

Estimating Planar Surface Orientation Using Bispectral Analysis

Hany Farid and Jana Košecká

Abstract—In this paper we propose a direct method for estimating the orientation of a plane from a single view under perspective projection. Assuming that the underlying planar texture has random phase, we show that the non-linearities introduced by perspective projection lead to higher-order correlations in the frequency domain. We also empirically show that these correlations are proportional to the orientation of the plane. Minimization of these correlations, using tools from polyspectral analysis, yields the orientation of the plane. We show the efficacy of this technique on synthetic and natural images.

I. INTRODUCTION

Many visual cues reveal the 3-D structure of a scene and relative pose of the camera with respect to the scene. While many of these cues are present in multiple views (disparity, optical flow), several others are already present in a single view. Gibson suggested, for example, the use of texture gradients and vanishing lines [7], while others proposed the use of shading gradients, e.g., [9], [16], [4]. Since then several different computational models based on these mechanisms have been proposed. These techniques differ in the assumptions made about the scene appearance and structure, the types of measurements used, and the computational methods employed for estimating shape or pose.

Feature-based techniques use elementary geometric features (points, lines, circles) and additional assumptions (orthogonality, co-planarity, eccentricity) to estimate the projective mapping between the world and the image plane, e.g., [1], [12]. These methods are most applicable when the desired features are easily extracted from the image.

The class of techniques termed “shape from texture” have also been popular, particularly when explicit features are not readily available in an

image [10]. One class of techniques directly use filtered outputs of the image intensities. As with the feature-based techniques certain assumptions are made regarding the spatial texture properties. The most standard assumption is one of homogeneity. With this assumption, the 3-D structure is estimated by first measuring local deformations of the texture, and then explicitly parameterizing them in terms of surface slant and tilt. In this setting both orthographic and perspective projections models have been considered, e.g., [19], [5]. Additional assumptions of isotropy, e.g., [6] and symmetry can also be made. In these cases the observed violation of the assumed statistical properties of texture elements, characterized by their gradient orientations, is used for estimating surface orientation. Frequency-based methods have also been extensively explored. These techniques explicitly model the effects of perspective projection on the frequency or phase of an image. The estimation stage typically employs measurements of instantaneous frequency [22], [23], model-based measurements of instantaneous phase [17], chirplets [14], or wavelets [8].

In this paper, we extend the frequency-based techniques by exploiting the fact that perspective projection introduces higher-order frequency correlations (beyond second-order). Like all texture-based techniques we make an assumption regarding the underlying texture. In our case, it is assumed that the texture has random phase. Shown in Fig. 1, for example, are several examples of texture with and without random phase – loosely speaking, random phase textures are those that do not contain spatially correlated or structured patterns. Under this assumption, we then show that when a textured plane is imaged under perspective projection, the resulting image contains third-order frequency correlations. The correlations are due to the non-linearities introduced by perspective projection. These correlations, measured using tools from polyspectral analysis,

Corresponding author: H. Farid, 6211 Sudikoff Lab, Computer Science Department, Dartmouth College, Hanover, NH 03755 USA (email: farid@cs.dartmouth.edu; tel/fax: 603.646.2761/603.646.1672)

