Scene dependence of the non-gaussian scaling properties of natural images

Ángel Nevado*, Antonio Turiel† and Néstor Parga‡
Departamento de Física Teórica
Universidad Autónoma de Madrid
Canto Blanco, 28049 Madrid, Spain

Abstract

We report results on the scaling properties of changes in contrast of natural images in different visual environments. This study confirms the existence, in a vast class of images, of a multiplicative process relating the variations in contrast seen at two different scales, as was found in [4]. But it also shows that the scaling exponents are not universal: Even if most images follow the same type of statistics, they do it with different values of the distribution parameters. Motivated by these results, we also present the analysis of a generative model of images that reproduces those properties and that has the correct power spectrum. Possible implications for visual processing are also discussed.

Submitted to *NETWORK*

*E-mail: anevado@delta.ft.uam.es
†E-mail: amturiel@delta.ft.uam.es
‡To whom correspondence should be addressed. E-mail: parga@delta.ft.uam.es

1 Introduction

As it has been suggested a long time ago [1], the early stages of the visual system must have evolved by adaptation to the statistics of the external stimuli. During this process the neurons in the visual pathway have developed their receptive fields in such a way that information about visual scenes is represented internally in an efficient way. The large amount of redundancy present in the external world is, at least in part, eliminated from the internal sensory code.

In order to achieve this quasi-optimal representation the visual system must have learnt the regularities present in the environment where the organism lived. If a given conjunction of some elementary features tend to appear together, a cell responding optimally to the combination of features is rather likely to exist.

Carrying out this program requires, as a first step, to perform an analysis of the statistical properties of the environment. In the case of visual scenes, a systematic study of this matter has began rather recently [2, 3, 4, 5, 6, 7]. Although the relevance of the second order statistics has been pointed out some time ago [8, 9], and a gaussian distribution for images has very often been used [10, 11, 12, 13] to make predictions on the receptive fields of cells in V1 and in previous stages of the visual pathway, there are many reasons to believe that this statistics leaves aside a vast number of qualitatively important properties. An indication of this is that once the image is decorrelated (i.e. the correlations between pairs of pixels are eliminated) the scene can still be recognized [14], something that is mainly due to the fact that the borders of the objects are still present.

As it was emphasized in [3, 4, 6], a better understanding of the statistics of images should be achieved before making predictions on the visual system. In fact, an analysis of the properties of changes in contrast in natural scenes revealed the existence of multiscaling properties: images do not have uniform scale properties, but they can be decomposed in sets of pixels such that only those pixels in a given set have similar scale properties. Interestingly enough, these properties can be explained [4] by means of a simple model, that obtains the statistics of changes in contrast at a scale r in terms of an independent multiplicative process applied to the changes occurring a a larger scale L.

The multiplicative process is a log-Poisson distribution. The events it generates represent sharp changes, or modulations, of the contrast gradient. It contains two parameters: the average number of events per unit of scale, $\bar{s}$, and the strength of each of these elementary changes, β. For the ensembles considered in [4, 6] these parameters took the values $\bar{s} \approx 1.$ and $\beta \approx 0.5.$ [1]

The set of images used for those studies was rather uniform, they were generally forest scenes, and their contrast distribution was also similar from image to image. This fact made the statistical analysis very stable, and the existence of the multiplicative relation between different scales was clear. It left however open the question of how much dependent those properties are on the chosen image ensemble. Their possible relevance for the development of the early visual system would be greatly

[1] Similar processes also occur in the physics of turbulent flows [15, 16, 17, 18].

diminished if they only held in a restricted set of images. On the other hand, if the same statistical properties of changes in contrast could be found in a large set of visual scenes, their importance would be more clear.

In this paper we address the issue of robustness of the properties of edges and softer textures analyzed in [4]. In order to establish their validity we use a larger dataset (more than a hundred times larger than the one used in our first analysis) and more heterogeneous images, containing many different environments.

Our result is both surprising and encouraging. The statistical analysis can be performed on an image by image basis. When this is done, one finds that most of them (375 out of 400) exhibit the scale properties found in the smaller dataset. Besides, the same log-Poisson model can be applied to explain the data. That is, for most images, contrast changes at different scales are related by the same type of multiplicative process. Nevertheless, the model parameter β depends on the image.

We have also carried out a similar statistical analysis for ensembles of images, where each ensemble was characterized by the type of scenes. Again, each of the ensembles analyzed presents the same multiscaling properties, although with different values of the parameter β. Finally, we have done the study over the whole set of images, and have again found that the same model is able to explain the observed quantities.

Given these results, our conjecture is that the visual system has adapted, during a long period of time, to the existence of these multiscale properties and, in particular, that its architecture has captured the existence of the multiplicative process. On the other hand, since different visual scenes differ in the value of the parameter β, the adaptation to this should occur more rapidly, and could be implemented, for instance, by gain control mechanisms.

Stochastic generative models of images can be useful to understand the role played by the statistical properties of images in the structure of receptive fields. These models become very simple if it is assumed that only the second order statistics is relevant; in this case the contrast distribution is just a gaussian with the correct correlation function between pairs of pixels. But the definition of a model able to describe correctly the statistics of changes in contrast requires a different approach. Here we have studied a model that reproduces the properties exhibited by the natural images analyzed in the first part of this work. As we will see, it also possesses the correct power spectrum. This generative model will also allow us to see more clearly the role of the log-Poisson parameters on the geometrical structure of the images.

The layout of this paper is as follows. The next section contains a brief theoretical background about the multiplicative log-Poisson process. The image datasets used in this work are described in Sec. 3, as well as the criteria that have been used to select them. The results about the statistics of changes in contrast and the existence of the multiplicative process in the various datasets are presented in Sec. 4. Sec. 5 is devoted to the presentation and analysis of a generative model of images that follows the non-gaussian statistical properties discussed in the previous sections. The last section is devoted to discussions, where the relevance of this work for the visual system is briefly discussed.

2 Multiscaling properties of images and the multiplicative process

In this section we give the basic theoretical background necessary for our numerical analysis. Since we are interested in quantifying changes of the contrast $C(\mathbf{x})$ at the point $\mathbf{x}$ and at a given scale r, it is natural to consider the following quantities:

$$\epsilon_{h,r}(\mathbf{x}) = \frac{1}{r} \int_{x_1}^{x_1+r} \left(\frac{\partial C(\mathbf{x}')}{\partial x_1'} \right)^2 \Bigg|_{\mathbf{x}'=\{x_1',x_2\}} dx_1' \tag{1}$$

and

$$\epsilon_{v,r}(\mathbf{x}) = \frac{1}{r} \int_{x_2}^{x_2+r} \left(\frac{\partial c(\mathbf{x}')}{\partial x_2'} \right)^2 \Bigg|_{\mathbf{x}'=\{x_1,x_2'\}} dx_2' \tag{2}$$

The variables $\epsilon_{h,r}(\mathbf{x})$ and $\epsilon_{v,r}(\mathbf{x})$ are defined at the position $\mathbf{x}$ and at a scale r. While the first takes contributions from edges transverse to a horizontal segment of size r, the second accumulates changes in contrast transverse to a vertical segment of the same size.

In [4] these variables have been interpreted as the *local linear edge variance* along the corresponding direction and at the scale r. Both of them give a measure of how far the changes in contrast are from being uniformly distributed in a segment, either horizontal or vertical, of size r. Although in this work we will discriminate between the vertical and horizontal directions, we mention that the behaviour of contrast changes at a scale r can also be studied by means of a bi-dimensional integral [6]:

$$\epsilon_r(\mathbf{x}) \ = \ \frac{1}{r^2} \int_{x_1-\frac{r}{2}}^{x_1+\frac{r}{2}} dx_1' \int_{x_2-\frac{r}{2}}^{x_2+\frac{r}{2}} dx_2' \quad | \nabla C(\mathbf{x}') | \quad . \tag{3}$$

Similarly to the other two variables, $\epsilon_r(\vec{x})$ quantifies the deviation from uniform of the distribution of $|\nabla C|$.

The question now arises of how the statistics of $\epsilon_{j,l}(\mathbf{x_0})$ (where $j = h, v$) at the scale l is related to the distribution of the variables at a larger scale L. This was discussed in [4, 6] for the marginal distribution of $\epsilon_{j,r}$. The answer is simple: they are related by a multiplicative process. Denoting by α_{lL} its associated stochastic variable, this means that

$$\epsilon_{j,l} \ \doteq \ \alpha_{lL} \, \epsilon_{j,L} \tag{4}$$

The random variable α_{lL} is independent of the ϵ's. The symbol $\doteq$ indicates that the equality holds in the distributional sense, that is, both sides of the equality have the same distribution.

The process for arbitrary changes in scale was derived in [4] where it was shown that the random variable α_{lL} follows a log-Poisson process [16]. This process can be justified as follows. As the scale is gradually reduced from L to l some changes in contrast are left outside the segment of size L. But some of these changes are special in that they give rise to a singular behaviour of the derivatives of the contrast, and in turn to a discontinuity in the ϵ's, which acquire a factor β each time that this occurs. $(0 \leq \beta < 1)$.

