

# **Blind Removal of Lens Distortion**

Hany Farid and Alin C. Popescu  
Department of Computer Science  
Dartmouth College  
Hanover NH 03755

Virtually all imaging devices introduce some amount of geometric lens distortion. This paper presents a technique for blindly removing these distortions in the absence of any calibration information or explicit knowledge of the imaging device. The basic approach exploits the fact that lens distortion introduces specific higher-order correlations in the frequency domain. These correlations can be detected using tools from polyspectral analysis. The amount of distortion is then estimated by minimizing these correlations.

# 1 Introduction


Virtually all medium- to low-grade imaging devices introduce some amount of geometric distortion. These distortions are often described with a one-parameter radially symmetric model [2, 8, 9]. Given an ideal undistorted image  $f_u(x, y)$ , the distorted image is denoted as  $f_d(\tilde{x}, \tilde{y})$ , where the distorted spatial parameters are given by:

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

where  $r^2 = x^2 + y^2$ , and  $\kappa$  controls the amount of distortion. Shown in Figure 1 are the results of distorting a rectilinear grid with positive and negative values of  $\kappa$ .

While these distortions may be artistically interesting it is often desirable to remove these geometric distortions for many applications in image processing and computer vision (e.g., structure estimation, image mosaicing). The amount of distortion is typically determined experimentally 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., [9]). But often such calibration is not available or direct access to the imaging device is not possible, for example when down-loading an image from the web. In addition, the distortion parameters can change as other imaging parameters are varied (e.g., focal length or zoom), thus requiring repeated calibration for all possible camera settings. An alternative calibration technique relies on the presence of straight lines in the scene (e.g., [1, 7]). These lines, mapped to curves in the image due to the distortion, are located or specified by the user. The distortions are estimated by finding the model parameters that map these curved lines to straight lines. While this technique is more flexible than those based on imaging a calibration target, it still relies on the scene containing extended straight lines.

In this paper a technique is presented for estimating the amount of lens distortion in the absence of any calibration information or scene content. The basic approach exploits the fact that



**Figure 1:** One-parameter radially symmetric lens distortion, Equation (1).

lens distortion introduces specific higher-order correlations in the frequency domain. These correlations can be detected using tools from polyspectral analysis. The amount of distortion is then determined by minimizing these correlations. These basic principles were used in a related paper in which we introduced a technique for the blind removal of luminance non-linearities [3].

Insight is gained into the proposed technique by first considering what effect a geometric distortion has on a one-dimensional signal. Consider, for example, a pure sinusoid with amplitude  $a$  and frequency  $b$ :

$$f_u(x) = a \cos(bx). \quad (2)$$

For purposes of exposition, consider a simplified version of the lens distortion given in Equation (1), where the spatial parameter is squared:

$$f_d(x) = a \cos(bx^2). \quad (3)$$

This signal is composed of a multitude of harmonics. This can be seen by considering its Fourier transform:

$$\begin{aligned} F_d(\omega) &= \int_{-\infty}^{\infty} f_d(x) e^{-i\omega x} dx \\ &= 2 \int_0^{\infty} a \cos(bx^2) \cos(\omega x) dx. \end{aligned} \quad (4)$$

Because the signal is symmetric (a cosine), the Fourier integral may be expressed from 0 to  $\infty$  and with respect to only the cosine basis (i.e., the sine component of the complex exponential integrates to zero). This integral has a closed form

solution [4] given by:

$$F_d(\omega) = 2a\sqrt{\frac{\pi}{2b}} \left[ \cos\left(\frac{\omega^2}{2b}\right) + \sin\left(\frac{\omega^2}{2b}\right) \right]. \quad (5)$$

Unlike the undistorted signal, with:

$$F_u(\omega) = \begin{cases} 1 & |\omega| = b \\ 0 & |\omega| \neq b \end{cases} \quad (6)$$

the Fourier transform of the distorted signal contains a multitude of harmonics. Moreover, the amplitude and phase of these harmonics are correlated to the original signal. Here the phases are trivially correlated as all frequencies are zero-phase. Nevertheless, if the initial signal consisted of multiple frequencies with non-zero phases, then the resulting distorted signal would have similar amplitude correlations and non-trivial phase correlations.