J. Kosecka, Department of Computer Science, George Mason University, Fairfax, VA 22030

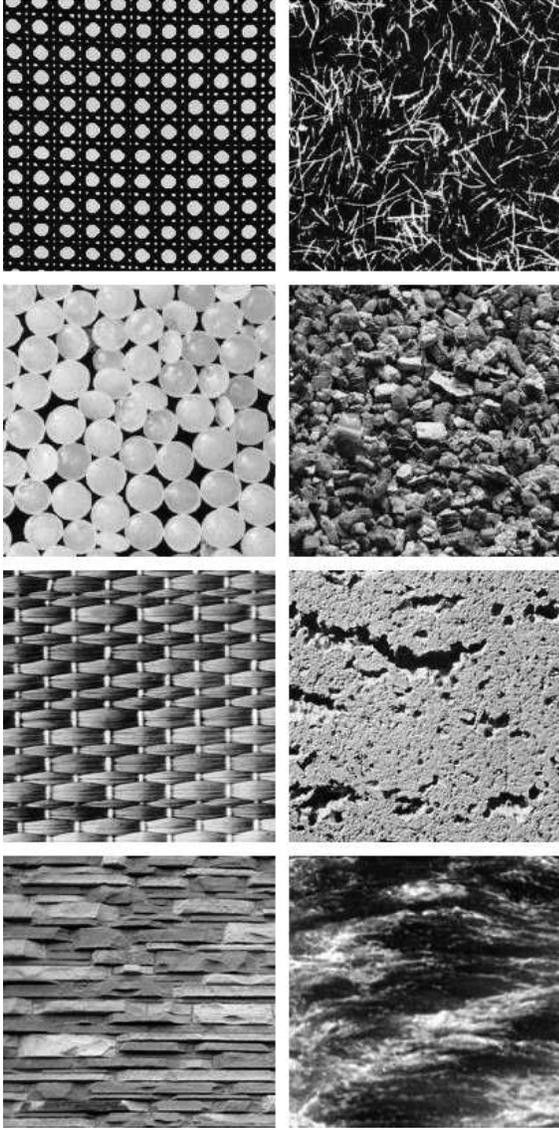


Fig. 1. Examples of texture with non-random (left) and random (right) phase.

are minimized to yield the orientation of the plane. This technique uses a similar approach as described in our earlier work, where we showed that non-linearities introduced by gamma correction and lens distortion could be blindly estimated and corrected for in the absence of any calibration information [2], [3].

II. PERSPECTIVE PROJECTION IN FREQUENCY DOMAIN

Consider, for purposes of illustration, a simple one-dimensional signal composed of two frequencies (for simplicity we will assume zero-phase):

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

Consider now this signal placed onto a one-dimensional plane oriented -45° from fronto-parallel. Assuming a perspective projection model, this signal can be written analytically as:

$$\begin{aligned} g(x) &= a_1 \cos\left(\omega_1 \frac{fx'}{z'}\right) + a_2 \cos\left(\omega_2 \frac{fx'}{z'}\right) \\ &= a_1 \cos\left(\omega_1 \frac{x'}{x'+2}\right) + a_2 \cos\left(\omega_2 \frac{x'}{x'+2}\right) \end{aligned} \quad (2)$$

where x denotes image coordinates and x' and z' denote camera coordinate systems. Without loss of generality, we assume that the focal length $f = 1$, and that the plane is a unit length along the optical axis away from the sensor, which, when coupled with a -45° oriented plane, yields the distance from the sensor, z' , to be simply $x' + 2$, Fig. 2. Example fronto-parallel and projected signals are shown in Fig. 2 (with frequencies $\omega_1 = 8\pi$ and $\omega_2 = 32\pi$, and amplitudes $a_1 = a_2 = 1$). Shown below each signal is the central portion of the magnitude of each signal's Fourier transform. Note that the original signal has exactly two frequencies (the Fourier magnitude is symmetric about its origin), and the projected signal contains a multitude of frequencies. Of particular interest to us are the frequency correlations introduced by the non-linear perspective projection. While in this simple example the power spectrum may reveal these correlations, when presented with more complex signals these correlations are less likely to be so evident. As in our previous work [2], [3] we turn to measuring higher-order correlations in the frequency domain. Specifically, we will show empirically that higher-order correlations between triples of frequencies - pairs of frequencies and their sum - are proportional to the amount of distortion introduced by perspective projection. We first discuss how these higher-order correlations are measured, and then show how the minimization of these correlations yields an estimate of surface orientation.

III. BISPECTRAL ANALYSIS

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

$$F(\omega) = \sum_{k=-\infty}^{\infty} f(k) e^{-i\omega k}. \quad (3)$$

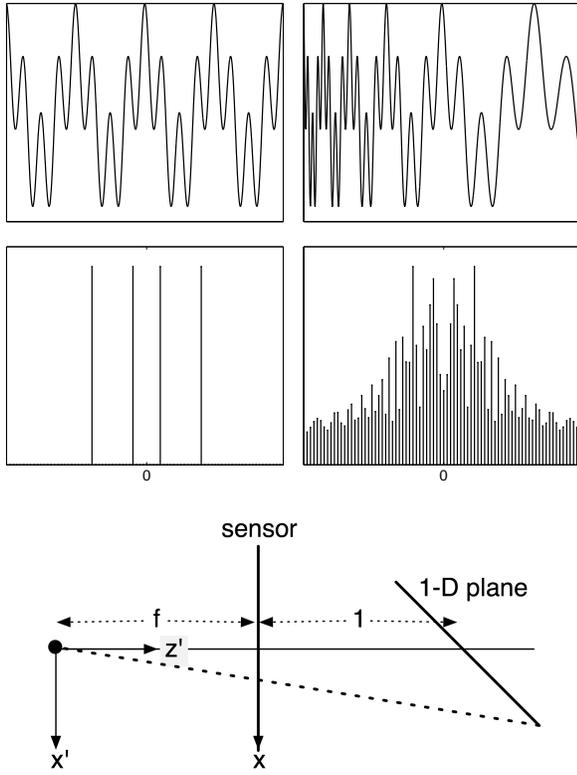


Fig. 2. Shown in the left column is a signal composed of two frequencies (top) and the central portion of the magnitude of its Fourier transform (bottom). Shown in the right column is the same signal placed onto a 1-D plane rotated -45° from fronto-parallel and projected under perspective projection. Note how this signal contains a multitude of frequencies. Shown below is the camera imaging geometry.

