Scales in Natural Images and a Consequence on their Bounded Variation Norm

Luis Alvarez[1], Yann Gousseau[2], and Jean-Michel Morel[2]

[1] Departamento de Informatica y Sistemas, Universidad de Las Palmas, Campus de Tafira, 35017 Las Palmas, Spain. luis@amihp710.dis.ulpgc.es
[2] Centre de Mathématiques et leurs Applications, ENS Cachan, 61 av. du Président Wilson, 94235 Cachan Cedex, France. gousseau@cmla.ens-cachan.fr, morel@cmla.ens-cachan.fr

Abstract. This paper introduces a new method for analyzing scaling phenomena in natural images, and draws some consequences as to whether natural images belong to the space of functions with bounded variation.

1 Introduction

A digital, gray level image may be seen as the realization of a random vector of size $H \times L$ taking values in a discrete set $V = 1, ..., G$. For typical values like $H = L = G = 256$, the number of possible realizations, $G^{HL} = 2^{524288}$, is huge. Obviously, "natural images", i.e. digital photographs of natural scenes, only form a small subset of all possible realizations. Looking at random realizations of such vectors is enough to be convinced of this fact. Natural images are highly improbable events. It is therefore interesting to look for statistical characteristics of such images: what are the relationships between gray level values at distant pixels? Is it possible to define a probability law for natural images? Moreover, statistics of texture images may be useful for synthesis purposes (see [9], [24], [23]).

Most of the statistical studies of natural images are concerned with first or second order statistics (through the power spectrum, the covariances, the cooccurrences) or with additive decompositions of images. The power spectrum $P(\omega, \nu)$ is known to be well approximated by a power function $\frac{C}{(\omega^2+\nu^2)^\gamma}$, where γ is an image dependant number usually close to 2 (see [5], [7]). The histogram of natural images has been found to have a peculiar, non-Gaussian shape (see [20], [10]). Nearest neighbors coocurences functions also exhibit non-Gaussian distributions (see [10]). Principal and independant component analysis on databases of such images yield localized and oriented images bases (see [17], [2]). We have a different approach, working in the image domain on items that can have a straightforward visual interpretation, and involve (relatively) long and high order interactions between pixels. We shall show that in natural images, there is a constant form for the size distribution of such items. The definitions of sizes we consider are of two types: area and boundary length. An experimental program which we performed on many photographs of very diverse natural scenes

indicates that the size distribution of homogeneous parts in images obeys a law

$$\text{Card}\{\text{Homogeneous regions with size s}\} = \frac{K}{s^\alpha},$$

where K is an image dependent constant. When the size s denotes the area, in most photographs, α is close to 2. We will define in Section 2 what we mean by homogeneous parts, the connected components of image domains where contrast does not exceed a certain threshold. Let us mention that power laws have been previously observed, e.g. for points statistics (see [18], [19]) or density of extrema in scale space (see [11]).

As a consequence of the size power law, some information can be obtained about the "natural" function space for images, as will be shown in Section 3: we focus our attention on the space BV of functions with bounded variation. We are in a position to tell when a given image is not in this space, provided the observed size distribution model remains true at smaller (not observable) scales as well.

2 Sizes of sections in natural images

2.1 The distribution of areas

We'll now make clearer what we mean by homogeneous region of an image. We begin by equalizing the image histogram, and uniformly quantify it in the following way. We consider a digital image I of size $H \times L$, with G integer gray levels, and write $I(i,j)$ for the gray level at pixel (i,j). Let k be an integer less than G. Let N_1 be the first integer such that more than $\frac{HL}{k}$ pixels have a gray level less than N_1, then N_2 the first integer such that more than $2\frac{HL}{k}$ have a gray level less than N_2, then $N_3, ..., N_k = G$ defined the same way, this sequence being possibly constant at some point. For l varying from 1 to k, let I_l be the binary image with $I_l(i,j) = 1$ if $I(i,j) \in [N_{l-1}, N_l)$ and $I_l(i,j) = 0$ otherwise. We call those images k-bilevels of I. Each bilevel image represents a quantization level of the equalized image.

Next, we look at the area histogram of the connected components of the bilevels. For s an integer varying from 0 to HL, let $f(s)$ be the number of connected components with area s of the set of 1's pixels, in any of the k-bilevels of I. We will both consider 4-connectivity (each pixel has 4 neighbors: up, down, right, left) and 8-connectivity (we add the diagonal neighbors, so that each pixel has 8 neighbors).

