

Shape from Periodic Texture Using the Eigenvectors of Local Affine Distortion

Eraldo Ribeiro and Edwin R. Hancock

Abstract—This paper shows how the local slant and tilt angles of regularly textured curved surfaces can be estimated directly, without the need for iterative numerical optimization. We work in the frequency domain and measure texture distortion using the affine distortion of the pattern of spectral peaks. The key theoretical contribution is to show that the directions of the eigenvectors of the affine distortion matrices can be used to estimate local slant and tilt angles of tangent planes to curved surfaces. In particular, the leading eigenvector points in the tilt direction. Although not as geometrically transparent, the direction of the second eigenvector can be used to estimate the slant direction. The required affine distortion matrices are computed using the correspondences between spectral peaks, established on the basis of their energy ordering. We apply the method to a variety of real-world and synthetic imagery.

Index Terms—Shape-from-texture, spectral analysis, affine distortion, eigen-analysis.

1 INTRODUCTION

THE recovery of surface shape using texture information is a process that is grounded in psychophysics [1]. Moreover, it has been identified by Marr [2] as being a potentially useful component of the $2\frac{1}{2}$ D sketch. Stated succinctly, the problem is as follows: Given a two-dimensional image of a curved textured surface, how can the three-dimensional shape of the viewed object be recovered? This is clearly an ill-posed problem. In order to be rendered tractable, restrictive simplifications must be made. These frequently hinge on the assumed periodicity, homogeneity [3], and isotropy of the underlying surface texture. When viewed from the standpoint of shape recovery, there are two distinct areas of activity in the literature. The first of these confines its attention to planar surfaces and focuses on the recovery of perspective geometry from texture gradient or vanishing point location [4], [5], [6], [7]. The second problem is that of interpreting the geometry of curved surfaces [8], [9], [10], [11], [12]. Without detailed knowledge of the viewing geometry or camera optics, the latter problem involves recovering local surface orientation from texture gradient.

This second problem is perhaps the most interesting from the standpoint of shape perception and is the focus of this paper. The problem can be approached in two ways. The first of these is to use presegmented structural texture primitives and to measure texture gradient using the variation in the size of the primitives. This approach will clearly be sensitive to the process used to segment the required texture primitives. However, the method can be used with aperiodic textures. The second commonly used method is to adopt a frequency domain representation of the underlying texture distribution. The frequency domain approach can be applied more directly to the raw image data and is potentially less sensitive to the segmentation process. However, the applicability of the method is restricted to periodic textures.

Although it is not dependant on the reliable detection of texture primitives or the measurement of texture gradient, the frequency

domain approach to curved shape-from-texture can be criticized on a number of grounds. First, in the methods of both Krumm and Shafer [13] and Malik and Rosenholtz [10], [11] the recovery of local surface orientation is based on either numerical optimization or exhaustive search. This is done to overcome the problems associated with the fact that no fronto-parallel sample of the texture is to hand. In addition, the method of Rosenholtz and Malik [10], [11] has difficulty in distinguishing between curved and planar surfaces, is sensitive to initial parameter values, and needs curvature to be specified as a parameter.

To overcome these shortcomings, our aim in this paper is to present a closed form method for recovering shape from curved or planar textured surfaces using frequency information. We follow Krumm and Shafer [13], Super and Bovik [5], [14], Malik and Rosenholtz [10], [11] by measuring the local texture variations due to the perspective using the frequency domain affine distortion of the pattern of spectral peaks. Our main contribution is to show that the directions of the eigenvectors of the affine distortion matrix can be used to directly estimate the slant and tilt angles of local tangent planes to curved surfaces. This result applies under the assumption that the underlying surface is painted with a uniform texture and is viewed under perspective projection onto the image plane. The texture is assumed to be homogeneous but not necessarily isotropic. Our main contribution is, therefore, to develop a direct and simple method for estimating slant and tilt angles.

The main advantage of the method presented in this paper over those of Super and Bovik [14], Krumm and Shafer [15], [4], [13] and Malik and Rosenholtz [11] is that it delivers slant and tilt angles by computing the eigenvectors of the affine distortion matrix for spectral peaks. These related methods, differ from our proposed method in the way in which the affine distortion of the texture spectrum is used to measure the local surface orientation of curved surfaces. Specifically, previous work in this area concentrates on using numerical optimization or exhaustive search to recover the local slant and tilt angles. By contrast, our method recovers the parameters in closed form. To provide some specific examples, Krumm and Shafer [13] numerically adjust the slant and tilt angles so as to minimize a “frontalization” error between the back-projected texture spectra. Although less direct than our method of orientation estimation, this backprojection method involves a feedback loop. The lack of feedback between the measured texture spectra and the estimated orientation field may be viewed as a weakness of our method. Super and Bovik [14], [5], on the other hand, invoke the homogeneity assumption and recover the surface orientation parameters which minimize the variance of the backprojected frequency vectors of the spectral peaks. Rosenholz and Malik [10], [11] have a more sophisticated method which involves using a singular value decomposition method to minimize a χ^2 measure of goodness of fit. The recovered parameters are the slant and tilt angles together with local curvature.

