IEEE Transactions on Image Processing, 10(2):XXX-XXX 2001.

Blind Inverse Gamma Correction

Hany Farid†

The luminance non-linearity introduced by many imaging devices can often be described by a simple point-wise operation (gamma correction). This paper presents a technique for blindly estimating the amount of gamma correction in the absence of any calibration information or knowledge of the imaging device. The basic approach exploits the fact that gamma correction introduces specific higher-order correlations in the frequency domain. These correlations can be detected using tools from polyspectral analysis. The amount of gamma correction is then estimated by minimizing these correlations.

Keywords: higher-order statistics, image processing, image restoration, blind inverse gamma correction

1 Introduction

The luminance non-linearity introduced by many imaging devices can often be described with a simple point-wise operation (gamma correction) of the form:

$$g(u) \quad = \quad u^{\gamma}, \tag{1}$$

where $u \in [0, 1]$ denotes the image pixel intensity. If the value of γ is known then inverting this process is trivial:

$$g^{-1}(u) \quad = \quad u^{1/\gamma}. \tag{2}$$

The value of γ is typically determined experimentally by passing a calibration target with a full range of known luminance values through the imaging system (e.g., a Macbeth chart [1]). But often such calibration is not available or direct access to the imaging device is not possible, for example when downloading an image from the web. In addition, most commercial digital cameras dynamically vary the amount of gamma. For many applications in digital photography, image processing, and computer vision it would be advantageous to remove these non-linearities prior to subsequent processing stages.

In this paper a technique is presented for estimating the amount of gamma correction in the absence of any calibration information or knowledge of the imaging device. We term this technique blind inverse gamma correction. The basic approach exploits the fact that gamma correction introduces specific higher-order correlations in the frequency domain. These correlations can be detected using tools from polyspectral analysis. The amount of gamma correction is then determined by minimizing these correlations.

Higher-order statistics have previously been used for various forms of image restoration: noise removal [2], deblurring [3, 4], and speckle removal [5]. See also [6, 7, 8] for general discussions on the use of higher-order statistics in image processing. Previous work in this area however has not addressed the issue of inverting the effects from gamma correction.

† 6211 Sudikoff Laboratory, Computer Science Department, Dartmouth College, Hanover, NH 03755.3510. tel: 603.646.2761 fax: 603.646.1672 email: farid@cs.dartmouth.edu.

Insight is gained into the proposed technique by considering a simple one-dimensional signal composed of a sum of two zero-phase sinusoids:

$$y(n) \quad = \quad a_1 \sin(\omega_1 n) + a_2 \sin(\omega_2 n). \tag{3}$$

When this signal is gamma corrected as in Equation (1) new harmonics are introduced with amplitudes and phase that are correlated to the original harmonics [9]. To see this more explicitly $g(\cdot)$ is rewritten in terms of it's Taylor series expansion:

$$g(u) = g(u_0) + \frac{g'(u_0)(u-u_0)}{1!} + \frac{g''(u_0)(u-u_0)^2}{2!} + ... \tag{4}$$

Considering only the first three terms of this expansion, dropping the various scalar constants, and rewriting using basic trigonometric identities yields:

$$\begin{aligned} g(y(n)) &= y(n) + y^2(n) \\ &= a_1 \sin(\omega_1 n) + a_2 \sin(\omega_2 n) \\ &+ \tfrac{1}{2}a_1^2(1 + \sin(2\omega_1 n)) + \tfrac{1}{2}a_2^2(1 + \sin(2\omega_2 n)) + 2a_1 a_2 \sin(\omega_1 + \omega_2) + 2a_1 a_2 \sin(\omega_1 - \omega_2). \end{aligned} \tag{5}$$

Notice the presence of several new harmonics, $2\omega_1$, $2\omega_2$, $\omega_1 + \omega_2$ and $\omega_1 - \omega_2$. Notice also that the amplitudes of these new harmonics are correlated to the amplitudes of the original harmonics. For example, the amplitude of $\omega_1 + \omega_2$ is proportional to the product of the amplitudes of ω_1 and ω_2. For purposes of exposition we have considered a simplified form of a signal and point-wise non-linearity. However, this model generalizes to more complex signals and point-wise non-linearities. Note also that a linear $g(u)$ has no effect on the higher-order correlations, i.e., the Taylor series expansion contains no terms beyond the first two linear terms.