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

$$P(\omega) = \mathcal{E} \{F(\omega)F^*(\omega)\}, \quad (4)$$

where $\mathcal{E}\{\cdot\}$ is the expected value operator, and $*$ denotes complex conjugate. However the power spectrum is blind to higher-order correlations introduced by non-linearities such as perspective projection. These correlations can however be estimated with higher-order spectra (e.g., [11], [15], [25]). For example the bispectrum estimates third-order correlations and is defined as:

$$B(\omega_1, \omega_2) = \mathcal{E} \{F(\omega_1)F(\omega_2)F^*(\omega_1 + \omega_2)\}. \quad (5)$$

Note that unlike the power spectrum the bispectrum of a real signal is complex-valued. The bispectrum measures the statistical dependence between frequency triples, (e.g., $[\omega_1, \omega_1, 2\omega_1]$ or $[\omega_1, \omega_2, \omega_1 + \omega_2]$). Specifically, if a triple of frequencies have independent phases, then the statistical averaging

by the expected value operator will cause the bispectrum to vanish. If, on the other hand, the triple of frequencies are non-linearly coupled, then the phases will be correlated and thus the statistical averaging will not cause the bispectrum to vanish. As a result, non-linearities can be revealed by analyzing the magnitude of the bispectrum.

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:

$$\hat{B}(\omega_1, \omega_2) = \frac{1}{N} \sum_{k=1}^N F_k(\omega_1)F_k(\omega_2)F_k^*(\omega_1 + \omega_2), \quad (6)$$

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 (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:

$$b^2(\omega_1, \omega_2) = \frac{|B(\omega_1, \omega_2)|^2}{\mathcal{E}\{|F(\omega_1)F(\omega_2)|^2\}\mathcal{E}\{|F(\omega_1 + \omega_2)|^2\}}. \quad (7)$$

It is straightforward 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:

$$\hat{b}(\omega_1, \omega_2) = \frac{|\frac{1}{N} \sum_k F_k(\omega_1)F_k(\omega_2)F_k^*(\omega_1 + \omega_2)|}{\sqrt{\frac{1}{N} \sum_k |F_k(\omega_1)F_k(\omega_2)|^2 \frac{1}{N} \sum_k |F_k(\omega_1 + \omega_2)|^2}}. \quad (8)$$

Note that the bicoherence is now a real-valued quantity. The bicoherence can be averaged across all frequencies to obtain a measure of overall correlation:

$$\frac{1}{N^2} \sum_{\omega_1=-N/2}^{N/2} \sum_{\omega_2=-N/2}^{N/2} \hat{b}\left(\frac{2\pi\omega_1}{N}, \frac{2\pi\omega_2}{N}\right). \quad (9)$$

This quantity is employed throughout this paper as a measure of higher-order correlations. In closing, we note that the bicoherence is invariant to linear transformation, thus making it applicable to a wide range of imaging conditions (e.g., blurring or scaling).

IV. ESTIMATING SURFACE ORIENTATION

The intuition behind our approach for estimating surface orientation is as follows. If a texture on a fronto-parallel plane has, as we assume, random phase, then the statistical averaging by the expected value operator will cause the bispectrum to vanish, Equation (5). If, on the other hand, the plane is slanted from fronto-parallel, then the non-linear perspective projection will cause phase coupling between the frequencies, and the statistical averaging will not cause the bispectrum to vanish. We will show empirically that the magnitude of the bispectrum is proportional to the deviation of the plane from fronto-parallel, and that the slant can therefore be estimated by minimizing the magnitude of the bispectrum.

We first derive an algorithm for estimating surface orientation for one-dimensional planes and then show how this algorithm can be extended to estimate the surface orientation of two-dimensional planes.