We computed the function f on many digital photographs. We did not attempt to use a single source of images; the digitized images either are scanned photographs or from a digital camera, with diverse optical systems and exposures. Those functions are of the form $f(s) = \frac{C}{s^\alpha}$, with C a constant and α a real number close to two, for values of s in a certain range and reasonable values of k (basically between 4 and 30). The observed fit is excellent, as can be seen from

Figure 2, which actually corresponds to one of the worst cases we observed. For fixed k, we consider the set of points

$$S = \{(\log(s), \log(f(s)), 0 \leq s \leq T_{\max}\},$$

where $T_{max}+1$ is the smallest value of s such that $f(s) = 0$. We perform a linear regression on this set S so as to find the straight line (in the log-log coordinates) $g(\log(i)) = A - \alpha \log(i)$ the closest to S in the least squares sense, and write E for the least squares error.

2.2 The distribution of areas in digital photographs

We present the results for two pictures having different scales and textures in Table 1. The value of α appears to be related to the amount of texture in the image; the more textured the image, the bigger the value of α. Typically, for photographs of natural scenes, the value of alpha is between 1.5 and 3 (the values close to 3 being reached for images as the baboon (Figure 1), which present textured areas), whereas for textures (e.g. from the Brodatz's album), it is typically between 2.5 and 3.5.

Fig. 1. baboon (512×512) and city (612×792) images

image	k	α	E	T_{max}	A
city	20	2.03	.32	184	11.7
city	16	1.94	.30	165	11.3
city	12	1.91	.42	202	11.1
city	8	1.80	.44	191	10.3

image	k	α	E	T_{max}	A
baboon	20	2.55	.30	70	11.7
baboon	16	2.38	.33	82	11.3
baboon	12	2.42	.47	78	11.4
baboon	8	2.35	.41	76	11.2

Table 1. different values of the quantization number k for the city and baboon images, 8-connectivity. Area distribution is $f(s) = As^{-\alpha}$, T_{max} is the maximal considered area

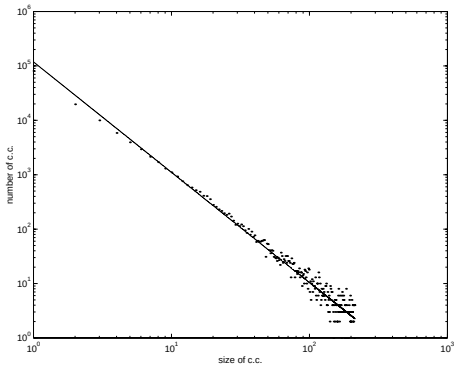

Fig. 2. function f (area distribution) for the city image (Figure 1), $k = 12$, 8-connectivity, $T_{max} = 202$

We also performed the linear regression on sets of points

$$S_{Tmin} = \{(\log(s), \log(f(s)), T_{min} \leq s \leq T_{max}\}$$

for various values of $T_{\min}$ to show that the fit of S to the power law was not forced by small areas only, and moreover that if the contribution of E mainly comes from the large areas, the value of α computed with those large areas was close to the initial value. The results for the image of the city are shown in Table 2. Those results about the stability of the slope of the regression across scales are of great importance in view of the hypothesis to be made in Section 3.

image	T_{min}	α	E	A
city	5	1.97	.30	11.6
city	10	1.98	.30	11.6
city	20	1.96	.31	11.5
city	40	1.91	.36	11.3

Table 2. different values of T_{min} for the city image, $k = 16$, 8-connectivity

2.3 The distribution of boundary lengths in digital photographs

We performed exactly the same analysis on the boundary lengths of connected components of bilevels as we did before on areas of those components. As a discrete definition of the length of a discrete connected set S (8-connectivity), we chose to count the pixels not belonging to S that are neighbors of some pixel

of S in the 4-connectivity sense. There are many other ways to define discrete boundary length. We tried several other methods that gave basically the same results as the one we detail here. The notations k, E, A, T_{min} and T_{max} refer to the same quantities as before; β now stands for the exponent of the power law whose fit to the boundary length distribution is the best in the least square sense. We chose $T_{min} = 10$, because some small values for the boundary lengths are attained only for regions touching the border of the image. The fit to the power law is again very good, and the exponent β is usually between 2 and 3. We present the results for the images of the city and the baboon in Table 3 . We note that $\beta \simeq 2\alpha - 1$ accounts for connected components of bilevel sets satisfying on the average a decent isoperimetric ratio, $c \leq \frac{area^{\frac{1}{2}}}{boundary\ length} \leq C$. This is not the case in general, except for some images of textures.