In what follows we will show that this observation is not limited to the specific choice of signal or distortion. We will also show empirically that when an image is geometrically distorted, higher-order correlations in the frequency domain increase proportional to the amount of distortion. As such, the amount of distortion can be determined by simply minimizing these correlations. We first show how tools from polyspectral analysis can be used to measure these higher-order correlations, and then show the efficacy of this technique to the blind removal of lens distortion in synthetic and natural images.

## 2 Bispectral Analysis

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

$$F(\omega) = \sum_{k=-\infty}^{\infty} f(k)e^{-i\omega k}. \quad (7)$$

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

$$P(\omega) = \mathcal{E} \{F(\omega)F^*(\omega)\}, \quad (8)$$

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 (1). These correlations can however be estimated with higher-order spectra (see [6] for a thorough survey). For example the bispectrum estimates third-order correlations and is defined as:

$$B(\omega_1, \omega_2) = \mathcal{E} \{F(\omega_1)F(\omega_2)F^*(\omega_1 + \omega_2)\}. \quad (9)$$

Note that unlike the power spectrum the bispectrum of a real signal is complex-valued.

The bispectrum reveals correlations between harmonically related frequencies, for example,  $[\omega_1, \omega_1, 2\omega_1]$  or  $[\omega_1, \omega_2, \omega_1 + \omega_2]$ . If it is assumed that the signal  $f(x)$  is ergodic, then the bispectrum can be estimated by dividing  $f(x)$  into  $N$  (possibly overlapping) segments, computing Fourier transforms of each segment, and then averaging the individual estimates:

$$\hat{B}(\omega_1, \omega_2) = \frac{1}{N} \sum_{k=1}^N F_k(\omega_1)F_k(\omega_2)F_k^*(\omega_1 + \omega_2), \quad (10)$$

where  $F_k(\cdot)$  denotes the Fourier transform of the  $k^{th}$  segment. This arithmetic average estimator is unbiased and of minimum variance. However, it has the undesired property that its variance at each bi-frequency  $(\omega_1, \omega_2)$  depends on  $P(\omega_1)$ ,  $P(\omega_2)$ , and  $P(\omega_1 + \omega_2)$  (see e.g., [5]). We desire an estimator whose variance is independent of the bi-frequency. To this end, we employ the bicoherence, a normalized bispectrum, defined as:


$$b^2(\omega_1, \omega_2) = \frac{|B(\omega_1, \omega_2)|^2}{\mathcal{E}\{|F(\omega_1)F(\omega_2)|^2\}\mathcal{E}\{|F(\omega_1 + \omega_2)|^2\}}. \quad (11)$$

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:

$$\hat{b}(\omega_1, \omega_2) = \frac{|\frac{1}{N} \sum_k F_k(\omega_1)F_k(\omega_2)F_k^*(\omega_1 + \omega_2)|}{\sqrt{\frac{1}{N} \sum_k |F_k(\omega_1)F_k(\omega_2)|^2} \sqrt{\frac{1}{N} \sum_k |F_k(\omega_1 + \omega_2)|^2}}. \quad (12)$$

Note that the bicoherence is now a real-valued quantity.

Shown in Figure 2 is an example of the sensitivity of the bicoherence to higher-order correlations that are invisible to the power spectrum.



**Figure 2:** Top: the normalized power spectrum and bicoherence for a signal with random amplitudes and phases. Bottom: the same signal with one frequency,  $\omega_3 = \omega_1 + \omega_2$ , whose amplitude and phase are correlated to  $\omega_1$  and  $\omega_2$ . The horizontal axis of the bicoherence corresponds to  $\omega_1$ , and the vertical to  $\omega_2$ . The origin is in the center, and the axis range from  $[-\pi, \pi]$ .

