

Illuminance flow over anisotropic surfaces

Stefan Karlsson,* Sylvia Pont, and Jan Koenderink

Department Physics of Man, Helmholtz Institute, 5, 3584 CC Utrecht, The Netherlands

**Corresponding author: S.Karlsson@phys.uu.nl*

Received April 27, 2007; revised October 26, 2007; accepted October 29, 2007;
posted November 6, 2007 (Doc. ID 82339); published January 7, 2008

A theory is presented to analyze images of anisotropic fine-scale surfaces. We investigate the estimates of illuminance flow by using structure tensors. For anisotropic surfaces, both the gradient-based and the Hessian-based tensors will yield deviations from the true illumination orientation. Our theory predicts this deviation. To show the use of this theory, an algorithm is derived that uses both tensors simultaneously to compensate for small amounts of anisotropy. Experimental results with rendered surfaces are shown to conform well to our theory. © 2008 Optical Society of America

OCIS codes: 150.2950, 330.7310, 100.2960.

1. INTRODUCTION

The light field [1] in which objects are situated is important in machine vision. Often slight changes in illumination produce large changes in the image of an object, especially with regard to mesoscale surface variations across the object (so-called 3D or surface texture).

A great deal of work has been done regarding methods for handling image variation due to translation, rotation, and scale [2–4]. Only in the past decade or so has the subject of lighting variation been given widespread attention within the machine vision field [5,6]. The focus has been mainly on arriving at a set of lighting-invariant features. The traditional approach has been to vary the light field and gather statistics that are insensitive to light change (often resulting in edges and corners in the albedo map).

In this paper, we investigate illuminated, Lambertian anisotropic surface textures of uniform albedo and try to infer the so-called illuminance flow [7,8], defined as the set of vectors resulting from projecting the light vector into the tangent planes of a surface. The illuminance flow over the surface of an object constitutes a flow field. An example of this flow field (for a sphere) is shown in the left-hand pane of Fig. 1. The surface illuminance flow is a continuous set of 3D vectors (one for each surface patch that is illuminated) and does not depend on the camera position.

The geometry of an object exists on several scales. The global scale constitutes the shape of the object as a whole (e.g., a sphere as in Fig. 1), and the mesoscale makes up the surface texture. At even finer detail are the microscales with variations so fine that they influence the reflectance properties of the surface rather than the observable geometry. Where one scale starts and the other ends is usually not given and can vary locally within the image and depends on the camera position. The global-scale geometry defines a set of tangent planes into which the light vector is projected to form the illuminance flow.

In this paper, we consider the case where the object is a plane on the global scale and where the camera is in normal position (illustrated by the right-hand pane of Fig. 1).

Assuming a collimated light source, the illuminance flow is a unidirectional flow field with direction ϕ (the azimuthal angle of the incident light, as illustrated in Fig. 1). For this case, the problem is reduced to estimating the azimuthal direction ϕ from the image.

Any algorithm that can estimate ϕ (locally) can be used to estimate the illuminance flow in normal viewing. This puts our work close to surface texture classification and photometric stereo, where illuminant azimuth (or tilt) estimation plays an important part [9]. In this context, this paper presents a novel theory that generalizes the Kube and Pentland imaging model [10] by use of structure tensor statistics to account for surface anisotropy. After presenting our theory and a proposed estimator, we will relate it to a few algorithms of illuminant estimation from texture analysis (Section 4) and show that our theory is relevant to them all. With respect to illuminance flow, this paper is in line with the study by Koenderink and Pont [7], where image structure tensors were used to estimate the light direction. The current paper is concerned primarily with the extension of the theory to the case of anisotropic surface textures.