2 LITERATURE REVIEW

Before we proceed to detail our new shape-from-texture method, we pause to review the related literature. The topic of shape-from-texture has been studied in the computer vision literature for almost three decades. Early work hinged on the use of texture-gradients [16], [17]. Bajcsy and Lieberman were among the first to demonstrate the use of texture gradient as a depth cue [17]. Drawing on psychophysics, Stevens has provided an analysis of the information content of texture-gradient [18] and has shown how it can be used to recover surface orientation through the estimation of slant and tilt angles [19]. Witkin [20] has developed a statistical method for recovering local orientation and, hence, surface shape from natural two dimensional imagery. The method commences by assuming a uniform distribution of edge orientation, i.e., an isotropic texture,

- E.R. Hancock is with the Department of Computer Science, University of York, York YO1 5DD, UK. E-mail: erh@cs.york.ac.uk.
- E. Ribeiro is with Oxford Metrics, Botley, Oxford, UK.

Manuscript received 31 Jan. 2000; revised 7 Nov. 2000; accepted 29 May 2001.

Recommended for acceptance by M. Shah.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number 111351.

and models how this distribution transforms under perspective geometry. The shape of the edge orientation distribution can be used to estimate slant and tilt angles. Several authors have developed methods which assume that a fronto-parallel view of a planar sample of the texture is available. For instance, Ikeuchi [21] has shown how the process of estimating surface shape from regular texture patterns can be simplified using spherical projection. Aloimonos and Swain [22] have computed a texture analogue of the reflectance map from the area gradient of the fronto-parallel texture elements. The regularized variational method of Ikeuchi and Horn [23] for shape-from-shading is adapted to recover surface orientation. In this way, poor initial surface orientation estimates can be improved using iterative relaxation operations [21], [22]. However, the smoothing of the field of surface normals is based on a simple neighborhood averaging method.

Several authors have attempted to elucidate more general frameworks for shape-from-texture. For instance, Blake and Marinos [3] have critically assessed Witkin's [20] edge isotropy assumption. They have developed an optimization-based method for recovering surface orientation from second-order texture moments. Kanatani and Chou [24] have further extended Witkin's [20] work by modeling the statistical transformation of texture density. They consider the effects of perspective geometry when isotropy and homogeneity assumptions apply. Garding [8], [9] has developed an elegant framework based on differential geometry. This simplifies the analysis of perspective geometry and provides a mathematical setting in which curvature can be directly recovered from weakly isotropic textures.

More recently, the use of spectral information has been explored as an alternative to texture-gradient information. The use of frequency domain or spectral measurements represents a way of overcoming some of the restrictive requirements imposed by the need to work with accurately determined texture gradients. Moreover, it increases the range of natural textures that can be accommodated. Frequency properties can be captured using a number of different representations including the Fourier transform, Gabor wavelets, and the Wigner distribution [25], [5], [10], [26], [27]. The frequency domain method has been used extensively in the recovery of the perspective pose of texture planes. In an important series of papers, Krumm and Shafer [15], [4], [13] recover slant and tilt angles for periodic surface textures. They commence by introducing the idea of using measurements of the affine distortion to recover the parameters of perspective projection from spectral peaks [15]. Parameter estimation is based on exhaustive numerical search. The method is later extended to provide a means of segmenting multiple planar patches from textured images [4]. Super and Bovik [5], [14] have an equivalent method which focuses on linearizing the Jacobian of the perspective transformation in the frequency domain. The method shares with that of Krumm and Shafer the feature of using exhaustive numerical search.

Turning our attention to curved surfaces, Krumm and Shafer [13] have shown how to relate the frequency distortion between local planar patches under an affine transform. Malik and Rosenholtz [10], [11] use Garding's framework [8] to estimate local orientation and curvature. Whereas Garding uses departures from isotropy as the texture measurement, Malik and Rosenholtz use Krumm and Shafer's ideas to construct affine distortion measures. The texture measurements are based on frequency domain derivatives. The method uses numerical minimization to recover the five parameters needed to estimate orientation and curvature. The method requires a good initial estimate of the magnitude of the curvature. It can be viewed as minimizing a measure of spectral back-projection error.

