Single-Image Superresolution of Natural Stochastic Textures Based on Fractional Brownian Motion

Ido Zachevsky and Yehoshua Y. (Josh) Zeevi

***Abstract*— Texture enhancement presents an ongoing challenge, in spite of the considerable progress made in recent years. Whereas most of the effort has been devoted so far to enhancement of regular textures, stochastic textures that are encountered in most natural images, still pose an outstanding problem. The purpose of enhancement of stochastic textures is to recover details, which were lost during the acquisition of the image. In this paper, a texture model, based on fractional Brownian motion (fBm), is proposed. The model is global and does not entail using image patches. The fBm is a self-similar stochastic process. Self-similarity is known to characterize a large class of natural textures. The fBm-based model is evaluated and a single-image regularized superresolution algorithm is derived. The proposed algorithm is useful for enhancement of a wide range of textures. Its performance is compared with single-image superresolution methods and its advantages are highlighted.**

***Index Terms*— Stochastic texture enhancement, superresolution, self-similarity, fractional Brownian motion.**

I. Introduction

SINGLE-IMAGE superresolution (SR) has attracted considerable attention in recent years [1]–[10]. This is a challenging task, since the original image needs to be recovered using only the degraded and subsampled image. Traditional image enhancement approaches for denoising and deblurring perform contour emphasis and result in sharper images. Consequently, these approaches often yield an unnatural cartoon-like image, compromising on the quality of some textures. In some cases, textures are even eliminated by these approaches. This compromise in image fidelity highlights the observation that textures are an important ingredient of image structure, that must be considered in the context of image enhancement tasks.

A. Enhancement Techniques

Image enhancement algorithms, used in deblurring and denoising methods, generally attempt to solve the following inverse problem [11]:

$$\hat{X} = \arg\min_{X\in\mathcal{X}} \|X * h - Y\|_2^2 + \lambda g(X), \qquad \text{(I.1)}$$

for a degraded image, $Y \in \mathcal{X}$, where $\mathcal{X} \subseteq \mathbb{R}^{N\times N}$, $X \in \mathcal{X}$ is a possible solution and $h \in \mathbb{R}^{n\times n}$ is a blurring kernel. The parameter λ is a Lagrange coefficient and $g(X)$ is a regularization function. Since h is usually a lowpass-type filter, and the measurement is noisy, the problem is ill-posed, and a regularization function is needed.

In Wiener filtering, or other linear methods, the regularizing function, $g(X)$, is a quadratic function of the image gradient. In $\mathcal{L}_1$-based methods, such as total variation (TV), $g(X)$ is the $\mathcal{L}_1$ norm of the image gradient.

Partial differential equations (PDE)-based approaches generally use $g(X) = G(|\nabla X|^2)$. A gradient descent minimization then yields the following scheme [11]:

$$X_t = -\tilde{h} * (X * h - Y) + \lambda\nabla(G'(|\nabla X|^2)\nabla X), \qquad \text{(I.2)}$$

where $\tilde{h}(\eta_1, \eta_2) = h(-\eta_1, -\eta_2)$ and $G'(|\nabla X|^2)$ is the diffusivity or edge detection function. This function is chosen so that high gradients are preserved and low gradients are smoothed.

When these methods are applied to images that are comprised of stochastic textures, they do not yield the desired results. This is due to the common assumption that low gradient areas in an image are originated by noise or optimization artifacts (such as ringing or aliasing) and not by a valuable texture. This, in turn, is due to the assumption that images reside in a bounded variation (BV) space, which often lends itself to the wrong choice of the regularization function. The BV assumption has been challenged in recent years [12] and deblurring schemes, developed under more suitable spaces, have yielded more successful results [13].

B. Texture Representation and Enhancement

Some of the aforementioned methods can be adopted for texture preservation. In [14], a potential function has been incorporated into the diffusion equation which effectively prevents smoothing of specific texture details in an image. In sparseness-based approaches (not discussed in this work; for a review see [15]), a separate dictionary is used for handling textures.

Textures, in general, can be divided into two main types: Regular, or structured, and stochastic [16], [17]. One can define the former as spatially-replicated instances of a single or several repetitive patterns. An example of a regular texture is a brick wall. To compare with, stochastic textures do not contain a specific pattern. Instead, they are considered to be realizations of random processes. This type of textures cannot be modelled in a similar manner to regular textures. It is

Manuscript received August 16, 2013; revised December 29, 2013 and March 10, 2014; accepted March 12, 2014. Date of publication March 18, 2014; date of current version April 3, 2014. This work was supported by the Technion Ollendorff Minerva Center for Vision and Image Sciences. The work of Y. Y. Zeevi was supported by the Minerva Foundation. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Charles Creusere.

The authors are with the Department of Electrical Engineering, Technion-Israel Institute of Technology, Haifa 32000, Israel (e-mail: ido@tx.technion.ac.il; zeevi@ee.technion.ac.il).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIP.2014.2312284

important to note that this division is by no means a dichotomy, as natural textures depict the entire range between regular and stochastic. The texture spectrum, as defined in [16], contains textures in varying complexity and regularity.

As the two types of textures are visually and conceptually different, different techniques are used in order to enhance them. Most of the effort in texture enhancement has been devoted, even inadvertently, to the regular textures. Regular textures contain replicated versions of a single or a few basic patterns, in varying amounts of distortion. One can, therefore, use enhancement methods based on a search for similar shapes, in higher resolution, within the same image [9] and/or in other images bearing some similarity to the target image [1].

Numerous methods, that are applied primarily to images containing cartoon-type skeletons but also on textures, are based on a machine learning approach. These methods attempt to build up and exploit a database of natural images and/or textures. This database is then used for training with sets of low and high resolution image patches, and the final image is then obtained by predicting the high resolution patch according to the training database [18]. It is important to note that regular textures still obey, in general, the model of bounded variation, and can therefore be enhanced using known approaches.

C. Stochastic Textures

Unlike regular textures, stochastic textures are not characterized by repetitive patterns. They are, instead, defined by their statistical properties. This type of textures exhibits statistical properties such as non-local [19], long-range dependencies and self-similarity, as their pixel distribution remains the same across scales, up to a scaling parameter [20]–[23]. Enhancement of such textures can hardly be achieved by using example-based methods that have been demonstrated to be successful on regular textures, as the stochastic textures do not contain basic patterns but are rather governed by the statistical properties of a latent random process.

Sparseness-based or Gaussian mixture model-based approaches perform enhancements locally (on patches) [24]. In the former, the resulting dictionary usually appears to be similar to Fourier, discrete cosine transform or overcomplete Wavelet bases, with discontinuities that match edges in images as well. These models assume that an image can be modelled locally, whereas in stochastic textures, there is an advantage to a long-range dependence or global model, exploiting correlations between as many pixels as possible.

Extensive work has been done deriving a model for natural images [25], [26]. These studies have substantiated the notion that the distribution of natural images is highly kurtotic and non-Gaussian. This can be seen in the Wavelet domain by observing the first- and second-order empirical distributions of their coefficients. However, for stochastic textures, the Gaussian assumption is indeed relevant, a behaviour that had been overlooked by considering an ensemble of a numerous natural images, where the least of which are distinctly stochastic textures.

While regular textures can be enhanced by using methods of edge enhancement (considering, for example, the enhancement of the edges of a brick in a brick wall texture), in the stochastic case such edges do not exist. Attempting to apply edge enhancement to such a texture, may, in some cases, create a staircasing effect, while smoothing out the fine details in the neighborhood of the newly-created edge.

A different approach for regular and stochastic texture enhancement is the texture synthesis, in which a sample-patch is used in order to create a newly formed image of larger size and the same visual appearance as the original [16], [17], [27]. While such methods show successful results in visual resemblance to the original, they are less effective in deconvolution problems such as superresolution, in which a high resolution estimate has to represent the input low resolution image. Further, such synthesis, based on local dependencies, may fail to capture the global statistical structure of the texture, in case of stochastic textures.

There also exist methods which combine example-based techniques with texture synthesis. In [9], the example database is built from the degraded image itself, and similar patches in a search window are used for texture synthesis. While the former reconstructs edges and corners, the latter reconstructs textural details.

We present a model for stochastic textures. This model is based on fractional Brownian motion (fBm); a Gaussian random process which exhibits properties that characterize stochastic textures [21]. This process is used to regenerate the lost high frequency content based on the structure of a given degraded image. Realizations of the model are displayed, and an optimization scheme is derived and applied to single-image superresolution.

The basic model is suitable for isotropic textures. In order to be suitable for a broader class of textures, tensorial PDE-based regularization [28] is incorporated into Eq. (I.2) to yield:

$$X_t = -\tilde{h} * (X * h - Y) + \lambda \nabla (D(\nabla X) \nabla X), \qquad \text{(I.3)}$$

where $D(\nabla X)$ denotes the tensor diffusivity.

We have previously reported [29] of an algorithm which performs fBm-based deblurring. In this work we substantiate the results and use the principles of the fBm to yield an image model and a true superresolution algorithm.

II. Problem Statement

The following form of the superresolution problem is considered: A high-resolution (HR) image is degraded by a blurring filter, representing, for example, the PSF of an optical sensor. It is subsequently subsampled. Noise is then additively mixed with the blurred and subsampled image to create the available low-resolution (LR) image. Let $X(\eta_1, \eta_2)$ and $Y(\eta_1, \eta_2)$ denote the original (HR) image and observed (LR) noisy image, respectively. The imaging model can be represented as follows:

$$Y(\eta_1, \eta_2) = \mathcal{D}\left((X * b)(\eta_1, \eta_2)\right) + N(\eta_1, \eta_2), \qquad \text{(II.1)}$$

where $\mathcal{D}$ is the subsampling operator, $b(\eta_1, \eta_2)$ is a noninvertible blur kernel and $N(\eta_1, \eta_2)$ is an independent additive white Gaussian noise. In the case of SR, we assume that the noise has low variance, unlike denoising problems where

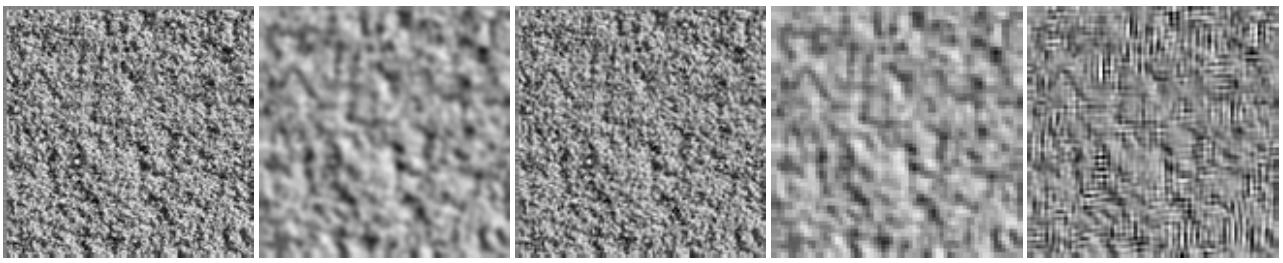

(a) (b) (c) (d) (e)

Fig. 1. Applying deblurring algorithms on blurred and decimated images: (a) The original image. (b) and (c) Blurred and BM3D-deblurrred image respectively [33]; (d) and (e) Blurred and subsampled image and BM3D-deblurrred image, respectively.

the noise is substantial. The blur kernel is assumed to have limited spatial support, and therefore the decimation operator introduces aliasing.

