



IRWIN AND JOAN JACOBS
CENTER FOR COMMUNICATION AND INFORMATION TECHNOLOGIES

Statistics of Natural Stochastic Textures

**Ido Zachevsky and
Yehoshua Y. Zeevi**

CCIT Report #862
June 2014

 Electronics
Computers
Communications

DEPARTMENT OF ELECTRICAL ENGINEERING
TECHNION - ISRAEL INSTITUTE OF TECHNOLOGY, HAIFA 32000, ISRAEL



Statistics of Natural Stochastic Textures*

Ido Zachevsky and Yehoshua Y. Zeevi
Technion, Israel Institute of Technology
Haifa, 32000, Israel

ido@tx.technion.ac.il, zeevi@ee.technion.ac.il

Abstract

Statistics of natural images have become an important subject of research in recent years. The highly kurtotic, non-Gaussian, statistics known to be characteristic of many natural images are exploited in various image processing tasks. We focus on Natural Stochastic Textures (NST), an important subset of natural images. In the first part of this report, we show their abundance in natural images and explore their statistical properties: Gaussianity, statistical self-similarity, stationary increments and long-range dependencies. These properties and, in particular, the Gaussianity, stand in contrast to known properties of the wider class of natural images. In the second part of this report, the importance of the Fourier phase in NST representation and processing is explored. A magnitude-phase variational flow is presented. This flow allows restoration/enhancement of images, combining statistical properties with the phase structure of a images.

1 Introduction

Statistics of natural images have been the subject of intensive studies in recent years [1–3]. With the increased use of statistical image enhancement algorithms, suitable priors play a crucial role in the enhancement and restoration of images, especially in ill-posed problems where only severely degraded images are available.

Various studies have consistently shown that natural images exhibit non-Gaussian behaviour. This has been observed by inspecting the 1D, 2D, or joint histogram of the wavelet coefficients of an image [4,5]. These histograms, evaluated on numerous images, provide a good indication of their statistical structure. Natural images exhibit highly kurtotic, non-Gaussian, behaviour, indicated by heavy tails in both 1D and 2D (joint) distributions [6]. Many models capture this behaviour successfully, such as Gaussian scale mixtures (GSM) [4,5] or generalized normal [7].

*Chapter 2 of this report is based on a paper accepted to ICASSP-2014, and chapter 3 is based on a paper submitted to ICIP-2014.

While previous studies address the entire range of natural images, we consider Natural Stochastic Textures (NST), that are abundant in natural images [8]. Many natural textures, such as sand, gravel, grass, grove and others exhibit fine details that are severely degraded by cameras' PSFs, sampling and noise. Such textures are not well represented by current models, and therefore enhancement algorithms, such as superresolution or denoising [9,10], do not perform well on NST. Particularly, some NST share similar spectra and/or statistical properties with that of noise, and are regarded as such under common image models. This problem is intensified when L_2 - or L_1 -based methods are implemented, since these reward smooth or piecewise-smooth images. This is addressed in recent studies attempting to figure out whether images are of bounded variation space [11].

NST are endowed with several important statistical properties, which differentiate them from the wider class of natural images. They are non-stationary, Gaussian, self-similar, and often exhibit long-range dependencies (LRD). As will be described later, these properties are approximately valid as some of them are contradictory, when turning to the exact mathematical definitions. Further, due to the fact that NST are very complex processes, to which no exact analytical description has been found so far, exact proof that these properties are found in NST does not exist. It will, however, be shown that this set of features is consistent with texture properties, and fits these images better than known models.

Despite of said properties, some NST cannot be completely described as realizations of random processes. This is due to the fact that many such images do contain skeleton-type structures. These structures, while being conceptually similar to edges in a cartoon-type or a general natural image, are not well described by natural image models. These properties will be validated via the phase information characteristic of an image.

In order to enhance an NST that contains edge structure as well, we propose a variational scheme, which exploits the Fourier phase of a signal. Using this scheme, statistical properties can be employed in a prior model, while the structural segment is not affected.

2 Statistics of Natural Stochastic Textures

The main properties of NST are Gaussianity and statistical self-similarity [6,9]. The former is in contrast with the well-established non-Gaussianity of structured natural images (rich in edge information) [6]. Another important statistical property of NST is the long-range dependence (LRD). While LRD is well-known in the context of stochastic processes [12,13], it has not yet been exploited to its full extent in the context of NST [6].

2.1 Evaluating natural image statistics

A common method of inspecting the statistics of natural images is the following [4,6]. An ensemble of images is analyzed in a transform domain. Wavelets are usually used, as a type of derivative filters, to capture the changes in images. Commonly used are steerable pyramidal wavelets, that allow different orientations to be analyzed, as well as easy reconstruction.

Once the transform coefficients are available, their 1D histogram, or empirical distribution, is calculated. This allows for an approximation to the ensemble's marginal statistics. 2D, or joint, histograms, are used to approximate joint statistics, where the pairs of coefficients analyzed are from different scales, orientations, or other types of coefficients, according to the chosen wavelet decomposition.

This method is not sufficient to prove that the distribution of an image takes a predetermined form. Even by taking large enough ensembles or image sizes, these are only first- and second-order statistics. Obtaining third-order statistics is a difficult task, when a limited amount of images is available. However, this method is sufficient when considering enhancement algorithms that use prior models that employ statistics up to second order. Additionally, we are able to observe expected properties that substantiate specific models even if only low-order statistics are at hand.

In our experiments, we use images from three databases: VisTex, Brodatz and McGill [14–16]. The two former databases are widely used in image processing, and the latter contains a significant amount of high-quality images, which are important for the study of NST. The steerable pyramid wavelets thus allow for different orientations as well as different scales (see [4]; other wavelets yield similar results). The images are generally not segmented and may contain several types of textures or a combination of cartoon-type and texture.

2.2 Gaussianity

Contrary to findings indicating that natural images are in general non-Gaussian, NST exhibit Gaussian statistics. In [6] and elsewhere, natural images were found to be of highly distributed kurtotic values. One of the proposed marginal distributions is the generalized normal distribution (or generalized Laplace distribution), which is characterized by the following pdf:

$$p(x) = c(\alpha, \beta) \exp(-(|x - \mu|/\alpha)^\beta), \quad (1)$$

with the normalizing constant, $c(\alpha, \beta) = \frac{\beta}{2\alpha\Gamma(1/\beta)}$. The μ and α parameters correspond to the mean and variance respectively, where the variance is defined explicitly by α . The β parameter determines the kurtosis of the distribution, whose excess kurtosis is defined by:

$$K = \frac{\Gamma(5/\beta)\Gamma(1/\beta)}{\Gamma(3/\beta)^2} - 3. \quad (2)$$

$\beta = 2$ corresponds to the normal distribution with zero excess kurtosis. The range $0 < \beta < 2$ corresponds to distributions with high kurtosis. Eq. (2) shows that the excess kurtosis is defined nonlinearly with β ; values of $\beta > 1$ correspond to much smaller kurtosis relative to $\beta < 1$.

Analyzing with this distribution assumes that the coefficients are independent. This is generally not the case, and more advanced methods exist [5]. Nevertheless, the generalized normal distribution has been successfully used as a distribution for wavelet coefficients of empirical distributions of natural images for $\beta \in [0.5, 1]$ [6, 7], which correspond to leptokurtic behaviour. We propose the normal distribution model for the class of stochastic textures, obtained for $\beta = 2$ in Eq. (1). Out of the numerous methods that determine the more suitable distribution, we use the Kullback-Leibler (KL) divergence, $D_{KL}(f_1||f_2) = \int_{\mathbb{R}} f_1(x) \log\left(\frac{f_1(x)}{f_2(x)}\right) dx$, where $f_1(x)$ is the empirical density and $f_2(x)$ is the density according to the evaluated model [17].

