

Scale Invariance and Noise in Natural Images

Daniel Zoran

Interdisciplinary Center for Neural Computation
The Hebrew University of Jerusalem, Israel

daniez@cs.huji.ac.il

Yair Weiss

School of computer science and Engineering
The Hebrew University of Jerusalem, Israel

<http://www.cs.huji.ac.il/~yweiss/>

Abstract

Natural images are known to have scale invariant statistics. While some earlier studies have reported the kurtosis of marginal bandpass filter response distributions to be constant throughout scales, other studies have reported that the kurtosis values are lower for high frequency filters than for lower frequency ones. In this work we propose a resolution for this discrepancy and suggest that this change in kurtosis values is due to noise present in the image. We suggest that this effect is consistent with a clean, natural image corrupted by white noise. We propose a model for this effect, and use it to estimate noise standard deviation in corrupted natural images. In particular, our results suggest that classical benchmark images used in low-level vision are actually noisy and can be cleaned up. Our results on noise estimation on two sets of 50 and a 100 natural images are significantly better than the state-of-the-art.

1. Introduction and Related Work

1.1. Scale Invariance in Natural Images

One of the most striking properties of natural image statistics is their scale invariance [14]. The most notable scale invariant property is the power-law spectrum. When decomposing an image to its local bandpass filter components, the power, or the variance of coefficient distributions decays as a power-law of the form: $P(\omega) = \frac{A}{|\omega|^{2-\eta}}$ where η is usually a small number and ω is the magnitude of the spatial frequency [15]. This property is very robust and holds across different images and scenes. Various other properties of natural images have been shown to be scale invariance as described in [13, 14].

Natural images, in addition to having scale invariant statistics, are also extremely non-gaussian. The distributions of the different Fourier coefficients, for example, have very large peaks, heavy tails and are highly kurtotic. These distributions can be generally well fitted with a generalized gaussian distribution, which captures this distinctive

shape [1]. The kurtosis of a generalized Gaussian distribution is directly dependent on its shape parameter α . Assuming x is generalized Gaussian distributed such that $x \sim GG(\mu, \sigma^2, \alpha)$ where μ is the mean, σ^2 the variance and α is the shape parameter, the kurtosis of x is:

$$\kappa_x(\alpha) = \frac{\Gamma(\frac{1}{\alpha})\Gamma(\frac{5}{\alpha})}{\Gamma(\frac{3}{\alpha})^2} \quad (1)$$

Where Γ is the standard gamma function [4]. As can be seen, the kurtosis is inversely related to the shape parameter α . For natural images, α is usually rather small, having values of between 0.5 and 1 [15].

In light of this, one would expect to see some sort of scale invariance in the kurtosis of marginal coefficient distributions, specifically, a reasonable assumption would be that the kurtosis should be constant throughout scales. This, however, is not always the case. There are inconclusive evidence to whether the kurtosis values change with the scale of the measured filter response distribution. In [6] it has been reported that the kurtosis is constant throughout scales for DCT filters marginal distributions, whereas in [1] it has been reported that the kurtosis changes with scale. Specifically, in [1] it is reported that for higher frequencies, the kurtosis values are lower than for the low frequency ones.

Figure 1 shows two different natural images, one is a natural scene captured in good light conditions and the other is the ubiquitous “Lena” image, which also depicts a natural scene. As can be seen in the figure, while the kurtosis values for the natural image are more or less constant throughout scales, the Lena image displays changes in kurtosis values in different scales. Higher frequency filter responses have less kurtotic distributions than lower frequency ones.

