

Blind Removal of Image Non-Linearities

Hany Farid and Alin C. Popescu
Department of Computer Science
Dartmouth College
Hanover, NH 03755
{farid,popescu}@cs.dartmouth.edu

Abstract

This paper presents a technique for blindly removing image non-linearities in the absence of any calibration information or explicit knowledge of the imaging device. The basic approach exploits the fact that a non-linearity introduces specific higher-order correlations in the frequency domain (beyond second-order). These correlations can be detected using tools from polyspectral analysis. The non-linearities can then be estimated and removed by simply minimizing these correlations.

1. Introduction

Most imaging devices introduce some form of luminance and geometric non-linearities (e.g., gamma correction and lens distortion). For many applications in image processing, digital photography and computer vision it is advantageous to remove these non-linearities prior to subsequent processing. For example, geometric non-linearities (e.g., lens distortion) may interfere with structure/shape estimation and image mosaicing. Luminance non-linearities (e.g., gamma correction) may interfere with shape from shading and photometric stereo.

Typically, luminance non-linearities are estimated by passing a calibration target with a full range of known luminance values through the imaging device (e.g., a Macbeth chart [1]). The resulting measurements are used to determine the functional form of the non-linearity. Similarly, geometric non-linearities are estimated by imaging a calibration target with known fiducial points. The deviation of these points from their original positions is used to estimate the amount of distortion (e.g., [13]).

We are grateful for the support from a National Science Foundation CAREER Award (IIS-99-83806), a departmental National Science Foundation Infrastructure Grant (EIA-98-02068), and a Department of Justice Grant (2000-DT-CX-K001).

If a device's non-linearity is constant, then the device can occasionally be calibrated and the non-linearities removed by passing the images through the inverse models. However, if direct access to the device is not possible or if it is impractical to place a calibration target in the scene, then this procedure becomes considerably more burdensome. In addition, image non-linearities may change dynamically with the camera and scene dynamics. For example, most commercial digital cameras dynamically alter the luminance non-linearity depending on the scene luminances. Similarly, changes in focal length and zoom may alter the amount of lens distortion. Under these conditions a single off-line calibration is insufficient: the device must be re-calibrated with every image.

We propose a technique for blindly removing image non-linearities in the absence of any calibration information or explicit knowledge of the imaging device. The basic approach exploits the fact that non-linearities introduce specific higher-order correlations in the frequency domain (beyond second-order). These correlations can be estimated using tools from polyspectral analysis. The non-linearities can then be estimated and removed by simply minimizing these correlations.

Polyspectral analysis and higher-order statistics have previously been used for various forms of image restoration: noise removal [8], deblurring [14, 4], and speckle removal [9]. See also [5, 11, 6] for general discussions on the use of higher-order statistics in image processing. Polyspectral analysis has been used to analyze non-linearities in signals (e.g, speech [2]).

Insight is gained into the proposed technique by considering first the effect of luminance non-linearities on a simple one-dimensional signal composed of a sum of two zero-phase sinusoids:

$$f(x) = a_1 \sin(\omega_1 x) + a_2 \sin(\omega_2 x). \quad (1)$$

When this signal is passed through a point-wise luminance non-linearity new harmonics are introduced with amplitudes that are correlated to the original harmonics [3]. To

see this more explicitly consider an arbitrary non-linear point-wise function expressed in terms of its Taylor series expansion:

$$g(u) = g(u_0) + \frac{g'(u_0)(u - u_0)}{1!} + \frac{g''(u_0)(u - u_0)^2}{2!} + \dots \quad (2)$$

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

$$\begin{aligned} g(f(x)) &\approx f(x) + f^2(x) \\ &= a_1 \sin(\omega_1 x) + a_2 \sin(\omega_2 x) \\ &\quad + \frac{1}{2} a_1^2 (1 + \sin(2\omega_1 x)) + \frac{1}{2} a_2^2 (1 + \sin(2\omega_2 x)) \\ &\quad + 2a_1 a_2 \sin((\omega_1 + \omega_2)x) + 2a_1 a_2 \sin((\omega_1 - \omega_2)x). \end{aligned} \quad (3)$$

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 . This observation generalizes to arbitrary signals/images and non-linearities.