We will show empirically that when an image is gamma corrected, higher-order correlations in the frequency domain monotonically increase proportional to the deviation of γ from unity. As such, the value of γ can be determined by simply minimizing these correlations. We first show how tools from polyspectral analysis can be used to detect the presence of these higher-order correlations, and then show the efficacy of this technique to the blind inverse gamma correction of synthetic and natural images.

2 Bispectral Analysis

Consider a one-dimensional signal $y(n)$, and its Fourier transform:

$$Y(\omega) \quad = \quad \sum_{k=-\infty}^{\infty} y(k) e^{-i\omega k}. \tag{6}$$

It is common practice to use the power spectrum to estimate second-order correlations:

$$P(\omega) \quad = \quad \mathcal{E}\left\{Y(\omega)Y^*(\omega)\right\}, \tag{7}$$

where $\mathcal{E}\{\cdot\}$ is the expected value operator, and * denotes complex conjugate. However the power spectrum is blind to higher-order correlations of the sort introduced by a non-linearity, Equation (5). These correlations can however be estimated with higher-order spectra (see [10] for a thorough survery). For example the bispectrum estimates third-order correlations and is defined as:

$$B(\omega_1, \omega_2) \quad = \quad \mathcal{E}\left\{Y(\omega_1)Y(\omega_2)Y^*(\omega_1 + \omega_2)\right\}. \tag{8}$$

Note that unlike the power spectrum the bispectrum of a real signal is complex-valued. Comparing the bispectrum with Equation (5) we can see intuitively that the bispectrum reveals the sorts of higher-order correlations introduced by a non-linearity. That is, correlations between harmonically related frequencies, for example, $[\omega_1, \omega_1, 2\omega_1]$ or $[\omega_1, \omega_2, \omega_1 + \omega_2]$. The bispectrum can be estimated by dividing the signal, $y(n)$, into N (possibly overlapping) segments, computing Fourier transforms of each segment, and then averaging the individual estimates:

$$\hat{B}(\omega_1, \omega_2) = \frac{1}{N} \sum_{k=1}^{N} Y_k(\omega_1) Y_k(\omega_2) Y_k^*(\omega_1 + \omega_2), \tag{9}$$

where $Y_k(\cdot)$ denotes the Fourier transform of the k^{th} segment. This arithmetic average estimator is unbiased and of minimum variance. However, this estimator has the undesired property that its variance at each bi-frequency (ω_1, ω_2) depends on $P(\omega_1)$, $P(\omega_2)$, and $P(\omega_1 + \omega_2)$ (see e.g., [11]). We desire an estimator whose variance is independent of the bi-frequency. To this end, we employ the bicoherence, a normalized bispectrum, defined as:

$$b^2(\omega_1, \omega_2) = \frac{|B(\omega_1, \omega_2)|^2}{\mathcal{E}\{|Y(\omega_1)Y(\omega_2)|^2\}\mathcal{E}\{|Y(\omega_1 + \omega_2)|^2\}}. \tag{10}$$

It is straight-forward to show using the Schwartz inequality that this quantity is guaranteed to have values in the range $[0, 1]$. As with the bispectrum, the bicoherence can be estimated as:

$$\hat{b}(\omega_1, \omega_2) = \frac{|\frac{1}{N}\sum_k Y_k(\omega_1) Y_k(\omega_2) Y_k^*(\omega_1 + \omega_2)|}{\sqrt{\frac{1}{N}\sum_k |Y_k(\omega_1)Y_k(\omega_2)|^2 \frac{1}{N}\sum_k |Y_k(\omega_1 + \omega_2)|^2}}. \tag{11}$$

Note that the bicoherence is now a real-valued quantity. This quantity is used throughout this paper to measure higher-order correlations.

