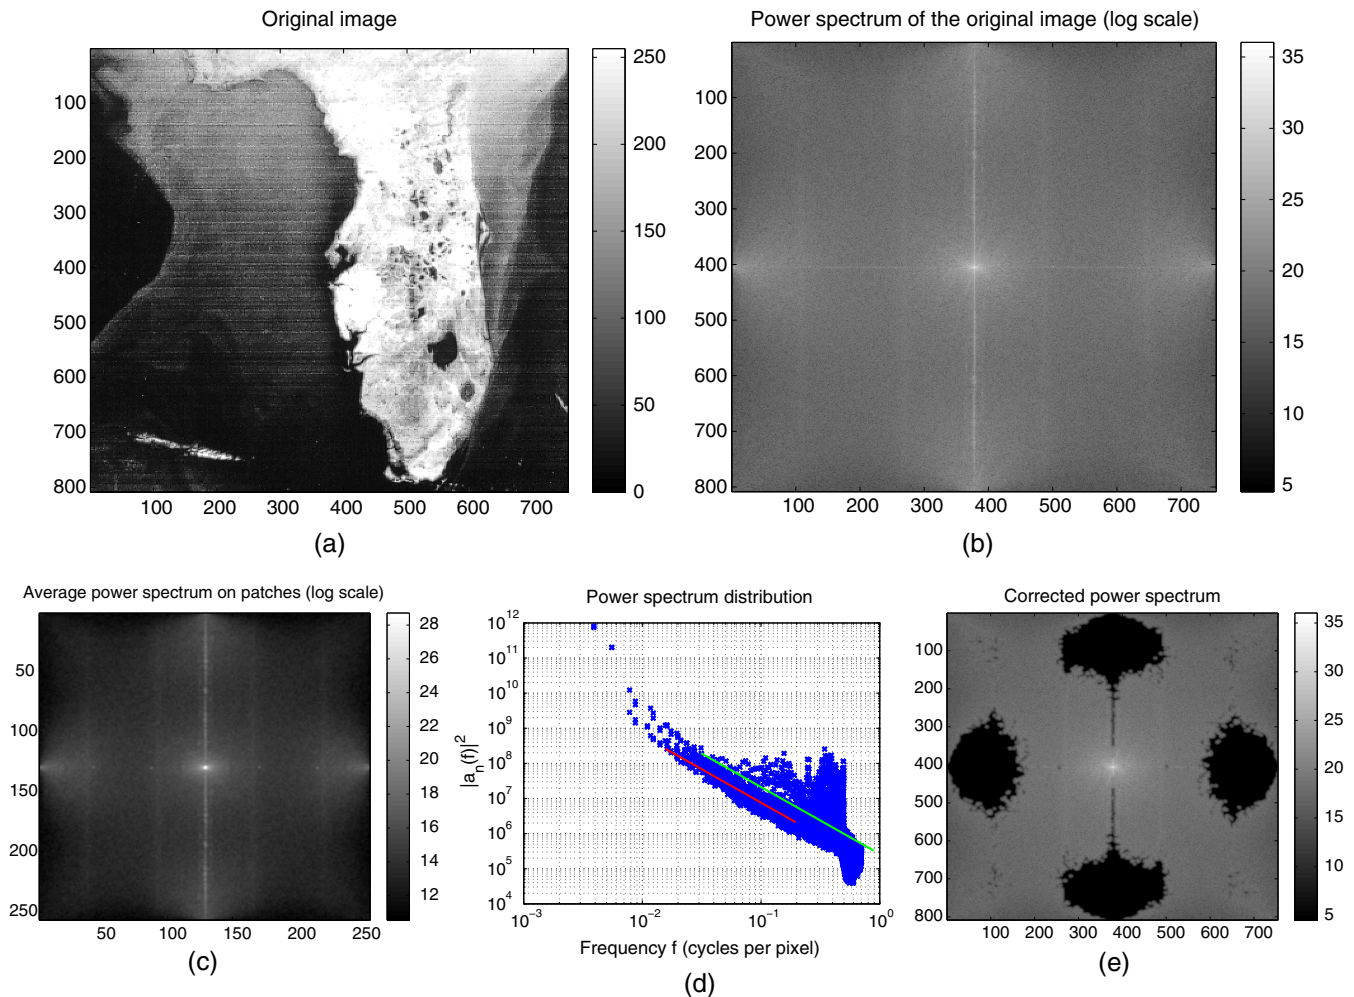


Fig. 12 Florida experiment (3). (a) Noisy image. (b) Power spectrum. (c) Average power spectrum. (d) Power spectrum coefficient distribution. (e) Corrected power spectrum.

The only remaining free parameter is the size L of the patches. Typically, periodic noise in an image with a width of around 500 to 1000 pixels is satisfactorily removed with an L around 100 pixels. The discussion concerning Step 1 should be borne in mind when setting the value of L .

Remark: The boundary effects in the DFT motivate us to multiply the patches by a window decreasing to 0 at its boundaries. The DCT does not suffer from this drawback since it is equivalent to computing the DFT of an even extension of the initial image along both the x and y axes. However, the drawback is to mix information from symmetric directions. We have observed that the proposed algorithm performs equally well with DCT instead of DFT, without the need to multiply the patches with a decreasing window, although the fitting of the power spectrum with the power law is slightly altered. In this paper, we use the DFT-based version.

3.3 Illustrative Example

Figure 1(a) shows a 600×581 image i obtained by the Mariner 4 probe. It is impaired by quasiperiodic noise.