In the following sections we will show empirically that when an image is passed through a non-linearity, higher-order correlations in the frequency domain increase proportional to the ‘‘magnitude’’ of the non-linearity. As a result, non-linearities can be estimated and removed by simply minimizing these correlations. This procedure requires no calibration information or explicit knowledge of the imaging device. We first show how tools from polyspectral analysis can be used to estimate these higher-order correlations, and then show the efficacy of the blind removal of luminance and geometric non-linearities in both synthetic and natural images.

2. Bispectral Analysis

Consider a one-dimensional signal $f(x)$, and its Fourier transform:

$$F(\omega) = \sum_{k=-\infty}^{\infty} f(k) e^{-i\omega k}. \quad (4)$$

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

$$P(\omega) = \mathcal{E}\{F(\omega)F^*(\omega)\}, \quad (5)$$

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 (3). These correlations can however be estimated with higher-order spectra (see [10] for a thorough

survey). For example the bispectrum estimates third-order correlations and is defined as:

$$B(\omega_1, \omega_2) = \mathcal{E}\{F(\omega_1)F(\omega_2)F^*(\omega_1 + \omega_2)\}. \quad (6)$$

Note that unlike the power spectrum the bispectrum of a real signal is complex-valued. Comparing the bispectrum with Equation (3) 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, $f(x)$, into N (possibly overlapping) segments, computing Fourier transforms of each segment, and then averaging the individual estimates:

$$\hat{B}(\omega_1, \omega_2) = \frac{1}{N} \sum_{k=1}^N F_k(\omega_1)F_k(\omega_2)F_k^*(\omega_1 + \omega_2), \quad (7)$$

where $F_k(\cdot)$ 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., [7]). We desire an estimator whose variance is independent of the bi-frequency. To this end, we employ the bicoherence, a normalized bispectrum, defined as:

$$b^2(\omega_1, \omega_2) = \frac{|B(\omega_1, \omega_2)|^2}{\mathcal{E}\{|F(\omega_1)F(\omega_2)|^2\}\mathcal{E}\{|F(\omega_1 + \omega_2)|^2\}}. \quad (8)$$

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:

$$\hat{b}(\omega_1, \omega_2) = \frac{|\frac{1}{N} \sum_k F_k(\omega_1)F_k(\omega_2)F_k^*(\omega_1 + \omega_2)|}{\sqrt{\frac{1}{N} \sum_k |F_k(\omega_1)F_k(\omega_2)|^2 \frac{1}{N} \sum_k |F_k(\omega_1 + \omega_2)|^2}}. \quad (9)$$

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

3. Luminance Non-Linearities

Shown in Figure 1 is a 1-D signal $f(x)$, the log of its normalized power spectrum $\log(P(\omega))$ and its bicoherence (normalized bispectrum) $\hat{b}(\omega_1, \omega_2)$. Also shown is the same signal passed through a point-wise luminance non-linearity $f^{1.5}(x)$ and its power and bispectral response. Notice that while the non-linearity leaves the power spectrum largely unchanged there is a significant increase in the bispectral response. It is the introduction of these correlations that we will exploit to estimate and remove image non-linearities.