3 Results

For simplicity, it is assumed that gamma correction can be modeled with the one-parameter family of functions: $g(u) = u^\gamma$, where $u \in [0, 1]$ denotes the image pixel intensity. Note that the scaling of intensities into this range has no effect on the bicoherence [10]. Given only the gamma corrected image, our task is to determine the value of γ. To accomplish this we apply a range of inverse gamma values $g^{-1}(u) = u^{1/\gamma}$ to the image and select the value of γ that minimizes the following measure of third-order correlations:

$$\sum_{\omega_1=-\pi}^{\pi} \sum_{\omega_2=-\pi}^{\pi} |\hat{b}(\omega_1, \omega_2)| \tag{12}$$

where $\hat{b}(\omega_1, \omega_2)$ is the bicoherence defined in Equation (11). In practice this simple search strategy is effective because the function being minimized is typically well-behaved, i.e., contains a single minimum. This is illustrated in Figure 3 where the bicoherence is plotted as a function of varying inverse gamma values. In this example a 1-D fractal signal $y(n)$ with 512 samples, a $1/\omega$ power spectrum and random phase was subjected to a gamma of $\gamma = 1.5$. Note that the bicoherence reaches a unique minimum at $\gamma^{-1} = 1.5$.

To avoid the memory and computational demands of computing an image's full four-dimensional bicoherence, our analysis is restricted to the one-dimensional horizontal scan lines of an image (vertical scan lines could have equally well been used). This is reasonable since gamma correction is a point-wise non-linearity and the correlations introduced in 1-D will be similar to those in 2-D. The amount of

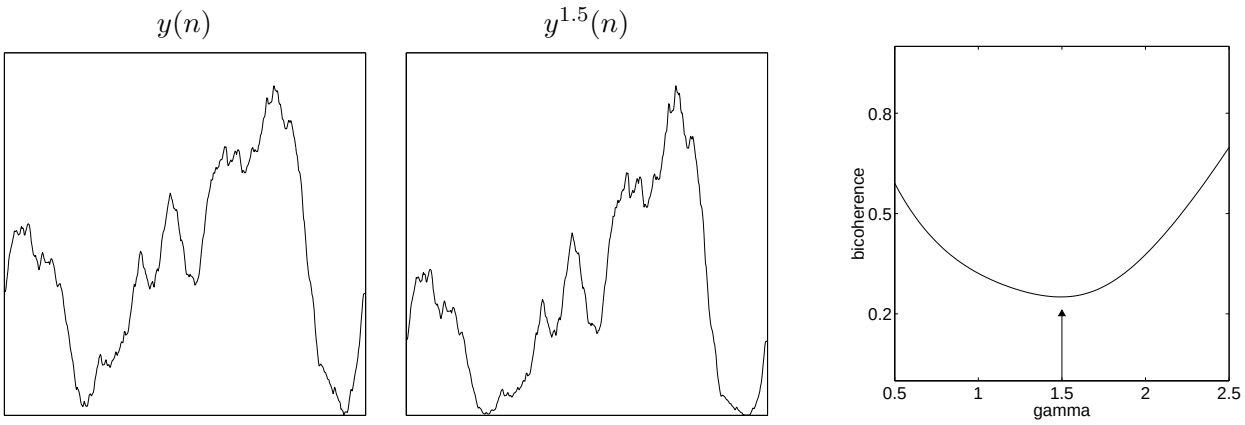

Figure 1: Shown on the left is a 1-D fractal signal $y(n)$ and this signal after gamma correction $y^{1.5}(n)$. Shown on the right is the average bicoherence (Equation (12)) computed for a range of inverse gamma values, $[y^{1.5}(n)]^{\gamma^{-1}}$. The bicoherence reaches a unique minimum at $\gamma^{-1} = 1/1.5$.

gamma is blindly estimated for an *image* by averaging over the estimates for each horizontal scan line (or a subset of them). The amount of gamma for a *scan line* is estimated by a brute force search for the inverse gamma value that minimizes the average bicoherence. In the results reported here, the bicoherence for each 1-D image slice $y(n)$ is computed by dividing the signal into overlapping segments of length 64 with an overlap of 32. A 128-point windowed DFT $Y_k(\omega)$ is estimated [1] for each segment from which the bicoherence is estimated, $\hat{b}(\omega_1, \omega_2)$ as in Equation (11). There is a natural tradeoff between segment length and the number of samples from which to average. We have empirically found that these parameters offer a good compromise, however their precise choice is not critical to the estimation results.