When the light field [1] has a dominant directional component (a collimated beam being the extreme case) anisotropy will be imposed in the image of the surface. In some cases this can be used to estimate the illuminance flow. It has been shown [7] that using the structure tensor of either the Hessian or the gradient yields good estimates. However, formally, this applies only under a set of assumptions, namely,

- Isotropic surface
- Lambertian reflectance
- Uniform albedo
- Low relief
- Normal view
- Light field [1] of dominant directional component (e.g., a collimated beam)

Experimental results have been presented for real-world surfaces [7]; the structure tensors yielded promis-

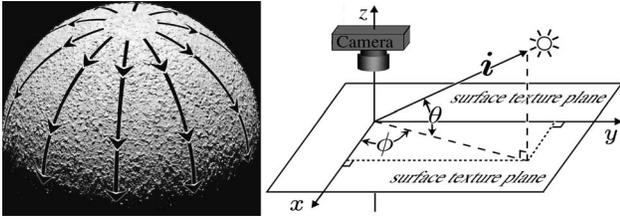


Fig. 1. Left, illustration of illuminance flow on a sphere. Two scales of interest are the mesoscale, surface variation due to texture, and object scale, defining the geometry of the object (here a sphere). Right, imaging geometry. We assume an object that is flat on the global scale for this paper. θ and ϕ define the direction of illumination.

ing results, with orientation estimates within a few degrees of the vertical orientation. The illumination elevation was not estimated due to the bas-relief ambiguity [11].

All assumptions are important in their own right. Some work has been done on extending the theory to account for nonnormal views [8] and nonuniform albedo [12]. In this paper, we focus on the isotropy assumption, which, as far as we know, has not been given any attention so far. For anisotropic surfaces, both the gradient- and the Hessian-based tensors will yield deviations from the true orientation of the illuminance flow. The theory presented here predicts these deviations, which are shown to conform well to experiments on rendered surfaces.

To show the use of this framework, an algorithm is proposed that uses both tensors simultaneously. Extensive experimental results on rendered surfaces are presented, as well as a few estimations on real-world textures.

2. THEORETICAL BACKGROUND

Derivatives will be denoted

$$h_{pq} = \frac{\partial^{p+q} h(x,y)}{\partial x^p \partial y^q},$$

the gradient $\mathbf{g}_h = \{h_{10}, h_{01}\}^T$, and the Hessian

$$H_h = \begin{pmatrix} h_{20} & h_{11} \\ h_{11} & h_{02} \end{pmatrix}.$$

$\hat{h}(\mathbf{k})$ is the Fourier transform of $h(\mathbf{r})$ (spatial coordinates $\mathbf{r} = \{x, y\}^T$ and frequency coordinates $\mathbf{k} = \{u, v\}^T$), and its square magnitude is referred to as the power spectrum, denoted $\hat{\rho}_h(\mathbf{k}) = |\hat{h}(\mathbf{k})|^2$.

$\langle h \rangle$ will denote averaging over the function $h(\mathbf{r})$, and m_h^{pq} denotes moments of the power spectrum defined as

$$m_h^{pq} = \langle u^p v^q \hat{\rho}_h \rangle. \quad (1)$$

For correlations between derivatives we have (because $\langle h_{p_1 q_1} \rangle = 0$)

$$\langle h_{p_1 q_1} h_{p_2 q_2} \rangle = m_h^{(p_1+p_2, q_1+q_2)} (-1)^{(p_1-p_2+q_1-q_2)/2}. \quad (2)$$

Equation (2) is derived easily by use of the following properties of the Fourier transform:

$$h_{pq} \Leftrightarrow (-ju)^p (-jv)^q \hat{h}(\mathbf{k}),$$

$$\langle f(\mathbf{r})g(\mathbf{r}) \rangle = \langle \hat{f}(\mathbf{k})\hat{g}^*(\mathbf{k}) \rangle,$$

where $*$ indicates complex conjugate. Because $\hat{\rho}_h$ is symmetric [due to Hermitian symmetry $\hat{h}(\mathbf{k}) = \hat{h}(-\mathbf{k})^*$] all moments for which $p+q$ is an odd integer will equal zero.

Gaussian Surfaces and the Structure Tensors. Henceforth $h(\mathbf{r})$ will denote a Gaussian surface, which is a 2D Gaussian (stochastic) process with several convenient properties [13,14] making it a good candidate for a surface model. It is (strongly) stationary with a constant autocorrelation function $\rho_h(\mathbf{r})$ within its ensemble. The covariance function $\text{cov}_h(\mathbf{r})$ is related to $\rho_h(\mathbf{r})$ by $\text{cov}_h(\mathbf{r}) = \sigma^2 \rho_h(\mathbf{r})$, where σ is the standard deviation of surface height. The spread of $\rho_h(\mathbf{r})$ is denoted s , i.e., its width as measured by second-order moments. s is a measure of scale, and the quantity σ/s equals the standard deviation of the derivative processes $h_{10}(\mathbf{r})$ and $h_{01}(\mathbf{r})$.