The superresolution (SR) problem is severely ill-posed due to the decimation operator. For this reason, the effect of decimation is often ignored in SR studies. The single-frame SR problem is a special case of the classical multi-frame SR problem [30]–[32]. In the latter, several degraded images are available, each containing unique details of the single original image, acquired by sub-pixel shifts. A high-resolution reconstruction is obtained in this case by recovering complementary information from all the measurements.

In contrast, in single-frame SR, known in the literature also as upscaling, only a single measurement is available. The single-frame SR problem can be formally stated as follows:

$$\hat{X}(\eta_1, \eta_2) = \arg\min_{X \in \mathcal{X}} \| Y(\eta_1, \eta_2) - \mathcal{D}\left((X * b)(\eta_1, \eta_2)\right) \|_2. \tag{II.2}$$

The SR image, $\hat{X}(\eta_1, \eta_2)$, thus obtained is the best one in that it yields the smallest $\mathcal{L}_2$ error relative to the original image (ground truth). The SR problem appears, at first look, to be equivalent to a deconvolution problem, in which a SR image, $\hat{X}(\eta_1, \eta_2)$, is to be recovered from the blurred and noisy image, $Y(\eta_1, \eta_2)$. It is important to emphasize that this is by no means the case. In the case of SR, the decimation operator introduces aliasing artifacts due to the loss of details in the subsampling.

Deblurring algorithms are not effective in the case of SR of textures. For example (Fig. 1), while BM3D-deblurring [33] performs deblurring successfully on blurred images, when decimation is introduced the results deteriorate. This is emphasized in the case of textures.

There are studies that first solve the deblurring problem, and then apply interpolation to "inverse" the effect of the decimation operator. This approach is applicable only when the blur filter acts as an anti-aliasing filter. Blur filters considered in the present study, and in SR problems in general, have, however, a small spatial support, and cannot be considered as anti-aliasing filters. Therefore, our emphasis is on the *model* for images, based on which, the missing data can be reconstructed, to yield a high fidelity estimate of the original image.

III. Fractional Brownian Motion

The fractional Brownian motion (fBm)-a Gaussian random process-was introduced by Mandelbrot and Van Ness as a model suitable for natural images [21]. The fBm generalizes the well-known Brownian motion in that the increments are stationary but not independent. It is defined, in one dimension, as a Gaussian process with zero mean and the following autocorrelation function:

$$E\left[B_H(t)B_H(s)\right] = \frac{\sigma^2}{2}\left(|t|^{2H} + |s|^{2H} - |t-s|^{2H}\right), \tag{III.1}$$

where

$$\sigma^2 = \frac{\sigma_W^2}{2}\frac{\cos(\pi H)}{\pi H}\Gamma(1-2H), \tag{III.2}$$

σ_W^2 is a known variance, and the Hurst parameter, $H \in (0, 1)$, controls the regularity of the process. This is a non-stationary process with stationary increments - a property to be exploited later in efficient synthesis. The first sample is usually set to zero, $B(0) = 0$, further indicating that it cannot be stationary (unless it is zero everywhere). For Hurst parameter values of $H \in \left(0, \frac{1}{2}\right)$, this process exhibits negative correlation between samples (anti-persistence), and for $H \in \left(\frac{1}{2}, 1\right)$, it exhibits positive correlation and long range dependencies. A special case is $H = \frac{1}{2}$, for which this process becomes the well-known Brownian motion, or Wiener process. This process exhibits two important properties, known to characterize natural images, in the context of Mandelbrot's work on fractals [21]. The first one is long-range dependencies between samples, where for $H > 0.5$ the sum of the correlations of the increments diverges. The second property of the fBm, which is exploited in this study, is the statistical self-similarity, defined as:

$$B_H(at) \stackrel{d}{=} |a|^H B_H(t), \tag{III.3}$$

for a positive number a, where the superscript d stands for equality in distribution. This equality indicates that the sample distribution across different scales is varied only by a constant depending on the scale, a, and the Hurst parameter. These two properties highlight the relevance of this process to natural textures, as the latter often exhibit similarities between adjacent as well as distant pixels [23].

A. Synthesis in Two Dimensions

Since the fBm process is a Gaussian process with known covariance function, one can explicitly synthesize a realization in the discrete domain [34]. Recall the following covariance property for multivariate random variables:

$$\text{cov}(LZ) = L\Sigma_Z L^T, \tag{III.4}$$

where Z is a random vector with autocorrelation matrix Σ_Z, and L is a matrix. In order to synthesize the fBm in this manner, one needs to first build the autocorrelation matrix according to Eq. (III.1). Then, the Cholesky decomposition is used to obtain a matrix, L, such that $M = LL^T$, where M is the fBm autocorrelation matrix. This decomposition is possible since it is positive-definite. Then, upon multiplying the lower-triangular matrix L by a multivariate random vector Z with a unity covariance matrix, the covariance of the resulting vector $B = LZ$ fits fBm.

This simple algorithm can be extended to two dimensions, by representing (stacking) a 2D image as a vector and constructing the autocorrelation matrix with the respective 2D dependencies. At this stage it is also worth to note that this process is isotropic in the statistical sense. Let $p = (x_1, y_1)^T$ and $q = (x_2, y_2)^T$ be two points. Then, the autocorrelation of the fBm in 2D is defined as follows:

$$E\left[B_H(p)B_H(q)\right] = \frac{\sigma^2}{2}\left(\|p\|^{2H} + \|q\|^{2H} - \|p-q\|^{2H}\right). \tag{III.5}$$

Let M be a rotation matrix, indicating $|M| = 1$ and $M^{-1} = M^T$. It is straightforward to check that for an arbitrary point $l = (x_0, y_0)^T$, $\|Ml\| = \|l\|$ and therefore $E[B_H(Mp)B_H(Mq)] = E[B_H(p)B_H(q)]$. This indicates the autocorrelation is invariant under rotations and the process is isotropic.

While the aforementioned method can produce the exact 2D fBm process, it is very inefficient. This is due to the dimensional requirement of the covariance matrix, which dictates a size of $N^2 \times N^2$ for an image size of $N \times N$ and time complexity of $O(N^6)$ due to the Cholesky decomposition. The space and time complexity render this method to become impractical even in application to moderately-sized images.

It is possible, however, to use more efficient methods of synthesizing nearly-exact realizations of the 2D fBm. We adopt the method proposed by Kaplan and Kuo [35], which implements Fourier synthesis, and yields accurate realizations with time complexity of $O(N^2 log_2(N))$ and space complexity of $O(N^2)$.

This method utilizes the stationary increments of the fBm and builds the realization by first calculating the autocorrelations of the increments in the $\underline{x}$, $\underline{y}$, and $(\underline{x}, \underline{y})$ directions, synthesizing the increments in the frequency domain, and then summing them to produce the final result. Two realizations of the 2D fBm process, for two typical values of H, are displayed in Fig. 2. A low value of H is better fitted to high-frequency content relevant to stochastic textures.

B. Synthesis of Non-Stationary Fields With Stationary Increments

The fBm can be considered as a special case of a family of non-stationary processes with stationary increments. As such, it is a simple case, as it is statistically isotropic and is being governed by a single parameter, H. In [36], the 2D fBm synthesis algorithm of Kaplan and Kuo is generalized for any such field. The synthesized fields are derived with reference to an initial white noise image, $W(\eta_1, \eta_2)$, and a structure function, $\phi(\eta_1, \eta_2)$, which defines the autocorrelation of the fBm, $F(\eta_1, \eta_2)$, by the following equation:

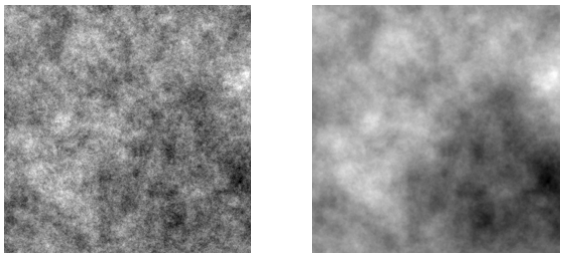

(a) (b)

Fig. 2. 2D fBm realization for two values of H: (a) A value of $H = 0.1$, depicting negative correlation between adjacent pixels. (b) A value of $H = 0.6$, depicting high correlation between adjacent pixels.

$$E[F(\eta_1, \eta_2)F(\eta_1', \eta_2')] = -E[F(0,0)^2] + \frac{1}{2}(\phi(\eta_1, \eta_2) + \phi(\eta_1', \eta_2') - \phi(\eta_1 - \eta_1', \eta_2 - \eta_2')), \tag{III.6}$$

where $F(0, 0)$ is set to zero for the synthesis process. This structure function defines the autocorrelation of the increments of the desired field. For the fBm case, we obtain:

$$\phi(x, y) = C(x^2 + y^2)^H, \quad H \in (0, 1), \tag{III.7}$$

for a suitable normalizing constant, C, and a proper Hurst parameter, H. In this manner, anisotropic fBm fields can be synthesized by choosing a different structure function - one that depends on more than a single Hurst parameter. For the exact synthesis algorithm, as well as other limitations of the structure function, see [36]. We take it one step further and propose an adaptive structure function to be used in image enhancement of textured images.

C. Remarks

We assume that the details missing in degraded textures had, originally, dependencies similar to those characteristic of an fBm. Therefore, using a single realization of a proper fBm as an initial image, we should be able to restore missing high-resolution details by fitting such a realization onto a degraded image. Since we are interested in high frequencies, the common values for the Hurst parameter should be in the range $\left(0, \frac{1}{2}\right)$.