3 SPECTRAL DISTORTION

This paper is concerned with recovering a dense map of surface orientations for surfaces which are uniformly painted with periodic textures. Our approach is a spectral one which is couched in the Fourier domain. We make use of an analysis of spectral distortion under perspective geometry extensively developed by Super and Bovik [5], Krumm and Shafer [4], [13], and by Malik and Rosenholtz [10], [11] among others. This analysis simplifies the full perspective geometry of texture planes using a local affine approximation. Suppose that \mathbf{U}_i represents the frequency vector associated with a spectral peak detected at the point with position-vector \mathbf{X}_i on a texture plane. Further, let \mathbf{U}_i and \mathbf{X}_i represent the corresponding frequency vector and position vector when this texture plane undergoes perspective projection onto the image plane. If the perspective projection can be locally approximated by an affine distortion $T_A(\mathbf{X}_i)$, then the relationship between the texture-plane and image-plane frequency vectors is $\mathbf{U}_i = T_A(\mathbf{X}_i)^{-T} \mathbf{U}_t$. The affine distortion matrix is given by

$$T_A(\mathbf{X}_i) = \frac{\Omega}{hf \cos \sigma} \begin{bmatrix} x_i \sin \sigma + f \cos \tau \cos \sigma & -f \sin \tau \\ y_i \sin \sigma + f \sin \tau \cos \sigma & f \cos \tau \end{bmatrix}, \quad (1)$$

where f is the focal length of the camera, σ is the slant angle of the texture-plane, τ is the tilt angle of the texture-plane, and $\Omega = f \cos \sigma + \sin \sigma(x_i \cos \tau + y_i \sin \tau)$.

We aim to exploit this property to recover shape-from-texture. Provided that the texture distribution painted on a curved surface is homogeneous, observed changes in the image plane texture pattern can be attributed to variations in surface orientation. Our aim is to compute the slant and tilt angles of local tangent planes to a textured surface using the observed distortions of the texture spectrum across the image plane. To do this, we measure affine distortion between corresponding spectral peaks.

Consider the point S on the curved texture surface. Suppose that the neighborhood of this point can be approximated by a local planar patch. This planar patch undergoes perspective projection onto the image plane. Further suppose we sample the texture projection of the local planar patch at two neighboring points A and B laying on the image plane. The coordinate vectors of the two points are, respectively, $\mathbf{X}_A = (x, y)^T$ and $\mathbf{X}_B = (x + \Delta x, y + \Delta y)^T$, where Δx and Δy are the image-plane displacements between the two points.

Suppose that the local planar patch on the texture surface has a spectral peak with frequency vector $U_S = (u_s, v_s)^T$. On the image plane, the corresponding frequency vectors for the spectral peaks at the points X_A and X_B are respectively $U_A = (u_A, v_A)^T$ and $U_B = (u_B, v_B)^T$. Using the Fourier domain affine projection property described above, the texture-surface peak frequencies are related to the image plane peak frequencies via the equations $U_A = (T_A(X_A)^{-1})^T U_S$ and $U_B = (T_A(X_B)^{-1})^T U_S$, where $T_A(X_A)$ is the local affine approximation to the perspective projection of the planar surface patch at the point A and $T_A(X_B)$ is the corresponding affine projection matrix at the point B . The consequence of this property is that the frequency vectors for the two corresponding spectral peaks on the image-plane are related to one another via the local distortion $U_B = (T_A(X_A)T_A(X_B)^{-1})^T U_A$. As a result, the texture-surface spectral distortion matrix $\Phi = (T_A(X_A)T_A(X_B)^{-1})^T$ is a 2x2 matrix. This matrix relates the affine distortion of the image plane frequency vectors to the 3D orientation parameters of the local planar patch on the surface. Moreover, it does not require a sample of the fronto-parallel texture since we have eliminated U_S . Substituting for the affine approximation to the perspective transformation from (1), the required matrix is given in terms of the slant and tilt angles as

$$\Phi = \frac{\Omega(\mathbf{A})}{\Omega^2(\mathbf{B})} \begin{bmatrix} \Omega(\mathbf{A}) + \Delta y \sin \sigma \sin \tau & -\Delta y \sin \sigma \cos \tau \\ -\Delta x \sin \sigma \cos \tau & \Omega(\mathbf{A}) + \Delta x \sin \sigma \sin \tau \end{bmatrix}, \quad (2)$$

where $\Omega(A) = f \cos \sigma + \sin \sigma(x \cos \tau + y \sin \tau)$ and

$$\Omega(B) = f \cos \sigma + \sin \sigma((x + \Delta x) \cos \tau + (y + \Delta y) \sin \tau).$$

The above matrix represents the distortion of the spectrum sampled at the location B with respect to the sample at the location A . In the next section, we show how to solve directly for the parameters of surface orientation, i.e., the slant and tilt angles, using the eigen-structure of the transformation matrix Φ .