The power spectrum $\hat{\rho}_h(\mathbf{k})$ is related to $\rho_h(\mathbf{r})$ through a Fourier transform, and thus $\hat{\rho}_h(\mathbf{k})$ is also constant within the ensemble, as are all the moments m_h^{pq} .

We will assume a Gaussian-shaped $\rho_h(\mathbf{r}) = \exp(-\mathbf{r}^T G_h \mathbf{r})$, yielding an equivalent Gaussian for the power spectrum: $\hat{\rho}_h(\mathbf{k}) = C \exp(-\mathbf{k}^T G_h^{-1} \mathbf{k})$, where C is a normalizing constant and G_h is the gradient-based structure tensor, a symmetric, positive semidefinite matrix consisting of the second-order moments:

$$G_h = \begin{pmatrix} m_h^{20} & m_h^{11} \\ m_h^{11} & m_h^{02} \end{pmatrix} = \langle \mathbf{g}_h \mathbf{g}_h^T \rangle. \quad (3)$$

The assumption that the autocorrelation function is Gaussian shaped is of no critical importance; it is only for convenience. All power spectra of interest have well-defined second-order moments, and thus also a structure tensor, though not necessarily explicit in closed form as with the Gaussian.

For an isotropic surface, $\rho_h(\mathbf{r})$ is isotropic and G_h is a scaled identity matrix (i.e., $m_h^{20} = m_h^{02}$ and $m_h^{11} = 0$). For an anisotropic surface the eigenvalues and eigenvectors of G_h yield the amount and orientation of anisotropy. We denote the maximum and minimum eigenvalues and their corresponding eigenvectors as $\lambda_{G_h}^{max}$, $\lambda_{G_h}^{min}$, $\mathbf{v}_{G_h}^{max}$, and $\mathbf{v}_{G_h}^{min}$. The gradient confidence is a measure of anisotropy defined as

$$\xi\{G_h\} = \frac{\lambda_{G_h}^{max} - \lambda_{G_h}^{min}}{\lambda_{G_h}^{max} + \lambda_{G_h}^{min}} \in [0, 1].$$

$\xi=0$ occurs when $h(\mathbf{r})$ is a completely isotropic surface, whereas $\xi=1$ indicates degeneracy (the power spectrum is reduced to a line and the surface contains only isolines). The two cases $\xi=0$ and $\xi=1$ will be referred to as the extrema, or extremal cases.

We are interested in the orientation of anisotropy. Here, we choose to make the distinction between orientation and direction, the latter being free to assume any angle, i.e., any real number modulo 360° . Orientation, on the other hand, is a quantity that treats directions that differ by 180° equally, defined by taking any directional angle modulo 180 {a two-to-one mapping $[0, 360) \rightarrow [0, 180)$ }.

The structure tensor was originally introduced to image analysis as a method for orientation estimation [15]. The orientation angle of anisotropy is given by the highest eigenvalue eigenvector and is denoted

$$\mu\{G_h\} = \angle \mathbf{v}_{G_h}^{max} \in [0^\circ, 180^\circ).$$

Closely related to G_h is the Hessian-based structure tensor defined as

$$T_h = \begin{pmatrix} m_h^{40} + m_h^{22} & m_h^{31} + m_h^{13} \\ m_h^{31} + m_h^{13} & m_h^{04} + m_h^{22} \end{pmatrix} = \langle H_h H_h^T \rangle. \quad (4)$$

A measure of anisotropy (similar to $\xi\{G_h\}$) is the Hessian confidence: $\xi\{T_h\}$. The orientation of anisotropy is given by $\mu\{T_h\}$.