IV. Phase of the Frequency Response

The importance of phase, and of "local phase", in signal and image processing is well-established [37], [38]. It has been shown that for natural images, the important information of the image is, in fact, stored in the phase rather than in the magnitude of the frequency domain representation of the image [38], [39]. The magnitude contains information about the frequencies present in the image, which are common to a large class of natural images. The phase contains information about the spatial relationship of these frequencies in a specific image. Let $X(\eta_1, \eta_2)$ and $Y(\eta_1, \eta_2)$ be two images with

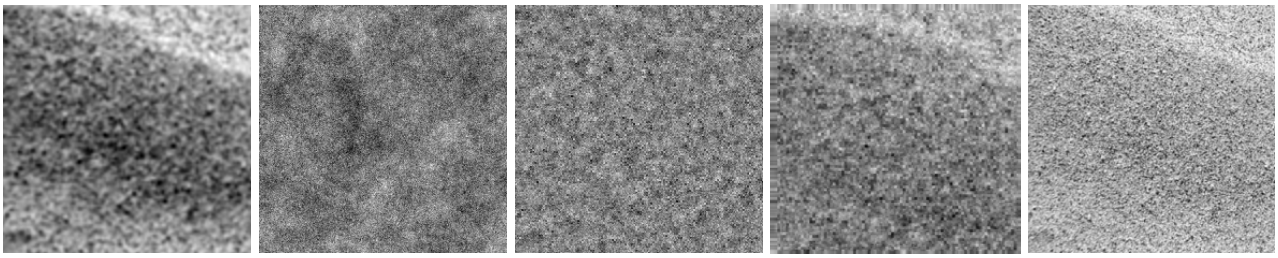

(a) (b) (c) (d) (e)

Fig. 3. Artificial texture, generated by means of the model outlined in section V, with a natural texture image as $X_{LP}(\eta_1, \eta_2)$. (a) Low-pass image, $X_{LP}(\eta_1, \eta_2)$, generated by blur and subsampling of a real image. (b) 2D fBm image, with Hurst parameter, $H = 0.1$. (c) $X_{HP}(\eta_1, \eta_2)$, the 2D fBm image obtained after phase matching for $X_{LP}(\eta_1, \eta_2)$. (d) The artificial texture image, $X(\eta_1, \eta_2)$. (e) The original image used for generating $X_{LP}(\eta_1, \eta_2)$, for comparison. Note the similar details obtained in this image and the one in which the fBm model has been applied.

respective Fourier transforms $\tilde{X}(\tilde{\eta_1}, \tilde{\eta_2})$ and $\tilde{Y}(\tilde{\eta_1}, \tilde{\eta_2})$. The Fourier transform, $\tilde{M}(\tilde{\eta_1}, \tilde{\eta_2})$, of a signal, $M(\eta_1, \eta_2)$, can be decomposed into magnitude and phase components:

$$\tilde{M}(\tilde{\eta_1}, \tilde{\eta_2}) = |\tilde{M}(\tilde{\eta_1}, \tilde{\eta_2})| \exp\left(j \cdot \angle \tilde{M}(\tilde{\eta_1}, \tilde{\eta_2})\right). \quad \text{(IV.1)}$$

In this manner, one can define a frequency representation of an image, $\hat{\tilde{Y}}(\tilde{\eta_1}, \tilde{\eta_2})$, by using the absolute value of $\tilde{Y}(\tilde{\eta_1}, \tilde{\eta_2})$ and the phase of $\tilde{X}(\tilde{\eta_1}, \tilde{\eta_2})$.

The resulting image, $\hat{Y}(\eta_1, \eta_2)$, obtained by the inverse Fourier transform of $\hat{\tilde{Y}}(\tilde{\eta_1}, \tilde{\eta_2})$, will contain the frequency components present in $Y(\eta_1, \eta_2)$, but the phase, and therefore the spatial shifts, will be according to $X(\eta_1, \eta_2)$.

This simple property enables us to exploit a synthetic 2D fBm image for our purposes. The synthetic fBm is inherently random, as it is generated by an iid Multivariate normal matrix. However, by using the phase matching technique described above, one can benefit from the frequency spectrum characteristic of a 2D fBm image, while fitting the phase relationships of a specific image. We thereby effectively exploit the fBm correlations while imposing minor dependence on the initial white noise.

V. A Model for Self-Similar Textures

We now present a model for self-similar textures. This model is based on the high frequencies of the 2D fBm process and fits a variety of images corresponding to the fBm statistics. The model can be used to generate such textures as well as be used for image enhancement tasks, such as superresolution or denoising (although denoising is not covered by this study).

The model is defined as follows. Let $X_{LP}(\eta_1, \eta_2)$ be a low frequency image. This image can be generated using fBm models, autoregressive-moving-average models, or be based upon a natural image. Let $X_{HP}(\eta_1, \eta_2)$ be a high-frequency image, obtained as follows:

$$X_{HP}(\eta_1, \eta_2) = \mathcal{P}_{H,W(\eta_1,\eta_2)}\left(X_{LP}(\eta_1, \eta_2)\right), \quad \text{(V.1)}$$

where $\mathcal{P}_{H,W(\eta_1,\eta_2)}$ is an operator performing phase matching, as described earlier, for a 2D-fBm image generated according to the Hurst parameter, H, from the white noise image, $W(\eta_1, \eta_2)$. The value of H will typically be low, $H \leq 0.5$, to generate high frequencies. The texture image, $X(\eta_1, \eta_2)$, is then constructed from the superposition of both images as follows:

$$X(\eta_1, \eta_2) = X_{LP}(\eta_1, \eta_2) + (X_{HP} * h_{HP})(\eta_1, \eta_2) + V(\eta_1, \eta_2), \quad \text{(V.2)}$$

where $h_{HP}(\eta_1, \eta_2)$ is a high-pass filter and $V(\eta_1, \eta_2)$ is residual noise compensating for model inaccuracies.

We note that an fBm image can be derived as a special case of this model, since $X_{LP}(\eta_1, \eta_2)$ can be the lowpass version of the fBm generated by $\mathcal{P}_{H,W(\eta_1,\eta_2)}$. Obviously, natural image textures cannot be represented by this model without error. It is straightforward to check that given an image, $I(\eta_1, \eta_2)$, the energy content of the error, $V(\eta_1, \eta_2)$, will reside in the high-frequency range. Therefore, in order to assess the model in natural images, we have to consider the error in the high-frequency range.

While a naïve measure norm, such as $\mathcal{L}_2$ norm, may be suggested for the evaluation of this error, we do not use it. This is due to the well-known shortcomings associated with $\mathcal{L}_2$-based comparisons [40], which are especially emphasized in the case of stochastic textures. Instead, and lacking a better method, the results will be inspected both visually and by means of image statistics such as the histogram.

An example of the model is depicted in Fig. 3. $X_{LP}(\eta_1, \eta_2)$ [Fig. 3(a)] is generated by blur and subsampling of a real image texture. While the resulting image [Fig. 3(d)] is not identical to the original image [Fig. 3(e)], the high frequencies are visually similar. This encourages us to use this model in our texture superresolution algorithm.

Fig. 4 depicts two typical images, one suitable for the proposed model and the other one is not. In the first row we have a stochastic texture. Two high-passed versions are shown; Fig. 4(b) is the high-pass of the original texture ($I_{HP}(\eta_1, \eta_2)$), and Fig. 4(d) is the high-pass of the texture according to the proposed model, $X_I(\eta_1, \eta_2)$. These two high-pass images are then compared by their histograms [Fig. 4(e)]. We observe that the high-pass versions are both visually similar, and have similar histograms. Additionally, we observe that the histogram is of a Gaussian shape. We show in [41] that a necessary condition for an fBm image is to have a Gaussian-shaped 1D histogram, in the limit, as $H \to 0$ and $N \to \infty$. Despite this being only a necessary condition,

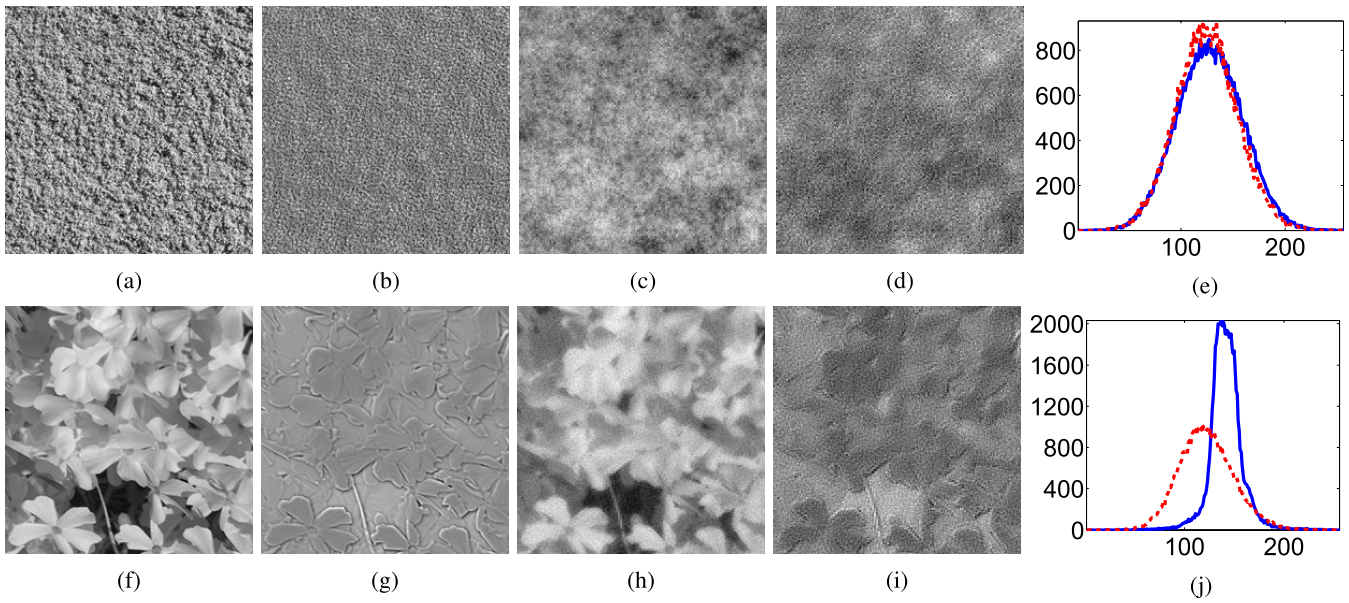

Fig. 4. Comparison of two natural textures, one suitable and one unsuitable for the model described in section V. (a) and (f) Original images of natural textures. (c) and (h) 2D fBm versions of the respective images. (b) and (g) High-pass versions of the original images, respectively. (d) and (i) High-pass versions of the 2D fBm images, respectively. (e) and (j) Histograms of the two high passed versions respectively. The high-pass version histogram of the original images is highlighted in red-dashed, and the histogram of the high pass versions of the 2D fBm is highlighted in blue.

we are encouraged to seek further indications of the relevance of this model to textured images. We further substantiate in the next subsection in the form of 2D histogram.

