Image Statistics for the British Aerospace Segmented Database

Jinggang Huang and David Mumford
Division of Applied Math
Box F, Brown University
Providence, RI02912

Abstract

British Aerospace has segmented a large database of images into 11 sets of regions representing distinct categories of visible objects and surfaces. We examine the nearest neighbor pixel statistics for each category and find major differences between their scaling properties and the shape of the full histogram.

1 Introduction

There has been enormous interest in the statistics of natural images recently. One remarkable property of natural images is that their statistics are nearly scale-invariant under block averaging. In terms of second order statistics, this means the power spectrum of natural images is proportional to $\frac{1}{frequency^2}$. But there is considerable variability in this scaling and in some investigations, the spectrum is better fitted by exponents other than 2. Going beyond second order statistics, one can study the full marginal distribution on the log of the ratio of intensities at adjacent pixels (assuming the images are calibrated, i.e. equal to the energy incident on each sensor). In various studies, these marginals have been computed and again they have shown both similarities and differences. British Aerospace has compiled a database of 214 calibrated RGB images, with 512x768 pixels each, which they have laboriously segmented into pixels representing 11 different parts of the scene (listed below). This allows us to make a precise comparison of scaling properties and of the nearest neighbor marginals for each category. We confirm the existence of systematic differences and make initial steps at quantifying the differences.

2 Observations with the unsegmented database

As mentioned, the database consists of calibrated RGB images. We work only with grey levels, hence we took the weighted average of the three components, with widely used weights: 0.299,0.587,0.114. Since the pixel values are positive and have the dimension of energy, we take logarithms, giving us images I which are dimension-free up to a single additive constant[1]. For each image I, we define the scaled down image $I^{(k)}$ by taking disjoint $k \times k$ blocks and computing the average intensity value of each block. The statistic we investigated here is the horizontal derivative, which for scale k, is simply: $D^{(k)} = I^{(k)}(i, j+1) - I^{(k)}(i, j)$, for any possible i, j pairs. If natural images are scale invariant, $D^{(k)}$ should have the same distribution for all k.

Figure 1: Two images from the image data set, and their segmentations.

In this section, we compare the histogram of D at different scales for the full dataset. First, we claim that D is a very stable statistic, consistent through different data sets. Figure 2 shows the histogram of D from the BA Image Database and from another larger image database acquired by Van Hateren [1]. Although these two statistics came from differ-

[1] In fact, the images had a gamma correction by a square root, so we also multiply by 2.

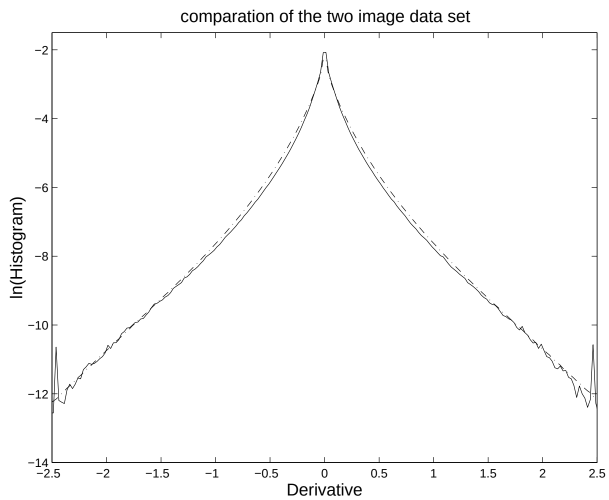

Figure 2: Comparison of ln(Histogram) for the derivative statistic from the British Aerospace Database (solid line) and van Hateren's database (dotted line), both at scale 2