Since this high-frequency noise may be smoothed out in the printed version of the paper, the reader is kindly asked to zoom in the pdf file. The power spectrum is shown in Fig. 1(b). It can be noted that the Fourier coefficients actually tend to decrease with f (the corners are darker than the middle), but the spikes corresponding to the noise periodic pattern can hardly be identified. The horizontal/vertical grid pattern is due to the regular marking on the sides of i . This illustrates the difficulty in automating the notch filter design from the power spectrum alone. It appears that the contribution of the periodic noise is hidden behind the contribution of contours or localized textures. This is confirmed by the distribution of the power spectrum coefficients against the frequency, which shows a decreasing trend with largely scattered points [cf. Fig. 1(c)]. Figure 1(d) shows the average power spectrum obtained by averaging the power spectrum of a set of $870 L \times L = 128 \times 128$ image patches [see Eq. (11)] regularly distributed over the whole image at $L/8 = 16$ pixels apart in the x and y directions. The power law is fitted to this average power spectrum through robust linear regression (we find $A = 11.27$ and $\alpha = 1.47$), see

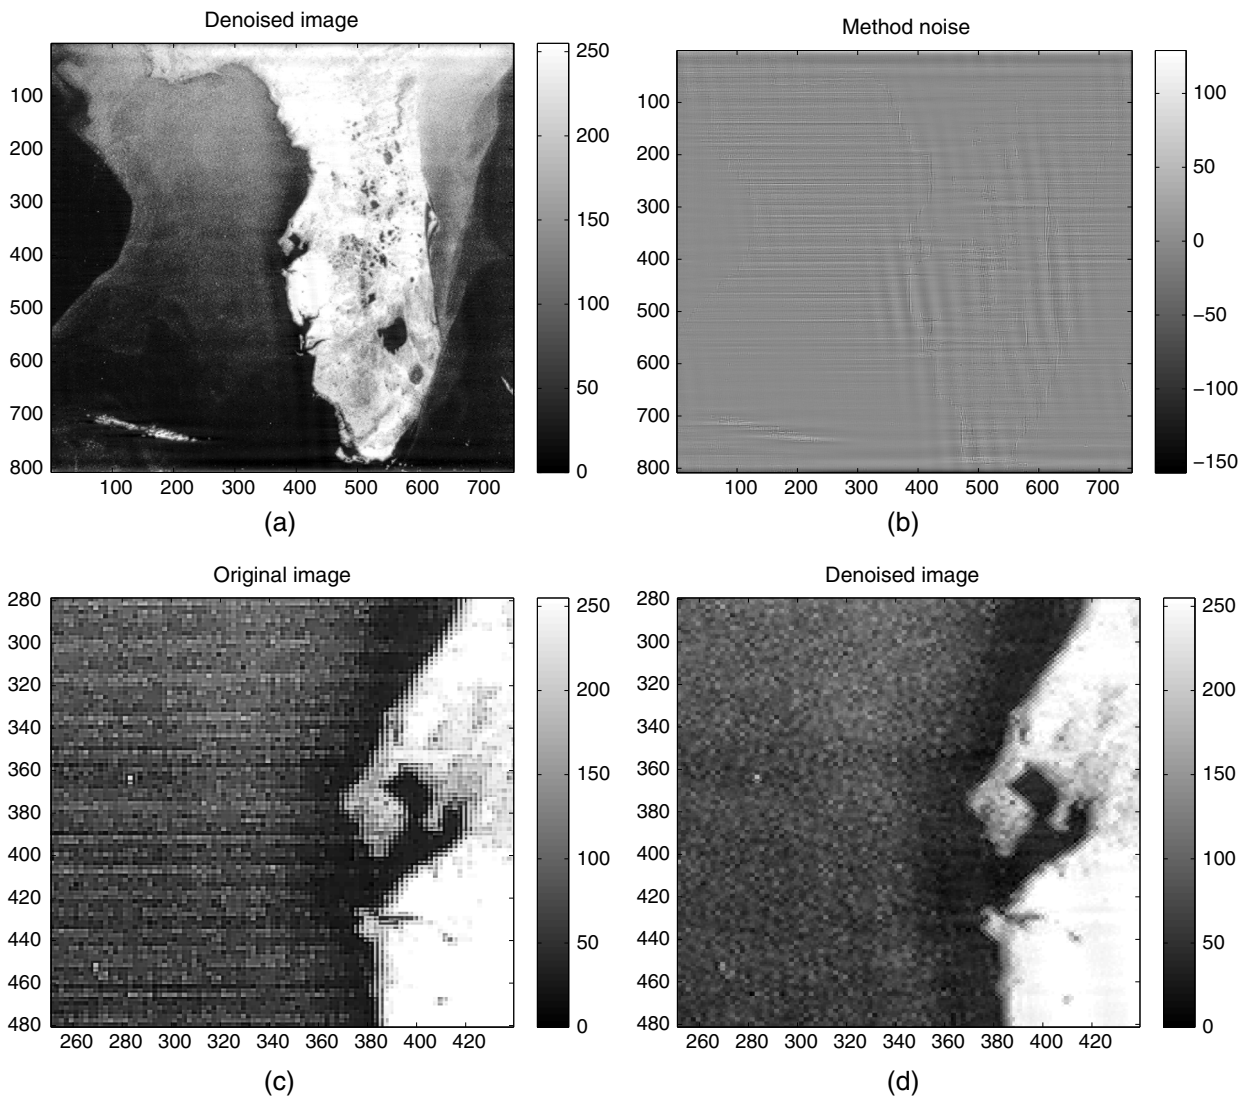


Fig. 13 Florida experiment (3). (a) Denoised image. (b) Estimation of the noise. (c) Close-up view of the noisy image. (d) Close-up view of the denoised image.

Fig. 1(e). Here, the red line corresponds to the linear trend (between f_0 and f_1 cycle per pixel), and the upper green line corresponds to the 3σ upper limit (beyond f_2). Outliers are above this latter straight line. Note that a few outliers are missed. We can see that a uniform threshold (i.e., constant over the frequency range) would not detect the spikes and would retain most of the low- to middle-frequency components. We also observe a fall-off in the higher frequencies.

The residuals between the average power spectrum and the expected linear relation are normalized as in the left-hand member of Eq. (10), giving the so-called normalized spectrum, depicted in Fig. 2(a). In this latter image, the spikes are well detached and easier to detect. We can also see that the power law actually depends on the orientation; though sufficient for our purpose, an isotropic modeling of the spectrum fall-off does not fit the actual distribution well. Here, the dependence is probably due to the orientation of the illumination source (from the bottom to the top of the image), which gives a dominant direction for the shadow distribution. Figure 2(b) shows the outlier localization in the average power spectrum, and the modified spectrum of the original image is shown in Fig. 2(c). (This is the original spectrum multiplied by the notch filter $1 - M_o$, cf. Step 4 in Sec. 3.1). The two peaks in the distribution of the coefficients

in Fig. 1(e) correspond, respectively, to the four spikes at the boundaries of the x and y axes, and to the four spikes in the corners. The low-frequency outliers detected on the vertical axis are the few “false positive” points over the green line just above f_2 in Fig. 1(e). The impact of such points is illustrated below.