A signal of length 4096 with random amplitude and phase is divided into  $N = 128$  overlapping segments of length 64 each. Shown in the top row of Figure 2 is the estimated power spectrum and the bicoherence estimated as specified in Equation (12). Shown below is the same signal where  $\omega_3 = \omega_1 + \omega_2$  has been coupled to  $\omega_1$  and  $\omega_2$ . That is,  $\omega_3$  has amplitude  $a_3 = a_1 \cdot a_2$  and phase  $\phi_3 = \phi_1 + \phi_2$ . Note that the remaining frequency content of the signal remains unchanged, but that the bicoherence is significantly more active (increasing from 0.08 to 0.20) at the bi-frequency  $\omega_1, \omega_2$ , as seen by the peaks in Figure 2. The multiple peaks are due to the inherent symmetries in the bicoherence.

As a measure of overall correlations, the bicoherence can be averaged across all frequencies:

$$\frac{1}{N^2} \sum_{\omega_1=-N/2}^{N/2} \sum_{\omega_2=-N/2}^{N/2} \hat{b}\left(\frac{2\pi\omega_1}{N}, \frac{2\pi\omega_2}{N}\right). \quad (13)$$

This quantity is employed throughout this paper as a measure of higher-order correlations.

### 3 Lens Distortions and Correlations

Shown in Figure 3 is a 1-D signal  $f_u(x)$ , of length 4096, with a  $1/\omega$  power spectrum and random phase. Also shown is the log of its normalized power spectrum  $P(w)$  and its bicoherence  $\hat{b}(\omega_1, \omega_2)$ . The bicoherence was estimated from 128 overlapping segments each of length 64 each. Also shown in Figure 3 is the same signal passed through a 1-D version of the lens distortion,  $f_d(x)$ , given in Equation (1):


$$f_d(x) = f_u(x(1 + \kappa x^2)) \quad (14)$$

where  $\kappa$  controls the amount of distortion. Notice that while the distortion leaves the power spectrum largely unchanged there is a significant increase in the bispectral response: the bicoherence averaged across all frequencies, Equation (13), nearly doubles from 0.08 to 0.14. This example illustrates that when an arbitrary signal is exposed to a geometric non-linearity, correlations between triples of harmonics are introduced.

For our purposes, what remains to be shown is that these correlations are proportional to the amount of distortion,  $\kappa$ . To illustrate this relationship a 1-D signal  $f_u(x)$  is subjected to a full range of distortions as in Equation (14). Shown in Figure 4 is the average bicoherence, Equation (13), plotted as a function of the amount of distortion. Notice that this function has a single minimum at  $\kappa = 0$ , i.e., no distortion.

These observations lead to a simple algorithm for blindly removing lens distortions. Beginning with a distorted signal:


1. select a range of possible  $\kappa$  values,



**Figure 3:** Shown in the left column is a fractal signal, the log of its normalized power spectrum and its bicoherence. Shown in the right column is a distorted version of the signal. While the distortion leaves the power spectrum largely unchanged there is a significant increase in the average bispectral response.

2. for each value of  $\kappa$  apply the inverse distortion to  $f_d$  yielding a provisional undistorted image  $f_{\kappa}$ ,
3. compute the bicoherence of  $f_{\kappa}$ ,
4. select the value of  $\kappa$  that minimizes the bicoherence averaged across all frequencies.
5. remove the distortion according to the inverse distortion model

This basic algorithm extends naturally to 2-D images. However in order to avoid the memory and computational demands of computing an



**Figure 4:** Shown is the bicoherence computed for a range of lens distortion ( $\kappa$ ). The bicoherence is minimal when  $\kappa = 0$ , i.e., no distortion.

image's full 4-D bicoherence, we limit our analysis to one-dimensional radial slices through the center of the image. This is reasonable assuming a radially symmetric distortion and that the distortion emanates from the center of the image. If the image center drifts, then a more complex three-parameter minimization would be required to jointly determine the image center and amount of distortion. The amount of distortion for an *image* is then estimated by averaging over the estimates from a subset of radial slices (e.g., every 10 degrees), as described above.