We will refer to this effect by saying that a modulation has occurred. We will assume that:

- modulations are independent events. The average number of them contained in the change of scale $(\ln L - \ln l))$ is denoted by s,
- the image ensemble is translational invariant,
- the multiplicative process is scale invariant.

Under a finite change from L to l, there will be n modulations with probability p_n that, given the first assumption, is Poisson:

$$p_n = \frac{s^n}{n!} e^{-s} \tag{5}$$

The value of α_{lL} after n modulations is proportional to β^n: $\alpha_{lL} = \beta^n M(l, L)$. The fact that the proportionality constant is not one can be understood by noticing that if no modulations occur (i.e., $n = 0$), $\epsilon_{j,l}$ keeps as much singular structure in the interval l as that present in $\epsilon_{j,L}$, which is distributed in the larger interval L.

The dependence of $M(l, L)$ with the scales l and L can be obtained by invoking a scale invariant multiplicative process. In this case, it can only depend on the ratio l/L, that is $M(l, L) = M(\frac{l}{L})$. A standard argument shows that it has to be a power law: if the change from L to l is done through an intermediate scale r, then because of the multiplicative character of the process we have $M(\frac{l}{L}) = M(\frac{l}{r})M(\frac{r}{L})$, which can only be satisfied if $M(\frac{l}{L}) = \left(\frac{l}{L}\right)^{-\Delta}$. The exponent Δ is another parameter of the model.

Taking all these arguments into account, the distribution of $\ln \alpha_{lL}$ can be expressed as:

$$\rho_{\alpha_{lL}}(\ln \alpha_{lL}) = e^{-s} \sum_{n=0}^{\infty} \frac{s^n}{n!} \delta \left(\ln \alpha_{lL} - n \ln \beta - \Delta \ln \left(\frac{L}{l} \right) \right) \quad . \tag{6}$$

Up to now we have considered s, β and Δ as independent parameters. However translational invariance fixes one of them. In fact, from the definition of the $\epsilon_{j,r}$'s one easily checks that its average over a translational invariant ensemble of images does not depend on r. In turn this implies that

$$\langle \alpha_{rL} \rangle = 1 \quad , \tag{7}$$

where $\langle \ldots \rangle$ indicates the average over an ensemble of images. But then, imposing this condition over the value of this average obtained from the distribution in eq. (6) one has:

$$s = \frac{\Delta}{1 - \beta} \ln \frac{L}{r} \quad , \tag{8}$$

from where the average number of modulations per unit of change in scale is $\bar{s} = \frac{\Delta}{1-\beta}$.

The existence of a multiplicative process has direct consequences on the scaling properties of the moments of $\epsilon_{j,r}$. Let us denote these moments by $\langle \epsilon_{j,r}^p \rangle$. If the log-Poisson model holds, then from eq. (6) it is easy to check that the moments have a property called Self-Similarity (SS),

$$\langle \epsilon_{j,r}^p \rangle \;=\; A_p \, r^{\tau_p} \quad , \tag{9}$$

where the τ_p's are the SS exponents. Notice that this also implies that any moment can be expressed as a power of the moment of order q. Choosing $q = 2$ this means that

$$\langle \epsilon_{j,r}^p \rangle \;=\; A(p,2) \left[\langle \epsilon_r^2 \rangle \right]^{\rho(p,2)} \quad , \tag{10}$$

This relation could hold even when SS is not true. It is called Extended Self-Similarity [19]. The $\rho(p, 2)$'s are the ESS exponents and the $A(p, 2)$'s are geometrical factors. The computation of the moments of α_{lL} yields:

$$\rho(p,2) \;=\; \frac{p}{1-\beta} \;-\; \frac{1-\beta^p}{(1-\beta)^2} \quad . \tag{11}$$

Let us notice that although the model has two parameters, β and $\bar{s}$, the ESS exponents $\rho(p, 2)$ depend only on the modulation parameter β. There is a simple relation between τ_p and $\rho(p, 2)$:

$$\tau_p \;=\; -\bar{s}(1-\beta)^2 \, \rho(p,2) \quad , \tag{12}$$

Conversely, it can be proven that if SS and ESS hold and the exponents $\rho(p, 2)$ verify eq. (11), then ϵ_r can be described in terms of a multiplicative process (eq. (4)) of the log-Poisson type (eq. (6)). It is then enough to check eq. (11), from where the existence of a log-Poisson process is derived. In this work we have also preferred to compute the ESS exponents $\rho(p, 2)$ instead of the SS exponents τ_p because the first property usually holds for a larger range of values of r, what makes the numerical computation more stable.

Given the experimental ESS exponents $\rho(p, 2)$, it is easy to obtain the value of β by performing a least squares regression on eq. (11). To define the log-Poisson process completely one should still estimate the parameter $\bar{s}$, which can be done for instance by means of eq. (12). Unfortunately, $\bar{s}$ contains a dependence in $1/(1-\beta)^2$ that makes the propagation of errors important for β close to one, even if the error bars of τ_p and $\rho(p, 2)$ are small. We will then limit our analysis to the parameter β (Notice that this is enough to check the existence of a log-Poisson multiplicative process).

3 Choice of environments

In the present study we used a dataset of 200 images, having 1536×1024 pixels and 16 bits in luminance depth. These images have been kindly provided to us by Hans van Hateren [22] and were selected from his dataset of about 4200 images. To compare with previous studies we notice that the images used in [4] contained 45 natural scenes with a size of 256×256 pixels and a luminance depth of 13 bits. This means that our present statistical database has 130 times as many bits as the previous one.

The dataset has been considered either complete or divided into four different subsets with 50 images in each of them. These four ensembles were selected to meet the following requirements:

- Ensemble A:

  It contains natural scenes of trees and woods. The images do not have artificial objects or open skies. The woods are dense and neither shadows nor direct light rays are allowed.

- Ensemble B:

  It contains natural scenes. The rest of conditions imposed in the definition of ensemble A have not been required here.

- Ensemble C:

  Images containing both artificial and natural objects are included. We chose images that in a visual scan appear to have predominantly horizontal structures, i.e, the horizontal edges seem to be longer than the vertical ones. We are interested in testing whether this implies a difference between the horizontal and vertical statistics.

- Ensemble D:

  Out of the whole database of 4212 images in [22] we picked up 50 randomly. No image belonging to any of the other three ensembles was included here. The intention is to have a dataset as varied as possible.

Each ensemble is used to study the vertical and the horizontal statistics (that is the marginal statistics of $\epsilon_{h,r}(\mathbf{x})$ and $\epsilon_{v,r}(\mathbf{x})$, respectively), what gives a total of eight different elements.

4 Results

The main purpose of this section is to establish whether the statistical properties of SS and ESS, eqs.(9, 10), are present in the ensembles considered, and whether the ESS exponents can be predicted by the log-Poisson model, eq.(6). This will be done taking three different categories of ensembles. In Sec. 4.2 we consider the four ensembles defined in Sec. 3 for the horizontal and vertical directions. Next, in Sec. 4.3 we take as ensembles the individual images themselves, again for both directions. Finally in Sec. 4.4 we regard the 200 images as a single ensemble and we accumulate the horizontal and vertical statistics in the same distribution. Notice that the moments are self-averaging quantities but the ESS exponents, and therefore the parameter β of the probability distribution of eq. (6) are not, so we do not expect the parameter β of an ensemble to be the average of the β's of its images.

4.1 Methods

Here we describe the methods used to check the presence of SS, ESS and the log-Poisson model for given p-moments of ϵ_r. The statistical moments will be computed over the three different categories we are going to analyze, i.e., for each of the eight orientational ensembles, for single images and for the whole dataset. The items of a category giving rise to the different moments $\langle \epsilon_r^p \rangle$ will be referred to as elements. (Hence, we have eight elements in the first category, 400 in the second and just one in the third).

The existence of SS will be checked by verifying that $\ln < \epsilon_{j,r}^p >$ is linear in $\ln r$ for each element (the symbol $< . >$ denotes the ensemble average over the images belonging to that element). In order to present the results for different elements in a single graph and to evaluate the standard deviation of the data, we have proceeded as follows. Since each element has its own scaling exponents the curves have to be normalized; we first define,

$$S_p(x) = \ln < \epsilon_{j,r}^p > \quad , \tag{13}$$

where $x = \ln r$. When $S_p(x)$ vs x is fitted by a least-squares regression it yields the fitting line: $\hat{S}_p(x) = a + bx$. We now perform a linear transformation $\hat{S}_p \to \hat{S}'_p$ such that $\hat{S}'_p = x$ for each ensemble. The same transformation is applied to $S_p(x)$ to give

$$S_p(x) \to {S_p}'(x) = \frac{S_p(x) - a}{b} \quad . \tag{14}$$

The average over different ${S_p}'$ of a given category can now be carried out, yielding the quantities $\bar{S}'_p$. ESS is analyzed in a similar way; now we define

$$E_p(y) = \ln < \epsilon_{j,r}^p > \tag{15}$$

where $y = \ln < \epsilon_{j,r}^2 >$. The fitting line for each element is expressed as $\hat{E}_p(y) = c + d\, y$. However now not only the fit parameters c and d depend on the element considered, but also the variable y, for fixed r, varies from element to element. Therefore, before transforming all the fitting lines onto a common function, one has to normalize y properly. This is done as:

$$y \to y' = \frac{y - y_{min}}{y_{max} - y_{min}} \tag{16}$$

where y_{max} and y_{min} are respectively the maximum and minimum values of y for a given element. This maps y into the interval [0:1]. As before we apply to E_p the transformation that yields $\hat{E}_p \to \hat{E}'_p = y'$:

$$E_p \to {E'}_p(y') = \frac{E_p(y) - (a + b y_{min})}{b(y_{max} - y_{min})} \quad . \tag{17}$$

The experimental ESS exponents $\rho(p, 2)$ for each element in the category are fitted by a least-squares regression. An important result is that each element follows the log-Poisson model although the optimal value of β varies from element to element.