The output of the proposed algorithm is the denoised image \hat{i} , and the pseudoperiodic noise component \hat{n} is shown in Figs. 3(a) and 3(b). Following the terminology of Buades et al.,²⁰ the retrieved noise component is called method noise. It should ideally consist only of noise; no structure from the original image should be visible in this image. We can see that the periodic noise has been removed from \hat{i} . The pseudoperiodic noise \hat{n} is much more complicated than a simple sine wave. A close-up view of the original and denoised images is shown in (c) and (d). Of course, the white noise component w is not removed by the algorithm.

4 Experiments

A series of examples are now presented and discussed to illustrate the benefits and limitations of this approach. A synthetic noise is first considered to rely on a reference image when discussing the result of the process (Sec. 4.1). Various




Fig. 14 Mariner 6 experiment (4). (a) Noisy image. (b) Power spectrum. (c) Average power spectrum. (d) Power spectrum coefficient distribution. (e) Corrected power spectrum.

noisy images from the literature or the Internet are then processed (Sec. 4.2). A comparison with a blind source separation algorithm is discussed in Sec. 4.3. In the illustrations, the x component corresponds to the rows and the y component corresponds to the columns.

4.1 Synthetic Dataset

Experiments concerning synthetic data are discussed in this section. Synthetic periodic noise is added to an 8-bit noise-free image i_0 of size $X \times Y$. The periodic noise intensity is given by

$$n(x, y) = 50 \sin\left(2\pi \frac{p}{X} x\right) \sin\left(2\pi \frac{q}{Y} y\right), \quad (12)$$

where p and q are the parameters governing the frequency of the pseudoperiodic noise along the x and y axes, respectively. The unit of n is the gray level. The root-mean-square error (RMSE) between the noisy image $i = i_0 + n$ and the ground truth is equal to 25.

In the first experiment, the ground-truth image is the standard Mandrill image. Here $X = Y = 512$ and we set $L = 128$ [hence $N = 576$ patches, see Eq. (11)]. In this example, $\alpha = 1.59$ and $A = 11.29$. Figures 4 and 5 present

an example showing noise with a high-frequency y component and a low-frequency x component (here $p = 2$ and $q = 200$). The spikes caused by the periodic noise are not visible in the initial spectrum. However, most coefficients are correctly retrieved as upper outliers of the power law distribution. As in the preceding illustrative experiment, the few values just below the 3σ limit do not prevent the denoising algorithm from giving visually satisfactory results. The reason is that these values correspond to the neighborhood of the detected outliers, which is smoothed out by the Gaussian convolution in Step 4 of the algorithm. Note that the power spectrum distribution shows a constant value for high frequencies. It is probably the effect of the four which plays the role of strong white noise, in accordance with Sec. 2.2. We can define the RMSE calculated between the retrieved image \hat{i} and the ground truth (corresponding to the average error on the retrieved intensity of a pixel). Here, the RMSE is equal to 0.598, which is below the intensity quantization level (equal to one gray level) and far below the original error, which was 25.

We also illustrate the performance of the algorithm with respect to the noise component frequency. The frequency p (respectively q , in cycles per image) spans $[0, X/2]$ (respectively $[0, Y/2]$) in steps of 20 (hence




Fig. 15 Mariner 6 experiment (4). (a) Denoised image. (b) Estimation of the noise. (c) Close-up view of the noisy image. (d) Close-up view of the denoised image.

13×13 tested frequencies), except for values where $\sqrt{(p/X)^2 + (q/Y)^2} < f_2$, since frequencies that are too low are not considered as periodic noise components and are not eliminated by the algorithm. The RMSE is computed for each of these 13×13 frequency pairs. Figure 6 gives the RMSE map for the 13×13 frequencies (p in ordinate, q in abscissa). The top-left corner corresponds to the skipped low frequencies, and the bottom-right corner to the highest frequencies. We can see that, except for quite low frequencies, the RMSE is below 1, and that in any case it is much lower than the original error of 25.

In the second experiment, the ground-truth image is the standard Boat image. As in the previous experiment, $K = 512$ and $L = 128$. Here, $p = 40$ and $q = 100$. In this example, $\alpha = 2.52$ and $A = 8.73$. Figures 7 and 8 depict an example deliberately chosen for didactic purposes such that the noise components are correlated with the image components. Therefore, the notch filter also removes noise-free image parts. We can see in the close-up view that ringing artifacts appear along the masts parallel to the noise pattern. Such a situation is unavoidable in any notch filter-based approach. Reconstructing the missing coefficients by spectrum interpolation (i.e., spectrum “inpainting”) would be

required. To the best of our knowledge, this is still an open problem (see the recent PhD thesis in Ref. 21 and the experiments in Ref. 22). Here, the RMSE is equal to 1.801, which is about three times the RMSE obtained in the previous example. Figure 9 gives the error map as above. We can see that the RMSE is slightly larger than in the preceding example, precisely because of the correlation between the periodic noise and certain image patterns.

However, we did not observe such a situation in the experiments presented in the following section. It appears that in real images, the noise spectrum is mostly well separated from the image spectrum.

4.2 Real Dataset

For each experiment, we present the noisy image, its power spectrum, the average power spectrum of the patches, the average power spectrum distribution together with the fitting line (in red) and the 3σ upper limit (in green), the corrected power spectrum, the denoised image, the estimated noise, and a close-up view of the noisy and denoised images.

The values of parameter L and of the deduced number N of patches are given in Table 1, together with the values of A and α estimated on the average power spectrum coefficients.