In contrast with the texture of Fig. 4(a), Fig. 4(f) depicts a regular (structured) texture, and its high-pass versions is shown in Fig. 4(g). In this case, the two histograms [Fig. 4(j)] are distinctly different. We conclude that the model can not faithfully represent regular textures. This is due to the fact that the fBm cannot represent such images.

A. 2D Histograms

The 1D histogram provides an empirical estimate for image first order distribution. Consequently, it constitutes many-to-one mapping that is not sufficient for a suitable assessment of the model. We therefore present a second indicator, in the form of 2D histograms of adjacent pixels in an image, in the $\underline{x}$, $\underline{y}$ and diagonal orientations. The latter provides an estimate of the image second order distribution. The 2D histogram of an image $X(\eta_1, \eta_2)$ with 2D bins $\left\{S_{l_1,l_2}\right\}_{l_1,l_2}$ is defined as follows:

$$H(i,j) = \frac{1}{N} \sum_{(p_1,p_2)\in\mathcal{Q}} 1_{(X(p_1),X(p_2))\in\mathcal{S}_{i,j}}, \qquad \text{(V.3)}$$

where the set $\mathcal{Q}$ contains the locations of all adjacent pixels in $X(\eta_1, \eta_2)$, $p_i = (\eta_1^i, \eta_2^i)$ is a pixel location, and $N = |\mathcal{Q}|$. This is a discrete surface with two coordinates, representing the two gray levels of adjacent pixels, denoted by l_1 and l_2.

For smooth images, it is expected that $|l_1 - l_2| < \epsilon$ for a small ϵ. Most of the energy of the histogram resides close to the line $l_1 = l_2$. This is the case in the presence of edges as well, due to the relatively low number of edge pixel pairs with respect to smooth pixel pairs in natural images.

In the case of stochastic textures, suitable for the fBm model, adjacent pixels exhibit negative correlation. It is therefore expected that $|l_1 - l_2|$ will be large. For isotropic stochastic textures, the histogram shape is expected to approximate a 2D normal distribution, whereas for anisotropic stochastic textures, negative correlation is still exhibited, but the histogram shape deviates from the normal distribution.

Fig. 5 depicts three examples of characteristic images concerned with the proposed model, along with their 2D histograms and cross-cuts along the main and secondary diagonals. The first image [Fig. 5(a)] is clearly a non-stochastic texture. In this and other structured natural images, there is significant correlation between adjacent pixels as they are mostly smooth. Therefore, the energy of the 2D histogram is concentrated near the diagonal, $l_1 = l_2$. The cross-cut graph shows a non-Gaussian behaviour.

The second image [Fig. 5(b)] is of a stochastic texture, well fitting the proposed fBm model. In this case, we distinctly see a 2D normal distribution. The third image [Fig. 5(c)] is of a stochastic texture, which does not fit the model in its current form. For this type of anisotropic textures we propose a regularization that takes care of the anisotropy in the context of SR.

Anisotropic stochastic textures, while being similar to isotropic in large scales, are composed of small oriented shapes, which are anisotropic and do not exhibit self-similarity in arbitrarily small scales. The characteristic 2D histogram of such textures deviates from a 2D multivariate normal, but it is still sufficiently close enough to the Gaussian.

In Wavelet-based analysis of coefficients in adjacent scales or orientations, similar to what was done in [25], the same properties can be observed, indicating that stochastic textures indeed obey a Gaussian distribution.

VI. Superresolution Algorithm

A. Anisotropic Diffusion

A brief review of the anisotropic diffusion that will suffice for our application is provided. For a comprehensive

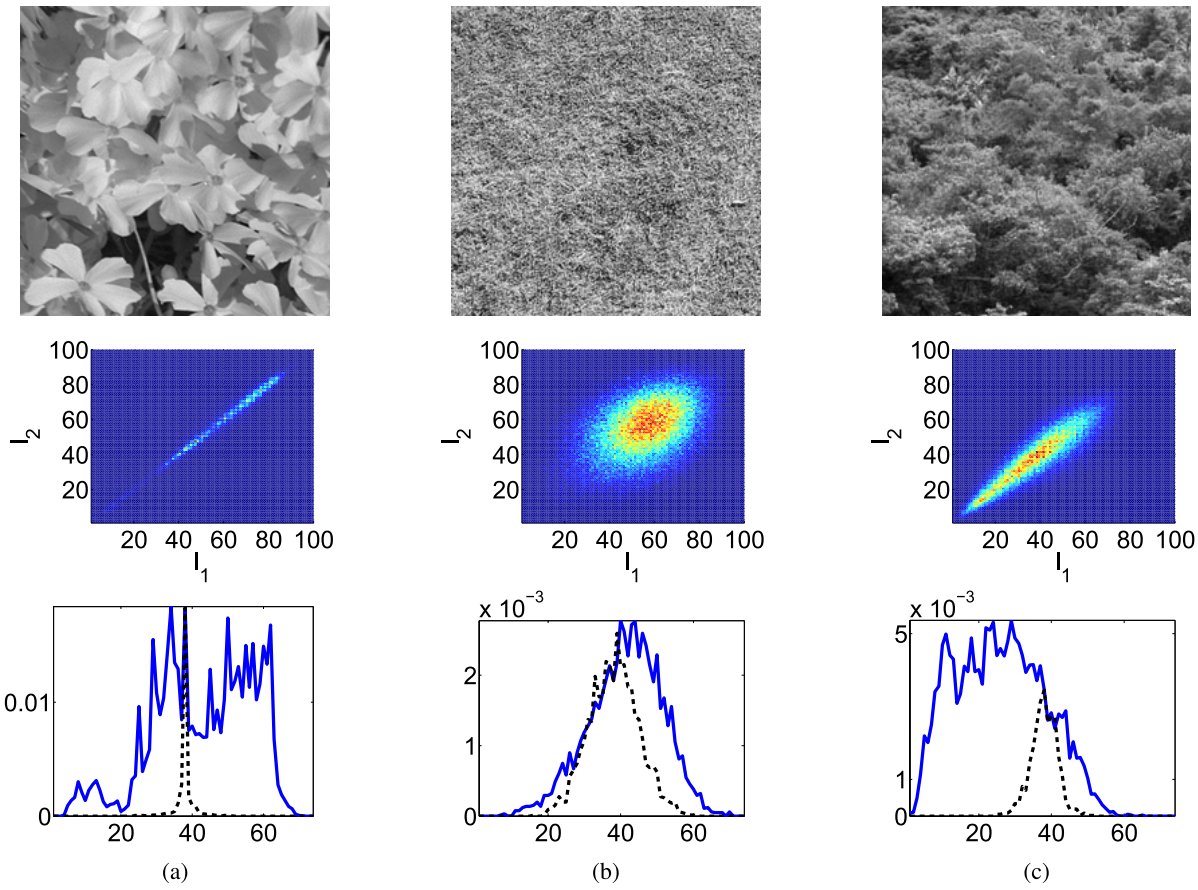

Fig. 5. 2D histograms of three characteristic textures: Original images (first row), 2D histograms (second row) and cross-cuts of the main and secondary diagonals of the histograms (third row), depicted in blue and black-dashed, respectively. (a) Non-stochastic texture: The 2D histogram data are clustered co-linearly along the line $l_1 = l_2$, due to the high correlation between adjacent pixels, and the cross-cuts show non-Gaussian behaviour. (b) Stochastic texture, suitable for the fBm model: The 2D histogram depicts a shape of a normal distribution, also apparent in the cross-cuts. (c) Anisotropic stochastic texture, which does not fit the current model: In this case, one obtains a non-Gaussian distribution. In both the histogram and the cross-cuts it is apparent that the shape is close to a Gaussian.

exposition see, for example, [28], [42]. Using PDE-based methods allows for adaptive filtering of an image, with low computational complexity. The following PDE equation suitable for image processing was introduced in this context by Perona and Malik [43]:

$$I_t = \nabla \cdot (g(\nabla I)\nabla I), \tag{VI.1}$$

with the initial condition $I|_{t=0} = I_0$, the diffusivity $g(s)$ being a decreasing function with the following properties: $g(0) = 1$, $g(s) \geq 0$ and $g(s) \to 0$ as $s \to \infty$. The diffusivity function is, in general, designed for the adaptive processing of images to allow high diffusion in low gradient areas, assumed to be noisy, and low diffusion in high gradient areas, indicating the presence of edges that should not be compromised.

Many choices have been suggested for this function. Perona and Malik (PM) proposed $g(s) = e^{-(s/K)^2}$ and $g(s) = \frac{1}{1+\left(\frac{s}{K}\right)^2}$, which are commonly used with a suitable constant, K. Also useful is the TV-based diffusivity, using the $\mathcal{L}_1$ norm: $g(s) = \frac{1}{|s|}$, or the regularized version, $g(s) = \frac{1}{\sqrt{\epsilon^2+s^2}}$. Gilboa et al. have even extended it to the negative regime [42].

This diffusion, although commonly referred to *anisotropic*, is in fact non-linear but isotropic. This has been noted by Weickert, who introduced a truly anisotropic diffusion process, commonly referred to as tensor diffusion:

$$I_t = \nabla \cdot (D(\nabla I)\nabla I), \tag{VI.2}$$

where $D \in \mathbb{R}^{2\times 2}$ is a tensor that is represented, using an eigenvalue decomposition, as follows:

$$D = (\omega_1, \omega_2)\begin{pmatrix}\lambda_1 & 0\\ 0 & \lambda_2\end{pmatrix}\begin{pmatrix}\omega_1\\ \omega_2\end{pmatrix}, \tag{VI.3}$$

where ω_1 and ω_2 are eigenvectors which satisfy:

$$\omega_1 \parallel \nabla I, \quad \omega_2 \perp \nabla I, \tag{VI.4}$$

and λ_1 and λ_2 are the corresponding eigenvalues. This formulation allows for different types of diffusion to be performed in different orientations within the image. In edge enhancing diffusion, for instance, only the diffusion coefficient perpendicular to the edge orientation will assume a significant value. This method further emphasizes edges while smoothing noisy image areas. Instead of a single diffusivity function, $g(x)$, two functions are used - one for each eigenvalue.