A. Synthetic Signals

Fractal signals were synthesized from a sum of sinusoids with amplitudes, $a_k = 1/k$, frequencies, $\omega_k = 2k\pi$, and random phases, $\phi_k \in [-\pi, \pi]$, as follows:

$$f(x) = \sum_{k=1}^n a_k \cos(\omega_k x + \phi_k), \quad (10)$$

with $n = 512$, and $x \in [-1, 1)$, and where the length of the signal is $2n$. The perspective projection of such a signal from an oriented one-dimensional plane (by an amount θ) can be expressed analytically. Without loss of generality, we assume that the plane is a unit length away from the sensor and that the focal length is one. The projected signal, $g(x)$, is synthesized directly as follows:

$$g(x) = \sum_{k=1}^n a_k \cos\left(\omega_k \frac{x'}{z'} + \phi_k\right), \quad (11)$$

where,

$$\begin{bmatrix} x' \\ z' \end{bmatrix} = \begin{bmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{bmatrix} \begin{bmatrix} x \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 2 \end{bmatrix}, \quad (12)$$

with x defined as above. Signals were synthesized in this manner to avoid possible interpolation artifacts that would be introduced by standard warping techniques.

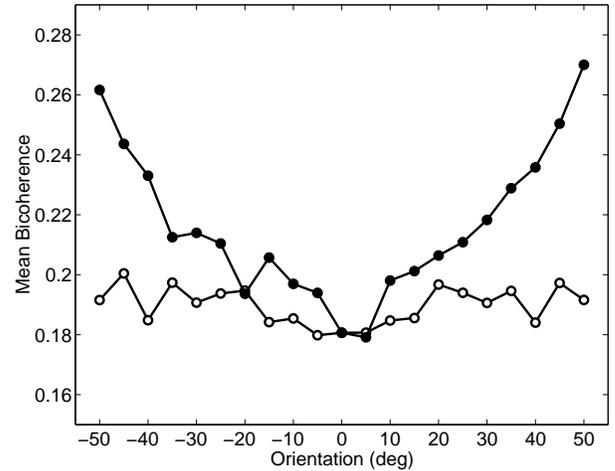
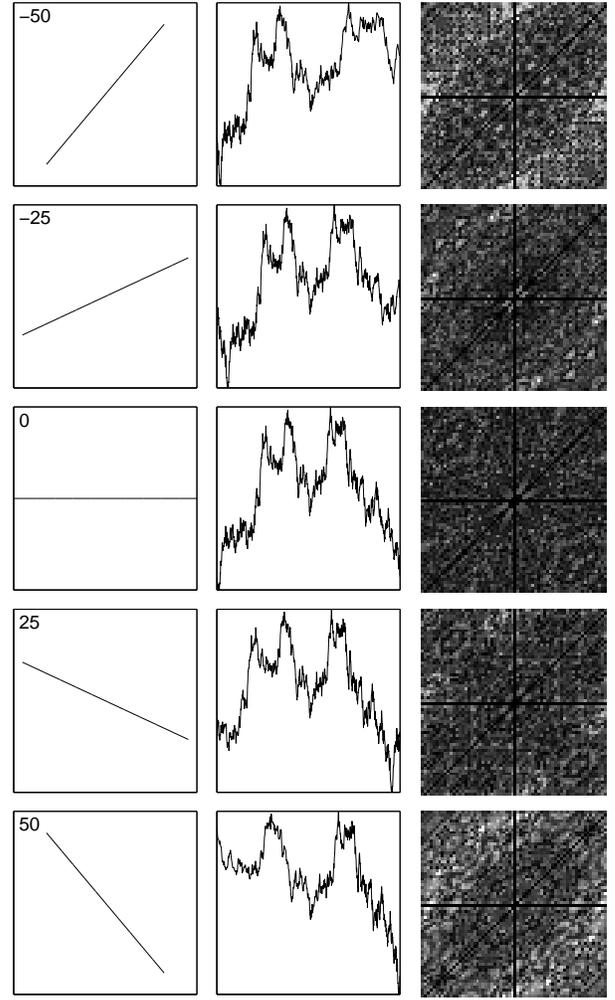


Fig. 3. Signals synthesized from a 1-D plane at several different orientations. Shown above are, from left to right, the 1-D plane, the projected signal, and the bicoherence of this signal. Shown below is the mean bicoherence, Equation (9), as a function of a full range of orientations. The filled circles correspond to the results from a perspective projection model, and the open circles to the same signal projected an orthographic projection. Note that under perspective projection, the mean bicoherence generally increases proportional to the deviation of the plane from horizontal (0°).