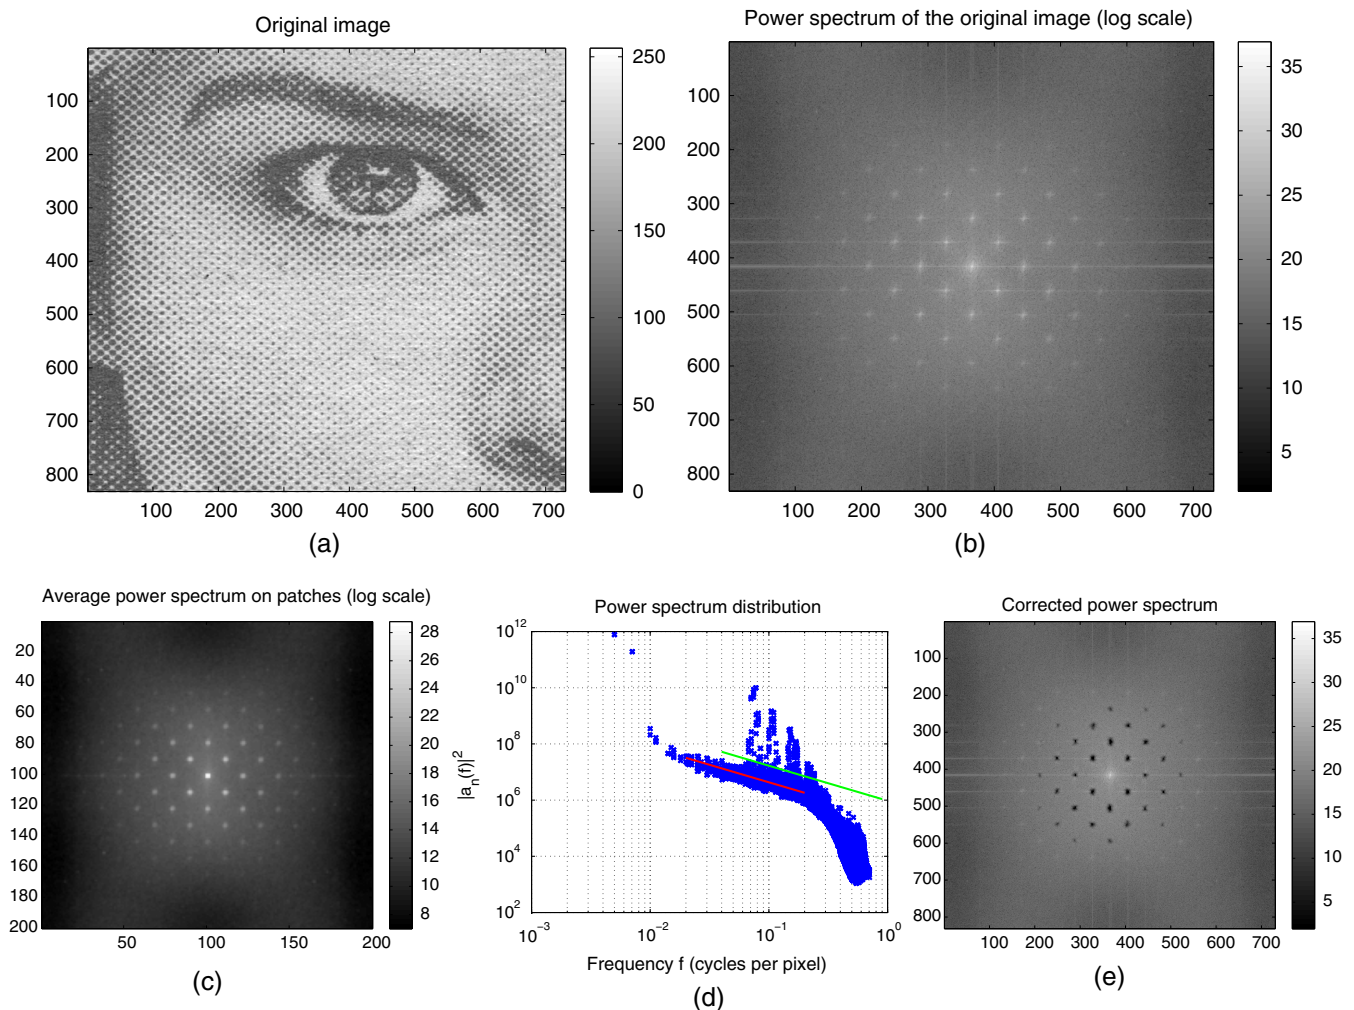


Fig. 16 Newspaper experiment (5). (a) Noisy image. (b) Power spectrum. (c) Average power spectrum. (d) Power spectrum coefficient distribution. (e) Corrected power spectrum.

Computation times are also mentioned. We can see that $\alpha \simeq 2$ (as expected for “generic” natural images) in the remote sensing images. Images essentially composed of dots (such as the challenging Apollo or Newspaper images) have a smaller α . The Bending image has a larger α . This image cannot be considered as a natural one, as explained briefly below, which explains its specific α and A values.

The original images were found on the Internet. The image for Experiment 1 comes from Ref. 23. Images for Experiments 2 to 5 come from Ref. 5. Additional experiments are available in a separate report.²⁴

Experiment 1 (Figs. 1 to 3) was fully discussed in Sec. 3.3.

In Experiment 2 (Figs. 10 to 11), the original image is composed of separate shapes (see close-up view) and is an extreme case of pseudoperiodic noise. The average power spectrum is dominated by the noise component and can only marginally be assimilated to the spectrum of a natural image. This explains the aspect of the power spectrum distribution which is quite widely spread. In spite of this, the automated process enables us to detect the noise component, and to reconstruct the image with only limited residual noise patterns. The complex noise pattern in the Fourier domain makes it all the more convenient to have an automated process. Note that a seemingly better reconstruction

is obtained in Ref. 5 by manually picking well-selected spikes.

Experiment 3 (Figs. 12 to 13) shows strong periodic striping in the sea in this satellite image of Florida. This explains the vertical straight lines in the spectrum. In addition to these lines, “blobs” at the center of the top/bottom and left/right boundaries of the images are also detected as outliers to the spectrum distribution. They are certainly explained by aliasing, another cause of spikes in the power spectrum which was not taken into account in this study. This is confirmed by the sharp, pixelized aspect of the coast in the close-up view. As a consequence, the method noise also contains information about the image edges. As expected, white noise is not discarded by the proposed algorithm and is still noticeable. Let us note that an approach to detecting aliasing based on statistical properties of the spectrum is available in Ref. 25.