In the results that follow in the next section, we assume a one-parameter radially symmetric distortion model. Denoting the desired undistorted image as  $f_u(x, y)$ , the distorted image is denoted as  $f_d(\tilde{x}, \tilde{y})$ , where

$$\tilde{x} = x(1 + \kappa r^2) \quad \text{and} \quad \tilde{y} = y(1 + \kappa r^2), \quad (15)$$

and  $r^2 = x^2 + y^2$ , and  $\kappa$  controls the amount of distortion. Given an estimate of the distortion, the image is undistorted by solving Equation (15) 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. In polar coordinates the undistorted image is denoted as  $f_u(r, \theta)$ , where

$$r = \sqrt{x^2 + y^2} \quad \text{and} \quad \theta = \tan^{-1}(y/x). \quad (16)$$

Similarly, the distorted image  $f_d(\tilde{x}, \tilde{y})$  in polar coordinates is  $f_d(\tilde{r}, \tilde{\theta})$ , where

$$\tilde{r} = \sqrt{\tilde{x}^2 + \tilde{y}^2} \quad \text{and} \quad \tilde{\theta} = \tan^{-1}(\tilde{y}/\tilde{x}). \quad (17)$$


Combining these parameters with Equations (15) and (16) yields

$$\tilde{r} = r(1 + \kappa r^2) \quad \text{and} \quad \tilde{\theta} = \tan^{-1}(y/x). \quad (18)$$

Note that since the distortion model is radially symmetric, only the radial component is effected. The undistorted radial parameter  $r$  is determined by solving the resulting cubic equation in Equation (18). These polar parameters are then converted back to rectangular coordinates and the distortion is inverted by warping the image  $f_d(\tilde{x}, \tilde{y})$  onto this new sampling lattice.

## 4 Results

In the results reported here, the bicoherence for each 1-D radial image slice is computed by dividing the signal into overlapping segments of length 64 with an overlap of 32. A 128-point DFT (windowed with a symmetric Hanning window)  $F_k(\omega)$  is estimated for each zero-mean segment from which the bicoherence is estimated as in Equation (12). 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, their precise choice, however, is not critical to the estimation results. Each equal length radial slice is obtained by bicubic interpolation. Running on a 933 MHz Pentium (under Linux), a  $512 \times 512$  image takes approximately 25 seconds to apply the inverse distortion model for a provisional estimate of the distortion, and compute the mean bicoherence of 90 1-D signals (every 2 degrees). The total runtime will depend on



**Figure 5:** Synthetic images with no distortion (center), negative (left) and positive (right) distortion.

the number of candidate distortion parameters tried.

Presented next are results on the blind estimation of lens distortion for synthetic and natural images.

### 4.1 Synthetic Images

Fractal images were synthesized from a sum of two-dimensional sinusoids with random orientation,  $\theta_n \in [-\pi, \pi]$ , random phase,  $\phi_n \in [-\pi, \pi]$ , amplitudes,  $a_n = 1/n$ , and frequencies,  $\omega_n = n\pi$ :

$$f_u(x, y) = \sum_{n=1}^N a_n \sin(\omega_n [\cos(\theta_n)x + \sin(\theta_n)y] + \phi_n), \quad (19)$$

These images were  $N \times N$  in size, with  $N = 512$ , and the horizontal ( $x$ ) and vertical ( $y$ ) coordinates normalized into the range  $[-1, 1]$ . The distortion of such an image by an amount  $\kappa$  was simulated from a similar sum of “distorted” sinusoids with the same orientations, phases, amplitudes, and frequencies:


$$f_d(\tilde{x}, \tilde{y}) = \sum_{n=1}^N a_n \sin(\omega_n [\cos(\theta_n)\tilde{x} + \sin(\theta_n)\tilde{y}] + \phi_n), \quad (20)$$

where,  $\tilde{x}$  and  $\tilde{y}$  are as in Equation (15). Shown in Figure 5 are examples of these images. The distorted images could have been synthesized by simply warping the undistorted image. This was not done in order to avoid any possible artifacts introduced by the required interpolation.