We first assume the image dataset fits a leptokurtic distribution, the generalized normal, and estimate its β value according to the mean of the sample kurtosis. The result is $\hat{\beta} = 0.711$. We then propose two distributions: $p_n(x)$, a normal distribution, and $p_g(x)$, a generalized normal, with $\beta = \hat{\beta}$. Equipped with the two distributions as possible models, the KL divergences for the empirical distribution and each of the models were calculated. Out of 1914 test images, 620 (32%) had lower KL divergence for the normal distribution, $p_n(x)$, indicating that the normal distribution better describes the data. The test images were obtained by dividing all images (of size greater than 256×256) in the VisTex database to 256×256 -sized images.

A second test was performed, with individual values of β , each estimated from an individual image. Inspecting all values for this parameter, 19% were above a threshold kurtosis value, which was chosen as the average between the Gaussian distribution kurtosis and the Laplace distribution kurtosis. The latter has the lowest kurtosis for known image models, with $\beta = 1$ and excess kurtosis of 3. In this case as well, we observe a significant number of images described better by the Gaussian distribution. Three individual images, along with their statistics, are displayed in Fig. 1, demonstrating the different statistics of NST in contrast with the statistics of general natural images.

2.3 Stationary increments

Stationary processes have been studied extensively in the context of theory of random processes [18]. A weak-sense stationary process is characterized by an autocorrelation function which depends only on the difference between lags:

$$E[X(t)X(s)] = E[X(t-s)X(0)] \triangleq R(\tau), \quad (3)$$

and has a constant expectation, $E[X(t)] = E[X(0)]$. In this sense, its distribution is invariant to translations. Estimation, analysis and synthesis are performed easily on

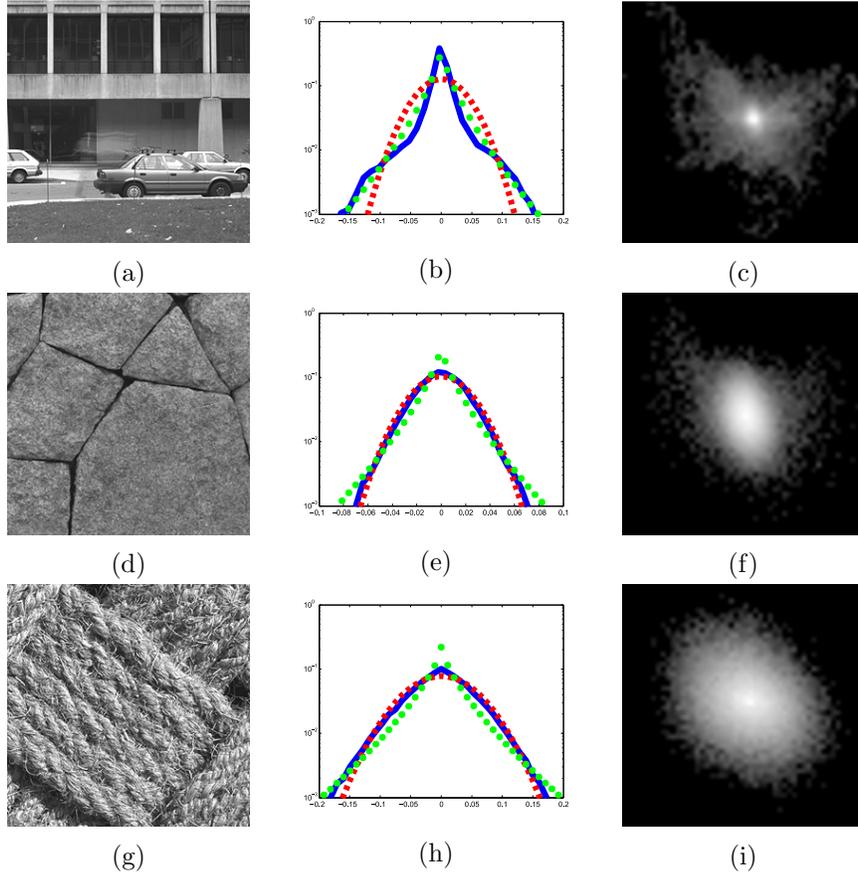


Figure 1: Examples of general image and NST, and their corresponding statistics: (i) A natural image. (b) Its marginal wavelet coefficient distribution, with the Gaussian model (red-dashed) and the leptokurtic model (dotted-green). (c) Joint (2D) distribution of the image shown in (i). The KL divergences for the Gaussian and leptokurtic models were 0.58 and 0.08 respectively, indicating non-Gaussianity. (d) NST: The KL divergences were 0.03 and 0.11 for the Gaussian and the leptokurtic models respectively. (g) Another NST: The KL divergences were 0.02 and 0.15, respectively. For NST, the Gaussian model describes the images better. (e) and (h) show the empirical, Gaussian and leptokurtic distribution fits (in blue, dashed-red and dotted-green, respectively). The Gaussianity of the NST is also apparent in the 2D histogram shapes in (f) and (i), compared with the non-Gaussian one, depicted in (c).

these processes. However, stationary processes are limited in the sense that signals with significant variability cannot be stationary; for example, a stationary signal cannot be self-similar (unless it is zero everywhere). It must also be homoskedastic, and therefore, cannot change in variability at different times (in 1D) or locations (in

2D). Therefore, considering the NST to be stationary renders it not to be fractal. We note that certain stationary processes are asymptotically self-similar.

However, even if the signal itself is not stationary, its increments can still possess this property. A process can be long-range dependent (LRD; discussed later) and self-similar, and its increments can be stationary [19]. Such a process may be endowed with interesting properties, relevant to image modelling, but yet possess some advantages of the ease of analysis and synthesis of stationary processes. Most of these properties are characteristic of weak-sense stationary processes (up to second-order statistics), and strict-sense-stationarity is not required.

The question then is how to assess whether a signal is stationary. This is a well-known problem in statistics [20, 21]. One notable method of stationarity testing (or stationarity assessment) is forming a hypothesis testing scheme, in which the null hypothesis, H_0 , is that a given process, $X(t)$, is non-stationary. If the null hypothesis is rejected, then the process is stationary.

In this framework of hypothesis testing [22], a model is assumed for the data, and it has different parameters for the two hypotheses, H_0 and H_1 . Then, a test statistic is proposed; usually asymptotically distributed in a known manner. Normal and Student's-t are notable examples. In known stationarity assessment tests, an autoregressive (AR) or autoregressive-moving average (ARMA) model is assumed. In this case, non-stationarity is synonymous with having a unit-root, and non-stationary processes are referred to as "unit-root processes" [23]. The existence of a unit-root is, therefore, the basis of these tests. A common such test is the augmented Dickey-Fuller (ADF) test [24].

The main concerns in applying these tests for NST are:

1. They are designated to time series, or 1D signals.
2. They assume an AR or ARMA model.

Due to these issues, we choose a different method to assess stationarity, based on extracting different segments of the image, and proposing a power spectral density, or autocorrelation function, to each segment [25–27]. The segments are usually overlapping or filtered via some smooth window function to create smooth transitions between different segments in the image. The spectral density can then be estimated either via a predetermined parameteric function, or via a non-parameteric estimation, fitted to the segment.

A given test statistic is then evaluated on these functions, and a hypothesis testing scheme is proposed. The basic assumption in this case is that the variance of the parameters will be low in case the data are stationary, and high otherwise. The null hypothesis, H_0 , is that the data are stationary. If the parameter set varies too much between different patches, then the null hypothesis is rejected and the data are assumed to be non-stationary.

There are several concerns with this method. First, it depends on the power spectral density function assumed under the null hypothesis. There are various

suggestions, such as the Matèrn class [28], used in [27], but certainly no single function can assess stationarity for a general 2D surface.

The distribution of the test statistic is another concern, as it should be determined under the null hypothesis in order for a threshold to be set correctly. This is done via setting the significance value of the test (usually under 10%) and inspecting the distribution of the test statistic under H_0 . When a simple model with known tests such as Student's t-test or F-test is assumed, this distribution is known. However, in more complex cases, this distribution should be estimated as well.

The third and final concern is the actual content of the images in question. If these images are not segmented, then their content will vary significantly across space and the test will obviously reject the null hypothesis.