4 EIGENSTRUCTURE OF THE AFFINE DISTORTION MATRIX

Let us consider the eigenvector equation $\Phi \mathbf{w}_i = \lambda_i \mathbf{w}_i$ for the distortion matrix Φ , where $\lambda_i, i = 1, 2$ are the eigenvalues of the distortion matrix Φ and \mathbf{w}_i are the corresponding eigenvectors. Since Φ is a 2x2 matrix the two eigenvalues are found by solving the quadratic eigenvalue equation $\det[\Phi - \lambda I] = 0$ where I is the 2x2 identity matrix. The explicit eigenvalue equation is

$$\lambda^2 - \text{Trace}(\Phi)\lambda + \det(\Phi) = 0, \quad (3)$$

where $\text{Trace}(\Phi)$ and $\det(\Phi)$ are the trace and determinant of Φ . Substituting for the elements of the transformation matrix Φ , we have

$$\lambda^2 - \left[\frac{\Omega(\mathbf{A})}{\Omega(\mathbf{B})} + \frac{\Omega^2(\mathbf{A})}{\Omega^2(\mathbf{B})} \right] \lambda + \left[\frac{\Omega(\mathbf{A})}{\Omega(\mathbf{B})} \times \frac{\Omega^2(\mathbf{A})}{\Omega^2(\mathbf{B})} \right] = 0. \quad (4)$$

The two eigenvalue solutions of the above quadratic equation are $\lambda_1 = \frac{\Omega^2(\mathbf{A})}{\Omega^2(\mathbf{B})}$ and $\lambda_2 = \frac{\Omega(\mathbf{A})}{\Omega(\mathbf{B})}$. The corresponding eigenvectors are

$$\mathbf{w}(\lambda_1) = [1, \tan \tau]^T$$

and

$$\mathbf{w}(\lambda_2) = \left[1, -\frac{\Delta x}{\Delta y} \right]^T.$$

As a result, we can directly determine the tilt angle from the vector components of the eigenvector associated with the eigenvalue λ_1 . The intuitive justification for this result is that under perspective projection the only direction which remains invariant at all locations on the image plane is the tilt direction. As a result, a frequency vector which is aligned in the tilt direction will maintain a constant angle, but it will change in magnitude according to the position on the image plane. In other words, the tilt direction is an eigenvector of the local affine transformation.

$$\tau = \arctan \left(\frac{\mathbf{w}_y(\lambda_1)}{\mathbf{w}_x(\lambda_1)} \right). \quad (5)$$

Once the tilt angle has been obtained, the slant angle can be recovered using the second eigenvalue via the relationship

$$\sigma = \arctan \left[\frac{f(\lambda_2 - 1)}{\zeta_1 + \zeta_2} \right], \quad (6)$$

where $\zeta_1 = (y(1 - \lambda_2) - \lambda_2 \Delta y) \sin \tau$ and

$$\zeta_2 = (x(1 - \lambda_2) - \lambda_2 \Delta x) \cos \tau.$$

With the slant and tilt angles to hand the surface normal \mathbf{n} may be computed.

5 COMPUTING LOCAL PLANAR ORIENTATION

In this section, we explain how to recover local planar surface orientation using our local affine distortion method. Our method shares with Krumm and Shafer [4] and Super and Bovik [5] the feature of using the affine distortion of spectra to estimate surface orientation. However, these two methods recover surface orientation by exhaustive spectral back-projection and error enumeration for all slant and tilt angles. The orientation is selected so as to numerically minimize a back-projection error. This is clearly a highly time consuming process due to the amount of search and numerical minimization required. Moreover, unless good initialization values are to hand, the method is prone to convergence to local optima.

Instead, our method solves for the local surface orientation parameters in closed-form. To directly recover the planar orientation angles, we use the eigenvectors of the spectral distortion matrix. We assume that the texture is homogeneous over the entire surface. The consequence of this assumption is that the spectral content of the texture does not change systematically over the curved texture surface. As a result of this assumption, the local spectral distortions measured on the image plane are attributable solely to changes in perspective. Shape effects such as changes in local surface orientation, which are attributable to surface curvature, must be assessed at a more global level.

In order to obtain a smooth spectral response, we use the Blackman-Tukey power spectrum estimator. This is defined to be the frequency response of the windowed autocorrelation function. We employ a triangular smoothing window $w(\mathbf{X})$ due to its well-documented spectral stability [28]. The spectral estimator is then

$$P(\mathbf{U}_i)^{BT} = \mathcal{F}\{r_{xx}(\mathbf{X}_i) \times w(\mathbf{X}_i)\}, \quad (7)$$

where r_{xx} is the estimated autocorrelation function of the image patch. We locate peaks in the output of the spectral estimator using a simple mode finding algorithm. This involves thresholding the spectral energy and locating the centroids of the regions of supra-threshold response. We use a spectral window of 32x32 pixels.