Shown in Fig. 3 is an example signal synthesized from a 1-D plane oriented at $-50^\circ, -25^\circ, 0^\circ, 25^\circ,$ and 50° . Shown in the right-most column of Fig. 3 is the estimated bicoherence - the bicoherence should be read similar to a 2-D Fourier transform with the origin in the center and the horizontal and vertical axes corresponding to frequencies ω_1 and ω_2 , respectively. Note the increase in overall magnitude as the plane deviates from fronto-parallel. Shown in the bottom panel of Fig. 3 is the mean bicoherence, Equation (9), as a function of a full range of orientations (filled circles). Also shown, for comparison, is the mean bicoherence when the fractal signal was synthesized under orthographic projection (open circles). Under orthographic projection, there is no systematic relationship between orientation and mean bicoherence. Under perspective projection, however, the mean bicoherence generally increases proportional to the deviation of the plane from horizontal (0°). This relationship suggests a simple algorithm for estimating orientation.

- 1) select a range of possible orientations;
- 2) for each orientation θ , assume that the signal lies on a 1-D plane oriented θ degrees away from fronto-parallel. Warp the signal under perspective projection to front-parallel, yielding a new signal f_θ ;
- 3) compute the bicoherence, Equation (8), of f_θ . Repeat for all of the orientations in the selected range;
- 4) select the value of θ that minimizes the mean bicoherence, Equation (9).

In subsequent sections, the range of possible orientations is fixed to be $[-30^\circ, 30^\circ]$ in steps of 5° . The signal is warped using bi-cubic interpolation and, because of the potential shortening of the signal due to the warping, only the central 512 samples are used to estimate the bicoherence. The bicoherence is computed by dividing the signal into overlapping segments of length 64 with an overlap of 32. A 64-point DFT (windowed with a symmetric Hanning window) is estimated for each zero-mean segment from which the bicoherence is estimated. These parameters were determined empirically, but in general we find that the results are not particularly sensitive to their choice. The value of θ that minimizes the mean bicoherence is taken to be the minimum of a second-order polynomial fit to the estimates across all orientations (this was done to yield more

reliable estimates in the face of potentially spurious estimates). And finally, estimates at either end of the orientation range, -30° or 30° , are discarded as outliers.

In our current MatLab implementation, and running on a 3GHz processor, the algorithm takes approximately 0.6 seconds to process a single 1024-length signal.

In the following section we show how this algorithm can be extended to estimate the orientation of two-dimensional planar surfaces, and show the efficacy of this algorithm on synthetic and natural images.

B. Synthetic Images

Similar to the synthetic signals of the previous section, fractal images were synthesized from a sum of two-dimensional sinusoids with amplitudes, $a_k = 1/k$, frequencies, $\omega_k = 2k\pi$, and random orientations, $\theta_k \in [-\pi, \pi]$, and random phases, $\phi_k \in [-\pi, \pi]$, as follows:

$$f(x, y) = \sum_{k=1}^n a_k \cos(\omega_k [\cos(\theta_k)x + \sin(\theta_k)y] + \phi_k), \quad (13)$$

where $n = 512$, $x \in [-1, 1)$, $y \in [-1, 1)$, and where the image is of size $2n \times 2n$. As before, the perspective projection of such an image from an oriented two-dimensional plane (by an amount θ_x and θ_y) can be expressed analytically. Without loss of generality, we assume that the plane is a unit length away from the sensor and that the focal length is one. The projected image, $g(x, y)$, is synthesized as follows:

$$g(x, y) = \sum_{k=1}^n a_k \cos\left(\omega_k \left[\cos(\theta_k)\frac{x'}{z'} + \sin(\theta_k)\frac{y'}{z'}\right] + \phi_k\right) \quad (14)$$

where,

$$\begin{bmatrix} x' \\ y' \\ z' \end{bmatrix} = R_y R_x \begin{bmatrix} x \\ y \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 2 \end{bmatrix}, \quad (15)$$

with x and y defined as above, and where the rotation matrices are defined as:

$$R_x = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\theta_x) & -\sin(\theta_x) \\ 0 & \sin(\theta_x) & \cos(\theta_x) \end{bmatrix} \quad (16)$$

$$R_y = \begin{bmatrix} \cos(\theta_y) & \sin(\theta_y) & 0 \\ 0 & 1 & 0 \\ -\sin(\theta_y) & \cos(\theta_y) & 0 \end{bmatrix} \quad (17)$$

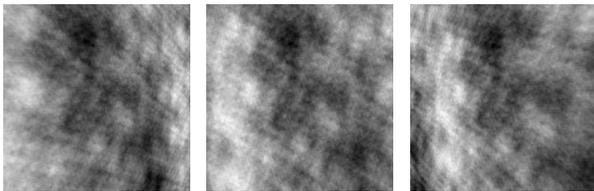


Fig. 4. Synthesized images from a 2-D plane rotated about vertical by -50° , 0° , and 50° .