In this work we propose a model which explains this discrepancy. We suggest that the kurtosis of marginal distributions in clean, natural images should be constant throughout scales, and that noise added to the image at various stages of production causes the kurtosis values to change, and violates the scale invariance principle. Figure 2 shows the result of adding noise at different standard deviations to a

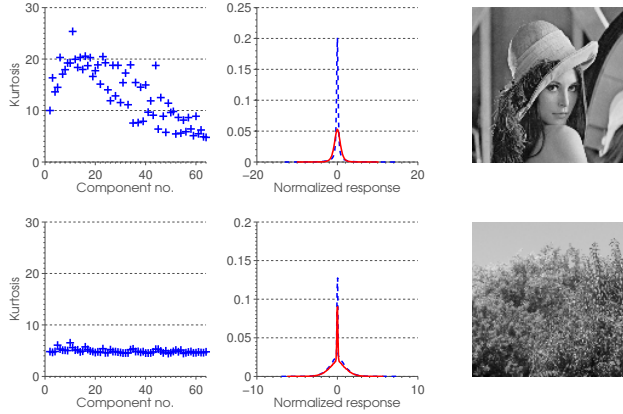


Figure 1: Kurtosis values for two different natural images. Top row on the left is the kurtosis profile - the kurtosis as a function of component number, or frequency, for the Lena image. Kurtosis values for higher frequencies have lower values than for low frequencies. In the middle of the row are the response histograms of two components for this image - the 12th (Low Frequency) and 55th (High Frequency) normalized by their variance. As can be seen, the lower frequency have a higher peak than the high frequency, making it more kurtotic. Bottom row shows the same plot for a clean natural image, taken at good light and downsampled. Kurtosis values are constant throughout scales. This can be seen in the response histogram as well. All results are for 8×8 DCT filters, both images are 512×512 pixels.




Figure 2: Result of adding noise at different standard deviation values to a clean image. The figure depicts the kurtosis of marginal filter response distribution for the DCT filter basis (8×8 pixels) as a function of spatial frequency. As the noise standard deviation rises, the kurtosis values drop, and the shape of the graph distorts with more change in the higher frequencies.

clean, natural image. As the noise standard deviation rises, kurtosis values drop, and more so for higher frequencies.

1.2. Estimating Noise Standard Deviation in Images

Many low-level computer vision algorithms combine the image evidence with a prior or regularization term. The rel-

ative weight of these two terms depends crucially on the observation noise and many computer vision algorithms assume this noise level is given as input to the algorithm [12, 11, 16, 18]. Different noise models are used in different algorithms but by far the most common model for noise is an additive, white Gaussian noise (sometimes referred to as AWGN). There has, however, been much effort to estimate the observation noise automatically. For the case of color images, Liu et al. [8] showed how an assumption of piecewise constant color allows estimating noise from a single image.

For gray level images, the MAD framework [19] uses the deviation from a smooth image model to estimate the noise. Specifically, two state-of-the-art methods [2, 10] take a similar approach. A Laplacian filter is convolved with the image, removing all second order dependencies between neighboring pixels. Since pixels in natural images have very high correlations between neighbours, this effectively removes most of the information in the original image. The only information that remains is at edges in the original image and the noise it self. Estimating the noise variance (or standard deviation) from the Laplacian image usually results in overestimation. This overestimation is due to edges contributing to the overall variance. To compensate for this, [2] apply a non-linear decay function over higher values of the block variance histogram in an iterative manner. [10] take a different approach - a Sobel edge detector is applied to the image, and using an adaptive threshold, edge pixels are marked and removed from the statistics. Both methods work quite well in general, but in images with prevalent edges, they overestimate the noise variance considerably.

The most similar approach to ours is that of Stefano et al. [3]. However, they assume a Laplacian distribution for the marginals, and do not assume anything about scale invariance or the relation between different wavelet coefficients. They do note, however, that their method works best for higher frequencies in which the SNR is lower - where the change in kurtosis is more pronounced.

We propose a method for noise estimation the relies on the assumption that kurtosis values of different scale filter distributions should not change with scale, and that any systematic change in these values is due to added noise. Our method performs much better on low-noise corrupted, elaborate natural image than current methods, and is comparable to other methods when the noise is higher, or the image simpler.