To represent the data from all the elements together we define a normalized $\rho'(p, 2)$ as:

$$\rho'(p,2) = \frac{\rho(p,2)}{\rho^{\beta}(p,2)}\rho^{\beta_0}(p,2) \tag{18}$$

where $\rho(p, 2)$ denotes the computed ESS exponents [2] obtained from each element, ρ^{β} denotes the exponents predicted by the log-Poisson model (11) using the value of β that best fits the data and $\rho^{\beta_0}(p, 2)$ are the exponents predicted by the log-Poisson model with a certain β_0 conventionally set equal to 0.5. (i.e., half of its maximum value).

4.2 Ensembles

We study the properties of the four ensembles in the two directions, which gives us eight elements altogether. Figs. 1Ia and 1Ib show the mean $\bar{S}'_p(x)$ and the standard deviation $\sigma_{S_p'(x)}$ for $p = 2$ and 7. It is observed that SS holds rather well, although there appear small deviations for the lowest and highest values of x.

The values of $E'_p(y')$ are plotted together for the eight elements as a function of y'. This time we cannot average over the values of $E'_p(y')$, since y' varies from element to element. Figs. 1IIa and 1IIb for $p = 2$ and 7, respectively, show that ESS holds better than SS.

Fig. 2 represents the mean value $\bar{\rho}'(p)$ over the eight ensembles as a function of p and its standard deviation. It is observed that the log-Poisson model is very well satisfied. In Table 1, the observed values of β are given. There are important numerical differences among them. As we will see in the next section this also occurs when the images are analyzed individually.

Ensemble and orientation	β
A horizontal	0.17
A vertical	0.022
B horizontal	0.27
B vertical	0.14
C horizontal	0.10
C vertical	0.10
D horizontal	0.28
D vertical	0.030

Table 1: Values of β for the four ensembles A-D and two orientations

[2] These exponents are obtained with a linear regression of the ESS curves for each moment. This regression is performed in the range $r = [8 : 25]$, where the ESS property is best realized. The function $\rho(p, 2)$ is computed up to $p = 7$. For higher order moments the properties of SS and ESS begin to appear less clearly. The reason for this is that the main contribution to these moments comes from the tail of the distribution of $\epsilon_{r,l}$, where the sampling is worse.

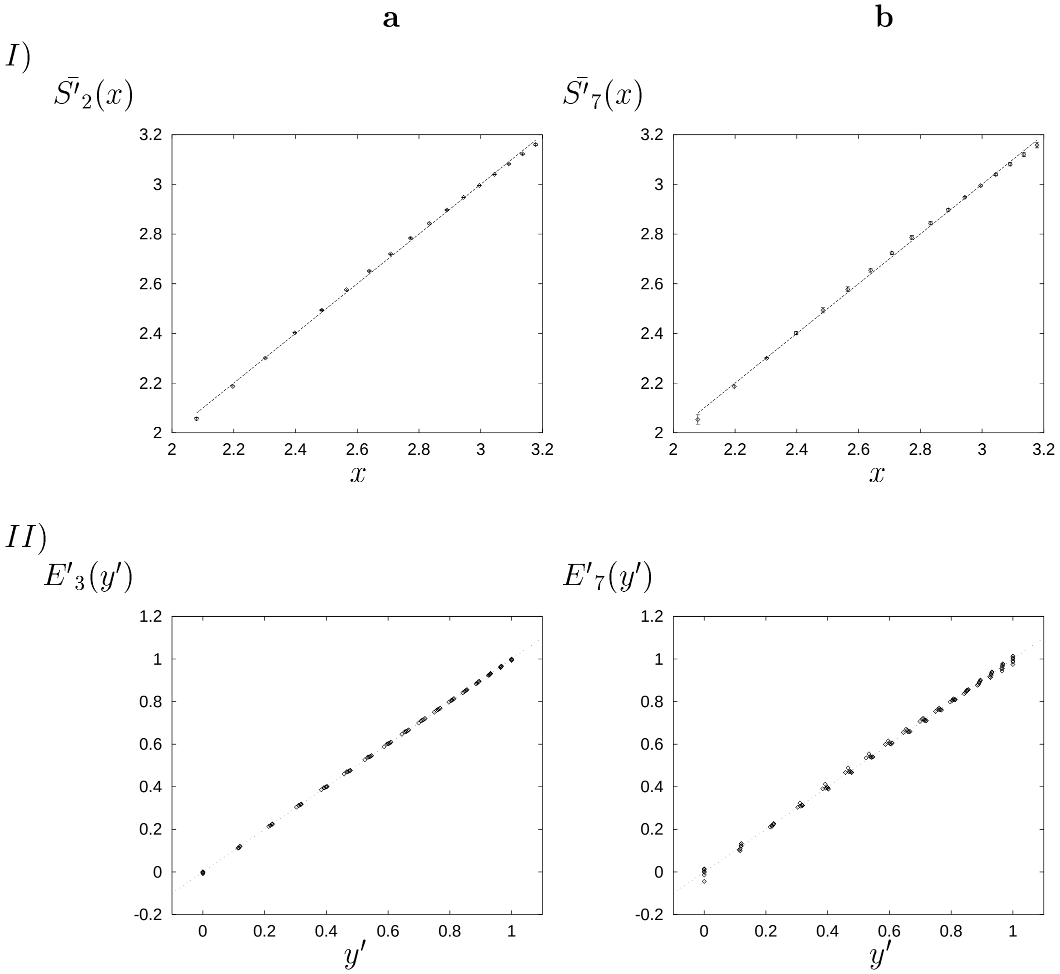

Figure 1: Test of SS and ESS for the eight ensembles. I) SS for the moments of order 2 (a) and 7 (b). The dashed line is $f(x) = x$. The diamonds represent the mean and the (small) bars the standard deviation over the eight elements. II) ESS for the moments of order 3 (a) and 7 (b). The dotted line is $f(y') = y'$

4.3 Single images

Again, our aim is to establish whether the statistical properties of SS and ESS are present but this time in single images, and whether the ESS exponents can be predicted by a log-Poisson model for each of them. Now the moments $< \epsilon^p_{r,l} >$ are obtained by averaging over the pixels of each individual image. Note that the resulting exponents $\rho(p, 2)$ of the images of a certain ensemble do not necessarily average to the ensemble exponents. This is checked for the set of 200 images for both the horizontal and vertical direction. From now on we shall refer to this set as containing 400 "images", implying that we consider the statistics from the horizontal and the vertical directions.

The results for SS, ESS and the ESS exponents are presented applying the procedures described in Sec. 4.1. Fig. 3) shows the data for SS and ESS while Fig. 4

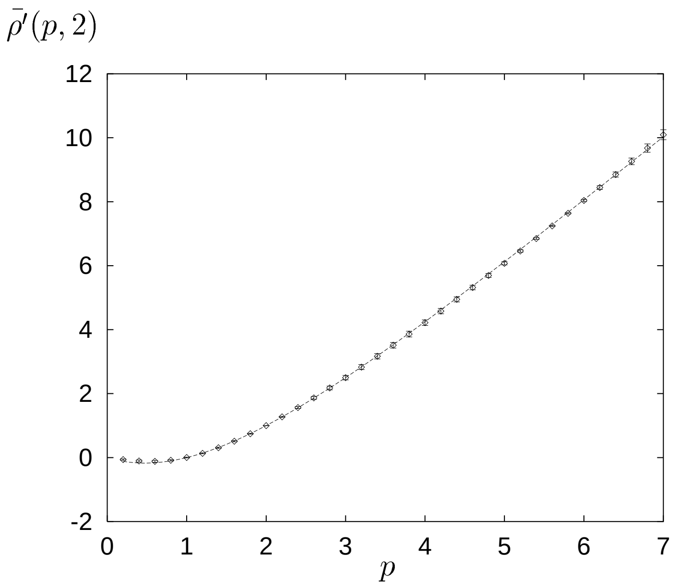

Figure 2: Test of the log-Poisson model for the eight ensembles. The diamonds represent the mean of $\rho'(p, 2)$ over the eight ensembles as a function of p and the error bars are twice the standard deviation. These error bars are small and difficult to appreciate. The dashed line represents $\rho^{0.5}(p, 2)$, the log-Poisson $\rho(p, 2)$ with $\beta = 0.5$.

includes the verification that the log-Poisson model is satisfied and the corresponding distribution of values of β.