The first step in orientation recovery is to estimate the affine distortion matrix which represents the transformation between different local texture regions on the image plane. These image texture regions are assumed to belong to a single local planar patch on the curved texture surface. We do this by selecting pairs of neighboring points on the image plane. At each point, there may be several clear spectral peaks. Since the affine distortion matrix Φ has four elements that need to be estimated, we need to know the correspondences between at least two different spectral peaks at the different locations. Suppose that $U_1^{p_1} = (u_1^{p_1}, v_1^{p_1})^T$ and $U_1^{p_2} = (u_1^{p_2}, v_1^{p_2})^T$ represent the frequency vectors for two distinct spectral peaks located at the point with coordinates $\mathbf{X}_1 = (x_1, y_1)^T$ on the image plane. The frequency vectors are used to construct the columns of a 2x2 spectral measurement matrix $V_1 = (U_1^{p_1} | U_1^{p_2})$. Further, suppose that $U_2^{p_1} = (u_2^{p_1}, v_2^{p_1})^T$ and $U_2^{p_2} = (u_2^{p_2}, v_2^{p_2})^T$ represent the frequency vectors for the corresponding spectral peaks at the point $\mathbf{X}_2 = (x_2, y_2)^T$. The corresponding spectral measurement matrix is $V_2 = (U_2^{p_1} | U_2^{p_2})$. Under the affine model presented in Section 3, the spectral measurement matrices are related via the equation $V_2 = \Phi V_1$. As a result the local estimate of the spectral distortion matrix is $\Phi = V_2 V_1^{-1}$.

In practice, we only make use of the most energetic peaks appearing in the power spectrum. That is to say, we do not consider the detailed distribution of frequencies. Our method requires that we supply correspondences between spectral peaks so that the distortion matrices can be estimated. There are a number of ways in which this can be effected. Here, we are primarily interested in demonstrating our new property of the affine distortion matrix to




Fig. 1. Affine distortion of the local spectra across the image plane. The arrows are the estimated normal vectors at the center of the local planes. The ellipses represent the affine distortion of circles under the local perspective geometry of the surface.

recover local surface orientation. For this reason, we draw on a heuristic which uses the energy amplitude of the peaks to establish the required correspondences. This is done by ordering the peaks according to their energy amplitude. The ordering of the amplitudes of peaks at different image locations establishes the required spectral correspondences. A more satisfactory method would be to use a large number of spectral peaks and to use a maximum-likelihood method to align the pattern of peaks and recover the associated affine distortion matrix. In fact, we have recently reported an expectation-maximization algorithm which can be used for aligning point-sets under an affine transformation when the correspondence information is not available and when there is point drop-out or contamination [29]. There is clearly scope for using this method in the spectral domain.

After estimating the affine transformation between two local spectral peaks, we can directly apply the eigenvector analysis described in Section 5 to estimate the tilt and the slant angles by using (5) and (6). However, there are some obstacles to the direct estimation of the local spectral distortion. The first of these is related to the choice of spectral scale. If the window size used in the spectral estimator is mismatched to the frequency bandwidth of the texture, then the power spectrum becomes defocused and poor peak localization results. The second problem arises if there is no significant affine distortion between the corresponding spectral peaks used in the analysis. In other words, in choosing the locations of the spectral samples we must strike a compromise. If the locations are too close to one-another then we risk poor orientation estimation since the local affine distortions due to global perspectivity are too small to detect. If, on the other hand, the points are too far apart or span a high curvature feature, then the local distortions in the power spectrum are likely to be due to significant changes in local surface orientation. We can overcome this latter effect through smoothing the field of local orientation estimates.

There are several artifacts of the texture images which may further limit the method described in this paper. The first of these is the effect of perspectivity on image brightness. To reduce this problem, we perform image equalization prior to spectral analysis. However, it is important to stress that the Fourier power spectrum is invariant to contrast reversal. As a result, our method is not sensitive to the perceived brightness or darkness dominance of the textures used in our experiments. Finally, certain textures may exhibit aliasing effects such as Moire fringes. These are particularly common in images of buildings or other objects which contain glass surfaces. For the data studied here, there are no such artifacts.

However, the presence of such effects may confound our shape-from-texture method.

In Fig. 1, we provide an illustration of the affine distortion of the local spectra across a plane viewed under perspective geometry and across a curved surface. At each location, we show the computed spectra. The brightness of the peaks is proportional to their energy. Notice that the energy ordering is preserved from location to location. It is also important to note that the distortions due to curvature are greater than those due to perspectivity.