2. Model

2.1. Noise and the Generalized Gaussian Distribution

We start by modeling the change in kurtosis of a generalized Gaussian distributed random variable due to added

Gaussian noise. Denote x a generalized Gaussian random variable such that $x \sim GG(\mu, \sigma_x^2, \alpha)$ and denote η an independent Gaussian random variable with zero mean and variance σ_n^2 . Denote y a random variable such that:

$$y = x + \eta$$

We wish to calculate the kurtosis κ_y of y . Going back to images, x would represent the original distribution for a local coefficient (for example) in the clean image, η the noise added and y the measured coefficient distribution from the noise corrupted image.

In this work, we refer to kurtosis as the fourth central moment normalized by the variance squared, or:

$$\kappa = \frac{\mu_4}{\sigma^4}$$

This is not the same as *excess kurtosis* which also includes a -3 term that makes the kurtosis of the Gaussian distribution zero. Due to the independence of noise, the variance of y is simply the sum of σ_{GG}^2 and σ_n^2 or:

$$\sigma_y^2 = \sigma_x^2 \left(1 + \frac{\sigma_n^2}{\sigma_x^2} \right)$$

The fourth central moment of y is easy to calculate using the cumulants and independence of noise:

$$\mu_4(y) = 3\sigma_x^4 \left(1 + \frac{\sigma_n^2}{\sigma_x^2} \right)^2 + \sigma_x^4 (\kappa_x(\alpha) - 3)$$

Where $\kappa_x(\alpha)$ is defined in Eq. 1. Finally, by normalizing with the squared variance calculated above we get:

$$\kappa_y = \frac{\kappa_x(\alpha) - 3}{\left(1 + \frac{\sigma_n^2}{\sigma_x^2} \right)^2} + 3 \quad (2)$$

Using this result we can predict the kurtosis of a marginal filter response distribution taken from a noise corrupted image, given the original image. This, however, is not usually the case as the original image is typically not available. In the next section we describe a method to estimate σ_n^2 using measurements of the noisy variable y and the assumption that κ_x is unknown, but constant throughout measurements of x for different scales.

2.2. Noise Estimation using Scale Invariance

At the base of our noise estimation procedure is the assumption that the original, uncorrupted image had scale invariant statistics. Specifically, we assume that the kurtosis of marginal filter response distributions for the original image is an unknown constant, and that adding noise to the image resulted in changes to kurtosis values throughout scales for the corrupted image.




Figure 3: Function space for Eq. 3. The minimum is the actual point found numerically. Noise added to the image had a standard deviation of 10.

The first step of the algorithm is gathering statistics over the noise corrupted image I_n . We convolve the image with each filter from the $N \times N$ DCT basis, to produce a response image for the y_i for the i -th filter. We estimate the variance and kurtosis for this response image to obtain $\hat{\sigma}_{y_i}^2$ and $\hat{\kappa}_{y_i}$. We do this procedure for every component i in the range $2..N^2$, hence ignoring the DC component.

Given these variance and kurtosis measures, we wish to estimate the variance of the added noise. This is done by finding the pair $\hat{\kappa}_x$, the kurtosis of the original uncorrupted image distribution and $\hat{\sigma}_n^2$ the variance of the noise, which minimizes:

$$\hat{\kappa}_x, \hat{\sigma}_n^2 = \arg \min_{\kappa_x, \sigma_n^2} \sum_{i=2}^{N^2} \left| \frac{\kappa_x - 3}{\left(1 + \frac{\sigma_n^2}{\hat{\sigma}_{y_i}^2 - \sigma_n^2} \right)^2} + 3 - \hat{\kappa}_{y_i} \right| \quad (3)$$

Figure 3 shows an example of what the function space look like, when the noise added to the image has a standard deviation of 10. As can be seen, there is a rather pronounced valley at the minimum point (which the was numerically found).