B. Texture-Based Tensor Diffusion

The tensor, D, is a function of the gradient of the image, ∇I. Due to the fact that textures contain small oriented elements,

the tensor diffusion is a natural choice for their enhancement. However, commonly used tensor diffusion schemes do not yield a successful enhancement of textures; while the edges are emphasized, the high frequency components of textural details are lost. This is due to their spectral resemblance of noise. Further, when attempting to enhance a blurred image, much of the high frequency details are in effect non-existing to begin with.

This has encouraged us to consider a different function, instead of ∇I, for the calculation of the tensor. This function is required to represent the desired properties of the texture, while still resembling the shape of the texture itself.

One cannot expect to represent a natural texture using a single parameter, H. As discussed in section III-B, it is possible to consider a structure function, $\phi(\eta_1, \eta_2)$, to create a non-stationary field which better represents the desired image. Instead of using a general function, we use a structure function generated from the degraded image itself. This yields an image which contains the details of the degraded image, along with correlations introduced according to the specific structure of the non-stationary field. We refer to the structure function derived from the degraded image as the empirical structure function (ESF).

The method to recover the ESF from a given, degraded, image is based on an inverse procedure to the method of obtaining the image from the structure function, devised in [36]. Let $Y(\eta_1, \eta_2)$ be a degraded image. The increments in the $\underline{x} = \eta_1$ and $\underline{y} = \eta_2$ orientations are defined as:

$$\begin{aligned} Y_{\Delta\eta_1}(\eta_1, \eta_2) &= Y(\eta_1, \eta_2) - Y(\eta_1 - \Delta_{\eta_1}, \eta_2), \\ Y_{\Delta\eta_2}(\eta_1, \eta_2) &= Y(\eta_1, \eta_2) - Y(\eta_1, \eta_2 - \Delta_{\eta_2}), \end{aligned} \quad \text{(VI.5)}$$

respectively, where Δ_{η_1} and Δ_{η_2} are small increments set to 1 in discrete schemes. The increments in the $(\underline{x}, \underline{y}) = (\eta_1, \eta_2)$ coordinates are defined as:

$$\begin{aligned} Y_{\Delta\eta_1,\Delta\eta_2}(\eta_1, \eta_2) &= Y(\eta_1, \eta_2) - Y(\eta_1 - \Delta_{\eta_1}, \eta_2) \\ &\quad - Y(\eta_1, \eta_2 - \Delta_{\eta_1}) + Y(\eta_1 - \Delta_{\eta_1}, \eta_2 - \Delta_{\eta_2}). \end{aligned} \quad \text{(VI.6)}$$

Let $R_{\eta_1}(\eta_1, \eta_2)$, $R_{\eta_2}(\eta_1, \eta_2)$ and $R_{\eta_1,\eta_2}(\eta_1, \eta_2)$ be the autocorrelation functions of the increments $Y_{\Delta\eta_1}$, $Y_{\Delta\eta_2}$ and $Y_{\Delta\eta_1,\Delta\eta_2}$, respectively. The autocorrelation functions for the 1D increments, $R_{\eta_1}(\eta_1, \eta_2)$ and $R_{\eta_2}(\eta_1, \eta_2)$, are derived from the structure function, $\phi(\eta_1, \eta_2)$, as follows:

$$\begin{aligned} R_{\eta_1}(\eta_1, \eta_2) &= \frac{1}{2}(\phi(\eta_1 + \Delta_{\eta_1}, \eta_2) + \phi(\eta_1 - \Delta_{\eta_1}, \eta_2) \\ &\quad - 2\phi(\eta_1, \eta_2)), \\ R_{\eta_2}(\eta_1, \eta_2) &= \frac{1}{2}(\phi(\eta_1, \eta_2 + \Delta_{\eta_2}) + \phi(\eta_1, \eta_2 - \Delta_{\eta_2}) \\ &\quad - 2\phi(\eta_1, \eta_2)), \end{aligned} \quad \text{(VI.7)}$$

and the autocorrelation for the 2D increments, $R_{\eta_1,\eta_2}(\eta_1, \eta_2)$, is accordingly derived as:

$$\begin{aligned} R_{\eta_1,\eta_2}(\eta_1, \eta_2) &= 2R_{\eta_1}(\eta_1, \eta_2) + 2R_{\eta_2}(\eta_1, \eta_2) \quad \text{(VI.8)} \\ &\quad - \frac{1}{2}(\phi(\eta_1 + \Delta_{\eta_1}, \eta_2 + \Delta_{\eta_2}) \\ &\quad + \phi(\eta_1 - \Delta_{\eta_1}, \eta_2 + \Delta_{\eta_2}) + \phi(\eta_1 + \Delta_{\eta_1}, \eta_2 - \Delta_{\eta_2}) \\ &\quad + \phi(\eta_1 - \Delta_{\eta_1}, \eta_2 - \Delta_{\eta_2})). \quad \text{(VI.9)} \end{aligned}$$

To obtain the empirical structure function, it is therefore required to invert the equations, and produce $\phi(\eta_1, \eta_2)$, given the increment autocorrelation functions of $Y(\eta_1, \eta_2)$. Substituting $\Delta_{\eta_1} = \Delta_{\eta_2} = 1$ in Eq. (VI.7), it follows that the 1D autocorrelation functions can be represented using convolution equations with derivative filters:

$$\begin{aligned} R_{\eta_1}(\eta_1, \eta_2) &= (\phi * f_d)(\eta_1, \eta_2), \\ R_{\eta_2}(\eta_1, \eta_2) &= (\phi * f_d^T)(\eta_1, \eta_2), \end{aligned} \quad \text{(VI.10)}$$

where $f_d = \frac{1}{2}(1, -2, 1)$. The 2D autocorrelation can be represented in a similar manner, using the following equation:

$$R_{\eta_1,\eta_2}(\eta_1, \eta_2) = (\phi * f_{d2})(\eta_1, \eta_2), \quad \text{(VI.11)}$$

where

$$f_{d2} = \frac{1}{2}\begin{pmatrix} -1 & 2 & -1 \\ 2 & -4 & 2 \\ -1 & 2 & -1 \end{pmatrix}. \quad \text{(VI.12)}$$

Obtaining the ESF from the degraded image, is therefore reduced to solving Eqs. (VI.10) and (VI.11). This can be formulated as the following least-squares (LS) problem (Appendix A):

$$\phi = \arg\min_x \|D_f x - r\|_2^2, \quad \text{(VI.13)}$$

where ϕ is the column-stack representation of the ESF, $\phi(\eta_1, \eta_2)$, D_f in a suitable matrix representation of f_d and f_{d2}, and r is a suitable column-stack representation of $R_x(\eta_1, \eta_2)$, $R_y(\eta_1, \eta_2)$ and $R_{x,y}(\eta_1, \eta_2)$. This is an ill-posed problem, due to rank deficiency of the derivative matrix, D_f. It is similar to problems encountered in gradient domain processing. This poses a challenge to a least-squares procedure and we currently do not employ regularization techniques, although these may be relevant in further studies. Additionally, unlike problems in which the vector, ϕ, needs to be recovered exactly, in this case only the derivatives (in the manner of the derivative filters, f_d and f_{d2}) of the ESF are required. The derivation of these matrices and vectors is addressed in further details in Appendix A.

An example is presented in Fig. 6, where two sets of images are shown. The first [Fig. 6(a) and (b)] contains a 2D fBm image, with $H = 0.1$, and the restored image. The latter is derived by extracting the ESF from the 2D fBm image by the process described earlier, and then reconstructing the field by the algorithm presented in [36]. Due to the ill-posedness of the LS problem, the two images do not look alike. The second set [Fig. 6(c) and (d)] depicts the respective images obtained from the first set, after a high-pass filtering was performed by subtracting the result from a Gaussian lowpass filter with $\sigma = 15$. Since only the high frequency range is lost by degradation, this range is of more importance. Indeed, the high pass versions of Fig. 6(a) and (b), depicted in Fig. 6(c) and (d) appear to be visually similar.

Using the ESF, it is possible to obtain an image, $Y_\phi(\eta_1, \eta_2)$, from the degraded image, $Y(\eta_1, \eta_2)$, by calculating the autocorrelation of the first- and second-order increments, solving the LS problem in Eq. (VI.13) to obtain a structure function $\phi(\eta_1, \eta_2)$, and using the synthesis algorithm in [36]. The resulting image is referred to as the empirical image.

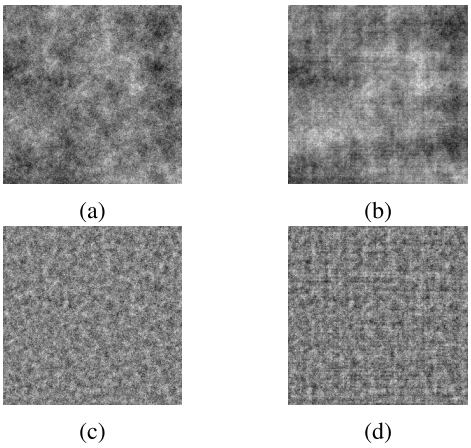

Fig. 6. Comparison of 2D fBm using the empirical structure function (ESF). (a): Original 2D fBm image. (b) Restored 2D fBm image, using ESF derived from (a). High pass versions of (a) and (b) are shown in (c) and (d), respectively.

C. Tensor Diffusion

We now consider the modifications required to enable the tensor diffusion to perform superresolution on natural textures. The tensor, $D(\nabla I)$, introduced earlier, is set instead to be $D(\nabla(I_t + \alpha Y_\phi(\eta_1, \eta_2)))$, where $Y_\phi(\eta_1, \eta_2)$ is the empirical image, and α is a weight parameter. This allows for the introduction of missing texture details, while still emphasizing the edges of a degraded texture image.