3.1 Synthetic Images

Fractal images of size 512 $\times$ 512 with a $1/\omega$ power spectrum and random phase are subjected to gamma correction with $\gamma \in [0.2, 2.8]$, Equation (1). The amount of gamma is estimated by averaging over the estimates from every sixteenth horizontal scan line. A range of possible gamma values from 0.1 to 3.8 are sampled in increments of 0.1. Shown in Figure 3.1 are the estimated values averaged across ten independent images. Also shown in Figure 3.1 is the estimated gamma plotted as a function of the actual gamma where each data point corresponds to one of the ten images. On average, the correct gamma is estimated within 5.3%

3.2 Natural Images

Shown in Figure 3.2 are twenty images taken from the database of [12]. The 8-bit images are 1024×1536 in size and calibrated to be linear in intensity. Each image was printed, digitally scanned with a flatbed scanner (8-bit, grayscale), and then subjected to a variety of gamma values in the range $[0.4, 2.2]$. Ground truth for the final gamma is determined by appending a small calibration strip to the bottom of the image. The actual gamma correction is determined from this calibration information, but only the original image is used in the blind estimation.

[1] The signal is windowed with a symmetric Hanning window prior to estimating the DFT.

actual	estimated gamma			
gamma	mean	std. dev.	min	max
0.20	0.22	0.01	0.20	0.24
0.40	0.44	0.04	0.40	0.56
0.60	0.65	0.05	0.60	0.77
0.80	0.84	0.03	0.78	0.88
1.00	1.05	0.04	0.99	1.11
1.20	1.26	0.04	1.21	1.31
1.40	1.42	0.04	1.35	1.47
1.60	1.67	0.06	1.56	1.76
1.80	1.92	0.11	1.73	2.12
2.00	2.06	0.07	1.98	2.19
2.20	2.29	0.04	2.22	2.36
2.40	2.44	0.12	2.19	2.58
2.60	2.62	0.12	2.34	2.78
2.80	2.83	0.06	2.69	2.92

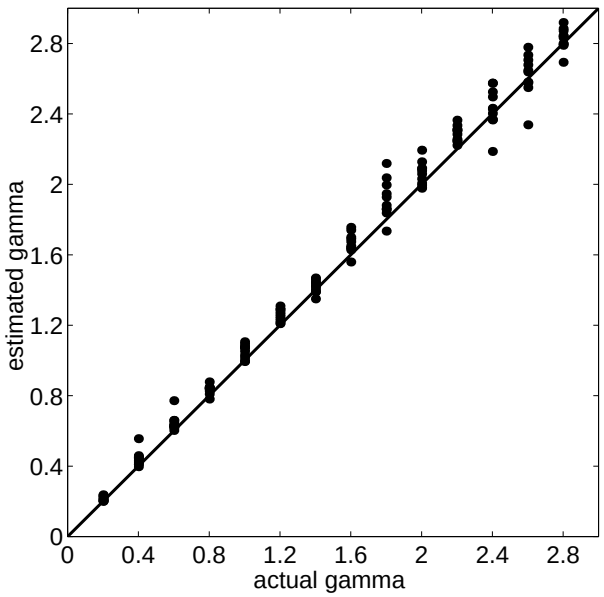

Figure 2: Shown on the left are estimated gamma values averaged over ten independent synthetically generated fractal images. Each data point on the right corresponds to one of the ten images. On average, the correct gamma is estimated within 5.3%.