2.3. Non Gaussian Noise

Although the above model assumes white Gaussian noise, it assumes that it is white and Gaussian in the filter domain only. Since many types of independent noise in the pixel domain will mix in to Gaussian noise in the frequency domain, this method work with other types of noise. The summation over the noise in pixels while calculating the response image for the DCT basis causes the distribution of the sum to be Gaussian - due to the central limit theorem and noise independence, an example can be shown in Figure 4, even though the noise is very non Gaussian in the pixel domain, it becomes Gaussian in the coefficient domain. In Section 3 we show that adding uniform noise to images, for example, does not change the performance of our method.

Other image corruption methods, such as JPEG compression artifacts, quantization noise and sensor noise were

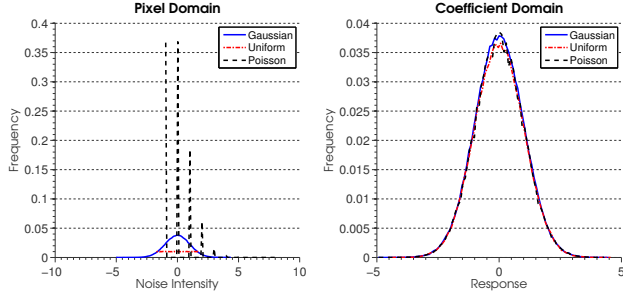


Figure 4: Response histograms of independent Gaussian, Uniform and Poisson pixel noise. Noise images were 512×512 pixels, having variance 1 and mean 0. Coefficient domain histograms made using an 8×8 DCT filter. It can be clearly seen that while the three look very different at the pixel domain, all three are mixed to be Gaussian with the same variance and mean in the coefficient domain.

tested. Results are shown in the next section.

3. Results

3.1. Methods

We first compared the proposed method with existing methods by synthetically adding noise to clean images. We minimized equation 3 using MATLAB’s `fminsearch` function.

3.2. Noise Estimation Results

When the original image is relatively simple, such as the one depicted in Figure 5 algorithms perform similarly, estimating the noise variance well for a range of values. However, when using a more complex image, with a lot of textured areas such as the one depicted in Figure 6, the picture change. All results are means over 3 noise realizations - standard deviations were not included in the graphs because they were too low for all methods to be visible in the graph.

Due to the large amount of edges in the image in Figure 6, the methods in [2, 10] over estimate the noise greatly. Our method performs better on the low to medium noise regime (σ_n between 1 to 15). As the noise levels rise, it is harder for our method to accurately estimate the noise. The reason for overestimation of the noise SD in other methods is that the textured areas have a lot of edges, causing the Laplacian filtering be less effective at separating the noise from the original image. Since the textured areas are from a natural image and have scale invariant statistics, our methods handles this quite well, and enables us to recover the noise SD more precisely. When the noise is sufficiently large, the variance contributed from the edges is small relatively to the noise variance, hence the similar performance at this regime. Table 1 summarizes results for all experi-

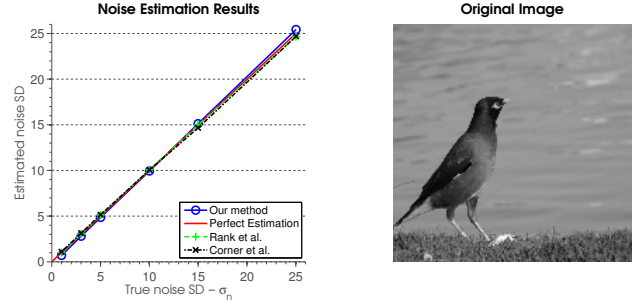


Figure 5: Noise estimation on a simple image, with no prevalent edges (BIRD). All methods estimate the noise correctly for a range of values.

ments. In the results table the estimation results for uniform noise can be observed. It seems that all methods handle this type of noise relatively well. Using different patch sizes for our method resulted in little difference in performance, it seems that the central limit theorem comes in to play even on small patch sizes (8×8 in this case, 64 pixels are more than enough).