The means $\bar{S'}_p$ and $\bar{\rho}(p, 2)$ are not over the 400 images but over the 375 images that best satisfy the log-Poisson model. Of the other 25 images (6.2%), 6 of them (1.5%), have been discarded because their β exceeds 1, the maximum admissible one. The rest have been discarded because the log-Poisson model fits their $\rho(p, 2)$ significantly less accurately than that of the 375 selected images. This can be seen in Figs. 4Ia and 4Ib The 375 images with smallest $\chi^2 = \sum_{p=1}^{n} \left(\rho(p, 2) - \rho^{\beta}(p, 2)\right)^2$ are a homogeneous set in the sense that discarding a few of the worst of them produces a small improvement of the overall fit represented in Fig. 4Ib. On the contrary, discarding the 25 worst images of the complete set produces a significant improvement of the overall fit.

The discarded images have some features in common. Almost all of them either are obscure and therefore lack definition or are scenes with little structure: sea, rivers, ground with little more than grass, buildings with simple repeating patterns.

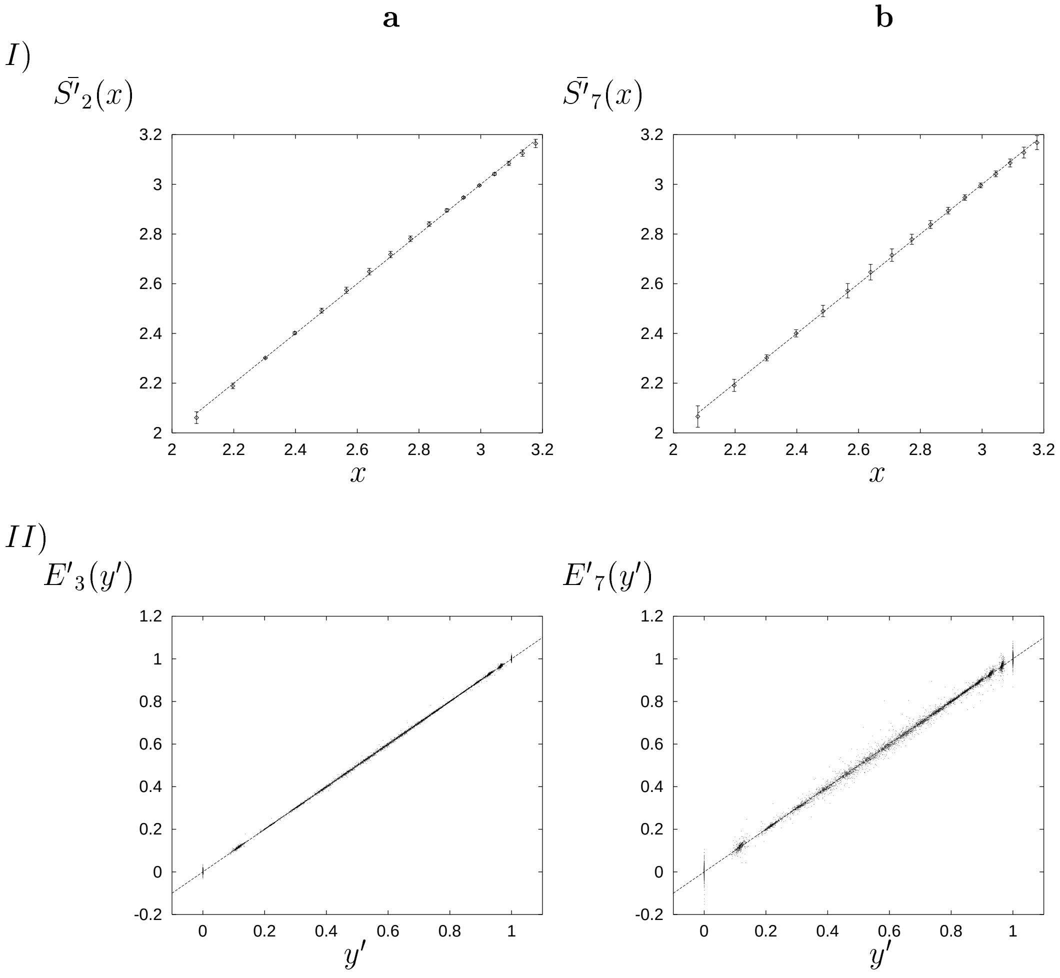

Figure 3: Test of SS and ESS for the 375 out of 400 chosen images. I) SS for the moments of order 2 (a) and 7 (b). The dashed line is $f(x) = x$. II) ESS for the moments of order 3 (a) and 7 (b). The dashed line is $f(y') = y'$

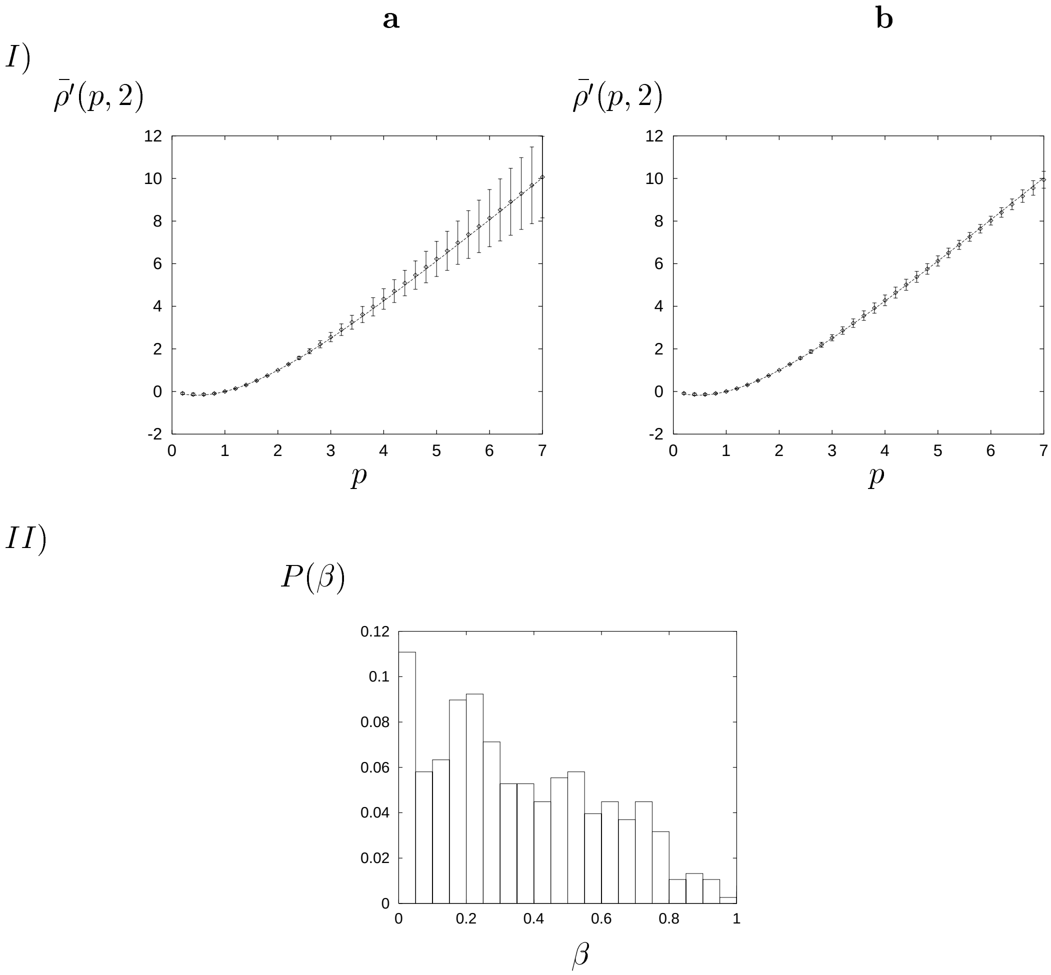

Figure 4: Test of the log-Poisson model for the 400 images. Ia) The mean of the $\rho'(p,2)$'s over the 400 images. Ib) The mean of the $\rho(p,2)$'s for the 375 best images. Again, for a) and b) the diamonds represent the mean of $\rho'(p,2)$ and the error bars are twice the standard deviation. The dashed line represents $\rho^{0.5}(p,2)$. II) Histogram of values of β for the 375 selected images.

4.4 Whole dataset

We now perform the same analysis on the global dataset containing the whole set of 200 images. These are again treated as 400 "images" by averaging together the moments for the horizontal and vertical variables ϵ. The results are presented in terms of the same quantities (eqs. (14) to (18)) used in Sec. 4.1 to allow for direct comparison. The value of β obtained from this set is 0.17. Fig. 5 shows the SS test for the moments of order two and seven and the ESS test for the moments of order three and seven. The fit of the ESS exponents with a log-Poisson curve is presented in Fig. 6 (Notice that, following the convention defined in eq. (18) the curve is referred to $\beta_0 = 0.5$). It can be seen that also for the whole dataset all these properties are well confirmed. Again ESS holds better than SS.