Shown in Figure 6 and summarized in Figure 7 are the results of blindly estimating the amount of lens distortion  $\kappa$ . In these simulations

actual $\kappa$	estimated $\kappa$			
	mean	s.d.	min	max
-0.60	-0.62	0.07	-0.76	-0.52
-0.50	-0.45	0.03	-0.52	-0.41
-0.40	-0.40	0.07	-0.53	-0.27
-0.30	-0.32	0.08	-0.41	-0.18
-0.20	-0.22	0.05	-0.29	-0.15
-0.10	-0.08	0.03	-0.13	-0.03
0.00	-0.01	0.03	-0.06	0.04
0.10	0.07	0.03	0.01	0.12
0.20	0.21	0.02	0.18	0.24
0.30	0.32	0.01	0.29	0.34
0.40	0.38	0.01	0.37	0.40

**Figure 6:** Shown are the blindly estimated distortion parameters (mean, standard deviation, and minimum and maximum values) averaged over ten independent synthetic images. On average, the correct value is estimated within 8% of the actual value. See also Figure 7.



**Figure 7:** Shown are the blindly estimated distortion parameters. Each data point corresponds to the average from ten synthetic images. See also Figure 6.

the bicoherence was estimated as described above. Values of  $\kappa$  from -0.8 to 0.6 in steps of 0.05 were sampled. The estimates for each value of  $\kappa \in [-0.6, 0.4]$  are averaged over ten independently generated images. On average, the correct value is estimated within 8% of the actual value.

Because of the unavoidable non-linear interpolation step involved in the warping during the model inversion, and extraction of 1-D radial slices, correlations are artificially introduced that confound those introduced by the lens distortion. As such, in all of our results the estimated distortion  $\kappa$  is related to the actual distortion  $\kappa'$  by the following empirically determined cubic relationship:

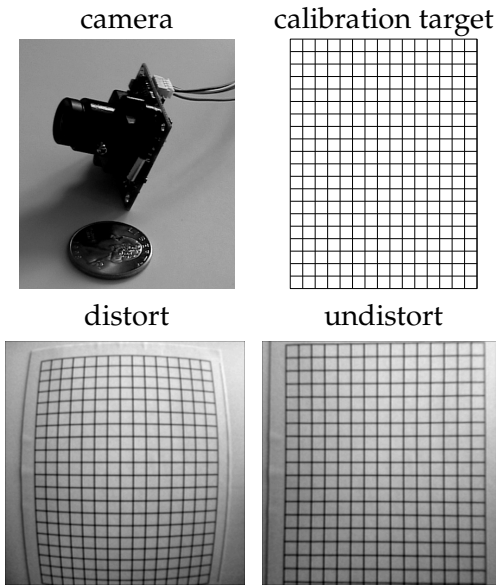
$$\kappa' = -1.5784\kappa^3 - 0.7752\kappa^2 + 1.6621\kappa - 0.0089 \quad (21)$$

This relationship holds for all the results presented here, but is dependent on the image size. That is, the image's spatial sampling lattice should be specified with respect to a  $512 \times 512$  image normalized into the range  $[-1, 1]$ .

## 4.2 Natural Images

Shown in Figure 8 is a low-grade camera used in our first experiment. The amount of distortion was estimated by imaging a calibration target. Shown in Figure 8 is an image of the calibration target before and after calibration. The amount of distortion was manually estimated to be  $\kappa = -0.16$ . Although the correction is not perfect, it does show that the one-parameter model can reasonably approximate the lens distortion from this and similar cameras.