Finally, we estimate noise in 50 images from the Van Hateren natural image database [17] as well as for a 100 images from the Berkeley database [9]. White Gaussian noise was added to each of the images and noise was estimated from them. It is obvious our method is on par or better than other methods for all noise levels, for both databases. Calculating the relative error for all noise levels over the van Hateren database we obtain a mean error rate of 3%, while other methods obtain 22% for Rank et al. and 28% for Corner et al. For the Berkeley our method obtains an average error that is 9%, while other methods obtain 49% (Rank et al.) and 65% (Corner et al.). Differences were even larger when neglecting the lowest noise level estimation (in which even a small error in estimation leads to a large relative error, results are then 1% for our method and 20% (Rank et al.) and 22% (Corner et al.)). Results are shown in Table 1.

3.3. Other types of noise and corruption scenarios

We tested the proposed method under several other image corruption scenarios. Under all scenarios under which there’s no clear parameter for the noise standard deviation σ_n , we calculate the standard deviation of the difference between the original image I and the corrupted image I_n such that:

$$\sigma_n = \langle |I - I_n| \rangle \quad (4)$$

Where the mean is over all the pixels in the images.

The first scenario we tested is first motion blurring an image, and then adding Gaussian noise to it. This scenario is important as correct estimation of noise prior to image deblurring is important to minimize common artifacts [7].

BIRD				FIELD			Van Hateren			Berkeley		
σ_n	$\hat{\sigma}_n$	$\hat{\sigma}_n^R$	$\hat{\sigma}_n^C$	$\hat{\sigma}_n$	$\hat{\sigma}_n^R$	$\hat{\sigma}_n^C$	$\hat{\sigma}_n$	$\hat{\sigma}_n^R$	$\hat{\sigma}_n^C$	$\hat{\sigma}_n$	$\hat{\sigma}_n^R$	$\hat{\sigma}_n^C$
1	0.68	1.01	1.1	2.06	7.16	8.3	0.89 ± 0.71	2.02 ± 1.19	2.25 ± 1.1	1.5 ± 1.8	2.9 ± 2.8	3.7 ± 2.8
3	2.8	3.05	3.13	3.67	8.04	8.76	2.92 ± 0.62	3.64 ± 0.89	3.75 ± 0.72	2.9 ± 1.8	4.5 ± 2.3	4.9 ± 2.3
5	4.83	5.11	5.14	5.52	9.43	9.56	4.96 ± 0.59	5.47 ± 0.72	5.55 ± 0.52	4.9 ± 1.9	6.4 ± 2	6.6 ± 1.9
10	9.98	10.11	10.08	10.35	13.56	12.89	10.13 ± 0.89	10.2 ± 0.49	10.2 ± 0.48	9.7 ± 1.8	11.1 ± 1.6	11.1 ± 1.3
15	15.15	14.97	14.98	15.34	17.9	17.05	15.32 ± 1.13	15 ± 0.38	14.87 ± 0.62	14.7 ± 1.8	15.9 ± 1.3	15.7 ± 1.1
25	25.47	24.75	23.73	25.57	26.8	26.27	26.33 ± 1.81	24.65 ± 0.28	23.9 ± 1.47	24.8 ± 2.0	25 ± 1.0	25 ± 1.2
$\langle \epsilon \rangle$	7%	1%	4%	24%	155%	176%	3%	22%	28%	9.8%	49%	65%

Table 1: Results summary for images and methods presented, for each image the first column is our method estimation $\hat{\sigma}_n$, the second Rank et al. $\hat{\sigma}_n^R$ and finally Corner et al. $\hat{\sigma}_n^C$. Last row is the average relative estimation error. On 50 images from the Van Hateren natural image database we obtain an average 3% error rate, while current state-of-the-art method obtain 22% and 28%. Results for the Berkeley database (100 images) are similar, our method out-performs current state-of-the-art methods.




Figure 6: Noise estimation on a more complex image having a lot edges and texture data (FIELD). With low noise standard deviation other methods perform poorly, overestimating the noise by a large percentage. Our method performs much better, estimating the noise standard deviation much better.