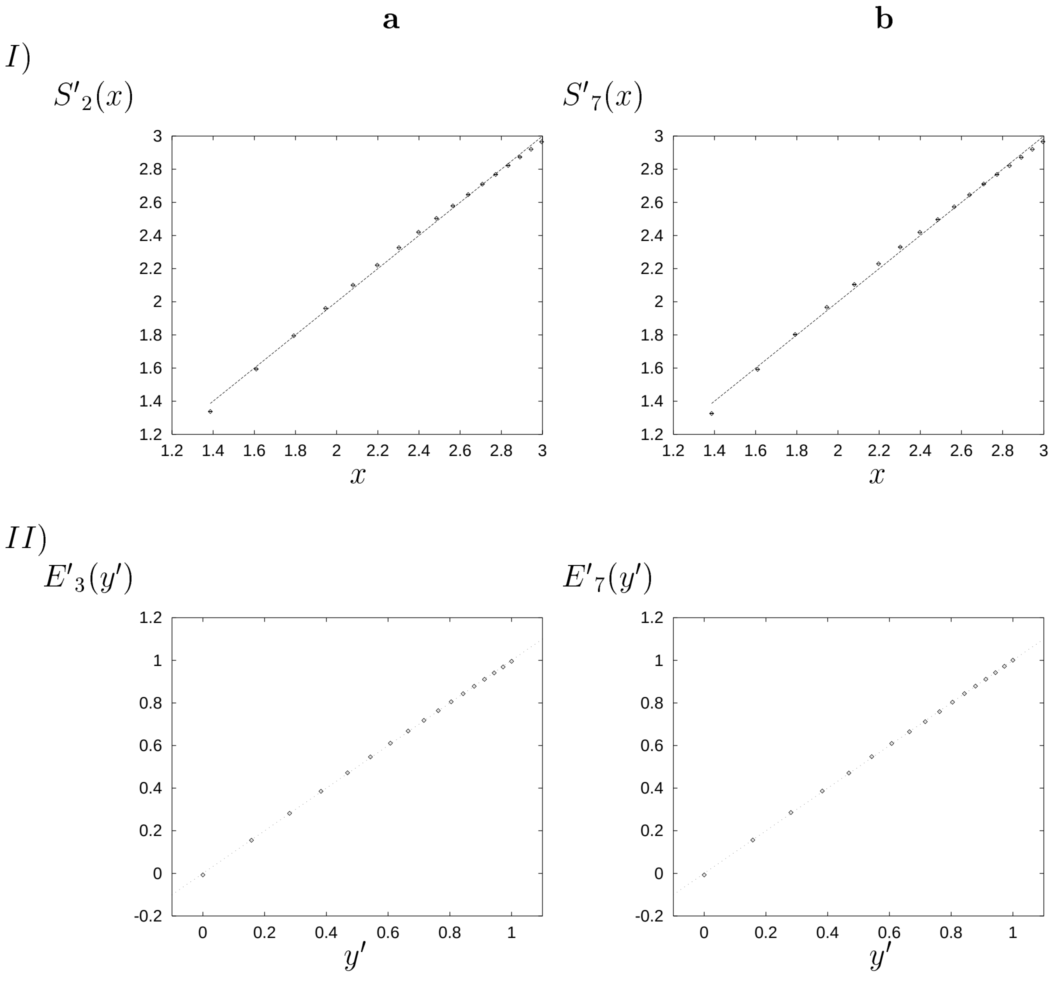

Figure 5: Test of SS and ESS for the whole image dataset (400 images). I) SS test for the moments of order two (a) and seven (b). The dashed line represents $f(x) = x$. II) ESS test for the moments of order three (a) and seven (b). The dotted line represents f(y')=y'.

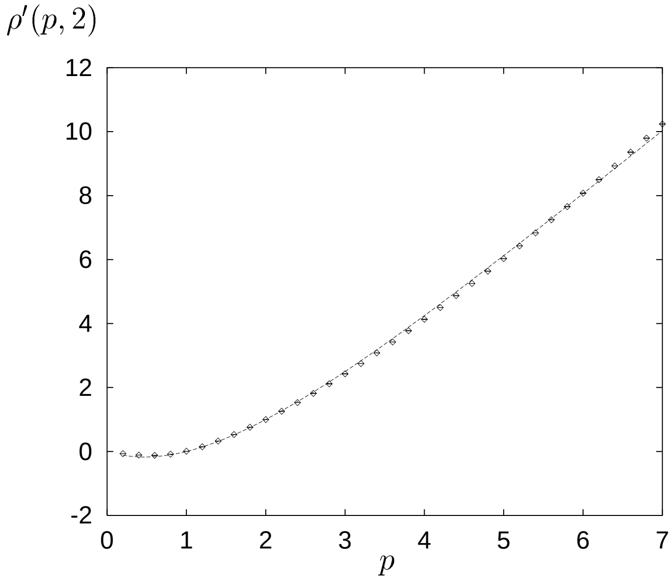

Figure 6: Test of the log-Poisson model for the whole dataset.

4.5 Analysis and geometrical interpretation

We can conclude from our experimental study of the three categories of ensembles that SS, ESS and the log-Poisson model are robust features of natural images. It is very remarkable that ensembles following log-Poisson laws of different parameters β can be combined together giving again a log-Poisson model. The multiplicative process neither gives a simple interpretation of this experimental fact by itself, nor relates the parameter β with geometrical properties of the images in a straightforward way. In this section we will attempt to provide some intuition about both questions.

Table 1 characterizes each of the eight elements of this category according to their values of β. First let us notice that in most of them there is an anisotropy between the observed values of β along the horizontal and vertical directions. For instance, in the ensemble A the horizontal β is much larger than the vertical one. A difference between the two β's is still present in the ensemble B, although it is smaller. This could be expected because the images contained in B are more isotropic than those in the ensemble A. Both are dominated by vertical statistics (they are made up of images of woods), which somehow is reflected in the different values observed for β along each direction. Naively, one would expect a difference of the opposite sign in the ensemble C, which was defined to be dominated by horizontal statistics. However, this particular ensemble reveals to be the most isotropic of the four. This apparent contradiction is explained by the fact that even when changes in contrast are normally smooth, they present very localized fluctuations that contribute to ϵ_r. In fact, a closer inspection of the images in ensemble C would show that although the contrast is more correlated along the horizontal direction, the edges appear very homogeneously, isotropically distributed. Finally, the ensemble D behaves similarly to A, except for the numerical values of β.

To explain the observed numerical values of β some intuition about the meaning of this parameter is needed. In all the following examples it will be assumed that SS is verified. Let us first consider a very simple image consisting of non-structured objects, that is, uniformly illuminated surfaces (A model that uses such type of images was proposed in [23]). Since the derivatives of the contrast are non-zero only at the set of points that contain the contours, all the contributions to the ϵ_r's come from the edges (each of them contributes an amount $1/r$). Given that almost every point in the image lies inside a particular object, ϵ_r computed at an arbitrary point at a distance l of the closest edge must vanish for $r < l$. That is, at $r = l$ a modulation of factor $\beta = 0$ occurs. This leads to a log-Poisson model with $\beta = 0$. The other extreme is realized in images for which the contrast does vary across the scene but in a smooth way. Since all the points give similar contributions the observed values of ϵ_r are order one. This leads to a log-Poisson process with $\beta = 1$.

It is remarkable that the case $\beta = 1$ can be understood in terms of images in which every point belongs to an "edge" (of a very smooth type, of course) and so the density of edges is maximal. On the other hand, the case $\beta = 0$ corresponds to very sharp but isolated transitions, for which the density of edges is (in the continuum) zero. Images with $0 < \beta < 1$ could be thought of as scenes consisting of objects giving rise to edges and edge-like transitions. If these edges were isolated structures, their contributions to ϵ_r would be $1/r$ what means that $\beta = 0$. To avoid this, those edges are necessarily not isolated along the spatial direction considered, with a particular density which is related to β by an increasing function (the smallest and the largest possible densities would result in the limiting values $\beta = 0$ and $\beta = 1$, respectively). Then, heuristically one can conclude that β is a measure of the density of the edges along the direction associated to ϵ_r. This discussion provides a geometrical interpretation (in terms of edges) of the parameter β, which was originally defined as the strength of a modulation in the multiplicative process (see Sec. 2).

Given this interpretation of β, it is expected that in images of the woods the horizontal β (which takes contributions from vertical edges) would be larger than the vertical one (which depends on the horizontal edges). This statement is in agreement with what is observed in the ensembles A and B. Besides, as ensemble B contains supplementary structure (i.e., blobs of light, shadows, clouds in the sky, etc) with respect to ensemble A, it is expected that both the horizontal and vertical β's increase as a result of the increasing spatial densities of edges (see Table 1). Finally, the very small values of the vertical β's in the ensembles A and D can be considered as a hint of the isolation of the horizontal edges. Ensemble A contains scenes taken in dense wood, in which horizontal branches are rather scarce and sufficiently apart from each other to be considered as isolated. Ensemble D mixes scenes very similar to those in A with town scenes (containing lots of almost flat surfaces) and with pictures of fields of tulips (having plenty of vertical edges and very scarce horizontal ones). This mixture gives a greater relative amount of edges in the horizontal direction and almost isolation along the vertical direction.