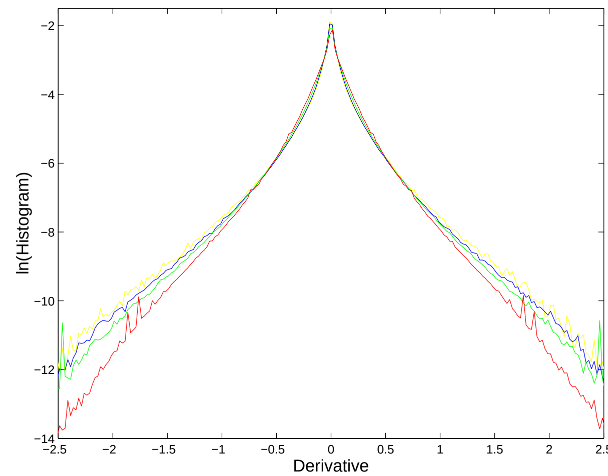

Figure 3: $ln(Histogram)$ of D at different scales: red = scale 1, green = scale 2, blue = scale 4, yellow = scale 8

ent groups, they match very closely. Figure 3 shows the log(histogram) of $D^{(k)}$, for $k = 1, 2, 4, 8$. We can see that, except for the tails, the histograms of $D^{(k)}$ match reasonably well over different scales. This is confirmed by the fact that their standard deviations are nearly equal, see table 2 category 'all';

3 Observations with the segmented database

Using the segmentations of the images, we can look closely at the scale invariance property on each category. Here are the categories and their frequencies:

Category	Frequency	Description
1	10.87	sky, cloud, mist
2	37.62	trees, grass, bush, soil, etc.
3	0.20	road surface marking
4	36.98	road surface
5	6.59	road border
6	3.91	building
7	2.27	bounding object
8	0.11	road sign
9	0.28	telegraph pole
10	0.53	shadow
11	0.64	car

Table 1: Categories and Their Weights

We found that the 'manmade' categories 6,8,9 and 11 have similar histograms and scaling behavior, so we put such categories together. Likewise, the 'vegetation' categories 2 and 7 behave similarly (note that category 7 contains vegetation like hedges, as well as fences, etc.). Figure 4 shows the histogram for these and the other large categories 1 and 4.

Some of these categories scale fairly well, e.g. the manmade one and the vegetation categories, while the sky and road categories do not scale at all well. To make this more precise, we calculate the standard deviation sd_ℓ of D at scale $level\ \ell = \log_2(k)$ for each category. These are given in the table 2 and shown

Category	$\ell = 0$	$\ell = 1$	$\ell = 2$	$\ell = 3$
1	0.07	0.04	0.03	0.03
2,7	0.35	0.34	0.31	0.28
4	0.14	0.12	0.09	0.08
6,8,9,11	0.30	0.35	0.39	0.41
all	0.26	0.26	0.26	0.27

Table 2: Standard Deviation of Different Categories at Several Scale Levels

graphically in figure 5. In figure 5, $\log_2$ of the standard deviation is plotted against ℓ, so the negative of the slope gives us an approximation to what the physicists call the 'anomalous dimension' η. In other words we are fitting $sd_\ell \approx 2^{-\eta\ell} sd_0$. Equivalently, this is fitting

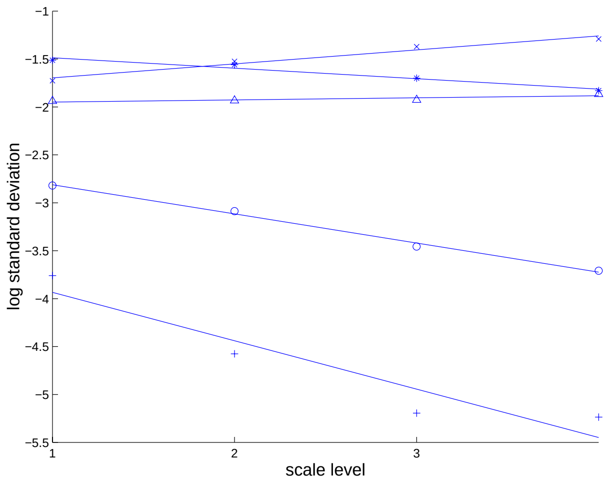

Figure 5: Plot of $\log_2$(standard deviation) vs. scale level and its linear fit in five cases: + = sky, * = vegetation categories, o = road, x = manmade categories and $\triangle$ = all pixels. The slopes are 0.15 for manmade, 0.02 for all, -0.11 for vegetation, -0.30 for road and -0.50 for sky.