The orientation estimates returned by the new shape-from-texture method are likely to be noisy and inconsistent when viewed from the perspective of local smoothness. In order to improve the consistency of our needle map and, hence, the surface shape description, we employ an iterative smoothing process to update the estimated normal vectors. Aloimonos and Swain [22] have used Horn's [30] local averaging method for this purpose. However, in order to avoid the oversmoothing of local surface detail associated with high curvature features, we use the robust vector field smoothing method of Worthington and Hancock [31]. Rather than using a quadratic penalty of the sort which underpins the method of Aloimonos and Swain [22], this method uses robust error kernels, to gauge the effect of the smoothness error. The reason for this is that the quadratic penalty grows indefinitely with increasing smoothness error. This can have the undesirable effect of oversmoothing genuine surface detail. Examples of such surface structures include ridges and ravines. By using robust error kernels, the Worthington and Hancock method moderates the effects of smoothing over regions of genuine surface detail and allow a more faithful topographic representation to be recovered [31].

6 EXPERIMENTS

We have experimented with both synthetic surfaces with known ground-truth and real-world images. The former are used to assess the accuracy of the method, while we use the latter to demonstrate the practical utility of the method.

6.1 Real World Textures

In this section, we experiment with real world textured surfaces. We have generated the images used in this study by moulding regularly textured sheets into curved surfaces. The images used in this study are shown in the first row of Fig. 2. There are two sets of images. The first two have been created by placing a tablecloth with a rectangular texture pattern on top of the corner of a box and a balloon. The second group of images have been created by




Fig. 2. Real curved surfaces. Top-row: original image; second-row: recovered needle map; third-row: smoothed needle map.

bending a regularly textured sheet of wrapping paper into various tubular shapes. The first of these is a cylinder, the second is a “wave,” while the final example is an irregular set of folds. The textures in all five images show strong perspective effects.

The remaining two rows of Fig. 2 show the initial needle-map and the final smoothed needle-map. There are clear regions of needle-map consistency. When robust smoothing is applied to the initial needle maps, then there is a significant improvement in the directional consistency of the needle directions. In the case of the ridge in the first image, the defining planes are uniform and the ridge-line is clearly segmented. The radial needle-map pattern emerges clearly in the case of the sphere. This radial pattern is also clear for the three “tubular” objects.

The main computational bottleneck in our reported algorithm is the estimation of the local power spectra. The computation of the initial surface orientation angles is fast, since it just involves calculating the eigen-values and eigen-vectors of a 2×2 matrix. The robust smoothing typically takes 10 iterations. Although our research code is not optimized for rapid execution, we provide some timing data for the processing of a 256×256 image on a Pentium III 900MHz processor. Here, the initial surface orientation estimation takes only 0.98 seconds, and the smoothing step takes five seconds. Leaving the power spectrum computation aside, the initial orientation estimation and robust smoothing steps involved in our method is faster than that reported by Krumm and Shafer in [15], [4], [13] since we do not need to perform iterative back-projection.

6.2 Sensitivity Study

In this section, we assess the ability of our shape-from-texture method to recover reliable slant and tilt information under increasing degradation of the texture regularity. Here, we have generated synthetic textured surfaces and have disturbed the texture regularity in a controlled manner.

We have investigated the effect of spatial domain texture irregularity by randomizing the positions of the texture primitives.

We do this by adding a random displacement sampled from a two-dimensional circularly symmetric Gaussian distribution of zero mean and known variance. We have projected these randomized textures onto a plane whose slant and tilt angles are both equal to 45 degrees. In Fig. 3, we visualize the estimated perspective pose of the plane by projecting a pattern of straight lines onto the image plane using the estimated slant and tilt angles. This pattern is composed of two families of parallel straight lines which are




Fig. 3. Estimated perspective pose.




Fig. 4. Circular primitive images: local power spectra.

perpendicular to one-another. One set of straight lines is parallel to the estimated vanishing line of the plane. From these projections, it is clear that the breaking point of the algorithm occurs once the ratio of the standard deviation of the displacement error to the interpoint distance becomes larger than 10 percent. In this example, it is interesting to notice that we still perceive the sensation of perspectivity from the images in the sequence. This may be attributable to the fact that the area gradient is invariant to the positional irregularity of the texture primitives [1], [18]. The random Gaussian noise in the locations of the texture primitives does not mimic perspective distortion. As a result, the density gradient of the primitives is not affected. However, our method is based on measurements that reflect the dominant directions of the texture. These dominant directions are severely affected by the positional noise.

The corresponding power spectra for each texture in the figure are shown in Fig. 4. The main effects to observe are that the individual peaks merge together, creating clouds of peaks and, that, the overall frequency structure of the texture spectra degrades. This collapse of the frequency content is a result of increasing the gap spacing between the texture primitives. The overlap of the texture primitives also contributes to the decrease of the total frequency energy.