Given that single images can also be described by the log-Poisson model it is natural to think that the log-Poisson model is the proper description for any single scene. Again the only parameter of the model, β, is a measure of the density of edges and edge-like structures in the image. The fact that grouping together the statistics of different single images also results in a log-Poisson model sugests that this description is a highly compatible, flexible one. In a sense, it seems that each image could be considered as a small part of a very large scene formed by patching together all the images of a given ensemble, and the type of statistical description is unaffected by moving across images. This "global image" (the considered ensemble) is characterized by a parameter β which is related to its density of edges (which is the average of the densities of the set of images in the ensemble).

5 Generative model of images

The fact that the SS exponents τ_p do not behave linearly with p implies that scale-invariance of natural images is rather subtle. It is related to the fact that for a given image not all the points transform in the same way under scale transformations [6]. The log-Poisson model (see Sec. 2), which explains the non-linear behaviour of τ_p, also predicts the existence of a multifractal structure: the image pixels can be decomposed in fractal components such that points in the same component have the same scaling properties.

Generative models of natural images based on their scale properties should take this important property into account. The intuitive model proposed in [23] puts the emphasis in reproducing the correct power spectrum, but it does it in such a way that all the points have the same scale properties. This is because all the changes in contrast are sharp, and equally singular, what gives rise to a single fractal component. An evaluation of the SS exponents τ_p would show that these are linear in p.

Here we address the issue of how to define a simple stochastic model for image generation which takes into account the properties discussed in the previous sections and that has the correct power spectrum. The model described below is based on a wavelet expansion of the contrast $C(\vec{x})$ itself. The multifractal structure is produced through a convenient stochastic process for the wavelet coefficients.

The model was proposed in [24] to study synthetic turbulence. We will express the contrast $C(\vec{x})$ of the image in terms of a dyadic wavelet set [25] $\Psi_{j,\vec{k}}(\vec{x}) \equiv \Psi(2^j\vec{x}-\vec{k})$, where the scale r is given by $r = 2^{-j}$ $(j \in \mathcal{Z})$, and the space is sampled at the points $\vec{x}_0 = 2^{-j}\vec{k}$, with $\vec{k} \equiv (k_1, k_2)$ (k_1 and k_2 are integer numbers taking values from 0 to $2^j - 1$).

$$C(\vec{x}) \;=\; \sum_{j=0}^{N} \sum_{k_1,k_2=0}^{2^j-1} \gamma_{j,\vec{k}} \Psi_{j,\vec{k}}(\vec{x}) \quad . \tag{19}$$

where N is the number of scales considered; thus, the number of pixels of the image is $2^N \times 2^N$ (each change in scale is of size 2). Synthetic images will be generated with $N = 8$, which gives rise to images with 256×256 pixels. The mother wavelet $\Psi(\vec{x})$ will be taken as the Laplacian of the gaussian. The wavelet coefficients $\gamma_{j,\vec{k}}$ are then generated following a hierarchical procedure which starts at the scale with the poorest resolution ($j=0$), which is described by the mother wavelet. Its coefficient, $\gamma_{0\vec{0}}$, will appear just as a normalization of the contrast level. It is then used to generate the coefficients $\gamma_{1,\vec{k}}$ ($k_1, k_2 = 0, 1$) of the wavelets located at the four points used to sample the next scale ($j = 1$). They are of the form $\gamma_{1,\vec{k}} = \alpha_{1\vec{k}}\gamma_{0,\vec{0}}$, where the $\alpha_{1\vec{k}}$'s are independent, identically distributed random variables which follow a given probability distribution (to be given below). This operation is repeated recursively at every new scale. At the scale $j-1$ each $\gamma_{j-1,\vec{k}}$ generates, by multiplication by the corresponding independent α's, four new coefficients at the scale j:

$$\gamma_{j\vec{k}} = \frac{1}{2}\,\alpha_{j\vec{k}}\gamma_{j-1,[\frac{\vec{k}}{2}]} \tag{20}$$

where $[\frac{\vec{k}}{2}]$ means the vector with components equal to the integer part, rounding down, of those of $\vec{k}$ divided by 2.

The projections of such $C(\vec{x})$ over appropriate wavelets of size r possess the property of SS. Let us consider first the dual wavelet $\tilde{\Psi}$ (that is, $\langle\tilde{\Psi}_{j\vec{k}}|\Psi_{j'\vec{k}'}\rangle = 2^{-2j}\delta_{jj'}\delta_{\vec{k}\vec{k}'}$). Defining the wavelet projection $T^r_{\tilde{\Psi}}C(\vec{x})$ as:

$$T^r_{\tilde{\Psi}}C(\vec{x}) = \frac{1}{r^2}\int d\vec{y}\;\; \tilde{\Psi}(\frac{\vec{x}-\vec{y}}{r})C(\vec{y}) \tag{21}$$

It follows from eq. (19) that:

$$\langle|T^r_{\tilde{\Psi}}C|^p\rangle = A^C_p\, r^{\tau^C_p} \tag{22}$$

where

$$\tau^C_p = p - \log_2 \overline{\alpha^p} \;, \tag{23}$$

where $\overline{\alpha^p}$ is the moment of order p of the distribution of α. Then, there exists at least the wavelet $\tilde{\Psi}$ for which the wavelet projections of C have SS, eq. (22) (in fact, all the wavelets vanishing the same number of moments define wavelet projections of C with SS of the same exponents τ^C_p). As shown in [6], this implies that also $\epsilon_{\vec{r}}$ verifies SS (i.e. it follows eq. (9)) and there is a simple relation between the SS exponents τ^C_p and τ_p given by:

$$\tau_p = \tau^C_p - p \tag{24}$$

and using eq. (23) we obtain $\tau_p = -\log_2 \overline{\alpha^p}$. It follows that in order to obtain SS exponents τ_p according to the log-Poisson model the random the $\alpha_{j\vec{k}}$ are computed according the log-Poisson distribution eq. (6) where the ratio $\frac{L}{r}$ is 2.

We have still to check that this model has the correct second order statistics [9]. Using a 512×512 image produced with the generative model we have verified that the power spectrum is indeed of the form $S(\vec{f}) \sim 1/f^2$, for an appropriately chosen Ψ's. The plot and the fit are exhibited in Fig. 7.

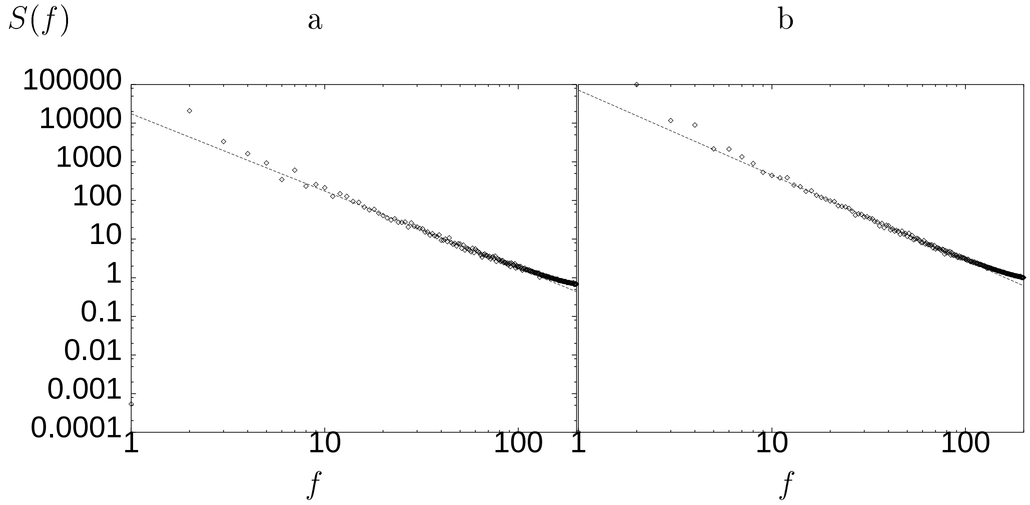

Figure 7: log-log plot of the power spectrum along the a) horizontal and b) vertical directions for a 512×512 image produced with the generative model. The straight lines correspond to $1/f^2$ fits; the horizontal axes are given in cycles *per* image.

5.1 Geometrical interpretation of the generative model

Table 2 presents a catalogue of images produced with the generative model for different values of the parameters β and $\bar{s}$. The images have been produced in such a way that those with identical $\bar{s}$ (shown in the same column) have been constructed with the same number of modulations for each $\eta_{j\vec{k}}$. That is, a given site $j\vec{k}$ has a random choice of $n_{j\vec{k}}$ modulations which is the same for all the images in the column. There is also a continuous change for fixed β as $\bar{s}$ varies, which is obtained keeping fixed the cumulative probabilities used to generate the number of modulations at a given site.