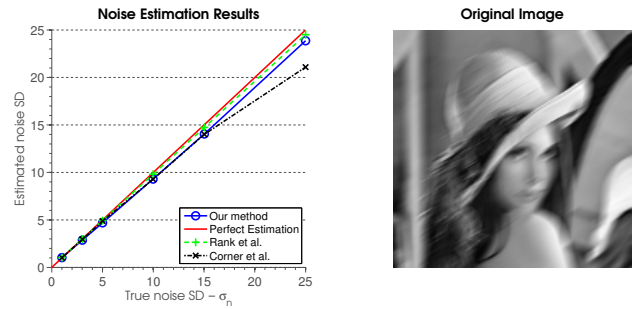


Figure 7: Noise estimation under motion blur. The image was first blurred with a motion blur mask and then corrupted by noise. Noise was then estimated by several methods. It seems that all methods estimate the noise well for a range of values.

Figure 7 shows the result of such a scenario - it seems that none of the methods is affected in any way by the blurring operation, and all methods perform well in estimating the noise standard deviation.

Second we tested whether corruption due to quantization can be estimated using this method. We quantized 256 gray scale levels images to 128, 64, 32, 16, 8 and 4 levels, noise standard deviation was estimated as in Equation 4. From the quantized image we tried to estimate this standard deviation. The proposed method works quite well, results can be seen in Figure 8.

Third, a more realistic noise scenario was used. As was done in [8], we obtained a Camera Response Function (or CRF) from [5]. With this CRF we inversely mapped an

intensity image into a lightness image, added noise to the lightness image and mapped the lightness image back to an noisy, saturated, intensity image. This effectively simulates the sensor noise of a digital camera for a single channel. Results of noise estimation can be seen in Figure 9. All methods slightly underestimate the noise in the image. The reason for this is the clipping of the noise due to the CRF's saturation reduces variance near the extreme values (near 0 and near 255).

Finally, we tested whether JPEG corruption can be estimated using the proposed method. A clean image was compressed in the JPEG format using different quality levels. The compressed images were then used to estimate the corruption in them and compared to the standard deviation of the difference image. Results were very poor for all meth-

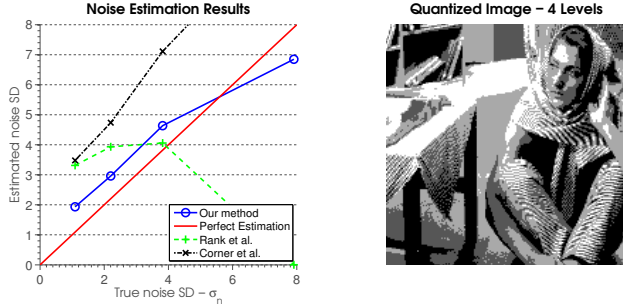


Figure 8: Noise estimation under quantization. The 256 gray scale level image was quantized to 128, 64, 32, 16, 8, and 4 levels. The standard deviation of the difference image between the original and quantized image was used as the true noise standard deviation. Estimation was done directly on the quantized image. Our method out-performs other methods on all quantization levels.

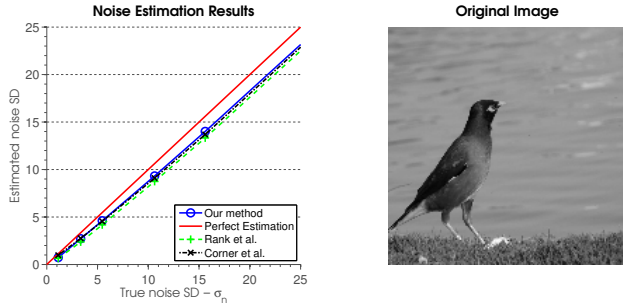


Figure 9: Noise estimation under simulated sensor noise. All methods slightly underestimate the noise in the image, due to the saturation caused by the CRF.

ods. It seems that this kind of artifacts disrupt the kurtosis of marginal filter distributions in such a way that this method can not handle properly. Specifically, since JPEG compression is primarily a frequency domain operation, the distributions change radically, having very extreme kurtosis values which are inconsistent with the above model.