All these concerns will be addressed hereinafter. Under the null hypothesis, the fractional Gaussian noise (fGn; [29]) model is assumed for the increments. This has several advantages: First, its autocorrelation function is known and has a simple form, depending only on a single parameter [19, 29]. Second, if an image is fBm, which has been shown to describe a wide range of textured images [10], then its first order increments are, indeed, fGn. Third, the test statistic under the null hypothesis can be determined numerically, via simulation of a large number of fBm images.

The hypothesis testing was performed as follows. The first order difference for each image was obtained, and divided into overlapping patches. Then, the fGn autocorrelation was assumed as a model:

$$\hat{R}(r) = ar + b, \quad (4)$$

where (a, b) are H -dependent parameters [19], and $r = \sqrt{x^2 + y^2}$ is the radius. This is possible since the distribution is isotropic. The parameters were estimated by regression to obtain (\hat{a}, \hat{b}) . The empirical variance of the estimates across patches was obtained, and the test statistic was chosen as follows:

$$T = \frac{\widehat{\text{var}}(\hat{b})}{\sqrt{\widehat{\text{var}}(\hat{b}_{fBm})}}. \quad (5)$$

The numerator is the variance of the estimator of \hat{b} . It is normalized by the standard deviation of the estimator for the same value of \hat{b} that was evaluated in synthesized fBm images. This normalization was chosen due to the fact that different values of b affect also the variance of the estimation in the optimal case in which the input images are indeed fGn; as b is directly related to the Hurst parameter, H , for b values indicative of high H , the variance will be larger, due to the smoother behaviour of the images (Fig. 2). Therefore, a high variance in case of a large H does not necessarily imply rejection of the null hypothesis. Although better estimation of the H parameter is attainable by using larger image and window sizes, the same conditions were applied in order to use properly for normalization with measurements of real images.

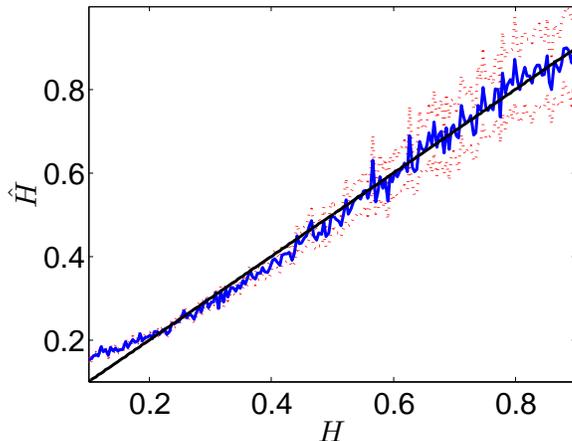


Figure 2: Estimated H value (blue) as a function of the true H value (black) in synthesized fBm. The variance of the estimator (dashed red) increases proportionally to H .

2.3.1 Obtaining the rejection threshold

The distribution of T under the null hypothesis is calculated as follows. 200 fBm images, spanning the entire range of $H \in (0, 1)$, are synthesized and the aforementioned process is performed. According to Eq. (5), T is then equal to the standard deviation of the estimator. This standard deviation is lower than the one to be obtained in real experiments, and any real experiment with the same standard deviation should retain the null hypothesis. To increase significance, the threshold is set so that the standard deviation will be less than 30% of the maximal standard deviation obtained here. This is depicted in Fig. 3; the threshold is on the lower range of standard deviations.

2.3.2 Results

The scheme was used on 1000 256×256 sized images, obtained from three databases - VisTex, Brodatz and McGill [14–16]. The patch size chosen was 32 with 10 pixel overlap. The results correspond to the main hypothesis and to visual inspection; the null hypothesis was retained by images of a single texture with fine detailed structure. Images containing several types of textures, or containing significant contours as well as other types of details, were rejected. Fig. 4 depicts several examples, and their respective increments are shown in Fig. 5.

It is important to note that we address textures of the same type as candidates for having stationary increments; images containing different types of textures are expected to be rejected. This scheme allows us to assess that the visual inspection and obtained results are consistent with formal testing of our hypothesis.

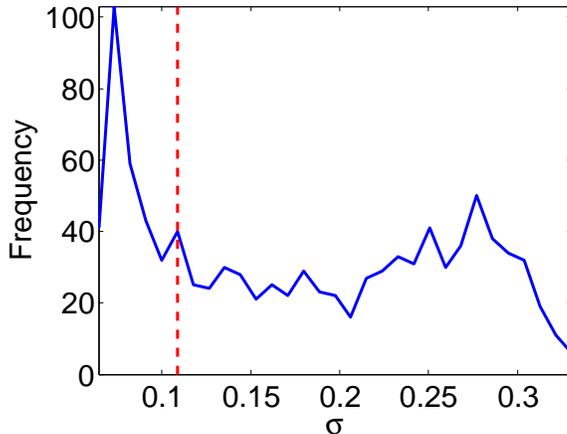


Figure 3: Histogram of standard deviations for synthesized fBm images. The dashed red line depicts the rejection threshold.

2.3.3 Remarks

The two main parameters arising from this formula are threshold and the patch sizes. As the textures transfer from having stationary increments with stochastic behaviour, to having a more regular structure, no single threshold is “best”. However, we try to set one in the reasonable range with high significance.

The second parameter, patch size, is also important. Small details may be missed due to averaging when considering too large patches, but patches too small have high variance in estimating the parameters. This presents a trade-off which needs to be resolved by assuming a typical detail size, or by employing a more complex, multi-scale-based analysis.

2.4 Self-similarity

Self-similarity is an important property of natural images. For Gaussian processes with stationary increments, it can be evaluated by performing log regression on the variance of the increments in the image domain. This is due to the fact that the only Gaussian process with stationary increments exhibiting self-similarity is the fractional Brownian motion (fBm), whose covariance function is known [29]. The log of the variance of the increments of order τ is given by:

$$y_H(\tau) = 2Hx(\tau) + b_H, \quad (6)$$

where $y_H(\tau) = \log \sigma^2(H, \tau)$ is the measurement for each τ , $x(\tau) = \log \tau$, and $b_H = \log \sigma_B^2(H)$ is a known parameter. H is the Hurst parameter, controlling the self-similarity of the process. Performing this regression on the images, found to be