Analogous to the $\xi\{G_h\}$ measure, $\xi\{T_h\}=0$ corresponds to the isotropic case, and $\xi\{T_h\}=1$ indicates degeneracy. However, for the nonextremal cases, the two measures differ in general (e.g., $\xi\{G_h\}=\frac{1}{2}$ and $\xi\{T_h\}=\frac{1}{2}$ have different interpretations). Also, important for our purposes, $\mu\{G_h\}$ and $\mu\{T_h\}$ will differ in general for the nonextremal cases. We will make frequent use of $\xi\{G_h\}$, and for the sake of convenience will write simply ξ , where it should be clear that $\xi=\xi\{G_h\}$, the confidence with respect to the gradient tensor of the surface $h(\mathbf{r})$, a well-defined and natural measure of the anisotropy of a surface.

3. ILLUMINATED GAUSSIAN SURFACES

In this section we will derive relationships between the hidden tensors of the surface and the observable ones of the image. The key property here is that surface anisotropy is encoded in the $\xi\{G_h\}$ statistic (equivalently, we could have used $\xi\{T_h\}$, yielding slightly more cumbersome expressions). Thus, the goal of this section is to express $\mu\{G_I\}$ and $\mu\{T_I\}$ in terms of $\xi\{G_h\}$ and ultimately to predict the deviations of the $\mu\{\cdot\}$ observables from the true orientation of illumination ϕ . The orientation of surface anisotropy, encoded by the hidden $\mu\{G_h\}$, will be addressed at the end of the section.

The unit normal of a surface can be written as

$$\mathbf{n} = \frac{\{-h_{10}, -h_{01}, 1\}^T}{\sqrt{h_{10}^2 + h_{01}^2 + 1}}.$$

The irradiance by a collimated beam is described by a single vector:

$$\mathbf{i} = c\{\cos \theta \cos \phi, \cos \theta \sin \phi, \sin \theta\},$$

where θ and ϕ are the elevation and azimuthal angles of the light source, respectively, and c is the intensity of the light (normal irradiance caused by the beam).

The assumption of a collimated beam is convenient for the derivation, but not critical since extended sources can be exactly accounted for by equivalent collimated sources if vignetting is not present. The following theory holds also for light fields [1] where a dominant directional component is present; the \mathbf{i} vector can be replaced by this dominant direction (for Lambertian surfaces this is equivalent, when self-shadowing is not present).

We will consider the image $I(\mathbf{r})$ that is formed by viewing the surface $h(\mathbf{r})$ normally and will assume that its in-

tensity is proportional to the irradiance of the surface. The image of a Lambertian low-relief surface will then be given by

$$I(h_{01}, h_{10}, \theta, \phi) = \mathbf{i} \cdot \mathbf{n} = c \frac{\sin \theta - \cos \theta (h_{10} \cos \phi + h_{01} \sin \phi)}{\sqrt{h_{10}^2 + h_{01}^2 + 1}}. \quad (5)$$

We assume a low-relief surface (which also entails no shadows or interreflections) with low values of h_{10} and h_{01} on average, making it reasonable to approximate Eq. (5) around $(h_{10}, h_{01})=(0, 0)$. A Taylor series approximation including the first-order terms of Eq. (5) yields

$$I(h_{01}, h_{10}, \theta, \phi) = c \sin \theta - c \cos \theta (h_{10} \cos \phi + h_{01} \sin \phi) + O(\mathbf{g}^T \mathbf{g}), \quad (6)$$

which is a weighted directional derivative of the surface plus a constant, $c \sin(\theta)$, of little interest. Equation (6) is a model of image formation first suggested by Kube and Pentland [10] and has been used in studies of texture segmentation [16,17] and light estimation [7,17].

The approximation is valid only under oblique lighting, which is evident from the fact that the approximation yields a constant image for $\theta=90^\circ$, and for lower values $\theta \rightarrow 0$ we enter the shadow region.

We consider the Taylor series expansion of the surface around any point of interest (expressed in a local frame where the origin has been aligned with the point of interest):

$$h(\mathbf{r}) = xh_{10} + yh_{01} + x^2h_{11}/2 + \dots \quad (7)$$

Now, we truncate Eq. (7) up to third-order terms and insert it into Eq. (6), yielding a closed-form approximation of the imaging process valid under the assumptions of low relief and oblique lighting of Gaussian surfaces:

$$I(\mathbf{r}, \theta, \phi) = c \left\{ \sin(\theta) - \frac{1}{2} \cos(\theta) [\sin(\phi)(h_{21}x^2 + 2h_{11}x + 2yh_{12}x + 2h_{01} + 2yh_{02} + y^2h_{03}) + \cos(\phi)(h_{30}x^2 + 2h_{20}x + 2yh_{21}x + 2h_{10} + 2yh_{11} + y^2h_{12})] \right\}. \quad (8)$$