image	k	β	E	T_{max}	A
city	20	2.42	.32	184	13.7
city	16	2.41	.35	184	13.7
city	12	2.28	.33	187	13.0
city	8	2.23	.48	192	12.4

image	k	β	E	T_{max}	A
baboon	20	3.02	.28	82	14.2
baboon	16	2.91	.35	81	13.9
baboon	12	2.93	.41	100	14.0
baboon	8	2.89	.33	96	13.8

Table 3. boundary lengths for the city and baboon images, with different values of the quantization number k

Let us mention that the length distribution of intersections of the homogeneous part of the image with lines (the so-called intercepts) also follows a power law. In a forecoming paper (see [1]), we use a morphological model, the dead-leaves model of G. Matheron (see [14]), as an object-based model for images. An image is defined as a sequential superposition of random objects. If we interpret the homogeneous parts as being the visible parts of objects after the occlusion process, it is possible to deduce the form of the length distribution of the intercepts from a power law distribution of the size of objects. This result is closely related to the ones of [19], where objects are defined in the image by visual segmentation, and where a power law is observed for the covariances.

2.4 Other types of images

In order to see whether the power law is in some sense characterizing digital photographs, we computed histograms of areas of bilevels for other types of images. We looked at noises images, white or correlated and text images.

White noise images, that is to say images in which the gray level values at distinct pixels are independent random variables, present an histogram of the form $f(s) = \exp(-Cs)$, with C a constant. We observed this fact on two different kinds of white noises: uniform and Gaussian. Text images produced by text editor do lead, as one would guess, to an histogram consisting of isolated peaks, whose height is not directly related to the value of s.

Then we looked at correlated noise. We performed convolutions between white noises and a Gaussian $\exp(-\frac{1}{2\sigma^2}(x^2+y^2))$, where σ is a variable parameter. This was done by multiplication in the frequency domain. Such a convolution can be seen as a crude approximation of the effect of an optical lens. The results we obtain for those images were similar to the ones for digital photographs of textures. We present the results obtained in the case of the uniform white noise in Table 4. We also tested the effect of the convolution with Bessel functions (Fourier transform of disks) and the results were very similar.

image convolved with a Gaussian	σ	α	E	T_{max}	A
noise convolved with a Gaussian	0.71	3.88	0.60	26	13.3
noise convolved with a Gaussian	0.85	3.57	0.50	30	13.1
noise convolved with a Gaussian	1.12	3.27	0.55	39	13.0

Table 4. Uniform white noise image, after convolution with a Gaussian of parameter σ, k=12, 8-connectivity

Those "non-natural images" lead to two remarks about the $\frac{1}{area^2}$ law. First, this law does not characterize natural images, even though a correlated noise looks similar to a natural texture. Secondly, the size law could be related to the way the optical photographic device captures the image, as suggested by the behavior of noises convolved with a Gaussian. More precisely, we observed that the convolution with a Gaussian increases the value of α for images where the initial α is small (such as text and synthetic images) whereas it tends to decrease its value when it is initially bigger than 2 (noises).

Another, and more satisfactory explanation of this power law is scale invariance. The assumption that natural images are scale invariant, so that all observed statistics should be scale (zoom) invariant, has been confirmed by the shape of the power spectrum mentioned in the introduction (see [7]), and also by the fact that some statistics are preserved when shrinking the image (see [20], [16]). Our experiments also confirm this assumption, since scale invariance yields the $\frac{1}{area^2}$ law. Indeed, if we suppose that the total area occupied by regions having an area between A and A' is the same as the total area occupied by regions with area between tA and tA', for all t, A, A', then the power law with exponent 2 is the only acceptable size distribution.

3 Size of sections and the BV norm of natural images

The aim of this section is to give a computational tool to decide whether an image can belong to the space BV of functions with bounded variations. The BV assumption for natural images is far ranging, from image restoration ([21], [22]) to image compression.