The geometrical characteristics of the images are clear from simple visual inspection of the figures in Table 2. The images in the column $\bar{s} = 0$ are all the same, as the possibility of having even one transition is neglected. These images are also obtained as the limit $\beta \to 1$ (this is because for $\beta = 1$ the transitions become irrelevant, as they do not change the value of ϵ_r). These two cases should correspond to a uniform contrast field (the non-uniformity in the borders are due to finite size effects). For non-trivial values of $\bar{s}$ and β one observes the following properties:

- as $\bar{s}$ increases keeping β fixed, the changes in contrast become more irregular and localized. In fact, the set of points where the changes in contrast are the sharpest can be proven to posses a fractal dimension $D_\infty = 2 - \bar{s}$ [4, 6].

- as β increases at fixed $\bar{s}$, contrast changes become smoother. This smoothing is achieved by increasing the number of points that are observed as more singular, that is, by increasing the density of edges.

These results have to be contrasted with the interpretation of β given in Sec. 4.5 according to which this parameter is a measure of the density of the sharpest transitions. In that analysis the relevance of the other parameter of the model, $\bar{s}$, was ignored because that discussion was limited to edges, which are curve-like sets (Notice however that $\bar{s}$ is only needed to compute the SS exponents but not to predict the ESS ones). This implies $\bar{s} = 1$, which may not be true for natural images. It is obvious that in the generative model $\bar{s}$ also characterizes the density of edges, which is a decreasing function of $\bar{s}$.

It is easy to provide an extended description of the geometry of the scenes in terms of both parameters. In this analysis, $\bar{s}$ gives a measure of the typical size of an elementary edge-like structure (which can be spread as a blob, more localized like a curve or a point, or even a structure of more irregular aspect), while β gives a measure of the amount of such elementary sets present in the image. Obviously, if the amount of them increases (β increases) the density grows, but it also does if the typical size increases ($\bar{s}$ decreases).

5.2 Numerical analysis of the generative model

The artificial images are constructed to have a certain theoretical β and $\bar{s}$. We study here the properties of 32 generated images of 256×256 pixels with $\beta = 0.40$. The histogram in Fig. 8 shows the distribution of the 64 computed $\beta's$ for the horizontal and vertical direction of each of the 32 images. The mean of the histogram is $\beta = 0.31$ with standard deviation $\sigma_\beta = 0.13$. The fluctuations around the mean and its bias with respect to the theoretical β are mainly due to the finite size of the images and not to the size of the dataset. To verify this we have randomly divided the dataset in two sets of 12 and 20 images and have checked that neither their means nor their standard deviations change appreciably.

On the other hand, averaging the moments of the 64 cases (which approximates very well the moments computed over a single 2048×2048 image) yields $\beta = 0.38$ (Fig. 9a), rather close to the theoretical 0.40.

What is interesting is that even when these fluctuations in β are due to finite size effects the multiplicative process for each image is preserved. This can be seen in Fig. 9b where the mean and variance of the normalized ESS exponent $\rho'(p, 2)$ (averaged over the 64 cases) is shown.

One can wonder whether the fluctuations in β observed in the true images could be thought of as finite size effects. This does not seem to be the case. As we have seen in Sec. 4.4 even the whole set of images follows the log-Poisson model with a given β. Considering this set as a single large image, each of its pieces (i.e. the single images) should follow the log-Poisson model with β's distributed as in Fig. 8. However, if the fluctuations of the artificial images were a good model for the fluctuations of the true ones, then the dispersion of the histogram in Fig. 8 should be larger than that of the histogram in Fig. 4III (because the synthetic images are smaller than the true ones).

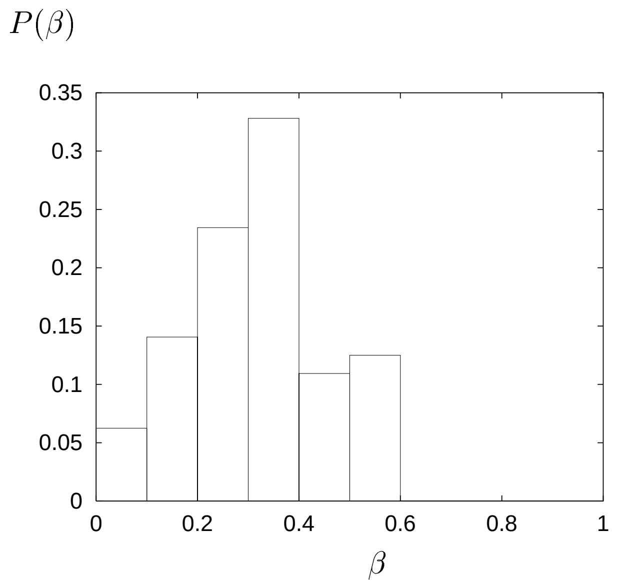

Figure 8: Generative model. Histogram of β for the 32 256×256 synthetic images (and two directions) with a theoretical $\beta = 0.40$

6 Discussion

One of the main conclusions of this work is that the multiscaling properties of changes in contrast are robust in natural scenes: Both SS and ESS are present in almost every single image considered, which are of very different type. In addition, a vast number of the images in our dataset (94%) can be described in terms of a multiplicative process of the log-Poisson type. The modulation parameter β is distributed over its admissible range, although the probability distribution decreases with increasing β.

Four different ensembles of images (described in Sec. 3) were considered in order to understand how the presence or absence of certain features affects the edge statistics. The moments of the relevant variables ϵ_r were then computed by averaging over all the images in the ensemble. It was found that SS, ESS and the same multiplicative process are also present in each of the four ensembles and for both vertical and horizontal statistics. Again, these eight cases differ only in the value of β. Analyzing the experimental values of β and simple theoretical situations, we have proposed a geometrical interpretation of β as a measure of the density of edges and similar structures present in the images. This interpretation allows for an understanding of the differences in the values of this parameter, which appear in Table 1.

The same properties (SS, ESS and log-Poisson) hold when the whole dataset is taken as the ensemble.

The generative model studied in this paper provides a simple way to produce images with the multiscaling properties of natural scenes. It could be used as a better description of the contrast distribution that contains not only the correct second order statistics but also a correct description of the non-gaussian aspects of changes

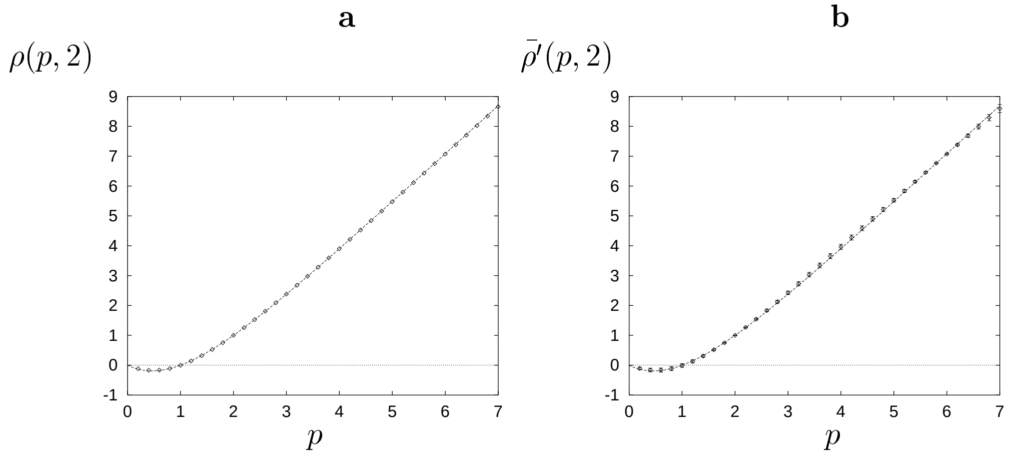

Figure 9: Test of the log-Poisson multiplicative process for the generative model. a) $\rho(p, 2)$ for the averaged moments of the whole ensemble of 32 256×256 images. The dashed line corresponds to the ESS exponents predicted by the log-Poisson model for $\beta = 0.38$. b) The mean of $\rho'(p, 2)$ over the 32 synthetic images and the two directions, where now $\rho'(p, 2)$ has been normalized to $\beta = 0.38$. The dashed line represents $\rho^{0.38}(p, 2)$, the log-Poisson $\rho(p, 2)$ with $\beta = 0.38$. The error bars are of the size of the data points.

in contrast. The model could then be used to study the consequences that these properties have on the early stages of the visual system [26].

The model is not based on objects. A quick glance to Table 2 shows cloudy structures that have little to do with natural images (except e.g. those of the sky). In particular their appearance is rather fractionary, lacking real contours. However the model could provide an indication of which aspects of the visual processing depend only on the multiscaling properties and which parts are more dependent on spatial correlations.

By analyzing the finite size effects of synthetic images, we have concluded that the value of β is an essential feature of each natural image, which cannot be regarded as a fluctuation from a universal value. The generative model was also used to illustrate the meaning of the parameters β and $\bar{s}$.