3.4. Uncovering the original Lena

One particularly interesting example is the famous “Lena” image. The Lena image is very old - photographed in 1972 and scanned in 1973. One can only assume that scanning and printing technology of these days were not of the highest grade, and noise has probably corrupted the image to some extent. Looking at the kurtosis values, this is clearly evident as was shown in Figure 1. It would be interesting to see if one can estimate the noise in the original Lena image, and maybe denoise it, uncovering the original, clean image.

Estimating the inherent noise in the original, uncorrupted

image by our method yields a standard deviation of $\hat{\sigma}_n = 2.08$. Other methods yield a bit more - this implies that Lena in itself is noisy, as the original kurtosis profile hints at. Taking this into account, it seems that in denoising experiments with Lena, which are very common [11] there is a maximal PSNR value that can be taken into account. When measuring the PSNR between the original image and the denoised image in a denoising experiment, any result above 41.76dB might not be possible without over fitting. It can certainly be the case that a denoising algorithm will discard some of the inherent noise in the original image, and that when measuring the PSNR value with the original image one will get a *lower* value for *cleaner* images. Of course, the inherent noise in Lena is not necessarily additive, white or Gaussian so the statement above should be taken with a grain of salt, but nevertheless, the noise is present.

Assuming that Lena is noisy, can we uncover the original, uncorrupted Lena image? We applied a simple Bayesian denoising scheme which assumes a Generalized Gaussian model for marginal filter distributions. We use our proposed noise estimation method to estimate the noise standard deviation in the image, and also the constant kurtosis value underlying in the original image. Using these two parameters, we estimated the MAP value of local DCT coefficients for all 8×8 patches of the images, and then averaged the inverse DCT results to obtain a denoised image. Results can be seen in Figure 10. Not surprisingly, estimating the noise standard deviation of the denoised image using our method yields a very low value (0.000001). This reason for this is obvious when looking at the kurtosis of the denoised image - as can be seen in Figure 11 it is almost constant throughout scales. Now we can measure the PSNR value between the original (which has noise in it) and the denoised version (which is assumed to be noise free) - this yield a PSNR of 42.1dB, consistent with what we suggest above. Denoising with a second algorithm, BRFOE by Weiss and Freeman [18] yielded similar results (see Figure 10). This is interesting because this denoising model does not assume scale invariance, and yet still, noise estimation by our method yields a very low value, since the kurtosis in the denoised image is rather constant.

4. Discussion and Future Work

In this work we describe and explain a baffling phenomena. When measuring the kurtosis of marginal filter response distributions in natural images, in many (but not all) natural images, values of kurtosis for lower frequency filters are higher than high frequency ones. This is in contrast to the scale-invariant nature of natural images. We argue that clean, natural images should have a constant kurtosis value throughout scales, and propose that deviations from this are due to noise inherent in the image. Using this assumption we show how the noise level in a corrupted image can be




Figure 10: Denoising results for the original Lena image using simple scale-invariant Bayesian and BRFOE denoising. On the top left is the original image (detail), on the top right is the denoised image (BRFOE). Details are fully preserved in the denoised image, but noise is much less apparent. MAP denoising is on bottom left. Difference image is scaled, and with BRFOE image.




Figure 11: Kurtosis plot for the denoised Lena image and original. An almost constant kurtosis throughout most of the scales is apparent in the denoised image. The dashed line shows the Kurtosis estimated by our noise estimation algorithm. Left is the result from scale-invariant Bayesian denoising, right is the BRFOE result.

accurately estimated, for a range of scene types and noise levels, and under different corruption scenarios.

A particularly intriguing example is the ubiquitous Lena image - a very common benchmark image - which we show here to be noisy. We showed that this image is noisy in its original form, and using it as "ground truth" in low-level vision experiments should be done with caution.