The space BV is the space of functions for which the sum of the perimeters of the level sets is finite. The space BV is of great importance in image modeling,

since such a simple image as a white disk on a black background is not in any Sobolev space, but belongs to BV. However, if the disk is replaced by an object whose boundary has an infinite length, such as a bidimensional Cantor set, then the corresponding function is no longer in BV. There is also another way for a function not to be in BV. Each of its level sets may be of finite perimeter, while the sum of those perimeters tends towards infinity. According to our analysis, this is the case with natural images, for which, in a sense, small objects are too numerous for the function to be in BV.

3.1 A lower bound for the BV norm

We consider $I \in BV(\Omega)$ a bounded image belonging to the space of functions with bounded variation ([25], [6]) on a domain (e.g. rectangular) $\Omega \subset I\!R^2$. For $\lambda \in I\!R$, define the level set of I with level λ by

$$\chi_\lambda I = \{x, I(x) \geq \lambda\}.$$

Recall that a function is of bounded variation if, for almost every $\lambda \in I\!R$, $\chi_\lambda I$ is a set with finite perimeter and, denoting by $\mathrm{per}(\chi_\lambda I)$ this perimeter (for a precise definition of the perimeter and the essential boundary we refer to [6]),

$$||I||_{BV} = \int_{I\!R} \mathrm{per}(\chi_\lambda I) d\lambda. \tag{1}$$

(By the coarea formula, see [6], we also have $||I||_{BV} = \int_\Omega |DI|$)

In addition, by the classical isoperimetric inequality, we have for every set O with finite perimeter,

$$\mathrm{per}(O) \geq 2\pi^{\frac{1}{2}} \nu(O)^{\frac{1}{2}}, \tag{2}$$

where $\nu(O)$ denotes the Lebesgue measure of O. In the following, we shall consider sections of the image. We always assume that the image I satisfies $0 \leq I(x) \leq C$. We first fix two parameters γ, λ, with $0 \leq \lambda \leq \gamma$. For any $n \in I\!N$, we consider the bilevel sets of I

$$\{x, \lambda + (n-1)\gamma \leq I(x) < \lambda + n\gamma\} = \chi_{\lambda+(n-1)\gamma} I \setminus \chi_{\lambda+n\gamma} I.$$

We call (γ, λ)-section of I any set which is a connected component of a bilevel set $\chi_{\lambda+(n-1)\gamma} I \setminus \chi_{\lambda+n\gamma} I$ for some n. We denote each one of them by $S_{\gamma,\lambda,i}$ for $i \in J(\gamma, \lambda)$, a set of indices. Notice that the (γ, λ)-sections are disjoint and their union is the image domain Ω,

$$\bigcup_{i \in J(\gamma,\lambda)} S_{\gamma,\lambda,i} = \Omega. \tag{3}$$

There are several ways to define the connected components of a set with finite perimeter, since such a set is defined up to a set with zero Lebesgue measure. We

denote by H^1 the one-dimensional Hausdorff measure, that is to say the length. In the following, we call Jordan curve a simple closed curve of $\mathbb{R}^2$, i.e. the range of a continuous map $c : [0,1] \to \mathbb{R}^2$, such that $c(s) \neq c(t)$ for all $0 < s < t < 1$, and $c(0) = c(1)$. A Jordan curve defines two and only two connected components (in the usual sense) of $\mathbb{R}^2 \setminus c([0,1])$, one bounded and one unbounded. We shall say that a Jordan curve separates two points x and y if they do not belong to the same connected component of $\mathbb{R}^2 \setminus c([0,1])$. One can prove ([8], [3]) that a definition of connected components for a set with finite perimeter permits the following statements :

Theorem 1 *(and definition)*
Let O be a set with finite perimeter.
(i) The essential boundary of O consists, up to a set of zero H^1-measure, of a countable set of noncrossing simple rectifiable closed curves c_j with finite length such that $per(O) = \sum_j H^1(c_j)$
(ii) Two points are in the same connected component of O if and only if for any representation of the essential boundary by a family of Jordan curves of the preceding kind, c_j, they are not separated by one of the c_j.
(iii) With this definition, the perimeter of a set with finite perimeter is the sum of the perimeters of its connected components.