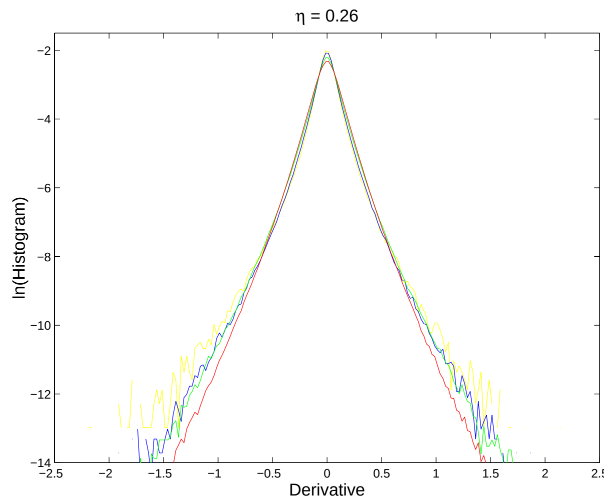

Figure 6: $ln(Histogram)$ of D for category 4 with $\eta = 0.26$.

a model for the second order statistics in which the power spectrum scales as $1/\text{freq}^{(2-2\eta)}$.

4 Discussion

Figure 3 shows that natural images as a single ensemble are very nearly scale invariant. When we look at different categories, however, we find major differences:

1. The vegetation category looks linear in the log plot of the histogram. It scales well although the power spectrum is modeled by $1/\text{freq}^{1.8}$ which is very close to what Ruderman and Bialek found [3]. The log(histogram) can be modeled by $C_1 - C_2|x|$, the 'double exponential' distribution.

2. The manmade category has a histogram with big 'shoulders' in the log plot. The center parts of the histograms match well for different scales, but the tails go up with increasing scale. We believe this phenomena is caused by large objects and their edges. Along an edge the total number of pixel pairs goes down by the factor of 2 when scaling, while the overall number of pairs goes down by the factor of 4. As a result, the frequency of edge pixels increases.

3. In category 1, the density of the distribution mainly concentrated around 0, and shifts further to the center with increasing scale. The scaling fit gives an power spectrum like $1/\text{freq}^{1.0}$.

4. In category 4, the log histogram is slightly concave, and scales badly. However, if we correct for the changing variance, using the assumption that its power spectrum is $1/\text{freq}^{(1.4)}$, we get a much improved fit as shown in figure 6.

To summarize, it appears that the multi-scale families of histograms for each category can be modeled with three parameters: a) their standard deviation, b) the anomolous dimension η and c) a 'shape' parameter for the histogram which we have elsewhere [2] identified as the parameter in an infinitely divisible family of probability distributions. These fits will be developed further later.

5 Acknowledgments

We'd like to thank Andy Wright for allowing us use the British Aerospace Image Database, and answering questions we have about it.

References

[1] J.H. van Hateren and A. van der Schaaf. "Independent Component Filters of Natural Images Compared with Simple Cells in Primary Visual Cortex". Proc.R.Soc.Lond. B 265:359-366,1998

[2] D.Mumford, B.Gidas and S-C Zhu, "Stochastic Models for Generic Images", in preparation.

