

Non-linear Filter Response Distributions of Natural Colour Images

Alexander Balinsky¹ and Nassir Mohammad^{1,2}

¹ School of Mathematics, Cardiff University, Cardiff, CF24 4AG, UK

² Hewlett Packard Laboratories, Bristol, BS34 8QZ, UK

{BalinskyA, MohammadN3}@Cardiff.ac.uk,

Nassir.Mohammad@hp.com

Abstract. We observe a non-Gaussian heavy tailed distribution for the non-linear filter

$$\gamma(U)(\mathbf{r}) = U(\mathbf{r}) - \sum_{\mathbf{s} \in N(\mathbf{r})} w(Y)_{\mathbf{r}\mathbf{s}} U(\mathbf{s}), \quad (1)$$

applied to the chromacity channel 'U' (and equivalently to 'V') on individual natural colour images in the colour space YUV. We fit a Generalised Gaussian Distribution (GGD) to the histogram of the filter response, and observe the shape parameter (α) to lie within the range $0 < \alpha < 2$, but rarely $\alpha > 1$.

Keywords: Non-Gaussian Statistics, Image Colorization, Non-Linear Filter Response, Natural Colour Statistics.

1 Introduction

Statistical analysis of natural luminance images have revealed an interesting property: non-Gaussian behaviour of image statistics, i.e. high kurtosis, heavy tails, and sharp central cusps (see e.g. [1], [2], [3], [4], [5], [6]). This property has been extensively studied via the emperical distributions on large databases of natural images, establishing image statistics, under common representations such as wavelets or subspace bases (PCA, ICA, Fishers etc.), as non-Gaussian. For example, a popular mechanism for decomposing natural images locally, in space and frequency, using wavelet transforms leads to coefficients that are quite non-Gaussian with the histograms displaying heavy tails and sharp cusps at the median [7]. In this study we show that this striking phenomenon readily follows across to natural colour images.

Given an RGB image we convert it to the colour space YUV. (The chromacity images U and V are similar and so we only explain our workings for the U component where the exact same procedure is repeated for the V component.) Our filter takes as input the chromacity channel U and the intensity image Y, the proposed filter is given below,

$$\gamma(U)(\mathbf{r}) = U(\mathbf{r}) - \sum_{\mathbf{s} \in N(\mathbf{r})} w(Y)_{\mathbf{r}\mathbf{s}} U(\mathbf{s}), \quad (2)$$

where \mathbf{r} represents a two dimensional point, $N(\mathbf{r})$ a neighborhood (e.g. 3x3 window) of points around \mathbf{r} , and $w(Y)_{\mathbf{rs}}$ a weighting function.

For our purpose we define two weights:

$$w(Y)_{\mathbf{rs}} \propto e^{-(Y(\mathbf{r})-Y(\mathbf{s}))/2\sigma_{\mathbf{r}}^2}, \quad (3)$$

and

$$w(Y)_{\mathbf{rs}} \propto 1 + \frac{1}{\sigma_{\mathbf{r}}^2}(Y(\mathbf{r}) - \mu_{\mathbf{r}})(Y(\mathbf{s}) - \mu_{\mathbf{r}}), \quad (4)$$

where $\mu_{\mathbf{r}}$ and $\sigma_{\mathbf{r}}^2$ are the mean and variance of the intensities in a window around \mathbf{r} .

The proposed filter thus takes a point \mathbf{r} in U and subtracts a weighted average of chromacity values in the neighborhood of \mathbf{r} . The $w(Y)_{\mathbf{rs}}$ is a weighting function that sums to one over s , large when $Y(\mathbf{r})$ is similar to $Y(\mathbf{s})$, and small when the two intensities are different. Filters with weights (3) and (4) have been used in [8] for the colorization problem. These types of filters are compatible with the hypothesis that the essential geometric contents of an image are contained in its level lines (see [9] for more details).