We denote by $J(n) \subset J(\gamma, \lambda)$ the set of indices of sections which are connected components of $\chi_{\lambda+(n-1)\gamma} I \setminus \chi_{\lambda+n\gamma} I$. As an obvious consequence of Proposition 1, we have

Corollary 1

$$per(\chi_{\lambda+(n-1)\gamma} I \setminus \chi_{\lambda+n\gamma} I) = \sum_{i \in J(n)} per(S_{\lambda,\gamma,i}).$$

When A is a set with finite perimeter, we have ([6])

$$\mathrm{per}(A) = ||\mathbb{1}_A||_{BV}.$$

Lemma 1 *If $B \subset A$ are two sets with finite perimeter, then*

$$per(A \setminus B) \leq per(A) + per(B).$$

Proof Indeed, by the subadditivity of the BV norm, we deduce from

$$\mathbb{1}_{A \setminus B} = \mathbb{1}_A - \mathbb{1}_B$$

that

$$\mathrm{per}(A \setminus B) \leq \mathrm{per}(A) + \mathrm{per}(B).$$

□

In the following theorem, we analyze the statistics of sizes of sections as follows. We fix γ, that is, the overall contrast of considered sections and for each

$0 \leq \lambda \leq \gamma$, we count all sections $S_{\gamma,\lambda,i}$ which have an area between s and $s+ds$. In other terms we consider the integer

$$\mathrm{Card}\{i, s \leq |S_{\gamma,\lambda,i}| \leq s+ds\}.$$

We average this number over all λ's in $[0,\gamma]$, and assume that this average number has a density $f(\gamma,s)$ with respect to s. In other terms,

$$\frac{1}{\gamma}\int_0^\gamma \mathrm{Card}\{i, s \leq |S_{\gamma,\lambda,i}| \leq s+ds\}d\lambda = f(\gamma,s)ds \tag{4}$$

Theorem 2 *Assume that there exists some $\gamma > 0$ such that (4) holds, i.e. the average number of sections with area s, for $0 \leq \lambda \leq \gamma$, has a density $f(\gamma,s)$. Then there is a constant c, not depending on I, such that*

$$||I||_{BV} \geq c\int_0^{\nu(\Omega)} s^{\frac{1}{2}} f(\gamma,s)ds. \tag{5}$$

Proof Applying Corollary 1 and Lemma 1

$$\begin{aligned}
||I||_{BV} &= \int_{I\!R} \mathrm{per}\{x, I(x) \geq \lambda\}d\lambda \\
&= \frac{1}{2}\Big(\int_{I\!R} \mathrm{per}\{x, I(x) \geq \lambda\}d\lambda + \int_{I\!R} \mathrm{per}\{x, I(x) \geq \lambda-\gamma\}d\lambda \\
&\geq \frac{1}{2}\int_{I\!R} \mathrm{per}(\chi_{\lambda-\gamma}I \setminus \chi_\lambda I)d\lambda \\
&= \frac{1}{2}\sum_{n\in\mathbf{Z}}\int_{n\gamma}^{(n+1)\gamma} \mathrm{per}(\chi_{\lambda-\gamma}I \setminus \chi_\lambda I)d\lambda \\
&= \frac{1}{2}\int_0^\gamma \sum_{n\in\mathbf{Z}} \mathrm{per}(\chi_{\lambda+(n-1)\gamma}I \setminus \chi_{\lambda+n\gamma}I)d\lambda \\
&= \frac{1}{2}\int_0^\gamma \sum_{i\in J(\gamma,\lambda)} \mathrm{per}(S_{\gamma,\lambda,i})d\lambda.
\end{aligned}$$

By isoperimetric inequality (2), we therefore obtain

$$||I||_{BV} \geq \pi^{\frac{1}{2}}\int_0^\gamma \sum_{i\in J(\gamma,\lambda)} |S_{\gamma,\lambda,i}|^{\frac{1}{2}}d\lambda.$$

Applying Fubini-Tonelli Theorem, some slicing and the assumption (4), we get

$$\begin{aligned}
||I||_{BV} &\geq \pi^{\frac{1}{2}}\int_0^\gamma d\lambda \int_0^{\nu(\Omega)} \mathrm{Card}\{i \in J(\gamma,\lambda),\ s \leq |S_{\gamma,\lambda,i}| \leq s+ds\}s^{\frac{1}{2}} \\
&= \pi^{\frac{1}{2}}\int_0^{\nu(\Omega)}\int_0^\gamma d\lambda \mathrm{Card}\{i \in J(\gamma,\lambda),\ s \leq |S_{\gamma,\lambda,i}| \leq s+ds\}s^{\frac{1}{2}} \\
&= \gamma\pi^{\frac{1}{2}}\int_0^{\nu(\Omega)} s^{\frac{1}{2}} f(\gamma,s)ds.
\end{aligned}$$