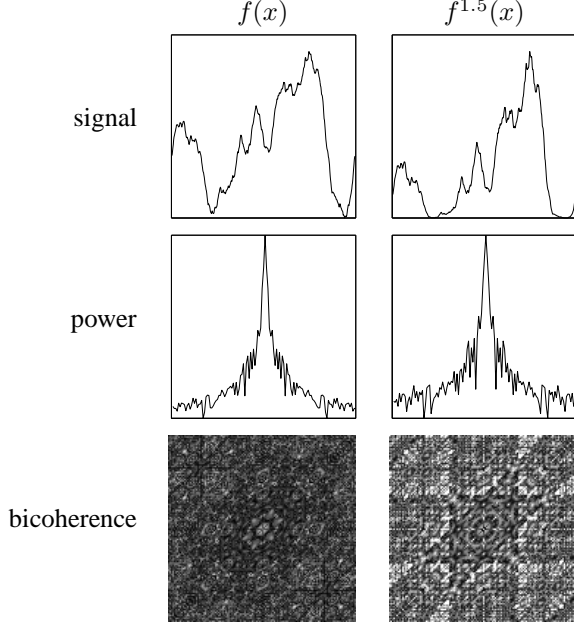


Figure 1: Shown on the left is a fractal signal and on the right the signal passed through a point-wise non-linearity. While the non-linearity leaves the power spectrum largely unchanged there is a significant increase in the bispectral response. The axis of the bicoherence corresponds to ω_1 and ω_2 . The origin is in the center, and the axis range from $[-\pi, \pi]$.

We begin by considering a parametric model of the luminance non-linearity. A generic one parameter gamma function is given by:

$$g(u) = u^\gamma, \quad (10)$$

where u denotes a pixel's luminance value, normalized into the range $[0, 1]$. The inverse function is given by: $g^{-1}(u) = u^{1/\gamma}$. Given an image that has been subjected to a non-linearity of this form, our task is determine the value of γ . This value is blindly estimated by sampling a range of inverse gamma values $1/\gamma$, applying the inverse function $g^{-1}(u)$ to the image and selecting the value of γ that minimizes the average bicoherence.

To avoid the memory and computational demands of computing an image's full four-dimensional bicoherence, our analysis is restricted to the 1-D horizontal scan lines of an image. This is reasonable since luminance non-linearities can typically be modeled as a point-wise operation so that the correlations introduced in 1-D will be similar to those in 2-D. The non-linearity $g(u)$ for a *scan line* is estimated by searching for the inverse function $g^{-1}(u)$ that

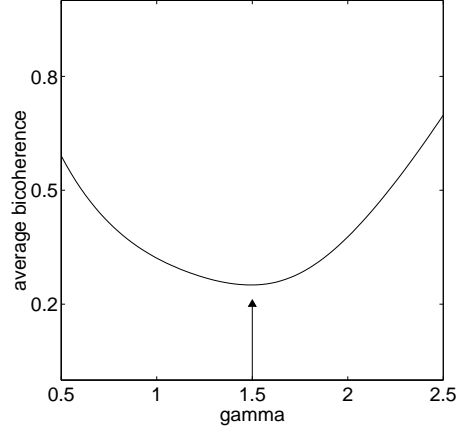


Figure 2: Shown is the bicoherence of a gamma corrected fractal image, $f^{1.5}(x, y)$, computed for a range of inverse gamma values, $[f^{1.5}(x, y)]^{\gamma^{-1}}$. The bicoherence reaches a unique minimum at $\gamma^{-1} = 1/1.5$.

minimizes the bicoherence averaged across all frequencies:

$$\sum_{\omega_1=-\pi}^{\pi} \sum_{\omega_2=-\pi}^{\pi} \hat{b}(\omega_1, \omega_2), \quad (11)$$

where $\hat{b}(\omega_1, \omega_2)$ is the bicoherence defined in Equation (9). The non-linearity for an *image* is estimated by averaging the estimates for each horizontal scan line (or a subset of them).

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 2 where the bicoherence is plotted as a function of varying inverse gamma values. In this example a 512×512 fractal image with a $1/\omega$ power spectrum and random phase is subjected to a non-linearity with $\gamma = 1.5$. Note that the bicoherence reaches a unique minimum at $\gamma^{-1} = 1/1.5$.

Shown in Figure 3 are ten images taken from the database of [12]. The 8-bit images, 1024×1536 in size, were originally calibrated to be linear in intensity. Each image was then printed (one non-linearity), digitally scanned (a second non-linearity) 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 is determined by appending a small calibration strip to the bottom of the image prior to printing. The actual non-linearity is determined from this calibration information, but only the original image is used in the blind estimation.