The amount of gamma is estimated for a horizontal scan line by applying a range of inverse gamma values between 0.1 and 3.6 in increments of 0.1 and selecting the value that minimizes the bicoherence, Equation (12). The horizontal scan lines were used as they afforded a longer signal from which to compute the bicoherence. The estimate for the entire image is determined by averaging the estimates from every sixteenth scan line. Shown in Figure 3.2 are the estimated gamma values. On average, the correct gamma is estimated within 7.5% of the actual value. Also shown in Figure 3.2 are results from a single image, showing the gamma corrected and inverse gamma corrected images.

4 Discussion

Gamma correction introduces specific higher-order correlations in the frequency domain. Using tools from polyspectral analysis, the amount of gamma can be blindly estimated by simply minimizing these correlations. In so doing, neither calibration information nor knowledge of the imaging device is required. Results from a range of natural images have shown that the amount of gamma correction can be accurately determined with errors averaging only 7.5%. Note, that because the bicoherence is unaffected by additive noise [10] the presence of noise should not adversely effect this technique. In addition, the bicoherence is unaffected by a linear transform, e.g., brightness/contrast changes or low-pass filtering [10].

There are several straight-forward extensions to this technique that may be easily realized. Although the results presented in this paper are for grayscale images, this technique is easily generalized to color images. For a multi-channel color image, the gamma for each channel can be individually estimated as described here. Since the bicoherence over varying inverse gamma values seems to be well-behaved, the brute force minimization should be replaced with a more efficient search (e.g., a gradient descent minimization).

image	gamma 0.42	0.80	1.10	1.63	2.11	image	gamma 0.42	0.80	1.10	1.63	2.11
	0.53	0.70	0.96	1.59	2.17		0.47	0.81	1.12	1.59	2.18
	0.46	0.73	1.06	1.48	2.13		0.42	0.72	1.00	1.56	2.16
	0.43	0.77	1.02	1.57	2.17		0.54	0.86	1.10	1.58	2.18
	0.46	0.80	1.07	1.57	2.18		0.46	0.79	1.08	1.49	2.15
	0.50	0.89	1.09	1.51	2.16		0.48	0.79	1.11	1.49	2.19
	0.44	0.71	0.98	1.36	2.08						
	0.49	0.82	1.14	1.55	2.18						
	0.46	0.73	1.02	1.54	2.16						
	0.49	0.86	1.19	1.60	2.22						

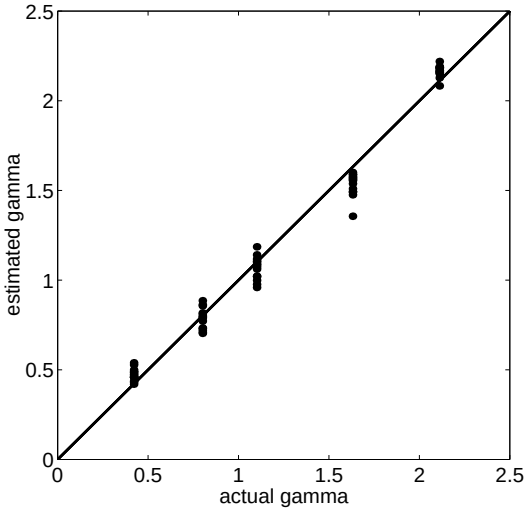

Figure 3: Each initially linear image is printed, digitally scanned and then synthetically gamma corrected with $\gamma \in [0.4, 2.2]$. Shown are the estimated gamma values determined by minimizing the bicoherence.