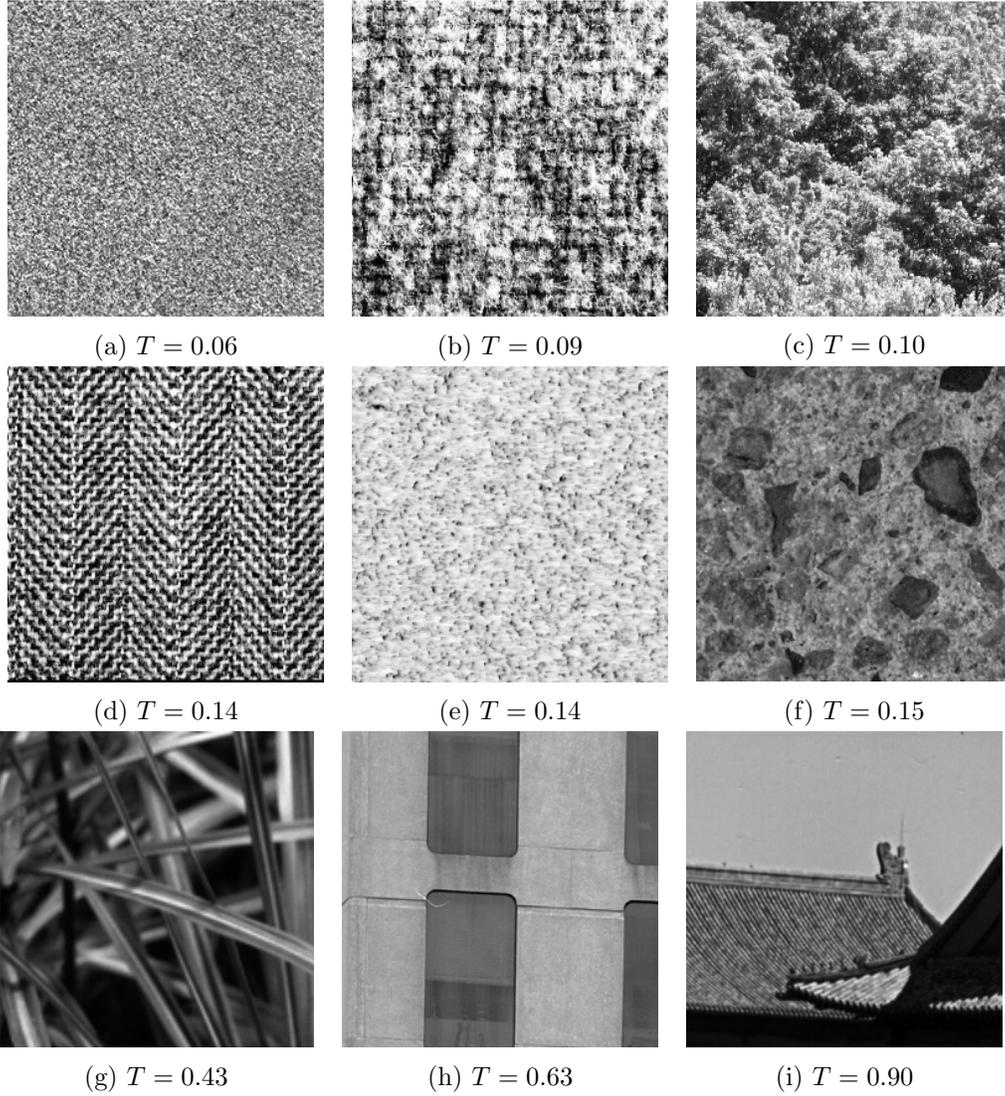


Figure 4: Example images. The threshold was $t = 0.116$. The increments for the images in the top row retained H_0 . The images in the second and third rows rejected H_0 . As higher variances indicate less stochastic behaviour, the images in the second row, with low test values, have nearly stationary increments, while the images in the third row do not. The images' respective increments are depicted in Fig. 5.

Gaussian in the analysis of Section 2.2, yielded an R^2 value with median of 0.96, which indicates a reliable result, reassuring that NST are Gaussian and self-similar. The median was chosen to minimize effects of outliers.

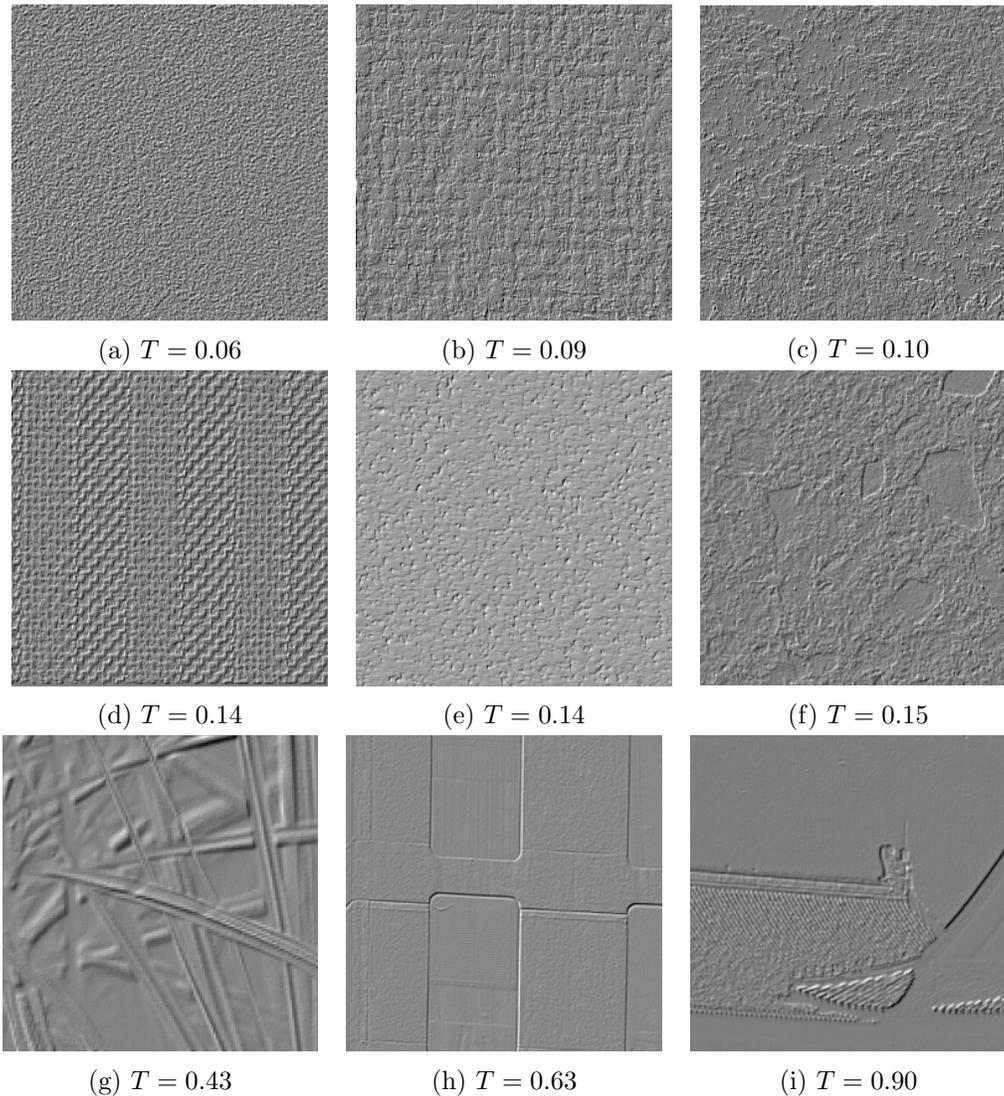


Figure 5: Increments for the example images shown in Fig. 4.

2.5 A model for NST

A possible prior model for NST, which emerges from the previous sections, is the fBm [10]. The three properties, of Gaussianity, self-similarity and stationary increments are the three sufficient ingredients for this process. FBm is a self-similar, isotropic, random process with stationary increments (see [19, 29] for an exposition of this process). While its generalized spectrum has been successfully in superresolution [10], and a direct approach was employed for denoising (Section 4.2; [9]), the development of a generalized model for a wide range of NST is in progress.

The emergence of fBm as a basis for an NST model is important for image enhancement, as NST often exhibit fine details which are easily corrupted and compromised in quality by blurring, decimation and noise in practical image acquisition processes. It is therefore imperative to have an accurate model, and treat NST in a separate domain than other types of textures or cartoon images.

Obtaining a practical image enhancement algorithm presents several preliminary issues. First, the direct application of fBm in the form of MAP or MMSE estimation requires inversion of a large matrix; a 256×256 sized image corresponds to a covariance matrix of size $2^{16} \times 2^{16}$, which is impractical. Second, for precise estimation purposes, image properties such as the Hurst parameter or the orientation (in terms of anisotropy or isotropy) may vary throughout space.

A third issue related to the estimation scheme itself. There are tasks in which MAP or MMSE estimators will not yield a good result. This is apparent in deblurring, where the noise variance is not significant, and MAP estimation yields a deconvolution scheme with little regard to the prior information. These issues are currently under investigation.

2.6 Long-range dependencies (LRD)

A weak-sense stationary random process exhibits LRD if the sum of its autocorrelation coefficients diverges [13]:

$$\sum_{k=1}^{\infty} R(k) = \infty, \quad (7)$$

where $R(k) = E[X(n)X(n+k)] = E[X(0)X(k)]$ is the autocorrelation function of $X(n)$. The interpretation of this property is high correlation between distant segments of the image. This is indeed found in textured images.