The superresolution algorithm is presented by considering the following energy functional, in column-stacked image representation:

$$E(X, \nabla X) = \int_\Omega (BX - Y)^2 + (\hat{X}_{HP} - H_{HP}X)^2 + \beta\Psi(|\nabla X + \alpha\nabla Y_\phi|^2)dxdy, \quad \text{(VI.14)}$$

where X is the SR image, Y is the degraded image, and B is a matrix performing blur and decimation. The second term penalizes deviations of the solution, X, from the fBm model of the image (discussed in section V), given by $\hat{X}$. This is derived as follows:

$$\begin{aligned} \|X - \hat{X}\|^2 &= \|X_{LP} + X_{HP} - \hat{X}_{LP} - \hat{X}_{HP}\|^2 \\ &\tilde{=} \|X_{HP} - \hat{X}_{HP}\|^2 \\ &= \|H_{HP}X - \hat{X}_{HP}\|^2, \end{aligned} \quad \text{(VI.15)}$$

where H_{HP} is a Gaussian high-pass filter described earlier, and $\hat{X}_{HP}$ is the high-frequency component, given by $\hat{X}_{HP} = \mathcal{P}_{H,W}(\hat{X}_{LP})$. X_{LP} and $\hat{X}_{LP}$ are assumed to be sufficiently close so that their difference is approximately zero. This is a valid assumption, as the imaging model does not degrade the low frequencies.

The solution for this equation satisfies the following Euler-Lagrange equation:

$$B^T(BX - Y) - \beta\nabla\cdot(\Psi'(|\nabla(X + \alpha Y_\phi)|^2)\nabla X) - 2\alpha\beta\nabla\cdot(\Psi'(|\nabla(X + \alpha Y_\phi)|^2)\nabla Y_\phi) = 0. \quad \text{(VI.16)}$$

This derivation is further addressed in Appendix B. A gradient descent scheme for the minimization of the energy functional, and the introduction of the tensor diffusion, $D(\nabla(X + Y_\phi))$, instead of the scalar diffusivity, $\Psi(|\nabla(X + Y_\phi)|^2)$, yield the following diffusion-reaction process:

$$\begin{aligned} X_t = &-2B^T(BX - Y) + 2H_{HP}^T(\hat{X}_{HP} - H_{HP}X) \\ &+\beta\nabla\cdot(D(\nabla(X + \alpha Y_\phi))\nabla X) \\ &+2\alpha\beta\nabla\cdot(D(\nabla(X + \alpha Y_\phi))\nabla Y_\phi). \end{aligned} \quad \text{(VI.17)}$$

This equation differs from the usual deblurring tensor diffusion with regard to the following.

1) As previously discussed, the tensor, D, is a function of two images. Based on this formulation, an extra facet emerges in the optimization.
2) The estimated high-frequency image, X_{HP}, is added, in order to recover the missing details.

These two arguments work in tandem; details are recovered by the texture model, and are diffused according to the correlations emerging from the empirical image in order to fit the current image. We show in [41] that both terms are required for adequate texture restoration.

The eigenvalue parallel to the gradient, λ_1, was set to a regularized $\mathcal{L}_1$ scalar diffusivity function:

$$\lambda_1(s) = \frac{1}{\sqrt{\epsilon^2 + s^2}}, \quad \text{(VI.18)}$$

and the perpendicular eigenvalue, λ_2, was set to 0.01.

D. Remarks

1) The empirical image, $Y_\phi(\eta_1, \eta_2)$, is initially derived from the degraded image, $Y(\eta_1, \eta_2)$. However, as the diffusion advances and the image is refined, it is beneficial to update $Y_\phi(\eta_1, \eta_2)$ as well. Due to the time consuming LS it entails, this is performed periodically after several iterations of the diffusion process.
2) The parameters of this algorithm are H, α, β and the number of diffusion iterations or stopping condition. H is estimated based on the degraded image itself (for more details see [41]). The other parameters have fixed values for all images. The diffusion process is completed when $H(i)$, estimated in the ith iteration, is equal to H.
3) The algorithm was implemented in MATLAB, using an explicit scheme for the tensor diffusion [28]. The empirical image was updated every 20 iterations. This process entails solving a sparse LS problem, which is more efficient than general LS. The average number of iterations required for obtaining the images (to be presented in Sec. VII, Fig. 7; the images size were 256×256 for the first two images and 128×128 for the last two) was 245 and the average processing time was 28.75 seconds. The same experiments, performed for image size of 512×512 pixels, took on the average 98.7 seconds. Processing time can be reduced by optimization of this code and by using more efficient diffusion schemes than the explicit one.
4) Given a statistical model, one can use MAP or MMSE estimation, with the fBm as a prior. However, neither MAP nor MMSE produce good results in the case

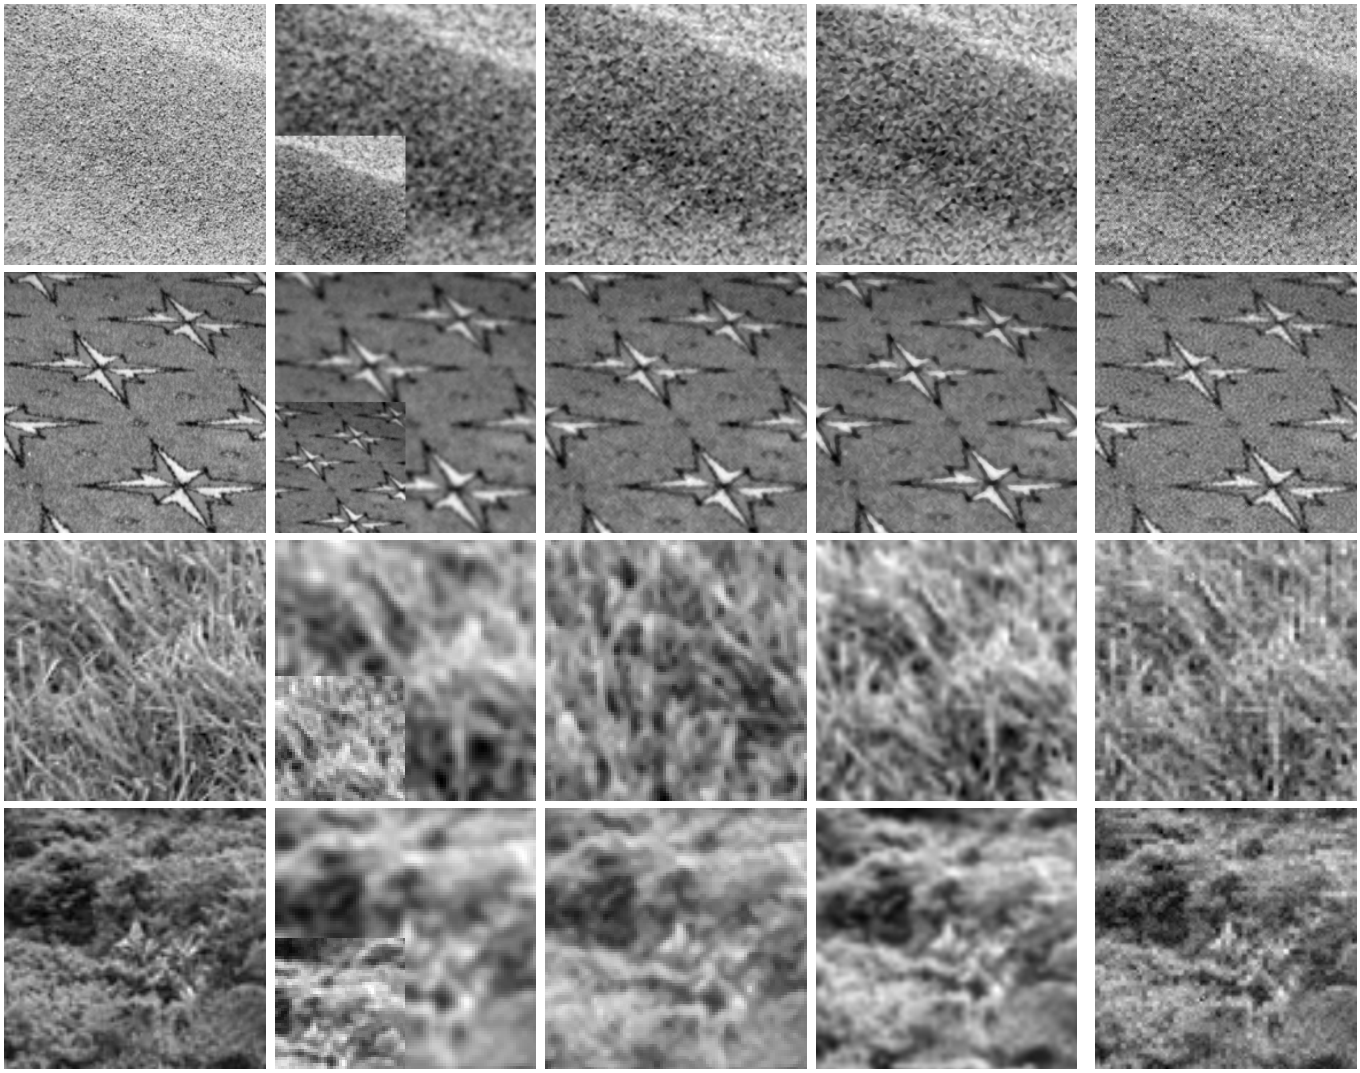

(a) (b) (c) (d) (e)

Fig. 7. Superresolution of textured images. The degraded images are obtained by $2 \downarrow 1$ subsampling in both axes after a Gaussian blur with $\sigma = 1.5$. (a) Original (ground-truth) images. (b) Bicubic interpolation with overlay of the input LR image. (c) Sparseness-based SR result [5]. (d) Example-based SR result [2]. (e) Proposed algorithm SR result.

of SR, due to the severe loss of details. This is a further indication that an $\mathcal{L}_2$-based error criteria is not suitable for textured images. In the case of denoising, however, substantial improvement was noted, which renders this type of optimization useful.

VII. Results

The proposed algorithm is implemented and used on stochastic textures. The decimation operator in Eq. (II.1) performs $\downarrow 2$ decimation in both dimensions, the blur kernel is a Gaussian with $\sigma = 1.5$ and effective support of 5×5 pixels. A small noise is added so that the BSNR is $40dB$. The contrast in all the examples is normalized after blurring. We present several types of textures from the texture "spectrum" (Fig. 7): isotropic stochastic (first row), near-regular (second row) and anisotropic stochastic (third and fourth rows). We find the latter the most challenging. The degraded versions of these images suffer from loss of high-frequency textural details, rather than edge and contour degradation. The results are compared with bicubic interpolation, example-based SR [2] and sparseness-based SR [5]. In all examples, the visual structure of the image, enhanced by the proposed algorithm, appears to resemble that of the original (ground-truth). While the high frequency spectrum is not identical to the original, the visually-appearing texture structure is restored.