The approximation in Eq. (8) can be used to derive expressions of the structure tensors of the image, which will equal

$$G_I(1, 1) = \frac{c^2}{2} \cos^2 \theta [m_h^{22} + m_h^{40} + (m_h^{40} - m_h^{22}) \cos 2\phi + 2m_h^{31} \sin 2\phi], \quad (9)$$

$$G_I(1, 2) = \frac{c^2}{2} \cos^2 \theta [m_h^{31} + m_h^{13} + (m_h^{31} - m_h^{13}) \cos 2\phi + 2m_h^{22} \sin 2\phi], \quad (10)$$

$$G_I(2,2) = \frac{c^2}{2} \cos^2 \theta [m_h^{22} + m_h^{04} + (m_h^{22} - m_h^{04}) \cos 2\phi + 2m_h^{13} \sin 2\phi], \quad (11)$$

$$T_I(1,1) = \frac{c^2}{2} \cos^2 \theta [m_h^{60} + 2m_h^{42} + m_h^{24} + (m_h^{60} - m_h^{24}) \cos 2\phi + 2(m_h^{51} + m_h^{33}) \sin 2\phi], \quad (12)$$

$$T_I(1,2) = \frac{c^2}{2} \cos^2 \theta [m_h^{51} + 2m_h^{33} + m_h^{15} + (m_h^{51} - m_h^{15}) \cos 2\phi + 2(m_h^{42} + m_h^{24}) \sin 2\phi], \quad (13)$$

$$T_I(2,2) = \frac{c^2}{2} \cos^2 \theta [m_h^{42} + 2m_h^{24} + m_h^{06} + (m_h^{42} - m_h^{06}) \cos 2\phi + 2(m_h^{15} + m_h^{33}) \sin 2\phi]. \quad (14)$$

The left-hand sides of Eqs. (9)–(14) are the elements of the structure tensors [Eqs. (3) and (4)] with respect to the image $I(\mathbf{r})$. The right-hand side contains power spectrum moments with respect to the surface $h(\mathbf{r})$.

Now, the moments of the surface will be affected by the anisotropy. We assume a Gaussian-shaped surface power spectrum $\hat{\rho}_h(\mathbf{k}) = C \exp(-\mathbf{k}^T G_h^{-1} \mathbf{k})$. Anisotropy is introduced as a nonzero $\xi\{G_h\}$. Assuming (for now) that the orientation of anisotropy is in the orientation of the x axis (0° or equivalently 180°), we can write the power spectrum as

$$\hat{\rho}_h(u,v) = g(u, s^{-1}) g\left(v, \frac{1 + \xi\{G_h\}}{1 - \xi\{G_h\}} s^{-1}\right),$$

where $g(x, a)$ is a 1D normalized Gaussian with standard deviation a . The parameter s remains a general measure of the scale of the surface.

Applying Eq. (1) will yield (for $\xi = \xi\{G_h\}$)

$$m_h^{40} = \frac{12\pi^2(\xi + 1)^5}{s^4(\xi - 1)^4},$$

$$m_h^{22} = \frac{4\pi^2(\xi + 1)^3}{s^4(\xi - 1)^2},$$

$$m_h^{04} = \frac{12\pi^2(\xi + 1)^5}{s^4(\xi - 1)^4},$$

$$m_h^{60} = \frac{60\pi^2(\xi + 1)}{s^6},$$

$$m_h^{42} = \frac{12\pi^2(\xi + 1)^3}{s^6(\xi - 1)^2},$$

$$m_h^{24} = \frac{12\pi^2(\xi + 1)^5}{s^6(\xi - 1)^4},$$

$$m_h^{06} = \frac{60\pi^2(\xi + 1)^7}{s^6(\xi - 1)^6},$$

The remaining moments ($m_h^{31}, m_h^{13}, m_h^{51}, m_h^{33}$, and m_h^{15}) all equal zero. Insertion into Eqs. (9)–(14) yields