We explain some further symbols that are used in this paper: Assuming X to be a random variable on \mathbb{R} with μ and σ^2 the mean and variance of X , respectively, we define:

$$k = \frac{\mathbb{E}(X - \mu)^4}{\sigma^4}, \quad S = \frac{\mathbb{E}(X - \mu)^3}{\sigma^3},$$

where k is the kurtosis, S is the skewness. For a normal distribution kurtosis and skewness take the values, 3 and 0, respectively. We also note here that our pictures of probability distributions are shown with the vertical scale not probability but *log of probability*. This is very important as it shows the non-Gaussian nature of these probability distributions more clearly, especially the nature of the tails. Subscript notation e.g. k_U and k_V denote statistic values for the 'U' and 'V' filtered components of an image, respectively.

2 Non-linear Filter Response of Individual Images

Figure 1 shows a sample of 8 natural colour images from our dataset of 25 images which are all bitmap uncompressed, captured using a Canon digital SLR camera, and were chosen to cover a wide spectrum of natural scenes in order to give some measure of robustness to our findings. We did not pay too much attention to the methods of capture, or any subsequent re-calibration as we wish to work with colour images captured via any mode, and believe that when images are considered to be natural this will have little effect on the general properties of the filter response.

Applying the non-linear filter on each of the colour channels, U and V , in the image outputs an intensity matrix on which we compute a histogram. We note that application of the filter is only possible within a boundary of the original



(a) balloons



(b) indoors



(c) houses



(d) sky



(e) objects



(f) seaside



(g) night



(h) nature

Fig. 1. Here we display a sample of 8 pictures taken from our dataset of 25 images. In order to give a measure of robustness to our findings we chose pictures covering a wide spectrum of natural scenes, ranging from natural landscapes to urban environments. Images shown here are all truecolour RGB obtained by a Canon digital SLR camera of varying resolutions in uncompressed bitmap format, and reduced to sizes in the region of 200x200 pixels using Adobe photoshop.

image, dependent on the size of the neighborhoods used in the filter construction. In our case the filter was not computed on a one pixel boundary of the image. Outputs of the filter on both the colour channels for two of our sample images, 'balloons' and 'objects', are shown in Figure 2 as grey-scale intensity images.

In [1] they consider the response of derivative filters on calibrated natural luminance images and model the histograms using the GGD. Similarly we fit the following GGD model to our data,

$$f(x) = \frac{1}{Z} e^{-|x/s|^\alpha}, \quad (5)$$




Fig. 2. Filter response of each of the colour channels, U and V, of two of our sample images, 'balloons' and 'objects', using the first weighting function (3)

where Z is a normalising constant so that the integral of $f(x)$ is 1, s the scale parameter and α the shape parameter, these are directly related to the variance and kurtosis by:

$$\sigma^2 = \frac{s^2 \Gamma(\frac{3}{\alpha})}{\Gamma^2(\frac{1}{\alpha})} \quad \text{and} \quad k = \frac{\Gamma(\frac{1}{\alpha}) \Gamma(\frac{5}{\alpha})}{\Gamma^2(\frac{3}{\alpha})}. \quad (6)$$

Special cases of this distribution occur when $\alpha = 1$ or $\alpha = 2$, giving the Laplacian or Gaussian distribution, respectively. We calculate the parameters of the model (numerically) directly from the variance and kurtosis using (6). In [1] they observe that such a calculated model is very close to the best fitting model obtained by minimisation of the mean square error.

Figures 3 and 4 show the histograms of the filter response (using weighting function (3)) on each chromacity channel, U,V, for two of our sample images, with the GGD fitting overlaid. The responses are typically concentrated around zero and highly non-Gaussian, exhibiting large kurtosis and heavy tails, as compared with the normal distribution.