□

We can repeat the preceding analysis by assuming now that

$$\frac{1}{\gamma}\int_0^{\gamma} \mathrm{Card}\{i, p \leq \mathrm{per}(S_{\gamma,\lambda,i}) \leq p+dp\}d\lambda = g(\gamma,p)dp. \tag{6}$$

Then we have the analog of Theorem 2 for the perimeters of sections:

Theorem 3 *Assume that there exists some $\gamma > 0$ such that (6) holds, i.e. the average number of sections with perimeter s, for $0 \leq \lambda \leq \gamma$, has a density $g(\gamma, p)$. Then*

$$||I||_{BV} \geq \frac{1}{2}\int_0^{+\infty} pg(\gamma,p)dp. \tag{7}$$

Proof The proof is essentially the same as for Theorem 2.

3.2 Application to natural images

In this section, we draw the consequences of Theorems 2 and 3 for the images analyzed in Section 2. According to the results of this section, we can assume that the considered images satisfy

$$f(\gamma, s) = \frac{C}{s^{\alpha}} \tag{8}$$

$$g(\gamma, p) = \frac{C}{p^{\beta}} \tag{9}$$

for some constants $\alpha > 0$, $\beta > 0$. This law has been experimentally checked for several values of $\gamma = \frac{256}{k}$, k ranging from 8 to 20. We also checked that the value of α was almost not modified when the bilevels were not defined from gray level 0, but from some gray level less that $\frac{256}{k}$ (that is to say, in the continuous model, for different values of λ). By Theorem 2 we have

$$||I||_{BV} \geq c\int_0^{\nu(\Omega)} \frac{Cs^{\frac{1}{2}}}{s^{\alpha}}ds = +\infty \text{ if } \alpha > \frac{3}{2}$$

and in the same way,

$$||I||_{BV} \geq c\int_0^{+\infty} \frac{Cp}{p^{\beta}}dp = +\infty \text{ if } \beta > 2.$$

thus if we admit that (8) and (9) indeed hold for natural images when $s \to 0$, as is indicated by the experiments of section 3.1, we obtain that the considered images are not in BV if $\alpha > \frac{3}{2}$, or $\beta > 2$. This strong assumption about the small scales behavior is motivated by the goodness of the fit at every scales and

by the stability of the fit with respect to T_{min}, see Section 2, Table 2. Notice, however, that $\alpha > 2$, which happens for several of the considered images, is not compatible with a finite image area, since then $\int \frac{s ds}{s^\alpha} = +\infty$. As suggested to us by Vicent Caselles and Stéphane Mallat, this raises the question of whether the area is correctly measured by covering pixels. In fact, if a region is very ragged, the cardinality of covering pixels may be related to its perimeter as well, in which case the estimate of $g(\gamma, s)$ is more reliable. This cardinality could also be related to a fractional Hausdorff measure.

We point out here that wavelet coefficients (see [15], [13] for an introduction to wavelet decompositions) also give a way to decide whether or not an image belongs to the space BV. Let (c_k) be the wavelets coefficients of the image I, ordered in a nonincreasing sequence. Let us suppose that the wavelets have compact supports. We say that the c_k's are in l^1 if $\sum |c_k| < +\infty$, and that they are in weak-l^1 if there exists a constant C such that $c_k \leq \frac{C}{k}$. Obviously l^1 is included in weak-l^1. It is quite easy to prove that if the c_k are in l^1, then I is in BV. In the other direction, Cohen and al., [4], recently proved that if I is in BV, then the c_k's are in weak-l^1. Thus it is possible to decide whether an image belongs or not to BV by looking at its wavelet coefficients decay, except if they decrease like $\frac{C}{k}$, which happens to be often the case ([12]). Moreover, it is worth noticing that the wavelet coefficients produced by the characteristic function of a simple shape already decay like $\frac{1}{k}$. We do not present here a precise comparison between the two criteria. Let us just mention that in the case of the baboon image (Figure 1), both methods agree: this image is not in BV. For the well-known image of Lena, our approach gives an α of 1.9 (for $k = 16$), which suggests Lena being out of BV, whereas from the wavelet approach, the image is in BV. In fact, according to our analysis, natural images are not in the space BV. Of course, one may objects the presence of an inner scale cut off, but our results indicate that the BV norm of continuous representations of natural images blows up as we consider smaller and smaller scales.