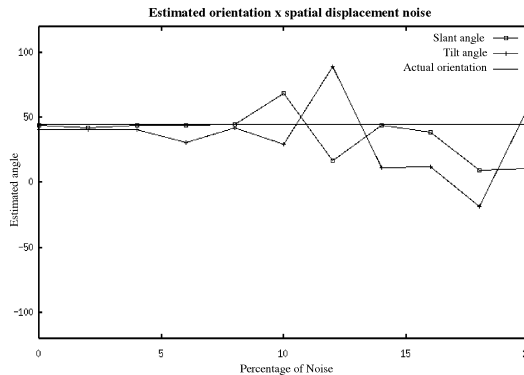


Fig. 5. Effect of texture irregularity on slant and tilt error.

In Fig. 5, we show the effect of the spatial domain irregularity in the positions of the texture primitives on the estimated slant and tilt angles. The plot shows the slant and tilt error as a function of the standard deviation of the random spatial displacement error. Here, the method recovers good estimates of the slant and tilt parameters provided that the ratio of the standard deviation of the displacement error to the interpoint distance does not exceed 10 percent.

The parameter of our spectral distortion method is the distance between the points used to estimate the affine distortion matrix on the image plane. As pointed out earlier, if this distance is too small then the affine distortion becomes undetectable. If, on the other hand, the distance is too large then we sample changes in surface orientation rather than perspective foreshortening. To investigate this effect in more detail, we have generated synthetic curved surfaces with known ground-truth slant and tilt angles. We have computed the regression line between the ground-truth and estimated orientation angles as the interpoint distance is increased.

In Fig. 6 for the smoothed and unsmoothed needle-maps, we plot the linear regression coefficients extracted from the scatter plots as a function of the interpoint distances. If the measurements are unbiased, then the linear regression coefficient should be unity. The main feature to note is that there is a critical value of the distance which results in a maximum value of the regression coefficient. For the smoothed needle-maps, the linear regression coefficient is closest to unity (0.97) when the interpoint distance is $r = 16$ pixels; this represents an improvement over the initial unsmoothed value of $\mu = 0.51$. For the unsmoothed needle-maps, the best regression coefficient (0.84) is obtained when $r = 48$ pixels; here the corresponding smoothed value is $\mu = 0.93$.

7 CONCLUSIONS

We have presented a new method for estimating the local orientation of tangent planes to curved textured surfaces. The




Fig. 6. Plot of the regression line correlation as a function of the neighborhood radius. (a) Smoothed needle map. (b) Unsmoothed needle map.

method commences by finding affine spectral distortions between neighboring points on the image plane. The directions of the eigenvalues of the local distortion matrices can be used to make closed form estimates of the slant and tilt directions. The initial orientation estimates returned by the new method are iteratively refined using a robust smoothing technique to produce a needle map of improved consistency.

The method is demonstrated on both synthetic imagery with known ground-truth and on real-world images of man-made textured surfaces. The method proves useful in the analysis of both planar and curved surfaces. Moreover, the extracted needle maps can be used to make reliable estimates of surface curvature information.

There are a number of ways in which the ideas presented in this paper could be extended and improved. First, our texture measures are relatively crude and could be refined to allow us to analyze textures that are regular but not necessarily periodic. Second, there is scope for improving the quality of the needle map through measuring the back-projection error associated with the smoothed surface normal directions. One possible approach to this problem would be to use a variation of the expectation-maximization algorithm to develop a statistical framework for local orientation estimation and backprojection error assessment. Finally, our current method makes use of a potentially fragile energy ordering heuristic to locate spectral correspondences. A more satisfactory solution to this problem would be to use more than two peaks and to find the set of correspondences which best align the pattern of peaks. We have recently developed a dual-step EM algorithm which can solve this correspondence problem for cluttered point-sets [29]. Our future plans involve using this algorithm for spectral correspondence.