A comparison of the performance of the proposed algorithm with that of a state-of-the-art method, based on self-similarity [1], is shown in Fig. 8. Similarly to the superior performance of the proposed algorithm demonstrated in the comparisons with other methods (Fig. 7), in this comparison too the proposed algorithm yields better results.[1]

We do not rely on PSNR or on other $\mathcal{L}_2$-based comparisons for the assessment of the algorithm's performance. A high PSNR value, in the range of $25dB$ and above, indicates

[1]Additional examples can be found at http://vision.technion.ac.il/demos/texture-sr/pde-based

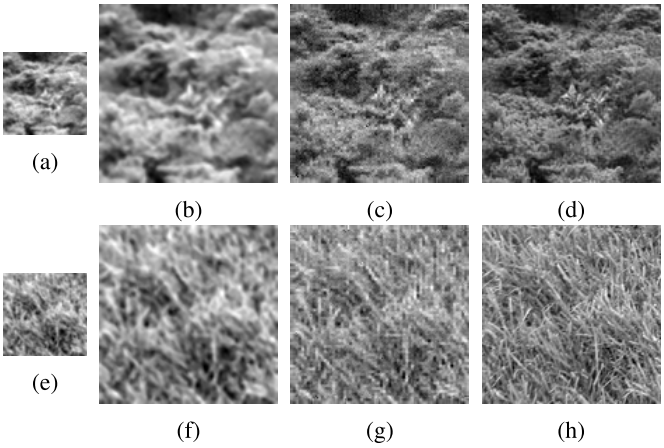

Fig. 8. The proposed algorithm and an alternate method. (a) and (e): LR image. (b) and (f): Self-similarity-based SR [1]. (c) and (g): fBm-PDE-based SR (proposed algorithm). (d) and (h) the original image (ground truth).

it is also a valid comparison method. This is the case in general images, containing sharp edges, but limited texture data. In the case of fully-textured images, the PSNR values are, however, significantly lower after degradation, and are confined to the range of 5–$20dB$. While other comparison methods, such as the structural similarity index (SSIM) [44], have been suggested, there do not exist widely-accepted criteria for performance evaluation, other than visual assessment (geometrically-based natural metrics, potentially-suitable for performance evaluation are under investigation) [40].

It is also important to note that in textured images, such as the ones considered here, one needs to look for the highest fidelity in comparison with the original image. This may be in contradiction to desired properties in other images. Sharp edges, for example, are desired in many applications, but in this case, a sharp edge in the SR image may in fact be undesired if the original image edges are not sharp.

The example-based SR results shown here depict the characteristic behaviour of the example-based SR methods; while emphasizing edges, improving on the result of the bicubic interpolation, it does not successfully restore missing fine details. This is apparent in all types of textures presented.

VIII. Discussion

The theoretical framework and algorithms presented in this study are concerned with superresolution of fully textured images, wherein the texture incorporates both stochastic and structured elements. The superresolution paradigm considered here is the so-called single-image superresolution, where only one image is available as an input. Considering first the more challenging aspect of the granularity and non-stationarity of structures often encountered in natural textures, a stochastic texture model has been developed, based on fBm. PDE-based regularization has been introduced in order to capture anisotropic texture details, and a diffusion-based single-image superresolution scheme was derived.

As is the case in similar underdetermined problems, the emphasis is on side information, inherent in the underlying image model. The results obtained in our study, encourage the use of global fBm-based model (rather than patch-based) for natural textured images, as a method for reconstruction of degraded textures.

The proposed model and concomitant algorithm are based on the empirical observation that stochastic textures are characterized by the property of self-similarity. An appropriate random process is estimated with reference to the existing low-resolution image. The initial restoration of missing details is based on an arbitrary realization of an fBm image. One may, therefore, expect different results for different evaluations. However, due to the phase matching and optimization, results for different random seeds yield almost identical results. In our current study, we attempt to remove the formal dependency on an initial arbitrary image, and obtain a model which depends on the fBm statistics.

A comparison with state-of-the-art example-based, single-image, superresolution algorithms highlights the main advantage inherent in the proposed algorithm: It reconstructs high frequency details that are otherwise missing, while the example-based algorithms emphasize edges but do not restore other, textural, missing details. On the other hand, as edges exist in some more complex textured images, possibly containing several types of textures, methods combining stochastic texture enhancement with edge enhancement are under study as well. This is performed by exploiting the Fourier phase of an image [45], [46].

Whereas the fBm has been widely used as a model of image structure, it is in fact most suitable for modelling natural textures, as this study indicates, but it is not congruous with image skeletal structures comprised of edges and contours. Further research is nonetheless called for in an attempt to expand the model to better model anisotropic textures as well, and to minimize thereby the need for regularization. Such a model may yield other enhancement algorithms suitable for a broader class of stochastic textures.

Despite of the above goal, yet to be accomplished, the proposed PDE-based regularization is interesting and important on its own merits. As such, it has to be further investigated. The empirical structure function is obtained via an ill-posed scheme, and better solutions for this problem may result in better understanding of textures and yield thereby better enhancement results.

The proposed model has been exploited for solving the SR problem. It can also be used for other image enhancement problems, such as denoising or in-painting. This is a challenge in the case of textures, due to the overlap in the frequency range with that of the noise, and due to the lack of local, small-scale, smoothness. It should be emphasized that existing denoising algorithms usually succeed in restoring edges and smooth segments, but not in the recovery of fine details. Preliminary results show that the fBm, used as a prior in MAP estimation, can effectively act as a regularizer which performs denoising on fBm-based images.

Appendix A
The ESF LS Problem

The LS problem, Eq. (VI.13), can be considered as follows. Let the filters f_d and f_{d2}, and the autocorrelation functions

$R_{\eta_1}(\eta_1,\eta_2)$, $R_{\eta_2}(\eta_1,\eta_2)$ and $R_{\eta_1,\eta_2}(\eta_1,\eta_2)$ be as presented in subsection VI-B. The matrix, D_f, in Eq. (VI.13) is derived by three vertically-stacked matrices, as follows:

$$D_f = \begin{pmatrix} D_{f,\eta_1} \\ D_{f,\eta_2} \\ D_{f2} \end{pmatrix}, \tag{A.1}$$

where D_{f,η_1} and D_{f,η_2} are the matrix representation of the filters f_d and f_d^T respectively, and D_{f2} is the matrix representation of the filter f_{d2}. The matrices perform convolution with symmetric boundaries.

In a respective manner, the vector r in Eq. (VI.13) is formed as follows:

$$r = \begin{pmatrix} r_{\eta_1} \\ r_{\eta_2} \\ r_{\eta_1,\eta_2} \end{pmatrix}, \tag{A.2}$$

where r_{η_1}, r_{η_2} and r_{η_1,η_2} are the column-stacked versions of the autocorrelation matrices $R_{\eta_1}(\eta_1,\eta_2)$, $R_{\eta_2}(\eta_1,\eta_2)$ and $R_{\eta_1,\eta_2}(\eta_1,\eta_2)$. The solution of the LS problem yields ϕ, a column-stacked vector, which is then converted back to matrix form in the size of the image.

Appendix B
Derivation of the Diffusion Equation

Let $E(X,\nabla X)$ be defined as the energy functional presented in Eq. (VI.14). The Euler-Lagrange equation for this functional is given by

$$\frac{\partial E(X,\nabla X)}{\partial X} - \sum_{i=1}^{2} \frac{\partial}{\partial \eta_i} \frac{\partial E(X,\nabla X)}{\partial X_{\eta_i}} = 0. \tag{B.1}$$

Let $E_r(X)$ denote the reaction, and $E_d(\nabla X)$ denote the diffusion term in $E(X,\nabla X)$, defined as follows:

$$\begin{aligned} E(X,\nabla X) &= \int_\Omega E_r(X) + E_d(\nabla X)dxdy, \\ E_r(X) &= (BX-Y)^2 + (\hat{X}_{HP} - H_{HP}X)^2, \\ E_d(\nabla X) &= \beta\Psi(|\nabla X + \alpha\nabla Y_\phi|^2). \end{aligned} \tag{B.2}$$

$\frac{\partial E_r(X)}{\partial X}$ is derived as follows:

$$\begin{aligned} \frac{\partial E_r(X)}{\partial X} &= 2B^T(BX-Y) - 2H_{HP}^T(\hat{X}_{HP} - H_{HP}X) \\ &= 2B^T(BX-Y) - 2H_{HP}^T\hat{X}_{HP} + 2H_{HP}^T H_{HP}X. \end{aligned}$$

Since a Gaussian filter is used, $H_{HP}^T = H_{HP}$.

The inner derivative for the diffusion term of index i in the second term of Eq. (B.1) is calculated as follows:

$$\begin{aligned} \frac{\partial E_d(\nabla X)}{\partial X_{\eta_i}} &= \beta\Psi'\left(\left|\nabla X + \alpha\nabla Y_\phi\right|^2\right) \frac{\partial}{\partial X_{\eta_i}} \left(\left|\nabla X + \alpha\nabla Y_\phi\right|^2\right) \\ &= 2\beta\Psi'\left(\left|\nabla X + \alpha\nabla Y_\phi\right|^2\right)\left(\nabla X + \alpha\nabla Y_\phi\right). \end{aligned}$$

Plugging this in the second term of Eq. (B.1) we thus obtain:

$$\begin{aligned} \sum_{i=1}^{2} \frac{\partial}{\partial \eta_i} \frac{\partial E_d(\nabla X)}{\partial X_{\eta_i}} &= \beta\nabla\cdot(\Psi'(|\nabla X + \alpha\nabla Y_\phi|^2)\nabla X) \\ &\quad + 2\alpha\beta\nabla\cdot(\Psi'(|\nabla X + \alpha\nabla Y_\phi|^2)\nabla Y_\phi). \end{aligned}$$

The rest of the derivatives are zero. Finally, substituting the derivatives of $E_r(X)$ and $E_d(\nabla X)$ we obtain the diffusion-reaction equation in Eq. (VI.16).

Acknowledgment

The comparison images in Fig. 8 were kindly provided by Dr. D. Glasner [1].

References

[1] D. Glasner, S. Bagon, and M. Irani, "Super-resolution from a single image," in *Proc. IEEE 12th Int. Conf. Comput. Vis.*, Sep. 2009, pp. 349–356.