Experiment 4 (Figs. 14 to 15) is another example of complex noise appearing in remote sensing imaging which is successfully removed. High-frequency components can be noted in the corners of the spectrum image.

Experiment 5 (Figs. 16 to 17) deals with the halftone printing used in newspapers. In this technique, gray levels are simulated by dots of various sizes distributed along a regular grid. Thus, it is possible to consider the image

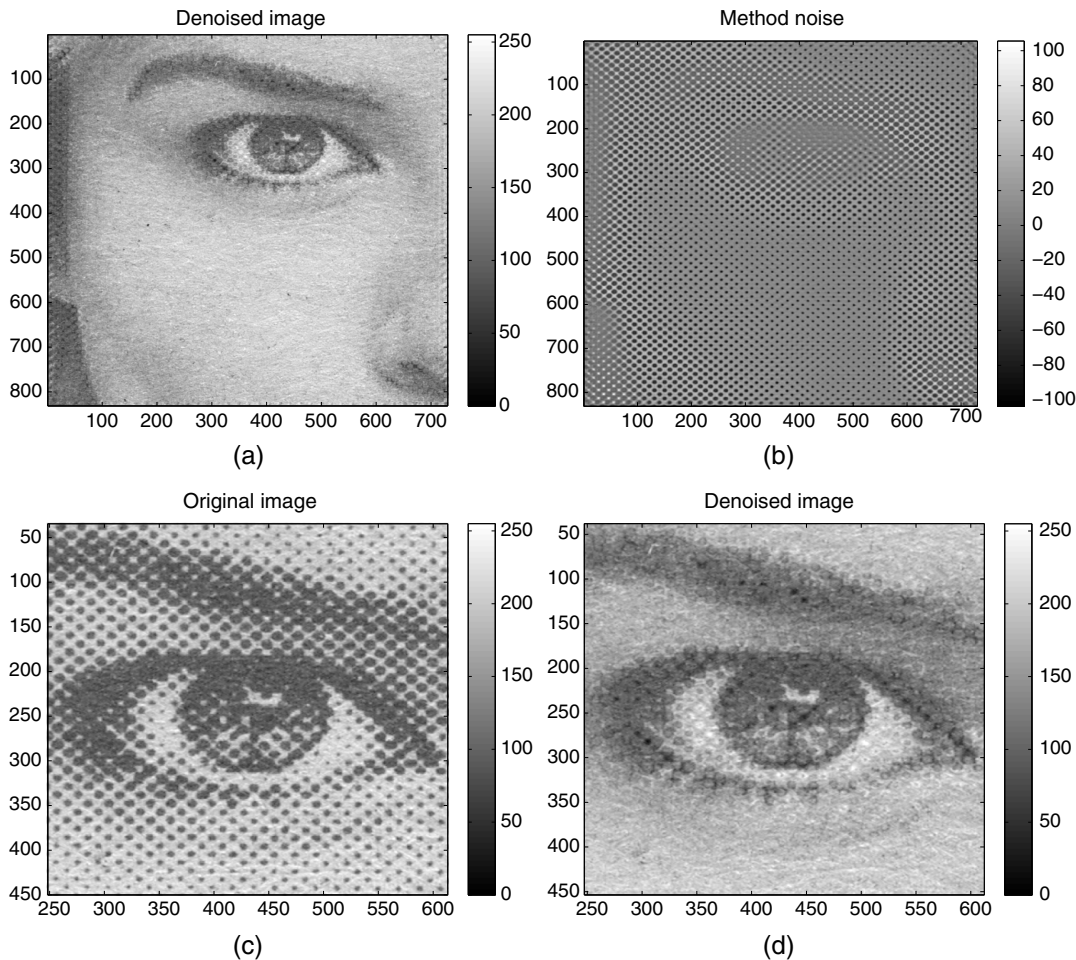


Fig. 17 Newspaper experiment (5). (a) Denoised image. (b) Estimation of the noise. (c) Close-up view of the noisy image. (d) Close-up view of the denoised image.

reconstruction as a pseudoperiodic noise removal problem. Although the linearity of the power spectrum distribution stops before the 0.2 cycle per pixel limit, we are still able to detect the spikes. We can see in the close-up view of the denoised image that the largest dots are still visible, and that subsequent processing is certainly required to obtain a better reconstruction.

Experiment 6 (Figs. 18 to 19) deals with data from experimental solid mechanics. This experiment is an application of the proposed algorithm to non-natural images. Figure 18(a) shows a shear strain map measured on the surface of a reinforced notched beam subjected to three-point bending. It is obtained using the so-called “grid method,” which consists of marking the surface with a regular pattern (a grid), taking pictures of this pattern during a test, and extracting the phases and their derivatives from the grid images using the windowed Fourier transform.²⁶ The displacement and strain maps are indeed respectively proportional to the phase and phase derivative changes between current and reference grid images. Low-frequency parasitic fringes like those shown in this figure sometimes appear on the maps. They are due to various causes, such as manufacturing defects of the grid. In the current case, these fringes are likely due to slight regular defects of the grid which are not

correctly compensated for by using the process described in Ref. 26. Of course, this image cannot be considered as a “natural” image. Here, the power spectrum distribution is satisfied up to 0.06 cycle per pixel, instead of the 0.2 value used in the algorithm in Sec. 3.1. “Blobs” can be seen in the denoised image. They are actually caused by sensor noise which propagates to the strain map^{27,28} as spatially correlated noise. We have shown here an example where the automated process provides satisfactory results. However, the power law assumption is, in general, not valid for strain maps. They still need hand-tuned dedicated approaches.²²

4.3 Comparison with Morphological Component Analysis

MCA^{10,11} implements blind source separation by decomposing a linear mixture of several (here two) signals, sparsely represented in a morphologically independent basis. Separating the periodic noise and the underlying natural image can be achieved by using a DCT representation for the periodic noise and a curvelet representation for the image, as suggested in Ref. 11.