REFERENCES

- [1] J.J. Gibson, *The Perception of the Visual World*. Boston: Houghton Mifflin, 1950.
- [2] D. Marr, *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Freeman, 1982.
- [3] A. Blake and C. Marinou, "Shape from Texture: Estimation, Isotropy and Moments," *Artificial Intelligence*, vol. 45, no. 3, pp. 323-380, 1990.
- [4] J. Krumm and S.A. Shafer, "Shape from Periodic Texture Using Spectrogram," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 284-289, 1992.
- [5] B.J. Super and A.C. Bovik, "Planar Surface Orientation from Texture Spatial Frequencies," *Pattern Recognition*, vol. 28, no. 5, pp. 729-743, 1995.
- [6] J.R. Kender, "Shape from Texture: An Aggregation Transform that Maps a Class of Texture into Surface Orientation," *Proc. Sixth Int'l Joint Conf. Artificial Intelligence*, pp. 475-480, 1979.
- [7] J.S. Kwon, H.K. Hong, and J.S. Choi, "Obtaining a 3-D Orientation of Projective Textures Using a Morphological Method," *Pattern Recognition*, vol. 29, pp. 725-732, 1996.
- [8] J. Garding, "Shape from Texture for Smooth Curved Surfaces," *Proc. European Conf. Computer Vision*, pp. 630-638, 1992.
- [9] J. Garding, "Shape from Texture for Smooth Curved Surfaces in Perspective Projection," *J. Math. Imaging and Vision*, vol. 2, pp. 329-352, 1992.
- [10] J. Malik and R. Rosenholtz, "A Differential Method for Computing Local Shape-from-Texture for Planar and Curved Surfaces," *Proc. IEEE Conf. Vision and Pattern Recognition*, pp. 267-273, 1993.
- [11] J. Malik and R. Rosenholtz, "Recovering Surface Curvature and Orientation from Texture Distortion, a Least Squares Algorithm and Sensitive Analysis," *Proc. European Conf. Computer Vision '94*, vol. 800, pp. 353-364, 1994.
- [12] B.J. Super and A.C. Bovik, "Shape from Texture Using Local Spectral Moments," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 17, no. 4, pp. 333-343, Apr. 1995.
- [13] J. Krumm and S.A. Shafer, "Texture Segmentation and Shape in the Same Image," *Proc. IEEE Int'l Conf. Computer Vision*, pp. 121-127, 1995.
- [14] B.J. Super and A.C. Bovik, "Filters for Directly Detecting Surface Orientation in an Image," *Proc. SPIE Conf. Visual Comm. and Image Processing*, pp. 144-155, 1992.
- [15] J. Krumm and S.A. Shafer, "Local Spatial Frequency Analysis of Image Texture," *Proc. IEEE Int'l Conf. Computer Vision*, pp. 359-358, 1990.
- [16] A. Rosenfeld, "A Note on Automatic Detection of Texture Gradients," *IEEE Trans. Computers*, vol. 24, pp. 988-991, 1975.
- [17] R. Bajcsy and L. Lieberman, "Texture Gradient as a Depth Cue," *Computer Graphics and Image Processing*, vol. 5, pp. 52-67, 1976.
- [18] K.A. Stevens, "The Information Content of Texture Gradients," *Biological Cybernetics*, vol. 42, pp. 95-105, 1981.
- [19] K. A. Stevens, "Slant-Tilt: The Visual Encoding of Surface Orientation," *Biological Cybernetics*, vol. 46, pp. 183-195, 1983.
- [20] A.P. Witkin, "Recovering Surface Shape and Orientation from Texture," *Artificial Intelligence*, vol. 17, pp. 17-45, 1981.
- [21] K. Ikeuchi, "Shape from Regular Patterns," *Artificial Intelligence*, vol. 22, pp. 49-75, 1984.
- [22] J. Aloimonos and M.J. Swain, "Shape from Texture," *Biological Cybernetics*, vol. 58, no. 5, pp. 345-360, 1988.
- [23] K. Ikeuchi and B.K.P. Horn, "Numerical Shape from Shading and Occluding Boundaries," *Artificial Intelligence*, vol. 17, pp. 141-184, 1981.
- [24] K. Kanatani and T. Chou, "Shape from Texture: General Principle," *Artificial Intelligence*, vol. 38, pp. 1-48, 1989.
- [25] L.G. Brown and H. Shvaytser, "Surface Orientation from Projective Foreshortening of Isotropic Texture Autocorrelation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 12, no. 6, pp. 584-588, June 1990.
- [26] K. Sakai and L.H. Finkel, "A Shape-from-Texture Algorithm Based on Human Visual Psychophysics," *Proc. IEEE Conf. Vision and Pattern Recognition*, pp. 527-532, 1994.
- [27] J.Y. Jau and R.T. Chin, "Shape from Texture Using Wigner Distribution," *Computer Vision, Graphics and Image Processing*, vol. 52, pp. 248-263, 1990.
- [28] S.M. Kay, *Modern Spectral Estimation: Theory and Application*. Prentice Hall, 1988.
- [29] A.D.J. Cross and E.R. Hancock, "Graph Matching with a Dual-Step EM Algorithm," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 20, no. 11, pp. 1236-1253, Nov. 1998.
- [30] B.K.P. Horn, *Robot Vision*, Mass.: MIT Press, 1986.
- [31] P.L. Worthington and E.R. Hancock, "New Constraints on Data-Closeness and Needle Map Consistency for Shape-from-Shading," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 21, no. 12, pp. 1250-1267, Dec. 1999.

► For more information on this or any other computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.