In the absence of this calibration information the amount of distortion was blindly estimated for each of the images in Figure 9. These images are  $640 \times 480$  pixels in size. In these experiments the bicoherence was estimated as described above. Values of  $\kappa$  from -0.5 to 0.1 in steps of 0.025 were sampled. The asymmetry in the sampling range was for computational efficiency, and reasonable in these examples with strictly negative lens distortions. The distortion, averaged across the four images (360 1-D radial slices, 90 per image) shown in Figure 9, is  $-0.15$




**Figure 8:** 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).

with a variance of 0.08. The distortion in each image was removed with this estimate.

Also shown in Figure 10 are results from images taken with a Nikon Coolpix 950 digital camera. These images are  $1600 \times 1200$  pixels in size. In these examples the distortion was experimentally determined to be  $-0.005$ : a small, but not insignificant amount of distortion. The blindly estimated distortion averaged from the four images shown in Figure 10 was  $-0.04$  with a variance of 0.03. With a distortion value close to zero, the error in the estimate is visually negligible, as can be seen in the resulting undistorted images.

Because of the individual variations from image to image, the blind estimation requires an average across several images. In our examples, we have found that as few as four images are sufficient. Note that this variation is consistent with the simulations shown in Figure 6, where, for example, the estimated parameters for  $\kappa =$



**Figure 9:** Shown are several distorted images (left) and the results of blindly estimating and removing the lens distortion (right).

0.0 ranged from  $-0.06$  to  $0.04$ .

As with the synthetic images, the estimated distortion parameter is related to the actual value as specified in Equation (21).

## 5 Discussion

Most imaging or recording devices introduce some amount of geometric lens distortion. While at times artistically pleasing, these distortions are often undesirable for a variety of applications in image processing and computer vision (e.g., structure from motion, image mosaicing).

The amount of lens distortion is typically de-





**Figure 10:** Shown are several distorted images (left) and the results of blindly estimating and removing the lens distortion (right).

terminated experimentally 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. This approach is complicated by the fact that the amount of distortion changes with varying camera settings (e.g., zoom or focal length). In addition this procedure is impossible in the absence of calibration information, for example, when downloading an image from the web.

In this paper we have presented a method for the blind removal of lens distortions in the absence of any calibration information or explicit knowledge of the imaging device. This method is based on the observation that a lens-distorted image contains specific higher-order correlations in the frequency domain. These correlations are detected using tools from polyspectral analysis. The distortion is estimated and removed by minimizing these correlations. We have experimentally verified this approach on a number of synthetic and natural images.

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. One such area may be in the consumer development of photographs taken with low-grade digital or disposable cameras. We are working to generalize these results to be used with higher-order lens distortion models. Such a system will require a multi-dimensional minimization of the same correlation measure over each of the model parameters. Such an approach will surely require a more adaptive minimization than the brute-force approach employed here. Finally, we are also working to incorporate our earlier work [3] on the blind removal of luminance non-linearities, for what we hope will be a complete system for the blind removal of image non-linearities.

## Acknowledgments

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

## References

- [1] F. Devernay and O. Faugeras. Automatic calibration and removal of distortion from scenes of structured environments. In *SPIE Conference on Investigative and Trial Image Processing*, San Diego, CA, 1995.
- [2] W. Faig. Calibration of close-range photogrammetric systems: Mathematical formulation. *Photogrammetric Eng. Remote Sensing*, 41(12):1479–1486, 1975.
- [3] H. Farid. Blind inverse gamma correction. *IEEE Transactions on Image Processing*, In press.
- [4] I.S. Gradshteyn and I.M. Ryzhik. *Table of Integrals, Series, and Products*. Academic Press, San Deige, CA, 1994.
- [5] 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.
- [6] 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.
- [7] R. Swaminatha and S.K. Nayar. Non-metric calibration of wide-angle lenses and poly-cameras. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 413–419, Fort Collins, CO, 1999.
- [8] R.Y. Tsai. A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses. *RA-3(4):323–344*, 1987.
- [9] J. Weng. Camera calibration with distortion models and accuracy evaluation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(10):965–979, 1992.