In the case of digital, finitely supported images, the LRD is determined according to the rate of decay of the partial sum of autocorrelation coefficients: If $R(m) = Cm^{-d}$ for $0 < d < 1$ and a constant C , the process exhibits LRD [12, 13].

In the case of non-stationary processes, determining LRD is, however, more complex [30]. Fortunately, the fBm model can be exploited in this case as well. One of the important properties of fBm is its stationary increments. Determining LRD is, therefore, reduced to checking the increments of the process. This is easily performed by the condition presented in Eq. (7).

Under the assumption of fBm, the image increments are stationary and constitute fGn [29], which has a known autocorrelation function and assumes the following general form (for a 1D process and $k \gg 1$):

$$R(k) = C|k|^{2H-2}, \quad (8)$$

where C is an appropriate constant and $H \in (0, 1)$ is the Hurst parameter, determining the self-similarity of the process. LRD in fBm and fGn is, therefore, obtained

for $H > 0.5$, and assessing LRD is reduced to estimating the power exponent in Eq. (8).

2.6.1 Evaluating LRD on real data

In order to evaluate LRD in images, radial autocorrelation is assumed for the increments (due to the isotropy of the fBm). A different exponent condition is obtained in 2D ([19]; see Fig. 6). The following model has been used for the increment autocorrelation:

$$R(k) = a|k|^b + c, \quad (9)$$

where (a, b, c) is the parameter vector.

A preliminary stage for evaluating LRD on textured images is to discriminate between NST and structured textures. This is performed similarly to [9], by testing the Gaussianity of an image. This stage is required as the LRD evaluation, described earlier, can only be performed on stochastic data. Otherwise, LRD may appear to be present even if the image is not at all a NST.

Regression was performed to find (a, b, c) of Eq. (9). 86% of the images yielded $R^2 > 0.9$, indicating a good fit for the model. $c \cong 0$ was found for all images. The results for the discriminating parameter, b , are depicted in Fig. 6 as a function of R^2 . Images with $R^2 \leq 0.9$ were discarded. From the rest, 28% were classified as Gaussian and 54% of these were classified as NST with LRD.

3 Magnitude and Phase in Natural Stochastic Textures

Although the stochastic structure of NST is indicative of random phase, certain NSTs incorporate phase information (Fig. 7). By optimizing the phase and the magnitude separately, we are able to optimize the image structure for a desired phase. This can be beneficial in deconvolution problems wherein the phase is easily attainable, as is the case in deblurring of images that were subjected to effects of Gaussian or motion filters, which are zero-phase.

The importance of phase in images is well-known [31,32]. The dominant features of natural images are edges and contours, which form the basis to segmentation and object recognition algorithms. Inspecting the frequency response of an image, these features are accounted for by the phase and not the magnitude. This is due to the fact that in order to create an edge, its Fourier components need to be in the same phase at the point of maximum contrast. This demand infers highly structured phase relationships.

Stochastic textures are, however, considered to be realizations of random processes [8,19]. As such, their phase is *random*. This is manifested by the fact that random deviations in the phase produce visually-similar images, and do not affect the image structure (visually speaking) in a significant manner. In images depicting

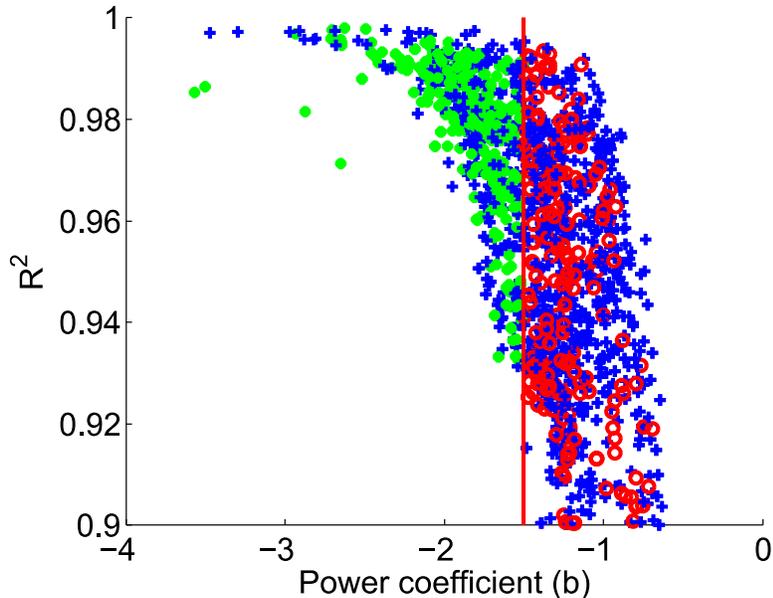


Figure 6: LRD regression results: The blue plus-signs represent non-NST images. The green dots represent non-LRD NSTs and the red circles represent NSTs with LRD. The red line indicated the boundary for LRD detection. 54% of the images satisfying Gaussianity in the database were NST with LRD.

NST, however, the same deviations severely distort the image due to smearing of edges (see Fig. 7).

Many textures, while exhibiting properties which render themselves to be NSTs, still contain some important edge structure. It is, therefore, required that modelling and enhancement be performed in a combined domain of stochastic and phase-based processing. This notion is the basis for proposing the scheme presented in Section 3.1.

3.1 Magnitude-Phase variational scheme

Consider the following image degradation model [33]:

$$Y(x, y) = (X * h)(x, y) + N(x, y), \quad (10)$$

where $X(x, y)$ is the original $M \times M$ image (ground truth), $N(x, y)$ is the noise, $h(x, y)$ is a filter and $Y(x, y)$ is the observed image (measurement). The goal of deconvolution algorithms is to obtain an estimate of $X(x, y)$, denoted by $\hat{X}(x, y)$, that minimizes the distance $d(X, \hat{X}) = \|X - \hat{X}\|_2$.

We propose an alternative, flow-based, variational scheme. This flow is initialized with $Y(x, y)$ and yields an estimate, $\hat{X}(x, y)$. Let $\tilde{X}(\omega_1, \omega_2) = \mathcal{F}\{X(x, y)\}(\omega_1, \omega_2)$

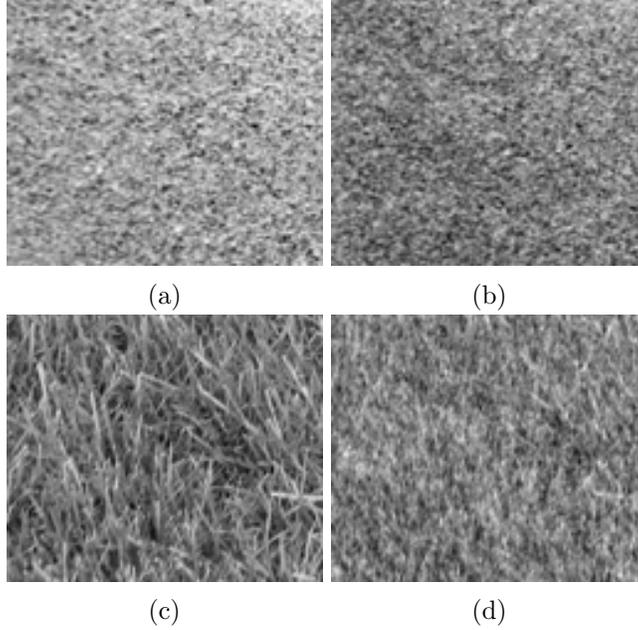


Figure 7: Importance of phase in NST. The phase of the NST displayed in (a) and (c) was distorted by Gaussian noise with $\sigma = 1$, to obtain the images of (b) and (d), respectively. In terms of visual details (other than brightness), the first image is not affected. This is indicative of a random phase. The second image is severely distorted, indicating a non-random, structured, phase.

denote the Fourier transform of $X(x, y)$. We introduce the following functional for a given Fourier transform of an image, $\tilde{I}(\omega_1, \omega_2)$:

$$E(\tilde{I}) = \iint \left(\|\tilde{I}\|^2 - |\tilde{X}_0|^2 \|^2 + \beta \|\angle \tilde{I} - \angle \tilde{X}_0\|^2 \right) dx dy, \quad (11)$$

where $|\tilde{X}_0|$ and $\angle \tilde{X}_0$ are the desired magnitude and phase respectively, to be addressed in Section 3.2. Let $F(\tilde{I})$ denote the Euler-Lagrange equation for this functional. The solution is found by gradient descent:

$$\tilde{I}_t = F(\tilde{I}), \quad \tilde{I}|_{t=0} = \tilde{Y}, \quad (12)$$

where the estimate $\hat{X}(x, y)$ is set to be $I|_{t_\infty}$ for the time of convergence, t_∞ . $\hat{X}(x, y)$ is the image with the closest magnitude and phase to $X_0(x, y)$ in the Fourier domain. This functional combines the magnitude and phase for optimization, where the magnitude is defined by $|\hat{X}|^2 = \hat{X}\hat{X}^*$, and X^* denotes complex conjugate. In the spacial domain, this is the empirical autocorrelation: $\mathcal{F}^{-1}\{\hat{X}\hat{X}^*\}(x, y) = (X * X')(x, y)$, where $X'(x, y) = X(-x, -y)$.