Future work will include several directions. The first is investigating what other types of image corruption come

into play when examining the kurtosis of marginal distributions. Second, handling non-white noise should be relatively simple as long as some model for the noise power spectrum is assumed. Finally, there's still the possibility that the kurtosis profile should have another kind of scale invariance, power-law or other, but not necessarily constant (which is a special case of power-law). Extending this method to include power-law is possible, but requires further work as merely adding another parameter to the minimization or model will not work.

References

- [1] M. Bethge. Factorial coding of natural images: how effective are linear models in removing higher-order dependencies? *23(6):1253–1268*, June 2006.
- [2] B. Corner, R. Narayanan, and S. Reichenbach. Noise estimation in remote sensing imagery using data masking. *International Journal of Remote Sensing*, 24(4):689–702, 2003.
- [3] A. De Stefano, P. R. White, and W. B. Collis. Training methods for image noise level estimation on wavelet components. *EURASIP J. Appl. Signal Process.*, 2004(1):2400–2407, 2004.
- [4] J. Domínguez-Molina, G. González-Farías, R. Rodríguez-Dagnino, and I. Monterrey. A practical procedure to estimate the shape parameter in the generalized Gaussian distribution. *technique report I-01-18_eng. pdf*, available through http://www.cimat.mx/reportes/enlinea/I-01-18_eng. pdf.
- [5] M. Grossberg and S. Nayar. What is the Space of Camera Response Functions? In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume II, pages 602–609, Jun 2003.
- [6] E. Lam and J. Goodman. A mathematical analysis of the dct coefficient distributions for images. *Image Processing, IEEE Transactions on*, 9(10):1661–1666, Oct 2000.
- [7] A. Levin, R. Fergus, F. Durand, and W. T. Freeman. Image and depth from a conventional camera with a coded aperture. In *SIGGRAPH '07: ACM SIGGRAPH 2007 papers*, page 70, New York, NY, USA, 2007. ACM.
- [8] C. Liu, R. Szeliski, S. Bing Kang, C. L. Zitnick, and W. T. Freeman. Automatic estimation and removal of noise from a single image. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(2):299–314, 2008.
- [9] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. 8th Int'l*

- Conf. Computer Vision*, volume 2, pages 416–423, July 2001.
- [10] K. Rank, M. Lendl, and R. Unbehauen. Estimation of image noise variance. In *Vision, Image and Signal Processing, IEE Proceedings-*, volume 146, pages 80–84, 1999.
 - [11] S. Roth and M. Black. Fields of Experts: A Framework for Learning Image Priors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, page 860. IEEE Computer Society; 1999, 2005.
 - [12] S. Roth and M. Black. On the Spatial Statistics of Optical Flow. *International Journal of Computer Vision*, 74(1):33–50, 2007.
 - [13] D. L. Ruderman. Origins of scaling in natural images. *Vision Research*, 37:3385–3398, 1997.
 - [14] D. L. Ruderman and W. Bialek. Statistics of natural images: Scaling in the woods. In *NIPS*, pages 551–558, 1993.
 - [15] A. Srivastava, A. B. Lee, E. P. Simoncelli, and S.-C. Zhu. On advances in statistical modeling of natural images. *J. Math. Imaging Vis.*, 18(1):17–33, 2003.
 - [16] M. Tappen and W. Freeman. Comparison of graph cuts with belief propagation for stereo, using identical MRF parameters. In *IEEE International Conference on Computer Vision*, volume 2, pages 900–906, 2003.
 - [17] J. van Hateren. Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings of the Royal Society B: Biological Sciences*, 265(1394):359–366, 1998.
 - [18] Y. Weiss and W. Freeman. What makes a good model of natural images? *Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on*, pages 1–8, June 2007.
 - [19] V. Zlokolica, A. Pizurica, and W. Philips. Noise Estimation for Video Processing Based on Spatio-Temporal Gradients. *IEEE SIGNAL PROCESSING LETTERS*, 13(6):337, 2006.