[3] D.L. Ruderman and W. Bialek. "Statistics of natural images: scaling in the woods", *Phys. Rev. Letter.*, 73:814-817,1994

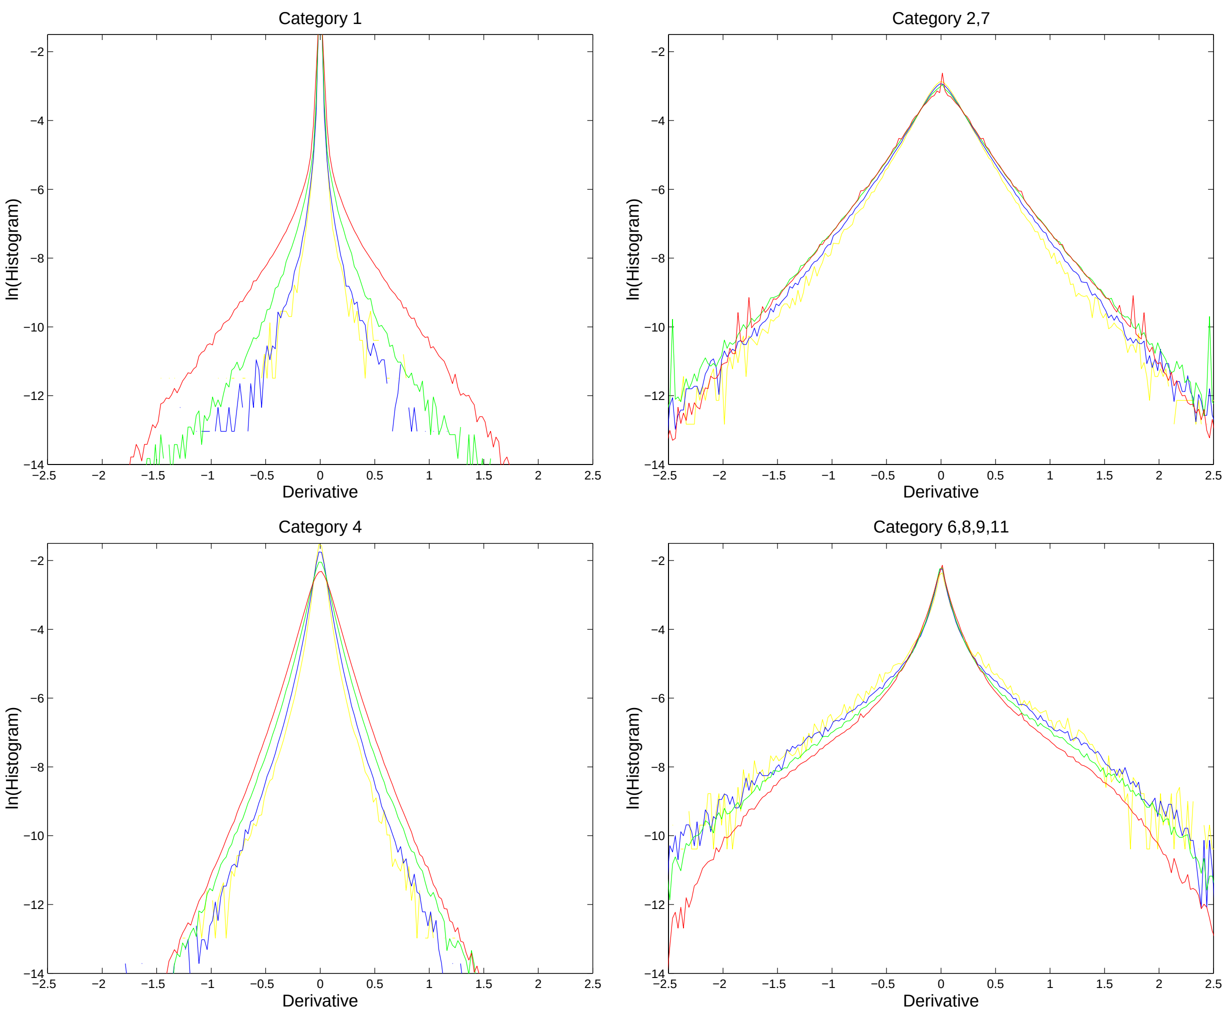

Figure 4: $ln(Histogram)$ of D at different scales and categories: red = scale 1, green = scale 2, blue = scale 4, yellow = scale 8