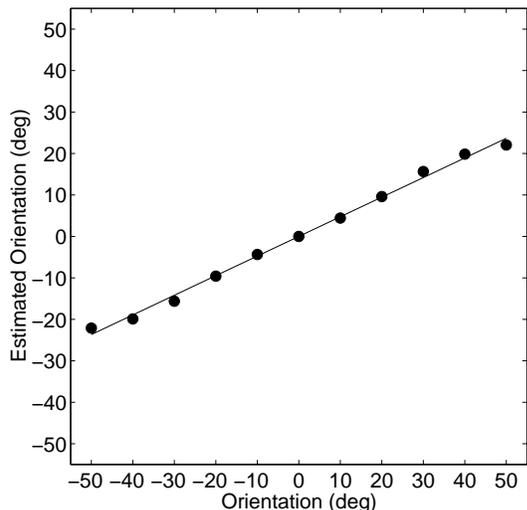


Fig. 5. Shown are the estimation results from synthetically generated images. Each data point corresponds to the average estimate from ten independently synthesized images (five of which correspond to estimating orientation about vertical, and five about horizontal).

Shown in Fig. 4 are example images from a plane rotated about vertical by -50° , 0° , and 50° .

In order to estimate the orientation about the vertical axis we independently estimate the orientation of each horizontal scan line, using the algorithm from the previous section, and average the results across the entire image. Similarly, in order to estimate the orientation about the horizontal axis we average the estimates from each vertical scan line. It is possible to perform a single two-dimensional search instead of the pair of one-dimensional searches. This would, of course, be more computationally demanding, and we have found that these two approaches yield very similar results.

Shown in Fig. 5 are the results of estimating the orientation from synthetically generated images rotated about either the vertical or horizontal axis. Each data point corresponds to the average estimate from ten independently synthesized images (five of which correspond to estimating orientation about vertical, and five about horizontal - these were aver-

	-30	-15	0	15	30
-30	-28.6 -30.9	-32.9 -15.1	-34.7 -0.9	-32.0 16.3	-29.0 33.6
-15	-13.3 -31.1	-15.2 -14.4	-14.7 -1.0	-14.5 15.4	-13.1 33.7
0	0.7 -31.1	0.9 -14.0	1.0 -1.0	0.7 15.2	1.0 34.1
15	13.1 -30.5	13.9 -14.6	14.2 -1.1	14.2 15.4	13.4 34.2
30	26.9 -30.9	29.0 -15.0	30.8 -1.0	29.8 15.8	27.0 33.2

TABLE I

SHOWN ARE THE ESTIMATION RESULTS (IN DEGREES) FROM SYNTHETICALLY GENERATED IMAGES. EACH ROW/COLUMN HEADING CORRESPONDS TO THE ACTUAL HORIZONTAL/VERTICAL ORIENTATION, AND EACH ENTRY CORRESPONDS TO THE HORIZONTAL/VERTICAL ESTIMATE, AVERAGED ACROSS TEN INDEPENDENTLY SYNTHESIZED IMAGES.

aged together since there were no significant differences between their estimates). Note that in all cases the estimated orientation is approximately one-half of the actual orientation (a fitted line yields a slope of 0.4734 and intercept of 0.0028). This bias in the estimation is due to the correlations introduced by the interpolation artifacts from warping each scan line prior to estimating the bicoherence [18] (bi-cubic interpolation was used in the warping). Since this warping is performed in the same way regardless of the underlying signal/image, the bias is consistent and can therefore be calibrated for.

In the results of Fig. 5 the synthetically generated images were rotated either about vertical or horizontal. In practice, of course, it is necessary to simultaneously measure both orientations. A full two-dimensional search over both orientations would be computationally costly. We have found, fortunately, that each orientation can be computed (as above) independent of one another. Shown in Fig. I are the results of simultaneously estimating both vertical and horizontal orientations. These results were averaged across ten independently synthesized images and passed through the calibration from Fig. 5. Across all orientations, the average estimation error is 1.4° with a standard deviation of 0.9° , and a maximum error of 4.7° .

In our current MatLab implementation, and running on a 3GHz processor, the algorithm takes approximately 19 minutes to process an image of size 1024×1024 pixels.