The bicoherence for a 1-D horizontal image slice is computed by dividing the signal into overlapping segments of length 64 with an overlap of 32. A 128-point windowed DFT is estimated for each segment, from which the bico-











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

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.

herence is estimated.¹ There is a natural tradeoff between segment length and the number of samples from which to average. We have found empirically that these parameters offer a good compromise, however their precise choice is not critical to the estimation results.

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 average bicoherence, Equation (11). The estimate for the entire image is determined by averaging the estimates from every sixteenth scan line. Shown in Figure 3 are the estimated gamma values. On average, the correct gamma is estimated within 7.5% of the actual value.

4. Geometric Non-Linearities

As with luminance non-linearities, geometric non-linearities (e.g., lens distortions) introduces specific higher-order correlations. This can be seen by first considering what effect a geometric non-linearity has on a 1-D signal. Consider, for example, a pure sinusoid: $f_u(x) = a_1 \cos(\omega_1 x)$, and for purposes of exposition, consider a simplified geometric distortion: $f_d(x) = a_1 \cos(\omega_1 x^2)$. It is straight-forward to show that, unlike the undistorted signal, this signal is composed of a multitude of (correlated) harmonics which yields an increase in the bicoherence. This observation generalizes to arbitrary signals/images and non-linearities.

As with the luminance non-linearities, the blind removal of geometric non-linearities requires a parameterized model. For the purposes of removing lens distortions, we consider a one-parameter radially symmetric distortion model. Denote the desired undistorted image as $f_u(x, y)$, and the distorted image as $f_d(\tilde{x}, \tilde{y})$, where:

$$\tilde{x} = x(1 + \kappa r^2) \quad \text{and} \quad \tilde{y} = y(1 + \kappa r^2), \quad (12)$$

and $r^2 = x^2 + y^2$, and κ controls the amount of distortion. Given only a distorted image our task is to determine the value of κ . To accomplish this we begin with a distorted image f_d , and sample a range of possible κ values, apply the inverse distortion to f_d yielding an undistorted image f_u , compute the average bicoherence of f_u , Equation (11), and select the value of κ that minimizes this quantity.

Given a guess as to the amount of distortion, κ , the distortion is inverted by solving Equation (12) for the original spatial coordinates x and y , and warping the distorted image onto this sampling lattice. Solving for the original spatial coordinates is done in polar coordinates where the solution takes on a particularly simple form.

To avoid the memory and computational demands of computing an image's full four-dimensional bicoherence,

¹The signal is first zero-meaned and then windowed with a symmetric Hanning window prior to estimating the DFT.

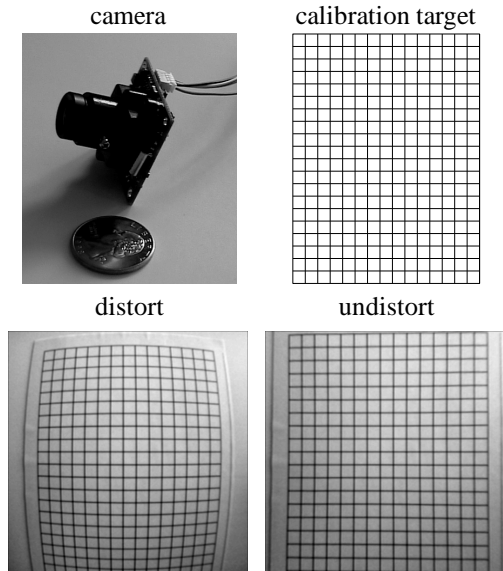


Figure 4: Shown along the top is a small low-grade camera, and a calibration target used to manually calibrate the lens distortion. Shown below is an image of the calibration target before (left) and after calibration (right).

the bicoherence is estimated from one-dimensional radial slices through the center of the image. This is reasonable if we assume a radially symmetric distortion and that the distortion emanates from the center of the image.² The amount of distortion for a *radial slice* is estimated by a brute force search for the inverse distortion that minimizes the average bicoherence, Equation (11). The amount of distortion for an *image* is then estimated by averaging the estimates from a subset of radial slices (every 10 degrees). The bicoherence was estimated as described in the previous section. Values of κ from -0.5 to 0.1 in steps of 0.02 were sampled. The asymmetry in the sampling range was for computational efficiency, and reasonable in our examples with strictly negative lens distortions.