[2] K. Kim and Y. Kwon, "Example-based learning for single-image super-resolution," in *Proc. Pattern Recognition*. Berlin, Germany: Springer-Verlag, 2008, pp. 456–465.

[3] L. C. Pickup, S. J. Roberts, and A. Zisserman, "A sampled texture prior for image super-resolution," in *Proc. Adv. Neural Inf. Process. Syst.*, 2003, pp. 1587–1594.

[4] D. Datsenko and M. Elad, "Example-based single document image super-resolution: A global map approach with outlier rejection," *Multidimensional Syst. Signal Process.*, vol. 18, no. 2–3, pp. 103–121, 2007.

[5] J. Yang, J. Wright, T. Huang, and Y. Ma, "Image super-resolution via sparse representation," *IEEE Trans. Image Process.*, vol. 19, no. 11, pp. 2861–2873, May 2010.

[6] K. I. Kim, and Y. Kwon, "Single-image super-resolution using sparse regression and natural image prior," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 6, pp. 1127–1133, Jun. 2010.

[7] B. Goldluecke and D. Cremers, "Superresolution texture maps for multiview reconstruction," in *Proc. IEEE 12th Int. Conf. Comput. Vis.*, Sep. 2009, pp. 1677–1684.

[8] Y.-W. Tai, S. Liu, M. S. Brown, and S. Lin, "Super resolution using edge prior and single image detail synthesis," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, Jun. 2010, pp. 2400–2407.

[9] C. Damkat, "Single image super-resolution using self-examples and texture synthesis," *Signal, Image Video Process.*, vol. 5, no. 3, pp. 343–352, Jan. 2011.

[10] J. Yang, J. Wright, Y. Ma, and T. Huang, "Image super-resolution as sparse representation of raw image patches," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2008, pp. 1–8.

[11] M. Welk, D. Theis, T. Brox, and J. Weickert, "PDE-based deconvolution with forward-backward diffusivities and diffusion tensors," *Scale Space PDE Methods Comput. Vis.*, pp. 585–597, 2005.

[12] Y. Gousseau and J.-M. Morel, "Are natural images of bounded variation?" *SIAM J. Math. Anal.*, vol. 33, no. 3, pp. 634–648, Jan. 2001.

[13] A. S. Carasso, "Singular integrals, image smoothness, and the recovery of texture in image deblurring," *SIAM J. Appl. Math.*, vol. 64, no. 5, pp. 1749–1774, 2004.

[14] O. Honigman and Y. Y. Zeevi, "Enhancement of textured images using complex diffusion incorporating Schrodinger's Potential," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, May 2006, pp. 633–636.

[15] M. Elad and M. A. T. Figueiredo, "On the role of sparse and redundant representations in image processing," *Proc. IEEE*, vol. 98, no. 6, pp. 972–982, Jun. 2010.

[16] W.-C. Lin, J. Hays, C. Wu, V. Kwatra, and Y. Liu, "A comparison study of four texture synthesis algorithms on regular and near-regular textures," Dept. Biol. Sci., Carnegie Mellon Univ., Pittsburgh, PA, USA, Tech. Rep. CMU-RI-TR-04-01, 2004.

[17] A. A. Efros and T. K. Leung, "Texture synthesis by non-parametric sampling," in *Proc. 7th IEEE Int. Conf. Comput. Vis.*, vol. 2. Sep. 1999, pp. 1033–1038.

[18] W. T. Freeman, T. R. Jones, and E. C. Pasztor, "Example-based super-resolution," *IEEE Comput. Graph. Appl.*, vol. 22, no. 2, pp. 56–65, Mar./Apr. 2002.

[19] G. Gilboa and S. Osher, "Nonlocal operators with applications to image processing," *Multiscale Model. Simul.*, vol. 7, no. 3, pp. 1005–1028, 2008.

[20] H.-O. Peitgen, D. Saupe, M. F. Barnsley, Y. Fisher, and M. McGuire, *The Science of Fractal Images*. New York, NY, USA: Springer-Verlag, 1988.

[21] B. B. Mandelbrot and J. W. Van Ness, "Fractional brownian motions, fractional noises and applications," *SIAM Rev.*, vol. 10, no. 4, pp. 422–437, 1968.

[22] B. Pesquet-Popescu and J. L. Vehel, "Stochastic fractal models for image processing," *IEEE Signal Process. Mag.*, vol. 19, no. 5, pp. 48–62, Sep. 2002.
[23] J. Keller, S. Chen, and R. Crownover, "Texture description and segmentation through fractal geometry," *Comput. Vis., Graph., Image Process.*, vol. 45, no. 2, pp. 150–166, 1989.
[24] D. Zoran and Y. Weiss, "From learning models of natural image patches to whole image restoration," in *Proc. Int. Conf. Comput. Vis.*, Nov. 2011, pp. 479–486.
[25] M. J. Wainwright and E. P. Simoncelli, "Scale mixtures of Gaussians and the statistics of natural images," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 12. Cambridge, MA, USA, May 1999, pp. 855–861.
[26] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli, "Image denoising using scale mixtures of Gaussians in the wavelet domain," *IEEE Trans. Image Process.*, vol. 12, no. 11, pp. 1338–1351, Jan. 2003.
[27] D. J. Heeger and J. R. Bergen, "Pyramid-based texture analysis/synthesis," in *Proc. 22nd Annu. Conf. Comput. Graph. Interact. Tech.*, vol. 3. New York, NY, USA, ACM, Aug. 1995, pp. 229–238.
[28] J. Weickert, *Anisotropic Diffusion in Image Processing*. Stuttgart, Germany: Teubner Stuttgart, 1998.
[29] I. Zachevsky and Y. Y. Zeevi, "Single-image superresolution of self-similar textures," in *Proc. IEEE Int. Conf. Image Process.*, Sep. 2013, pp. 952–956.
[30] M. Irani and S. Peleg, "Super resolution from image sequences," in *Proc. 10th Int. Conf. Pattern Recognit.*, vol. 2. Jun. 1990, pp. 115–120.
[31] M. Elad and A. Feuer, "Restoration of a single superresolution image from several blurred, noisy, and undersampled measured images," *IEEE Trans. Image Process.*, vol. 6, no. 12, pp. 1646–1658, Jan. 1997.
[32] J. Yang and T. Huang, "Image super-resolution: Historical overview and future challenges," in *Super-Resolution Imaging*. Boca Raton, FL, USA: CRC, 2010.
[33] A. Danielyan, V. Katkovnik, and K. Egiazarian, "BM3D frames and variational image deblurring," *IEEE Trans. Image Process.*, vol. 21, no. 4, pp. 1715–1728, Apr. 2012.
[34] S. Hoefer, H. Hannachi, M. Pandit, and R. Kumaresan, "Isotropic two-dimensional fractional Brownian motion and its application in ultrasonic analysis," in *Proc. 14th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, Nov. 1992, pp. 1267–1269.
[35] L. M. Kaplan and C.-C. Kuo, "An improved method for 2D self-similar image synthesis," *IEEE Trans. Image Process.*, vol. 5, no. 5, pp. 754–761, Jan. 1996.
[36] B. Pesquet-Popescu and P. Larzabal, "Synthesis of nonstationary fields with stationary increments," in *Proc. Image Process. Appl. Sixth Int. Conf.*, vol. 1. Jul. 1997, pp. 303–307.
[37] A. V. Oppenheim and J. S. Lim, "The importance of phase in signals," *Proc. IEEE*, vol. 69, no. 5, pp. 529–541, May 1981.
[38] J. Behar, M. Porat, and Y. Y. Zeevi, "Image reconstruction from localized phase," *IEEE Trans. Signal Process.*, vol. 40, no. 4, pp. 736–743, Apr. 1992.
[39] N. Skarbnik, C. Sagiv, and Y. Y. Zeevi, "Edge detection and skeletonization using quantized localized phase," in *Proc. Eur. Signal Process. Conf.*, Aug. 2009, pp. 1542–1546.
[40] Z. Wang and A. C. Bovik, "Mean squared error: Love it or leave it? A new look at Signal Fidelity Measures," *IEEE Signal Process. Mag.*, vol. 26, no. 1, pp. 98–117, Jan. 2009.
[41] I. Zachevsky and Y. Y. Zeevi, "Superresolution of self-similar textures," CCIT, Technion-Israel Inst. Technol., Tech. Rep. 838, 2013.
[42] G. Gilboa, N. Sochen, and Y. Y. Zeevi, "Forward-and-backward diffusion processes for adaptive image enhancement and denoising," *IEEE Trans. Image Process.*, vol. 11, no. 7, pp. 689–703, Jan. 2002.
[43] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 12, no. 7, pp. 629–639, Jul. 1990.
[44] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
[45] I. Zachevsky and Y. Y. Zeevi, "Combining long-range dependencies with phase information in natural stochastic texture enhancement," in *Proc. IEEE Int. Conf. Image Process.*, Paris, France, Sep. 2014.
[46] I. Zachevsky and Y. Y. Zeevi, "On the statistics of natural stochastic textures," CCIT, Technion-Israel Inst. Technol., Tech. Rep., 2014.

Ido Zachevsky is currently pursuing the Doctoral degree with the Department of Electrical Engineering, Technion - Israel Institute of Technology, where he was awarded the Ollendorff Fellowship. He served in the communication corps of the Israel Defense Forces and then received the Degree from the Technion. He was recently admitted to the departmental special program of direct track toward the Ph.D. degree.

Yehoshua Y. (Josh) Zeevi is a Professor of Electrical Engineering with the Technion - Israel Institute of Technology, where he was the Barbara and Norman Seiden Professor of Computer Sciences from 1988 to 2008, the Dean of the Faculty of Electrical Engineering from 1994 to 1999, and a member of the Board of Governors and the Technion Council at various times. He is currently the Head of the Ollendorff Minerva Center for Vision and Image Sciences. He is the co-inventor of over 30 patents and patent applications in vision and image technologies, and a co-author of over 300 publications in scientific journals and conference proceedings, and a co-editor of three books.

Dr. Zeevi received the Ph.D. degree from University of California, Berkeley. He was then a Vinton Hayes Fellow at Harvard University, where he continued to be a regular visitor for many years after joining the Technion. He was also a Visiting Professor with the Massachusetts Institute of Technology (on a joint appointment with Harvard University), Rutgers, Columbia University. He is the co-founder and Editor-in-Chief of the *Journal of Visual Communication and Image representation*.