C. Noise Sensitivity

We tested the sensitivity of our algorithm to additive noise. Noise was added to a synthetic image

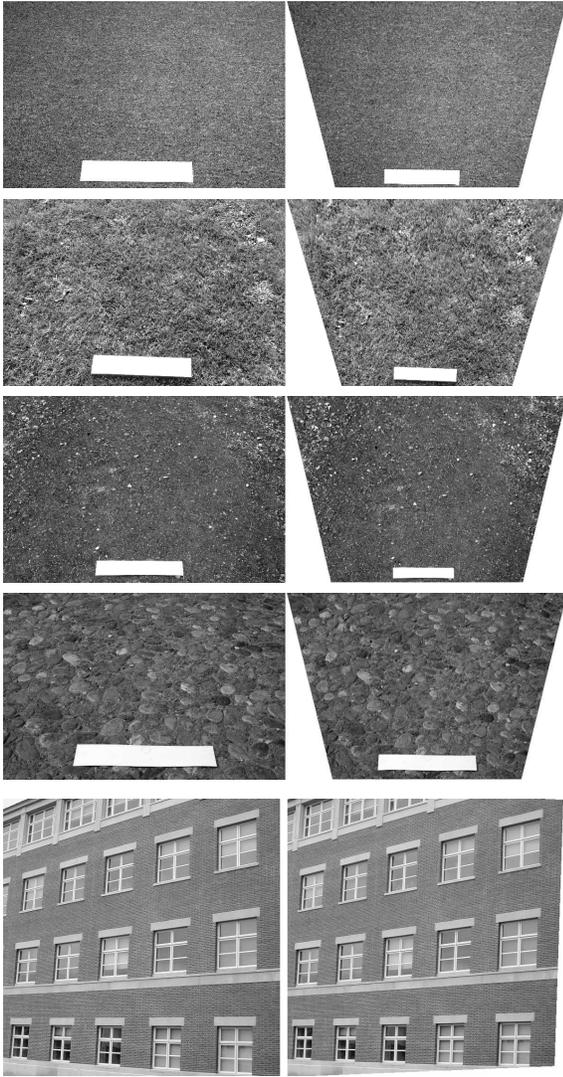


Fig. 6. Shown on the left is the original perspective distorted image. Shown on the right is the image rectified by the estimated surface orientation. The bottom panel demonstrates a failure of our algorithm.

with a horizontal and vertical orientation of -15° and 0° , Fig. 4. Uniform white noise, ranging from 15dB to 40dB, was added to the synthetic image. Here we report on the errors in estimating the horizontal orientation averaged over ten independently synthesized images. At noise levels of 40dB, 35dB, 30dB, 25dB, 20dB, and 15dB, the averaged estimates were -15.0 , -15.0 , -14.5 , -11.8 , -8.8 , and -4.5 , respectively. Note the graceful degradation as the signal-to-noise ratio (SNR) decreases. At low SNR, the additive noise, with random phase, dominates and the image begins to appear as fronto-parallel. The estimation is, nevertheless, reasonably robust in the presence of noise.

D. Natural Images

Shown in the top four panels of Fig. 6 are images of a carpet, grass, dirt road and stones (left). Also shown are the rectified images after correcting for the estimated horizontal orientation (right). The estimation was identical to that described in the previous section with all of the parameters of the algorithm held fixed. A small planar calibration target of known geometry was placed in each scene and used to determine ground-truth as described in [24]. With this technique, the actual orientation of each planar surface about the horizontal axis is 33, 32, 28 and 29 degrees, respectively, with a less than one degree rotation about the vertical axis. The estimated orientations using our bispectral method was 28, 33, 25 and 26 degrees, for an average error of 3 degrees. Note that the portion of the image containing the calibration target was not used in the estimation.

In the bottom panel of Fig. 6 is an example of where our algorithm fails to accurately estimate orientation. The reason for this failure is that the basic assumption of random phase, Section III, is violated in the periodic brick pattern. We verified this by generating synthetic images whose phases were correlated – the resulting estimation of orientation degraded as a function of the amount of correlation.

V. DISCUSSION

We have presented a direct method for estimating the orientation of a plane from a single view under perspective projection. This technique exploits higher-order correlations in the frequency domain that are introduced by perspective projection. These correlations, when minimized, yield the plane’s orientation.

Our proposed method adds to an existing body of literature in shape from texture. This technique has the advantage that it is applicable to images in which geometric features are not easily extracted. Like all single-view techniques, we must make some assumption regarding the underlying planar texture - in our case, we make an assumption that the texture has random phase. Specifically, if a texture on a fronto-parallel plane has random phase, then the statistical averaging by the expected value operator will cause the bispectrum to vanish. If the plane is slanted from fronto-parallel, however, the non-linear

perspective projection will cause phase coupling between the frequencies, and the averaging will not cause the bispectrum to vanish. We have empirically shown that the magnitude of the bispectrum is proportional to the deviation of the plane from fronto-parallel, and that the slant can be estimated by minimizing the magnitude of the bispectrum. If the phase of the original texture is not random, however, then the correlations due to the non-linear projection are confounded with the correlations due to the structural properties of the texture. As a result, the orientation cannot be reliably estimated, as shown in the bottom panel of Fig. 6.