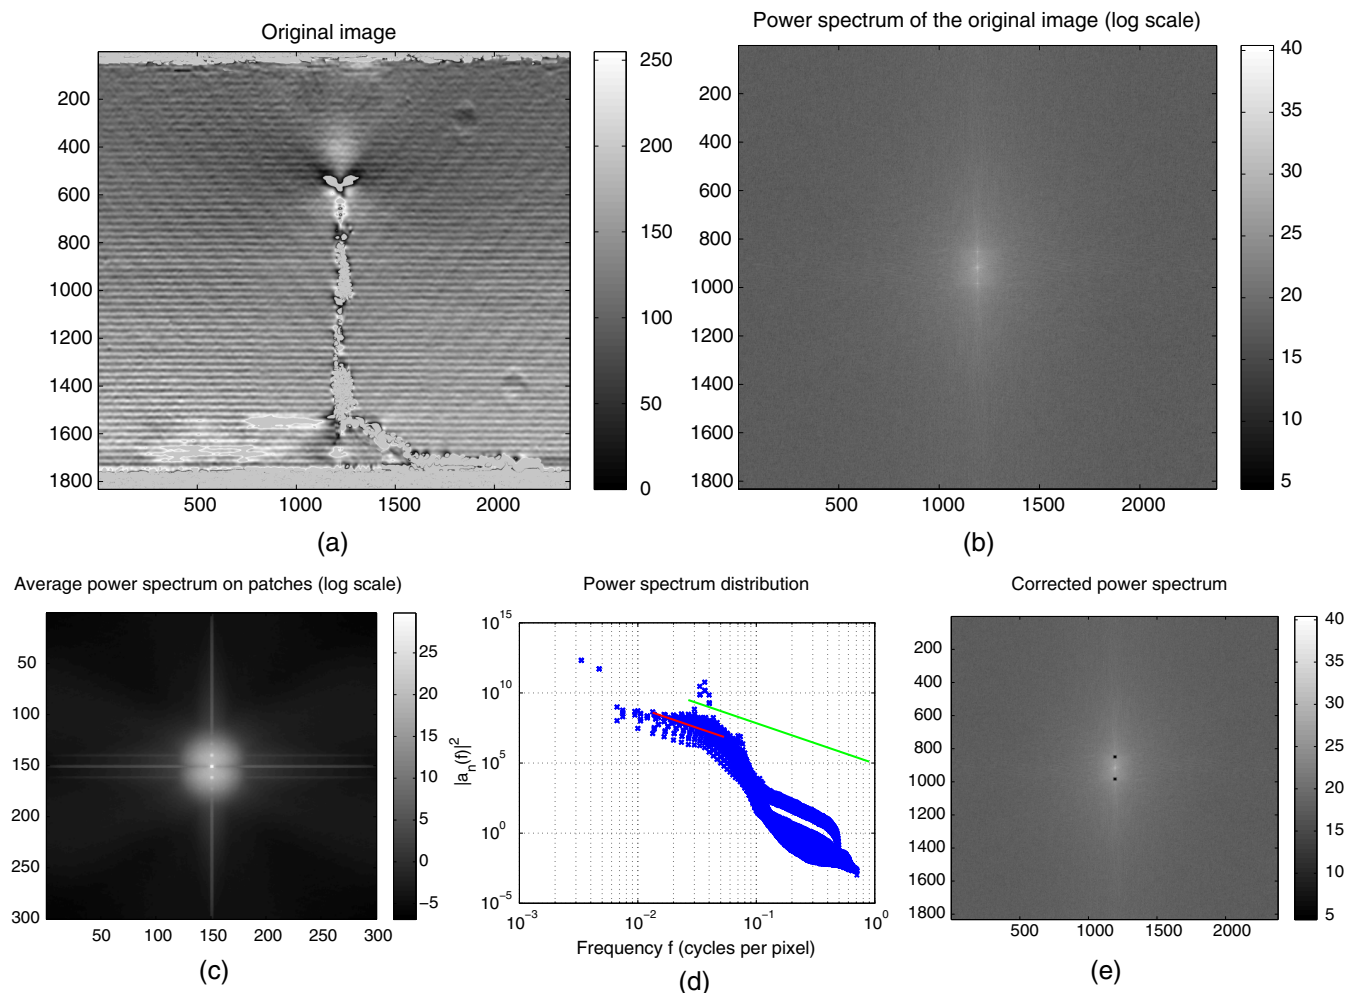


Fig. 18 Bending experiment (6). (a) Noisy image. (b) Power spectrum. (c) Average power spectrum. (d) Power spectrum coefficient distribution. (e) Corrected power spectrum.