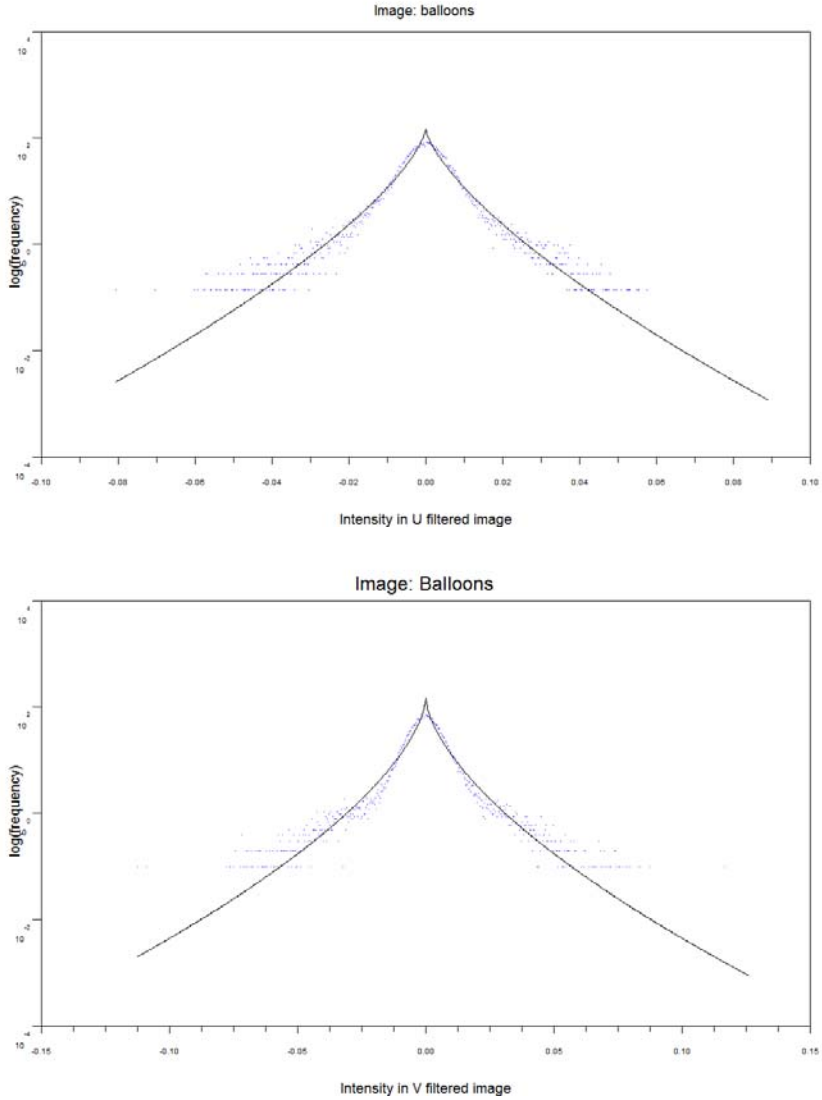


Fig. 3. Distribution of the filter response for both chromacity channels U and V for the image 'balloons' from figure 1 using the first weighting function (3)

Table 1 shows the associated parameters for each filtered image using the first weighting function (3). We observe that kurtosis is greater than that of the normal distribution for all the images considered. α is seen to lie within the range $[0, 1]$, with the parameter varying from image to image, the only exception being the distribution of the U-filtered response for image: 'indoors'. This was the only component of an image to show $\alpha > 1$ in our diverse dataset of images.




Fig. 4. Distribution of the filter response for both chromacity channels U and V for the image 'objects' from figure 1 using the first weighting function (3). The non-Gaussian; high kurtosis, heavy tailed distribution is clearly observed and is typical of the images used in our dataset.

Generally, responses also exhibit some degree of skewness and have very low variance.

Table 2 shows the statistics obtained by filtering the same set of 8 images, but using the second weighting function. Results are similar and show again that the filter response is highly non-Gaussian.