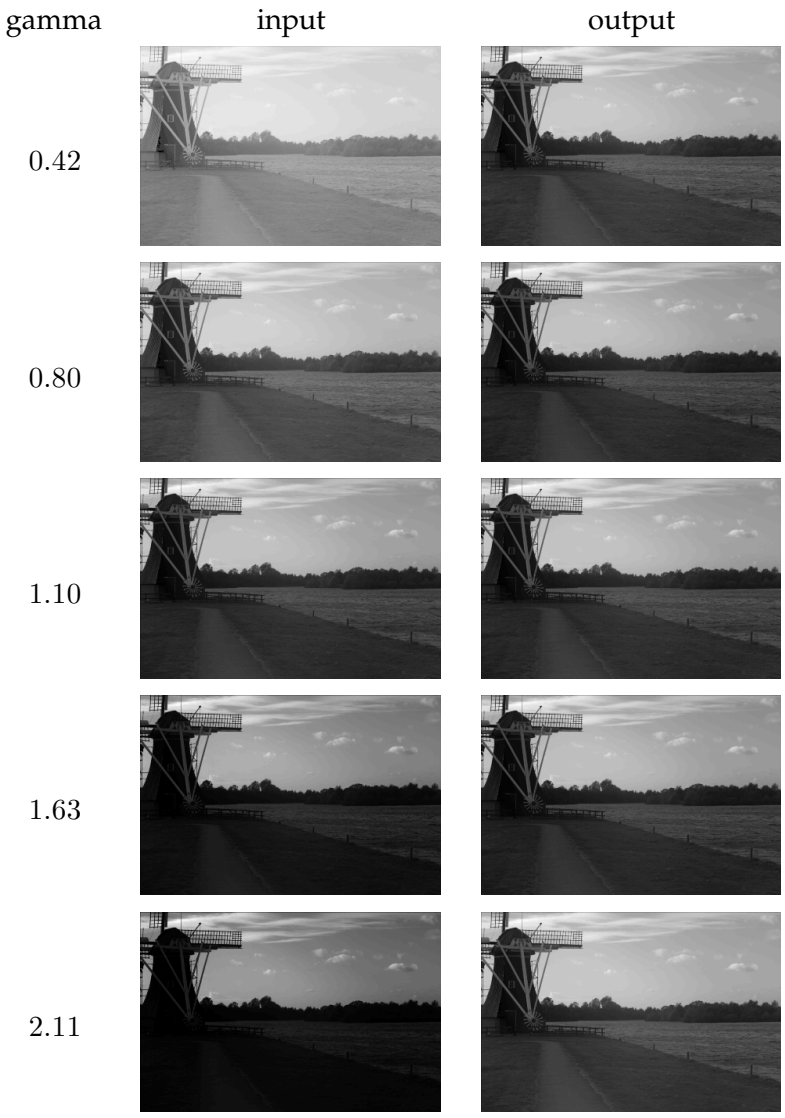

Figure 4: Shown in the left column is one image that has been subjected to varying degrees of gamma correction (see Figure 3.2). Shown on the right are the results of blindly inverting the gamma correction.

It should be pointed out that there are certain limitations and restrictions to these results. First, the simple one-parameter gamma function that is assumed is unlikely to accurately model a full range of devices (see e.g., [13]). This simple model was chosen in order to limit the minimization to one dimension. A higher parameter gamma function will require a correspondingly higher-dimensional minimization. We are currently investigating the stability of such an approach. Second, it is assumed that the gamma correction is applied uniformally throughout the image. A device that employs local gamma correction will require a more sophisticated inversion technique that allows for multiple spatially varying gamma values. We are also investigating the feasibility of such a system.

It is intriguing to observe that such "unnatural" operations as gamma correction can fundamentally change the statistics of an image while not dramatically altering the appearance of an image. Moreover, the class of functions that modulate the higher-order correlations is not limited to point-wise operations. Virtually any non-linearity will effect these correlations. This may prove to be particularly useful for understanding and classifying the statistics of natural images (e.g., [14, 15, 16, 17, 18, 19, 20])

Acknowledgments

Portions of this work were done while the author was at the Massachusetts Institute of Technology, where he was generously supported by Ted Adelson under National Institute of Health (NIH) Grant EY11005-04 and Multidisciplinary University Research Initiative (MURI) Grant N00014-95-1-0699. The author is currently at Dartmouth College where he is funded by a National Science Foundation CAREER Award (IIS-99-83806) and a departmental National Science Foundation Infrastructure Grant (EIA-98-02068). Thanks to Mary Bravo and Eero Simoncelli for their insightful and stimulating suggestions.