The advantage of implementing this functional is in its ability to optimize for a desired autocorrelation, representing stochastic properties, while the phase represents structural properties. Further, in the case of isotropic or motion blur filters, which are zero-phase, the phase of the observed image, $Y(x, y)$, is the same as that of the ground truth, $X(x, y)$, up to distortions caused by noise, as will be shown hereinafter. Therefore, for the Fourier phase, this constitutes a problem of denoising without deblurring.

Further inspecting Eq. (11), and using $\tilde{I} = A + jB$, where $A(\omega_1, \omega_2)$ and $B(\omega_1, \omega_2)$ are real-valued images and j denotes $\sqrt{-1}$, we obtain:

$$E(A, B) = \iint \left\| A^2 + B^2 - |\tilde{X}_0|^2 \right\|^2 + \beta \left\| \arctan 2 \left(\frac{B}{A} \right) - \angle \tilde{X}_0 \right\|^2 dx dy, \quad (13)$$

where $\arctan 2(\theta)$ is the signed arctangent function used for the phase of the Fourier transform. This functional is solved by using the Euler-Lagrange equations for $A(\omega_1, \omega_2)$ and $B(\omega_1, \omega_2)$. Optimization of these equations yields the following scheme:

$$\begin{aligned} A_t &= 2Af_1(A, B) + 2\beta \left(-\frac{B}{A^2 + B^2} \right) f_2(A, B) \\ B_t &= 2Af_1(A, B) + 2\beta \left(\frac{A}{A^2 + B^2} \right) f_2(A, B), \end{aligned} \quad (14)$$

where

$$f_1(A, B) \triangleq A^2 + B^2 - |\tilde{X}_0|^2 \quad (15)$$

$$f_2(A, B) \triangleq \arctan 2 \left(\frac{B}{A} \right) - \angle \tilde{X}_0. \quad (16)$$

The estimated image, $\hat{X}(x, y)$, is found by solving the inverse Fourier transform:

$$\hat{X}(x, y) = \mathcal{F}^{-1} \{ A(\omega_1, \omega_2)|_{t_\infty} + jB(\omega_1, \omega_2)|_{t_\infty} \}(x, y). \quad (17)$$

3.2 NST deblurring

For the following derivation we assume:

1. The noise variance is small relative to the signal: This is usually the case for NST deblurring, as even low variance noise is challenging. It will be demonstrated by the performance of a state-of-the-art algorithm.
2. The convolution filter is zero-phase: Many degradation effects are caused by zero-phase filters. However, the derivation holds for general filters as well, with slightly more complex derivations.

Recall Eq. (10):

$$Y(x, y) = (X * h)(x, y) + N(x, y), \quad (18)$$

which has the following Fourier transform:

$$\tilde{Y}(\omega_1, \omega_2) = \tilde{X}(\omega_1, \omega_2)\tilde{H}(\omega_1, \omega_2) + \tilde{N}(\omega_1, \omega_2). \quad (19)$$

In the noiseless case, we have $N(x, y) \equiv 0$. The magnitude yields the following equation:

$$\left| \tilde{Y}(\omega_1, \omega_2) \right| = \left| \tilde{X}(\omega_1, \omega_2) \right| \left| \tilde{H}(\omega_1, \omega_2) \right|, \quad (20)$$

and the phase yields the following equation:

$$\angle \tilde{Y}(\omega_1, \omega_2) = \angle \left\{ \tilde{X}(\omega_1, \omega_2)\tilde{H}(\omega_1, \omega_2) \right\} \quad (21)$$

$$= \angle \tilde{X}(\omega_1, \omega_2) + \angle \tilde{H}(\omega_1, \omega_2). \quad (22)$$

While this is not true in the noisy case, these equations approximately hold if we use the following Taylor expansion:

$$e^x = \sum_{n=0}^{\infty} \frac{1}{n!} x^n \Rightarrow e^{j\theta} = 1 + j\theta + O(\theta^2), \quad (23)$$

for a given angle, θ . Therefore, using $\alpha = |\alpha|e^{j\angle\alpha}$ for any $\alpha \in \mathbb{C}$, we obtain the following:

$$|\tilde{Y}(\omega_1, \omega_2)|e^{j\angle\tilde{Y}(\omega_1, \omega_2)} = |\tilde{W}(\omega_1, \omega_2)|e^{j\angle\tilde{W}(\omega_1, \omega_2)} + |\tilde{N}(\omega_1, \omega_2)|e^{j\angle\tilde{N}(\omega_1, \omega_2)}, \quad (24)$$

where $W(x, y) \triangleq (X * h)(x, y)$. Discarding the frequency variables, we obtain:

$$|\tilde{Y}| \left(1 + j\angle\tilde{Y} + O(\angle\tilde{Y}^2) \right) = |\tilde{W}| \left(1 + j\angle\tilde{W} + O(\angle\tilde{W}^2) \right) + \quad (25)$$

$$+ |\tilde{N}| \left(1 + j\angle\tilde{N} + O(\angle\tilde{N}^2) \right), \quad (26)$$

and discarding the second order terms, we obtain the following approximation:

$$|\tilde{Y}| + j|\tilde{Y}|\angle\tilde{Y} \cong |\tilde{W}| + |\tilde{W}|j\angle\tilde{W} + |\tilde{N}| + |\tilde{N}|j\angle\tilde{N}. \quad (27)$$

This equation can be separated into two equations, for the imaginary and real domains:

$$\begin{cases} |\tilde{Y}| \cong |\tilde{W}| + |\tilde{N}| \\ |\tilde{Y}|\angle\tilde{Y} \cong |\tilde{W}|\angle\tilde{W} + |\tilde{N}|\angle\tilde{N}. \end{cases} \quad (28)$$

If $N(x, y)$ is a white Gaussian noise image, then its Fourier transform will be distributed similarly. Additionally, if $h(x, y)$ is a zero-phase filter:

$$\angle\tilde{W} = \angle \left(\tilde{X} \cdot \tilde{H} \right) = \angle\tilde{X}. \quad (29)$$

Combining the set of equations in Eq. (28), we obtain the following:

$$|\tilde{X}| = |\tilde{H}|^{-1} (|\tilde{Y}| - |\tilde{N}|) \quad (30)$$

$$|\tilde{Y}| \angle \tilde{Y} = (|\tilde{Y}| - |\tilde{N}|) \angle \tilde{X} + |\tilde{N}| \angle \tilde{N} \quad (31)$$

$$\angle \tilde{X} = \frac{|\tilde{Y}| \angle \tilde{Y} - |\tilde{N}| \angle \tilde{N}}{|\tilde{Y}| - |\tilde{N}|} \quad (32)$$

$$= \frac{|\tilde{Y}|}{|\tilde{Y}| - |\tilde{N}|} \angle \tilde{Y} - \frac{|\tilde{N}|}{|\tilde{Y}| - |\tilde{N}|} \angle \tilde{N}. \quad (33)$$

If, additionally, the noise is of small variance, we obtain an estimate for the phase:

$$\angle \tilde{X}_0 = \angle \tilde{Y}, \quad (34)$$

where \tilde{X}_0 is the desired signal from Eq. (13). As an initial estimate for the magnitude, we have the following:

$$|\tilde{X}_1(\omega_1, \omega_2)| = \frac{|\tilde{Y}(\omega_1, \omega_2)| - |\tilde{N}(\omega_1, \omega_2)|}{|\tilde{H}(\omega_1, \omega_2)| + \epsilon}, \quad (35)$$

for a small $\epsilon > 0$, added for regularization in the case of $|\tilde{H}(\omega_1, \omega_2)| \ll 1$. Note that this is not \tilde{X}_0 , which will be found hereinafter by incorporating long-range dependencies to the model as well.