Our proposed technique builds on earlier work, where we showed that luminance non-linearities introduced by gamma correction, and geometric non-linearities introduced by lens distortion could be estimated and corrected for in the absence of any other information or calibration. With the recent interest in understanding and modeling the statistics of natural images [20], [26], [13], [21], these techniques may provide interesting insights into the statistical properties of images.

ACKNOWLEDGMENTS

Hany Farid was supported by an Alfred P. Sloan Fellowship, a National Science Foundation CAREER Award (IIS-99-83806), and a departmental National Science Foundation Infrastructure Grant (EIA-98-02068). Jana Košecká was supported by a National Science Foundation CAREER Award (IIS-03-4774).

REFERENCES

- [1] A. Criminisi, I. Reid, and A. Zisserman. Single view metrology. *International Journal of Computer Vision*, 40(2):11–20, 2000.
- [2] H. Farid. Blind inverse gamma correction. *IEEE Transactions on Image Processing*, 10(10):1428–1433, 2001.
- [3] H. Farid and A.C. Popescu. Blind removal of lens distortions. *Journal of the Optical Society of America*, 18(9):2072–2078, 2001.
- [4] R.T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading. *IEEE Transactions on Pattern and Machine Intelligence*, 10(7):439–451, 1988.
- [5] J. Garding. Shape from texture for smooth curved surfaces in perspective projection. *Journal of Mathematical Imaging and Vision*, 2(2):327–350, 1992.
- [6] J. Garding. Surface orientation and curvature from differential texture distortion. In *International Conference on Computer Vision*, 1995.
- [7] J. J. Gibson. *The perception of the visual world*. Houghton Mifflin, 1950.
- [8] T. Greiner and S. Das. Recovering orientation of a textured planar surface using wavelet transform. In *ICVGIP*, 2002.
- [9] B.K.P. Horn. *The Psychology of Machine Vision*, chapter Obtaining shape from shading information, pages 115–155. McGraw-Hill, 1975.
- [10] K. Kanatani and T. C. Chou. Shape from texture: General principle. *Artificial Intelligence*, 38:1–48, 1989.
- [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. Kosecka and W. Zhang. Extraction, matching and pose recovery based on dominant rectangular structures. In *International Conference on Computer Vision, Workshop on Higher Level Knowledge in Vision*, 2003.
- [13] G. Krieger, C. Zetsche, 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, 1997.
- [14] S. Mann and R. Picard. Video orbits of projective group: A simple approach to featureless estimation of parameters. *IEEE Transactions on Image Processing*, 6(9):1281–1295, 1997.
- [15] 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, 1991.
- [16] A. Pentland. Shading into texture. *Artificial Intelligence*, 29(2):147–170, 1986.
- [17] H. Permuter and J. Francos. Estimating the orientation of planar surfaces: Algorithms and bounds. *IEEE Transactions on Information Theory*, 46(6):1908–1920, 2000.
- [18] Alin C. Popescu. *Statistical Tools for Digital Image Forensics*. PhD thesis, Dartmouth College, Hanover, NH, December 2004. Dartmouth Computer Science Tech. Rep. TR2005-531.
- [19] R. Rozenholtz and J. Malik. Computing local surface orientation and shape from texture. *International Journal of Computer Vision*, 2(23):149–168, 1997.
- [20] D.L. Ruderman and W. Bialek. Statistics of natural image: Scaling in the woods. *Phys. Rev. Letters*, 73(6):814–817, 1994.
- [21] E.P. Simoncelli. Modeling the joint statistics of images in the wavelet domain. In *Proceedings of the 44th Annual Meeting*, volume 3813, 1999.
- [22] B. Super and A. Bovik. Shape from texture using local spectral moments. *IEEE Transactions on Pattern and Machine Intelligence*, 17(4):333–342, 1995.
- [23] B.J. Super and A.C. Bovik. Planar surface orientation from texture spatial frequencies. *Pattern Recognition*, 28(5):729–743, 1995.
- [24] W. Zhang and J. Kosecka. Extraction, matching and pose recovery based on dominant rectangular structures. *Computer Vision and Image Understanding*, 100:274–293, 2005.
- [25] G. Zhou and G. B. Giannakis. Polyspectral analysis of mixed processes and coupled harmonics. *IEEE Transactions on Information Theory*, 42(3):943–958, 1996.
- [26] 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.