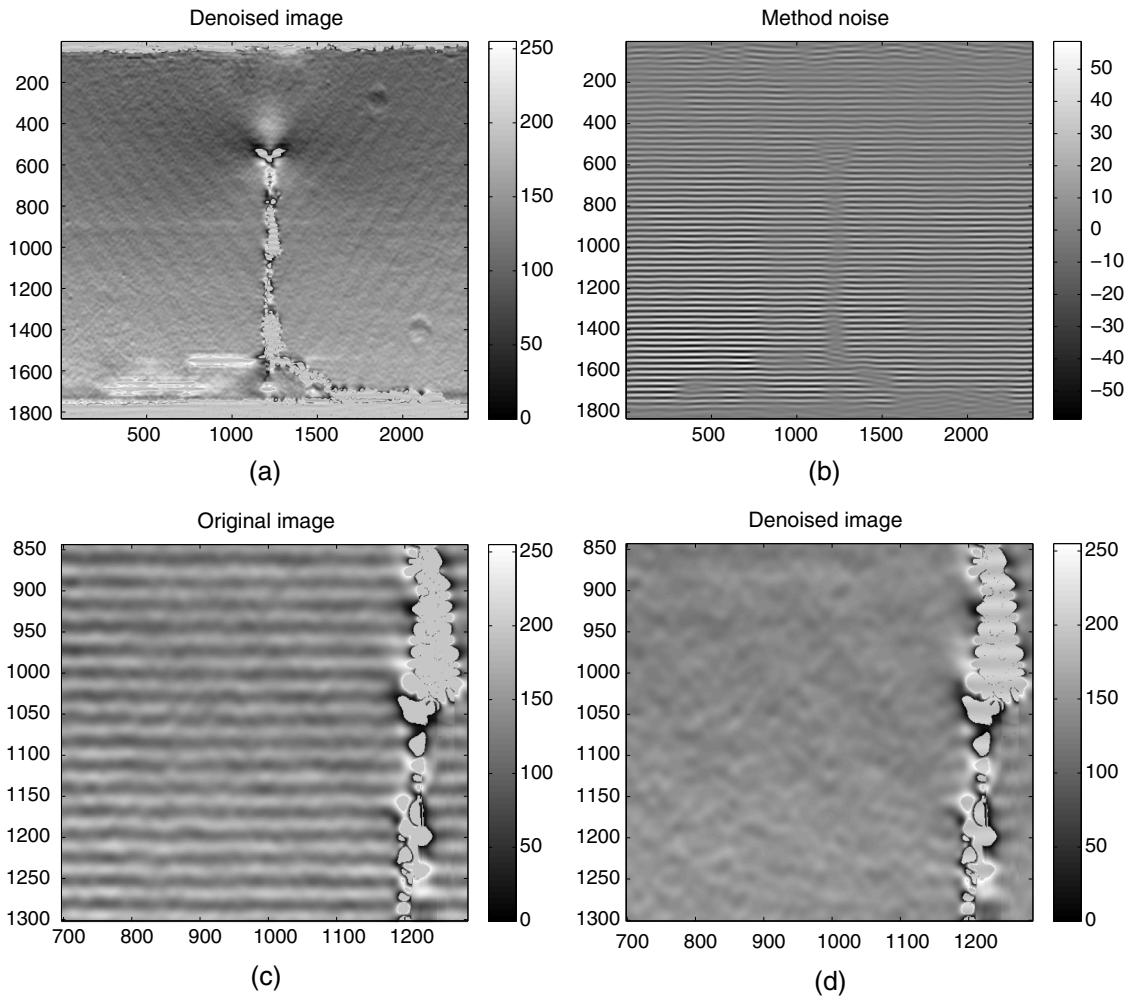


Fig. 19 Bending experiment (6). (a) Denoised image. (b) Estimation of the noise. (c) Close-up view of the noisy image. (d) Close-up view of the denoised image.

The results of this decomposition²⁹ for some of the images of Secs. 4.1 and 4.2 can be found in Figs. 20 to 22. The denoised image is the curvelet component, and the method noise is the DCT component. Computation times are reported in Table 2. Additional experiments are available in Ref. 24.

Even if the parameters would certainly benefit from careful setting, we can see that MCA does not perform as well as the proposed approach. Details of the image can be seen in the method noise (which means that they have been lost

from the denoised image), and some periodic noise can still be seen in the denoised image. The reason is that the sparsity assumption is probably too optimistic in the problem of interest, as could be expected from the shape of the corresponding notch filter in the Fourier domain. Furthermore, the retrieved denoised image appears smoother than with the proposed notch filter. In addition, MCA is 150 to 350 times slower than the proposed algorithm, despite the fact that MCA uses compiled functions (our code is native MATLAB).




Fig. 20 Morphological component analysis (MCA), Mandrill experiment. Denoised image, periodic noise component, close-up view of the denoised image. To be compared to Fig. 5.




Fig. 21 MCA, Florida experiment (3). Denoised image, periodic noise component, close-up view of the denoised image. To be compared to Fig. 13.

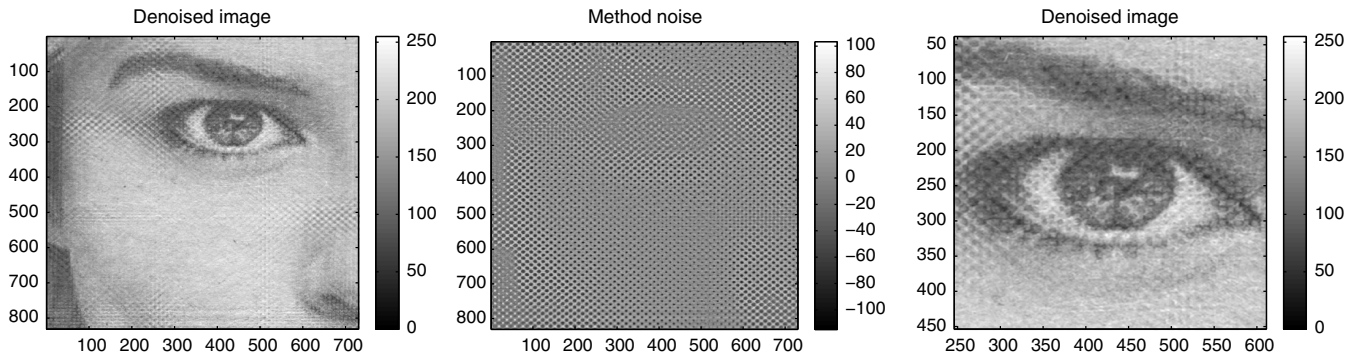


Fig. 22 MCA, Newspaper experiment (5). Denoised image, periodic noise component, close-up view of the denoised image. To be compared to Fig. 17.

Table 2 Morphological component analysis (MCA) algorithm: experiment number (as in Sec. 4.2); image name and corresponding figure; computation time (in seconds, Intel Xeon E3-1240 processor).

Number	Image	Time
—	Mandrill (Fig. 20)	30.2
3	Florida (Fig. 21)	134.4
5	Newspaper (Fig. 22)	134.6

5 Conclusion

This paper presents a novel automated approach to pseudo-periodic noise removal in natural images. It is based on the observation that the localization of the Fourier coefficients from the periodic noise are likely to be the only spikes present in the average power spectrum estimated on a set of patches regularly distributed in the image. The corresponding Fourier coefficients are detected as outliers of the power law distribution which is expected from a natural image. This gives an automated notch filter design, which can be used in turn to successfully remove pseudoperiodic noise. The resulting method is limited to “natural” images: this terminology includes remote sensing imaging but not, for example, strain maps in experimental mechanics which, in general, still need a dedicated approach.²² Aliasing is also potentially detected as pseudoperiodic noise if it affects a large part of the image. Moreover, a

comparison shows that the proposed method performs better than a blind source separation algorithm such as MCA. A perspective for this work would be to replace the parametric approach to detecting the spikes caused by noise in the Fourier domain by a nonparametric detection such as an a-contrario approach (see, e.g., Refs. 30 and 31) in the same spirit as in Ref. 25.