References

[1] Y.-C Chang and J.F. Reid. RGB calibration for analysis in machine vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 5(10):1414–1422, 1996.

[2] R.P. Kleihorst, R.L. Lagendiik, and J. Biemond. An adaptive order-statistic noise filter for gamma-corrected image sequences. *IEEE Transactions on Image Processing*, 6(10):1442–1446, 1997.

[3] X. You and G. Crebbin. Image blur identification by using higher order statistic techniques. In *Proceedings of the 3rd IEEE International Conference on Image Processing*, pages 77–80, Lausanne, Switzerland, 1996.

[4] M.A. Ibrahim, A.W.F Hussein, S.A. Mashali, and A.H. Mohamed. A blind image restoration system using higher-order statistics and Radon transform. In *Proceedings of SPIE - The International Society for Optical Engineering*, volume 3388, pages 173–181, Orlando, FL, USA, 1998.

[5] D.T. Kuan, A.A. Sawchuk, T.C. Strand, and P. Chavel. Adaptive restoration of images with speckle. In *Proceedings of SPIE - The International Society for Optical Engineering*, volume 359, pages 29–38, San Diego, CA, USA, 1983.

[6] G. Jacovitti. Applications of higher order statistics in image processing. In *Higher Order Statistics. Proceedings on the International Signal Processing Workshop*, pages 63–69, Chamrousse, France, 1992.

[7] B.H. Soong, J. Liu, and T.T. Goh. Using higher order statistics to reconstruct signals. In *Second International Conference on Automation, Robotics and Computer Vision*, volume 2, pages 1–5, Singapore, 1992.

[8] M.G. Kang and A.K. Katsaggelos. Deterministic estimation of the bispectrum and its application to image restoration. In *European Signal Processing Conference*, volume 2, pages 1449–1452, Trieste, Italy, 1996.

[9] R. P. Feynman, R. B. Leighton, and M. Sands. *The Feynman Lectures on Physics*. Addison-Wesley Publishing Company, 1977.

[10] J.M. Mendel. Tutorial on higher order statistics (spectra) in signal processing and system theory: theoretical results and some applications. *Proceedings of the IEEE*, 79:278–305, 1996.

[11] Y.C. Kim and E.J. Powers. Digital bispectral analysis and its applications to nonlinear wave interactions. *IEEE Transactions on Plasma Science*, PS-7(2):120–131, 1979.

[12] 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.

[13] T. Mitsunaga and S.K. Nayar. Radiometric self calibration. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, pages 374–380, Fort Collins, CO, USA, 1999.

[14] D. Kersten. Predictability and redundancy of natural images. *Journal of the Optical Society of America A*, 4(12):2395–2400, 1987.

[15] D.J. Field. Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A*, 4(12):2379–2394, 1987.

[16] J.G. Daugman. Entropy reduction and decorrelation in visual coding by oriented neural receptive fields. *IEEE Transaction on Biomedical Engineering*, 36(1):107–114, 1989.

[17] D.L. Ruderman and W. Bialek. Statistics of natural image: Scaling in the woods. *Phys. Rev. Letters*, 73(6):814–817, 1994.

[18] S.C. Zhu, Y. Wu, and D. Mumford. Filters, random fields and maximum entropy (frame) - towards the unified theory for texture modeling. In *IEEE Conference Computer Vision and Pattern Recognition*, pages 686–693, 1996.

[19] G. Krieger, C. Zetzsche, and E. Barth. Higher-order statistics of natural images and their exploitation by operators selective to intrinsic dimensionality. In *Proceedings of the IEEE Signal Processing Workshop on Higher-Order Statistics*, pages 147–151, Banff, Alta., Canada, 1997.

[20] E.P. Simoncelli. Modeling the joint statistics of images in the wavelet domain. In *Proceedings of the 44th Annual Meeting*, volume 3813, Denver, CO, USA, 1999.