3.3 Combining LRD

As noted in Section 2.6, NSTs often exhibit LRDs. Even more generally, their increments conform to a simple autocorrelation function, as indicated by high R^2 values. This further substantiates the fBm model. A generalized power spectrum can be defined analytically for fBm, using the H parameter of its increment process (the fGn) [19]. Let $|\tilde{X}_2(\omega_1, \omega_2)|$ denote this spectrum, which can be used as a regularizer to compensate for the degradations of the observed autocorrelation (Eq. (35)). We obtain $|\tilde{X}_0(\omega_1, \omega_2)|$ by averaging the two estimates:

$$|\tilde{X}_0(\omega_1, \omega_2)|^2 = \alpha |\tilde{X}_1(\omega_1, \omega_2)|^2 + (1 - \alpha) |\tilde{X}_2(\omega_1, \omega_2)|^2, \quad (36)$$

where $\alpha \in (0, 1)$ controls the regularization, $|\tilde{X}_1(\omega_1, \omega_2)|^2$ was defined in Eq. (35), and the value of H for the fGn can be estimated from blurry data [34].

4 Results

4.1 Deblurring

We use the scheme presented in Eq. (14) to deblur NST. For simplicity we use zero-phase filters, which still pose a challenge for deblurring [35]. In this case,

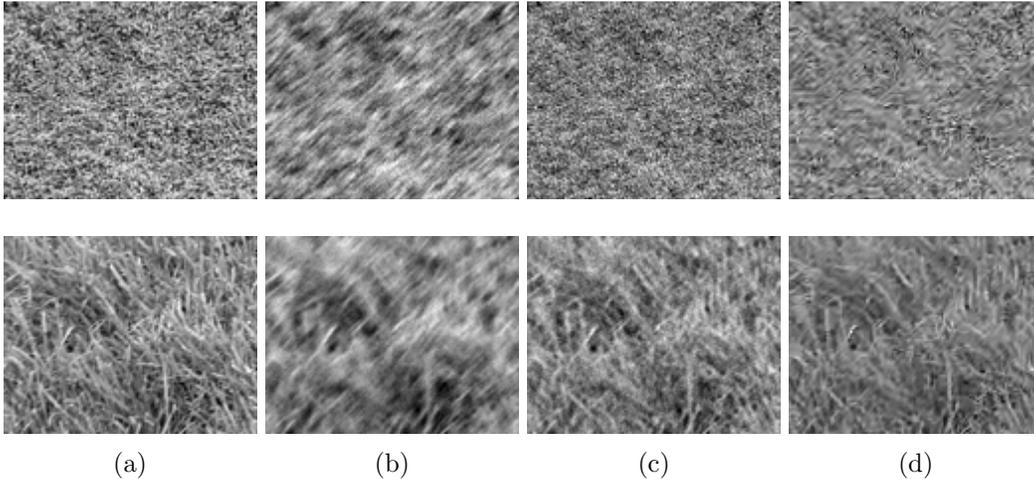


Figure 8: Deblurring examples: (a) Original images. (b) Blurred and noisy images. A 7×7 motion blur at an angular orientation of 45° was used, and additive white Gaussian noise with standard deviation of σ was added. Values of 2% and 1% of σ relative to the image were used for the top and bottom images, respectively. (c) The recovered images using the proposed method. (d) BM3D deblurring results [35].

the estimated phase is presented in Eq. (34). Small noise variance is used in the experiments, which nevertheless presents a challenge for deconvolution of textures (as the results will demonstrate).

The scheme was implemented directly, without code optimization. The step-size for the gradient descent was chosen to be small, $\Delta t = 5 \cdot 10^{-4}$, and the flow was stopped when a small enough change was obtained. The regularization parameter, ϵ , of Eq. (35) was set to 0.01. The parameter β was set so that the weight of both terms in Eq. (14) was similar in terms of the L_2 norm of the gradients. Due to the small step-size, several thousand iterations were required until convergence, but each step was performed with low computational complexity, as indicated by Eq. (14).

Further, in order to keep the optimization result in the real image domain, the following projection was performed every $T = 20$ iterations: $\tilde{X} := \mathcal{P}(\tilde{X})$, where $\mathcal{P}(X)$ projects the complex function X to the even function domain, rendering it to have a real inverse Fourier transform. Smaller values of T did not produce significant deviations in the results.

The scheme was evaluated by using several NSTs and compared with BM3D; a state-of-the-art deblurring algorithm [35]. A 7×7 -pixel motion blur with 45° angle was used, with noise standard deviation of 2% and 1% (Fig. 8). Both examples exhibit restoration of details that were missing due to blurring, comparable to restoration with the BM3D algorithm. We do not show PSNR values, as these do not correspond to visual appearance of the results and vary significantly for similar

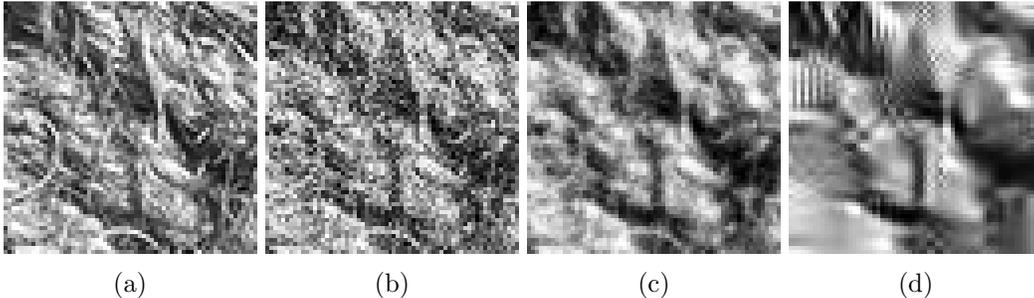


Figure 9: Example of denoising. (a): The original image. (b) Noisy and blurred image, obtained from (a) by blurring with a Gaussian blur with $\sigma = 0.5$ and contaminating with white Gaussian noise so that the blurred PSNR is 11dB. (c): MMSE fBm-based denoised image, 17.60dB PSNR. (d): BM3D denoised image, 14.73dB PSNR.

images [36]. For more deblurring results see <http://tx.technion.ac.il/~ido>.

4.2 Denoising

Denoising of NST is in particular challenging due to the fine structure of the textures. The fBm prior can be incorporated into any prior-based method. As a first attempt we chose to use optimal MMSE estimation, which yields a convenient and well-known linear estimator, due to the fact that both the image and the noise are Gaussian.

We choose 64×64 patches from images found to be Gaussian from the databases mentioned above and estimated their H parameter. Each image was blurred by a small Gaussian filter with a $\sigma = 0.5$. Noise was added so that the blurred-PSNR varied from 10dB to 20dB. The recovered PSNR was averaged for all images. Despite using a naive scheme, we see that in terms of both PSNR values and visual inspection (Fig. 9; for more examples see <http://tx.technion.ac.il/~ido>), the fBm-based denoised images yielded similar results to those of BM3D, a state-of-the-art denoising algorithm (Fig. 10) [37]. This is due to the underlying model, which recovered fine details rather than smoothing artifacts and preserving edges. It should be noted that the fBm-based denoising is by far more computationally efficient, as it is a linear estimator.