Acknowledgments

Pr. S. Amziane and Pr. E. Toussaint are gratefully acknowledged for their help in obtaining the strain maps shown in the Bending experiment. The authors of MCALab are gratefully acknowledged for making their software library available. This work is partially supported by GdR ISIS (TIMEX project).

References

1. R. Schowengerdt, *Remote Sensing: Models and Methods for Image Processing*, Academic Press, Orlando, Florida (2006).
2. M. Lebrun et al., “Secrets of image denoising cuisine,” *Acta Numerica* **21**(1), 475–576 (2012).
3. P. Milanfar, “A tour of modern image filtering: new insights and methods, both practical and theoretical,” *IEEE Signal Process. Mag.* **30**(1), 106–128 (2013).
4. T. C. Rindfleisch et al., “Digital processing of the Mariner 6 and 7 pictures,” *J. Geophys. Res.* **76**(2), 394417 (1971).
5. R. Gonzalez and R. Woods, *Digital Image Processing*, 3rd ed., Prentice-Hall, Upper Saddle River, New Jersey (2006).
6. I. Aizenberg and C. Butakoff, “A windowed Gaussian notch filter for quasi-periodic noise removal,” *Image Vis. Comput.* **26**(10), 1347–1353 (2008).
7. G. Hudhud and M. Turner, “Digital removal of power frequency artifacts using a Fourier space median filter,” *IEEE Signal Process. Lett.* **12**(8), 573–576 (2005).

8. J. Fehrenbach, P. Weiss, and C. Lorenzo, "Variational algorithms to remove stationary noise: applications to microscopy imaging," *IEEE Trans. Image Process.* **21**(10), 4420–4430 (2012).
9. M. Zibulevsky and B. Pearlmutter, "Blind source separation by sparse decomposition in a signal dictionary," *Neural Comput.* **13**(4), 863–882 (2001).
10. M. Fadili et al., "Image decomposition and separation using sparse representations: an overview," *Proc. IEEE* **98**(6), 983–994 (2010).
11. J.-L. Starck, M. Elad, and D. Donoho, "Image decomposition via the combination of sparse representations and a variational approach," *IEEE Trans. Image Process.* **14**(10), 1570–1582 (2005).
12. M. Cannon, A. Lehar, and F. Preston, "Background pattern removal by power spectral filtering," *Appl. Opt.* **22**(6), 777–779 (1983).
13. R. Srinivasan, M. Cannon, and J. White, "Landsat data destriping using power spectral filtering," *Opt. Eng.* **27**(11), 939–943 (1988).
14. D. Field, "Relations between the statistics of natural images and the response properties of cortical cells," *J. Opt. Soc. Am. A* **4**(12), 2379–2394 (1987).
15. A. Oliva and A. Torralba, "Modeling the shape of the scene: a holistic representation of the spatial envelope," *Int. J. Comput. Vis.* **42**(3), 145–175 (2001).
16. A. Torralba and A. Oliva, "Statistics of natural image categories," *Network* **14**(3), 391–412 (2003).
17. A. van der Schaaf and J. van Hateren, "Modelling the power spectra of natural images: statistics and information," *Vis. Res.* **36**(17), 2759–2770 (1996).
18. A. Hyvärinen, J. Hurri, and P. Hoyer, *Natural Image Statistics: A Probabilistic Approach to Early Computational Vision*, Springer, London, United Kingdom (2009).
19. <http://www.loria.fr/~7Esur/software/ARPENOS/> (January 2015).
20. A. Buades, B. Coll, and J.-M. Morel, "A review of image denoising algorithms, with a new one," *SIAM J. Multiscale Mode. Simul.* **4**(2), 490–530 (2005).
21. K. Jalalzai, "Regularization of inverse problems in image processing," PhD Thesis, École Polytechnique, Palaiseau, France (2012).
22. M. Grédiac, F. Sur, and B. Blaysat, "Removing quasi-periodic noise in strain maps by filtering in the Fourier domain," (2015). Submitted for publication.
23. <http://www.exploratorium.edu/mars/earlymissions.php> (January 2015).
24. F. Sur and M. Grédiac, "An automated approach to quasi-periodic noise removal in natural images," Research Report 8660, INRIA (2015).
25. B. Coulange and L. Moisan, "An aliasing detection algorithm based on suspicious colocalizations of Fourier coefficients," in *Proc. IEEE Int. Conf. on Image Processing (ICIP)*, pp. 2013–2016, IEEE, Hong Kong (2010).
26. C. Badulescu, M. Grédiac, and J.-D. Mathias, "Investigation of the grid method for accurate in-plane strain measurement," *Meas. Sci. Technol.* **20**(9), 095102 (2009).
27. F. Sur and M. Grédiac, "Towards deconvolution to enhance the grid method for in-plane strain measurement," *AIMS Inverse Probl. Imaging* **8**(1), 259–291 (2014).
28. M. Grédiac et al., "Using deconvolution to improve the metrological performance of the grid method," *Opt. Lasers Eng.* **51**(6), 716–734 (2013).
29. <https://fadili.users.greyc.fr/demos/WaveRestore/downloads/mcalab/Home.html> (January 2015).
30. F. Cao et al., "A theory of shape identification," in *Lecture Notes in Mathematics*, Vol. 1948, Springer, Berlin, Germany (2008).
31. A. Desolneux, L. Moisan, and J.-M. Morel, *From Gestalt Theory to Image Analysis: A Probabilistic Approach*, Interdisciplinary applied mathematics, Springer, Berlin, Germany (2008).

Frédéric Sur is an assistant professor of computer science at the Université de Lorraine, France. His current research interests include image restoration, image analysis, and computer vision. He co-authored a book on shape recognition (Springer, 2008).

Michel Grédiac is a professor at Blaise Pascal University, Clermont-Ferrand, France. He has been working for more than 25 years in the development and use of full-field measurement techniques in experimental solid mechanics. He has been an associate technical editor for *Experimental Mechanics* (Springer) since 2010.