4 Conclusions

We realized experimentally that the size distribution of homogeneous parts in digital natural images follows a power law. This power law confirms the scale invariance of natural images. Moreover, this enables us to show that, provided this power law is valid for small, non-observable scales, most natural images are not in the space BV of functions with bounded variations.

Acknowledgments

We would like to thank Vicent Caselles, Stuart Geman, Stéphane Mallat and David Mumford for several helpful comments and discussions. We also thank Les Treilles foundation which made some of these discussions possible.

References

1. L. Alvarez, Y. Gousseau, and J.-M. Morel. The size of objects in natural and artificial images. *Advances in Imaging and Electron Physics*, 111,1999.
2. A. J. Bell and T. J. Sejnowski. The "independent components" of natural scenes are edge filters. *Vision Research*, 37:3327–3338, 1997.
3. V. Caselles and J.-M. Morel. The connected components of sets of finite perimeters in the plane. Preprint, 1998.
4. A. Cohen, R. DeVore, P. Petrushev, and H. Xu. Non linear approximation and the space $BV(I\!R^2)$. Preprint, submitted to Amer. J. Math., 1998.
5. N. G. Deriugin. The power spectrum and the correlation function of the television signal. *Telecomunications*, 1:1–12, 1956.
6. L. C. Evans and R. F. Gariepy. *Measure Theory and Fine Properties of Functions.* Studies in Advanced Math. CRC Press, 1992.
7. D. J. Field. Relations between the statistics of natural images and the response properties of cortical cells. *Jour. Opt. Soc. Amer.*, A4:2379–2394, 1987.
8. E. De Giorgi. Su una teoria generale della misura $(r-1)$-dimensionale in uno spazio ad r dimensioni. *Ann. Mat. Pura ed Appl., IV. Ser.*, 36:191–213, 1954.
9. D. J. Heeger and J. R. Bergen. Pyramid based texture analysis/synthesis. *Computer Graphics Proceedings*, pages 229–238, 1995.
10. J. Huang and D. Mumford. Statistics of natural images and models. *To be published*, 1998.
11. T. Lindenberg. Detecting salient blob-like image structure and their scales with a scale-space primal sketch: A method for focus-of-attention. *International Journal of Computer Vision*, 11(3):283–318, 1993.
12. S. Mallat. Personal communication.
13. S Mallat. *A Wavelet Tour of Signal Processing.* Academic Press, 1997.
14. G. Matheron. Modèle séquentiel de partition aléatoire. Technical report, CMM, 1968.
15. Y. Meyer. *Wavelets: Algorithms and Applications.* SIAM, 1993.
16. D. Mumford, S. C. Zhu, and B. Gidas. Stochastic models for generic images. Working draft, 1997.
17. B. A. Olshausen and D. J. Field. Natural images statistics and efficient image coding. *Network: Computation in Neural Systems*, 7:333–339, 1996.
18. A. P. Pentland. Fractal based description of natural scenes. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 6(6):661–674, November 1984.
19. D. L. Ruderman. Origins of scaling in natural images. *Vision Research*, 37(23):3385–3398, 1997.
20. D. L. Ruderman and W. Bialek. Statistics of natural images: Scaling in the woods. *Physical Review Letters*, pages 814–817, 1994.
21. L. Rudin. *Images, Numerical Analysis of Singularities and Shock Filters.* PhD thesis, California Institute of Technology, 1987.
22. L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.
23. E. P. Simoncelli and J. Portilla. Texture characterisation via joint statistics of wavelet coefficient magnitudes. In *5th International Conference on Image Processing*, 1998.
24. S. W. Zhu, Y. Wu, and D. Mumford. Filters, random fields and maximum entropy (FRAME). *Int'l Journal of Computer Vision*, 27:1–20, 1998.
25. W. P. Ziemer. *Weakly Differentiable Functions.* Springer Verlag, 1989.