Table 1. Statistics of the non-linear filter response for our sample images using the first weighting function (3)

Image	U filtered response			V filtered response		
	α_U	k_U	S_U	α_V	k_V	S_V
balloons	0.695	11.23	-0.17	0.624	14.23	0.03
indoors	1.11	5.22	-0.05	0.619	14.45	0.16
houses	0.624	14.18	-0.39	0.633	13.74	0.68
sky	0.344	94.00	0.64	0.328	114.87	-2.58
objects	0.54	20.35	0.68	0.662	12.43	-0.08
seaside	0.539	20.44	0.63	0.491	26.60	0.09
night	0.944	6.52	0.03	0.561	18.37	0.13
nature	0.745	9.76	-0.11	0.826	8.11	0.26

Table 2. Statistics of the non-linear filter response for our sample images using the second weighting function (4)

Image	U filtered response			V filtered response		
	α_U	k_U	S_U	α_V	k_V	S_V
balloons	0.685	11.57	-0.19	0.624	14.19	-0.01
indoors	1.094	5.31	-0.07	0.599	15.62	0.16
houses	0.61	14.98	-0.41	0.607	15.14	0.76
sky	0.339	99.38	0.78	0.321	126.27	-2.48
objects	0.534	21.02	0.62	0.654	12.79	-0.07
seaside	0.54	20.39	0.60	0.489	26.84	0.14
night	0.931	6.66	0.07	0.556	18.85	0.18
nature	0.736	10.00	-0.07	0.811	8.37	0.26

We note here that the jpeg standard of image compression and storage is common place and hence we wanted to see how the filter holds under this form of compression. In order to do this we converted samples of the bitmap images from our dataset to the jpeg standard and filtered the images. Results (not shown) were again similar: non-Gaussian, heavy tailed distribution of the filter response on both chromacity channels.

3 Summary

In this paper we observe that conversion of an RGB natural colour image into the space YUV and subsequent performance of the non-linear filter independently on each of the chrominance chanel, U and V, results in a filter response that is highly non-Gaussian, exhibiting heavy tails and large kurtosis. We fit a Generalised Gaussian Distribution to the histogram of the filter response and obtain scale parameters α that vary from image to image, lying within the range $0 < \alpha < 2$, but rarely $\alpha > 1$. In order to show some measure of robustness we used a dataset of images that covered a diverse range of natural scenes, with the

distribution of the filter response always the same, i.e. non-Gaussian with heavy tails. In future work we intend to develop Bayesian analysis of the colorization problem using the GGD (5) as a regularization term. Our results indicate an interesting connection between 'compressive sensing' and the U,V components of natural images.

Acknowledgments

We would like to thank Stephen Pollard and Andrew Hunter from HP Labs, Bristol, UK for providing some of the images used in the study and for fruitful discussion. This work was supported in part by grants from the Engineering and Physical Sciences Research Council (EPSRC) and Hewlett Packard Labs, awarded through the Smith Institute Knowledge Transfer Network.

References

1. Huang, J., Mumford, D.: Statistics of Natural Images and Models. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition, Fort Collins, CO, vol. 1, pp. 541–547 (1999)
2. Mallat, S.G.: A Theory for Multiresolution Signal Decomposition: The Wavelet Representation. *IEEE Trans. Pattern Anal. Machine Intell.* 11, 674–693 (1989)
3. Mumford, D.: Empirical Investigations into the Statistics of Clutter and the Mathematical Models it Leads To, Lecture for the Review of ARO Metric Pattern Theory Collaborative, Brown Univ., Providence, RI (2000)
4. Wainwright, M.J., Simoncelli, E.P., Willsky, A.S.: Random Cascades on Wavelet Trees and their use in Analysing and Modelling Natural Images. *Appl. Comput. Harman, Anal.* 11, 89–123 (2001)
5. Field, D.J.: Relations between the Statistics of Natural Images and the Response Properties of Cortical Cells. *J. Opt. Amer.* 4(12), 2379–2394 (1987)
6. Olshausen, B.A., Field, D.J.: Natural Image Statistics and Efficient Coding. *New. Computation Neural syst.* 7 (1996)
7. Srivastava, A.: Stochastic Models for Capturing Image Variability. *IEEE Signal Processing Magazine* 19(5), 63–76 (2002)
8. Levin, A., Lischinski, D., Weiss, Y.: Colorization using optimization. *ACM Transactions on Graphics* 23(3), 689–694 (2004)
9. Caselles, V., Coll, B., Morel, J.M.: Geometry and color in natural images. *Journal of Mathematical Imaging and Vision* 16(2), 89–105 (2002)