5 Discussion

NST exhibit many statistical properties which differentiate them from general natural images. These statistical properties are only approximately characteristic of the image, in the sense that it is not possible to mathematically prove an image holds one property or the other. It is, however, possible to observe that the set of properties proposed here, fits NST *better* than other proposed models. This observation

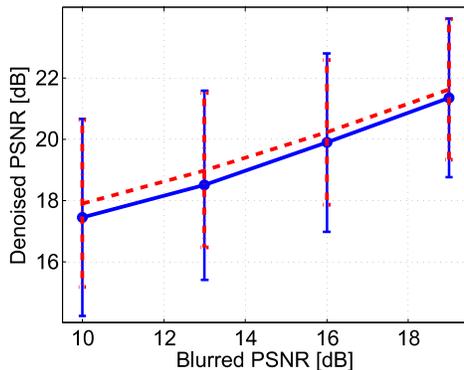


Figure 10: Denoising results. The estimated MMSE estimate, using fBm prior (dashed red line), yielded similar results to those of the BM3D denoising algorithm (blue) in terms of PSNR, when performed on NST. Each PSNR value is the mean for a test of 30 images. The error bars indicate the standard deviation.

can be performed by two means; first, testing the likelihood of one model versus another, by distribution distance measures, and second, comparing results based on either model.

The task of combining stochastic properties, such as long-range dependencies, with structural properties, remains challenging, in our continuing effort to obtain a complete model which captures these two types of properties in a closed form.

Acknowledgements

This research program has been supported by the Technion Ollendorff Minerva Center for Vision and Image Sciences, and by a special grant awarded to Y.Y.Z by the Minerva Foundation. Ido Zachevsky has been an Ollendorff Fellow.

References

- [1] A. B. Lee, D. Mumford, and J. Huang, “Occlusion Models for Natural Images: A Statistical Study of a Scale-Invariant Dead Leaves Model,” *IJCV*, vol. 41, pp. 35–59, 2001.
- [2] P. Moulin and J. Liu, “Analysis of Multiresolution Image Denoising Schemes Using Generalized Gaussian and Complexity Priors,” *Inf. Theory, IEEE Trans.*, vol. 45, no. 3, pp. 909–919, 1999.
- [3] D. Mumford and B. Gidas, “Stochastic models for generic images,” *Q. Appl. Math.*, vol. LIX, no. 1, pp. 85–111, 2001.

- [4] M. J. Wainwright and E. P. Simoncelli, “Scale Mixtures of Gaussians and the Statistics of Natural Images,” in *Adv. Neural Inf. Process. Syst.*, vol. 12. MIT Press, May 1999, pp. 855–861.
- [5] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli, “Image denoising using scale mixtures of Gaussians in the wavelet domain,” *Image Process. IEEE Trans.*, vol. 12, no. 11, pp. 1338–1351, Jan. 2003. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/18244692>
- [6] A. Srivastava, A. B. Lee, and E. P. Simoncelli, “On Advances in Statistical Modeling of Natural Images,” *J. Math. Imaging Vis.*, vol. 18, pp. 17–33, 2003.
- [7] E. P. Simoncelli and E. H. Adelson, “Noise removal via Bayesian wavelet coring,” in *Image Process. Int. Conf.*, 1996, pp. 379–382.
- [8] 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,” Carnegie Mellon University, Tech. Rep., 2004.
- [9] I. Zachevsky and Y. Y. Zeevi, “On the statistics of natural stochastic textures and their application in image processing,” in *ICASSP2014 (Accepted)*, Florence, Italy, 2014.
- [10] —, “Single-image superresolution of self-similar textures,” *IEEE Trans. Image Process. (Accepted)*.
- [11] 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. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/S0036141000371150>
- [12] G. Samorodnitsky, “Long memory and self-similar processes,” pp. 1–12, 2005.
- [13] —, “Long Range Dependence,” *Found. Trends Stoch. Syst.*, vol. 1, no. 3, pp. 163–257, 2006. [Online]. Available: <http://www.nowpublishers.com/product.aspx?product=STO&doi=0900000004>
- [14] R. Pickard, C. Grasznyk, S. Mann, J. Wachman, L. Pickard, and L. Campbell, “Vistex database,” *Media Lab., MIT, Cambridge, Massachusetts*, 1995.
- [15] P. Brodatz, *Textures: a photographic album for artists and designers*. Dover New York, 1966, vol. 66.
- [16] A. Olmos and F. A. A. Kingdom, “A biologically inspired algorithm for the recovery of shading and reflectance images,” *Perception*, vol. 33, pp. 1463—1473, 2004.
- [17] S. Kullback, *Information theory and statistics*. Courier Dover Publications, 1997.

- [18] A. M. Yaglom, *An introduction to the theory of stationary random functions*. Courier Dover Publications, 2004.
- [19] B. Pesquet-Popescu and J. L. Vehele, “Stochastic fractal models for image processing,” *IEEE Signal Process. Mag.*, vol. 19, no. 5, pp. 48–62, Sep. 2002. [Online]. Available: <http://ieeexplore.ieee.org/ielx5/79/22084/01028352.pdf?tp=&arnumber=1028352&isnumber=22084>
- [20] M. B. Priestley and T. Subba Rao, “A Test for Non-stationarity of Time-series,” no. 1, pp. 140–149, 1968.
- [21] J. J. Dolado, J. Gonzalo, and L. Mayoral, “A Fractional Dickey-Fuller Test.”
- [22] L. Wasserman, *All of statistics: a concise course in statistical inference*. Springer, 2004.
- [23] P. Perron, “The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis,” *Econometrica*, vol. 57, no. 6, pp. 1361–1401, 1989.
- [24] J. C. Ooms, “Bachelors Thesis: Assessing Stationarity in Univariate Time Series,” 2012.
- [25] D. Higdon, J. Swall, and J. Kern, “Non-Stationary Spatial Modeling,” no. 1992, 1998.
- [26] G. Nason, “A test for second-order stationarity and approximate confidence intervals for localized autocovariances for locally stationary time series .” pp. 1–36, 2012.
- [27] M. Fuentes, “A formal test for nonstationarity of spatial stochastic processes,” *J. Multivar. Anal.*, vol. 96, no. 1, pp. 30–54, Sep. 2005. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0047259X04001757>
- [28] B. Matérn, *Spatial variation*. Springer-Verlag New York, 1986, vol. 36.
- [29] 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.
- [30] A. Ayache, S. Cohen, and J. Vehele, “The covariance structure of multifractional Brownian motion, with application to long range dependence,” *2000 IEEE Int. Conf. Acoust. Speech, Signal Process. Proc.*, vol. 6, pp. 3810–3813, 2000. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=860233>
- [31] A. V. Oppenheim and J. S. Lim, “Importance of Phase in Signals,” *Proc. IEEE*, vol. 69, no. 5, pp. 529–541, 1981.

- [32] J. Behar, M. Porat, and Y. Y. Zeevi, “The importance of localized phase in vision and image representation,” *SPIE Vis. Commun. Image Process.*, vol. 1001, pp. 61–68, 1988. [Online]. Available: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=1254476>
- [33] M. R. Banham and A. I. Katsaggelos, “Digital Image Restoration,” *IEEE Signal Process. Mag.*, vol. 14, no. March, pp. 24—41, 1997.
- [34] I. Zachevsky and Y. Y. Zeevi, “Superresolution of self-similar textures,” *CCIT Rep. 838. EE Pub, Tech. Isr. Inst. Technol.*, vol. 1795, 2013.
- [35] 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. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/22128008>
- [36] Z. Wang and A. C. Bovik, “Mean squared error: Love It or Leave It?” *IEEE Signal Process. Mag.*, pp. 98–117, Jan. 2009.
- [37] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, “Image denoising with block-matching and 3D filtering,” *Proc. SPIE*, vol. 6064, pp. 606 414–606 414–12, Feb. 2006. [Online]. Available: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=728740>