The wavelet used in the generative model studied in this work was somewhat arbitrary. A better way to represent the multiscaling properties of natural images is to derive an intrinsic wavelet basis from the image dataset itself [26]. When this is done, it is found that the coefficients of the wavelet expansion are scale-independent and weakly space-dependent, thus providing a quasi-optimal code for the image. Their distribution for a given image is log-Poisson (eq. (6)), with exactly the same parameter β that would be obtained from the statistics of ϵ_r. This is a remarkable result, and supports our proposal that the visual system has adapted to the multiscaling properties of natural scenes and that its architecture contains, in particular, the log-Poisson multiplicative process.

Further equalization of these quasi-independent coefficients that follow the distribution given in eq. (6) could be an advantageous strategy for the visual system since it optimizes the mutual information between the source (visual stimuli) and the internal representation [27]. Since the probability distributions of the quasi-independent coefficients depend on β this proposed equalization depends on the image being processed. Thus the value of β may also play a role in the processing carried out by the visual system in its early stages.

To conclude, in this work we have established that the multiscaling properties of changes in contrast are a robust feature in natural images. The existence of this statistical regularity for single images has several implications for visual processing. It is relevant because it implies a reduction of the entropy of the distribution probability of the ensemble of natural images. As a new statistical regularity is found, the set of images that can be considered as natural becomes more restricted. The fact that this property is present not only in ensembles but also in single images greatly reduces the entropy of the ensemble of natural images. Therefore its existence should be taken into account when modeling visual processing in the brain.

Acknowledgments

We are thankful to Angela del Pozo for her help during the initial stages of the preparation of this work. We are grateful to Hans van Hateren, some of whose images we used for the statistical analysis. Antonio Turiel is financially supported by an FPI grant from the Comunidad Autónoma de Madrid, Spain. This work was funded by a Spanish grant PB96-0047.

References

[1] Barlow H. B., Possible principles underlying the transformation of sensory messages. in *Sensory Communication* (ed. Rosenblith W.) pp. 217. (M.I.T. Press, Cambridge MA, 1961).

[2] Ruderman D. and Bialek W., Statistics of natural images: Scaling in the woods *Phys. Rev. Lett.* **73**, 814 (1994)

[3] Ruderman D., *Network* **5**, 517-548 (1994)

[4] Turiel A., Mato G., Parga N. and Nadal J.-P. The self-similarity properties of natural images resemble those of turbulent flows. *Phys. Rev. Lett.* **80**, 1098-1101 (1998)

[5] Turiel A., Mato G., Parga N. and Nadal J.-P. Self-similarity properties of natural images. Proceedings of NIPS'97. Cambridge, MA: MIT Press (1998)

[6] Turiel A. and Parga N., The multi-fractal structure of contrast changes in natural images: from sharp edges to textures. To appear in *Neural Computation* (1999).

[7] Turiel A., Parga N., Ruderman D. and Cronin T.W., Multiscaling and information content of natural color images. Submitted to *J. Optical Soc. Am. A* (1999).

[8] Burton G.J. and Moorhead I.R., Color and spatial structure in natural scenes *Applied Optics* **26** 157-170 (1987).

[9] Field D. J., Relations between statistics of natural images and the response properties of cortical cells. *J. Opt. Soc. Am. A* **4** 2379-2394 (1987).

[10] R. Linsker. Self-organization in a perceptual network. *Computer*, **21** 105-117 (1988)

[11] J. J. Atick. Could information theory provide an ecological theory of sensory processing. *NETWORK*, **3** 213-251 (1992)

[12] J.H. van Hateren, Theoretical predictions of spatiotemporal receptive fields of fly lmcs, and experimental validation. *J. Comp. Physiology A*, **171** 157-170 (1992)

[13] Zhaoping Li and J. J. Atick. Towards a theory of the striate cortex *Neural Computation*, **6** 127-146 (1994)

[14] Field D. J., Scale-invariance and Self-similar "wavelet" transforms: an analysis of natural scenes and mammalian visual systems. in *Wavelets, Fractals, and Fourier Transforms* (ed. Farge M., Hunt J.C.R. and Vassilicos J.C.) pp. 151-193. (Clarendon Press, Oxford, 1993).

[15] Zhen-Su She and Leveque E., Universal scaling laws in fully developed turbulence. *Phys. Rev. Lett.* **72**, 336-339 (1994).

[16] Zhen-Su She and Waymire E.C., Quantized energy cascade and log-Poisson statistics in fully developed turbulence. *Phys. Rev. Lett.* **74**, 262-265 (1995).

[17] Castaign B., Gagne Y. and Hopfinger, Velocity probability density functions of high Reynolds number turbulence *Physica (Amsterdam)* **46D**, 177 (1990).

[18] Castaign B., The temperature of turbulent flows *J. Physique II (France)* **6**, 105-114 (1996).

[19] Benzi R., Ciliberto S., Baudet C., Ruiz Chavarria G. and Tripiccione C., *Europhys. Lett.* **24** 275-279 (1993)

[20] Benzi R., Ciliberto S., Tripiccione R., Baudet C., Massaioli and Succi, Extended self-similarity in turbulent flows *Phys. Rev. E***48**, R29 (1993)

[21] Benzi R., Ciliberto S., Baudet C. and Ruiz Chavarria G., *Physica D* **80** 385-398 (1995)

[22] van Hateren J.H. and van der Schaaf A., Independent component filters of natural images compared with simple cells in primary visual cortex. *Proc.R.Soc.Lond.* **B265**, 359-366 (1998)

[23] Ruderman D., Origin of scaling in natural images *Vision Research* **37**, 3385-3398 (1997)

[24] Benzi R., Biferale L., Ciliberto S., Struglia M.V. and Tripiccione R., Generalized scaling in fully developed turbulence, *Physica D* **96** 162-181 (1996)

[25] Daubechies I. *"Ten Lectures on Wavelets"*. CBMS-NSF Series in App. Math. *Capital City Press*. Montpelier, Vermont (1992)

[26] Turiel A. and Parga N. The multifractal wavelet filter of natural images. Submitted to NIPS'99.

[27] Nadal J.-P. and Parga N. Nonlinear neurons in the low-noise limit: a factorial code maximizes information transfer *Network* **5** 565-581 (1994)

	$\bar{s} = 0.0$ ($D_\infty = 2.0$)	$\bar{s} = 0.5$ ($D_\infty = 1.5$)	$\bar{s} = 1.0$ ($D_\infty = 1.0$)	$\bar{s} = 1.5$ ($D_\infty = 0.5$)	$\bar{s} = 2.0$ ($D_\infty = 0.0$)
$\beta = 0.0$	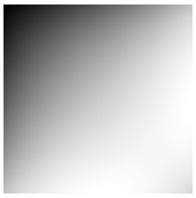	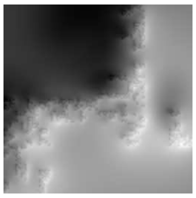	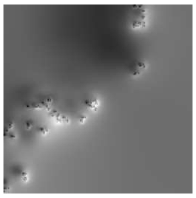	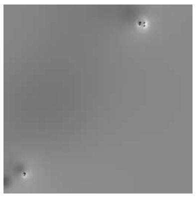	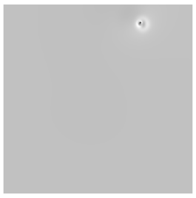
$\beta = 0.2$	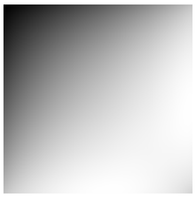			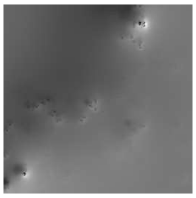	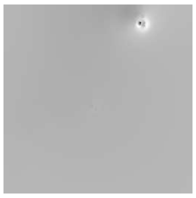
$\beta = 0.4$	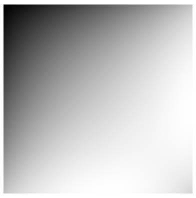				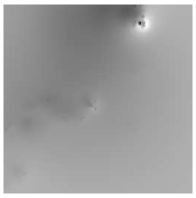
$\beta = 0.6$	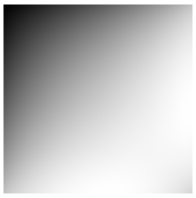				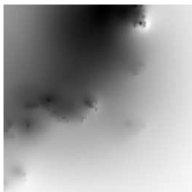
$\beta = 0.8$	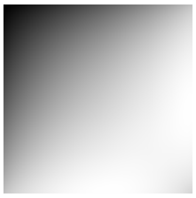	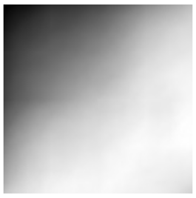	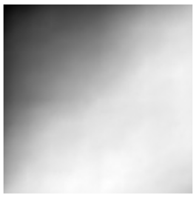		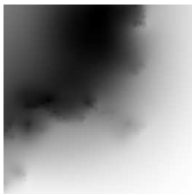

Table 2: Catalogue of synthetic images. The case $\beta = 1$ is not shown because the images are like those for $\bar{s} = 0$ (see text)