Shown in Figure 4 is a low-grade camera used in our experiments. Also shown in Figure 4 is an image of the calibration target before and after calibration. The amount of distortion was manually estimated to be $\kappa = -0.16$. In the absence of this calibration information the amount of distortion was blindly estimated for each of the images in Figure 5. The individual estimates are -0.20 , -0.02 , -0.10 , and -0.14 , with an average estimate of -0.13

Because of the unavoidable non-linear interpolation step involved in the warping during the model inversion, correlations are artificially introduced that confound those introduced by the lens distortion. As such, in all of our results

²If the image center drifts, then a more complex minimization is required to jointly determine the image center and amount of distortion.




Figure 5: Shown are several distorted images (left) and the results of blindly estimating and removing the lens distortion (right).

the estimated distortion κ is related to the actual distortion κ' by the following empirically determined cubic relationship:

$$\kappa' = -1.5784\kappa^3 - 0.7752\kappa^2 + 1.6621\kappa - 0.0089 \quad (13)$$

We have found that errors in the estimation of geometric non-linearities are typically higher than those of luminance non-linearities. We suspect that there are at least two reasons for this. First, the images used in these experiments are 480×640 , significantly smaller than those used in the removal of luminance non-linearities. The shorter signal length makes it more difficult to obtain an accurate estimate of the bicoherence. Second, geometric non-linearities introduce a multitude of correlations, while the bicoherence only measures correlations between harmonically related frequencies. In numerous simulations (not presented here), we were able to overcome the noisier estimates by simply averaging over five to ten images. As such, we show in Figure 5 the results of averaging the estimates from all four images. In practice, this should not be burdensome, as it is relatively easy to acquire a small number of images.

5. Discussion

Image non-linearities introduce specific higher-order correlations in the frequency domain. Using tools from polyspectral analysis, these non-linearities can be blindly estimated and removed by simply minimizing these correlations. In so doing, neither calibration information nor explicit knowledge of the imaging device is required. We did assume a parametric model of the non-linearity. We also assumed that the image non-linearities are applied uniformly throughout the image. A device that employs, for example, local gain control will require a more sophisticated inversion technique that allows for multiple spatially varying model parameters.

The blind estimation of luminance non-linearities has proven to be reasonably accurate for most applications. The advantage of this approach is that it does not require any calibration. The drawback is that the blind estimation requires a computationally more intensive estimation of the bicoherence across a range of possible model parameters. Since the minimization is well behaved this brute force search should be replaced with a more efficient search (e.g., a gradient descent minimization).

The accuracy of blindly estimating lens distortion is by no means comparable to that based on calibration. As such we don't expect that this approach will supplant other techniques in areas where a high degree of accuracy is required. Rather, we expect this approach to be useful in areas where only qualitative results are required.

It is intriguing to observe that the non-linearities in an imaging device can fundamentally change the statistics of an image while not dramatically altering the appearance of an image. This may prove to be useful for understanding and classifying the statistics of natural images.

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] J.W.A. Fackrell and S. McLaughlin. Detecting nonlinearities in speech sounds using the bicoherence. *Proceedings of the Institute of Acoustics*, 18(9):123–130, 1996.
- [3] R. P. Feynman, R. B. Leighton, and M. Sands. *The Feynman Lectures on Physics*. Addison-Wesley Publishing Company, 1977.
- [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] 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.
- [6] 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.
- [7] 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.
- [8] 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.
- [9] 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.
- [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] 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.
- [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] J. Weng. Camera calibration with distortion models and accuracy evaluation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(10):965–979, 1992.
- [14] 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.