$$G_I(1,1) = C_1 \frac{2(\xi^3 + 1) + (\xi + 1)((\xi - 4)\xi + 1) \cos 2\phi}{(\xi - 1)^2}, \quad (15)$$

$$G_I(1,2) = C_1 \frac{(\xi + 1)^3 \sin 2\phi}{(\xi - 1)^2}, \quad (16)$$

$$G_I(2,2) = C_1 \left[\frac{(\xi + 1)^3 2(\xi^2 + \xi + 1)}{(\xi - 1)^4} - \frac{(\xi(\xi + 4) + 1) \cos 2\phi}{(\xi - 1)^4} \right], \quad (17)$$

$$T_I(1,1) = C_2 \left[\frac{2\xi^5 - 2\xi^4 + 4\xi^3 + 4\xi^2 - 2\xi + 2}{(\xi - 1)^4} + \frac{(\xi^5 - 5\xi^4 - 5\xi + 1) \cos 2\phi}{(\xi - 1)^4} \right], \quad (18)$$

$$T_I(1,2) = C_2 \frac{(\xi + 1)^3 (\xi^2 + 1) \sin 2\phi}{(\xi - 1)^4}, \quad (19)$$

$$T_I(2,2) = C_2 \frac{a + 2 - (b + 1) \cos 2\phi}{(\xi - 1)^6},$$

$$a = 2(\xi^7 + 5\xi^6 + 13\xi^5 + 21\xi^4 + 21\xi^3 + 13\xi^2 + 5\xi),$$

$$b = (\xi^7 + 9\xi^6 + 27\xi^5 + 43\xi^4 + 43\xi^3 + 27\xi^2 + 9\xi), \quad (20)$$

for $C_1 = [4c^2\pi^2 \cos^2(\theta)]/s^4$ and $C_2 = [24c^2\pi^2 \cos^2(\theta)]/s^6$. We are interested in the eigenanalysis of the tensors. Multiplication by arbitrary positive constants does not affect the eigenvalues or vectors, so we can safely choose to ignore C_1 and C_2 —they become relevant only when we have $s = \infty$, $\theta = 90^\circ$, or $c = 0$, for which both C_1 and C_2 are zero. For $s = \infty$ the surface has degenerated to a plane (for which no variation is present in the image), for $\theta = 90^\circ$ the surface is illuminated strictly from above (no anisotropy is imposed by the illumination), and $c = 0$ indicates complete absence of light.

We now investigate the statistics that will be given by the eigenanalysis of the tensors. For the $\mu\{\cdot\}$ measures we get

$$\mu\{G_I\} = -\tan^{-1}\left(\frac{a}{b + \sqrt{c}}\right),$$

$$a = 2(\xi^2 - 1) \sin(2\phi),$$

$$b = (4\xi^2 + 2) \cos(2\phi) - 6\xi,$$

$$c = 10(\xi^2 + 4)\xi^2 + 4 + 6(\xi^2 + 2)\xi^2 \cos(4\phi) - 24(2\xi^3 + \xi)\cos(2\phi), \quad (21)$$

$$\mu\{T_I\} = -\tan^{-1}\left(\frac{a}{b + \sqrt{c}}\right),$$

$$a = 2(\xi^2 - 1)\sin(2\phi),$$

$$b = (8\xi^2 + 2)\cos(2\phi) - 2\xi^3 - 8\xi,$$

$$c = 4\xi^6 + 66\xi^4 + 76\xi^2 + 4 + 10(3\xi^2 + 2)\xi^2 \cos(4\phi) - 8(\xi^2 + 4)(4\xi^2 + 1)\xi \cos(2\phi). \quad (22)$$

The $\xi\{\cdot\}$ measures will yield

$$\xi\{G_I\} = \frac{\sqrt{a-b}}{c},$$

$$a = 5(\xi^2 + 4)\xi^2 + 2 + 3(\xi^2 + 2)\xi^2 \cos(4\phi),$$

$$b = 12(2\xi^3 + \xi)\cos(2\phi),$$

$$c = \sqrt{2}(\xi^2 - 3\xi \cos(2\phi) + 2),$$

$$\xi\{T_I\} = \frac{\sqrt{a-b}}{c},$$

$$a = 2\xi^6 + 33\xi^4 + 38\xi^2 + 2 + 5(3\xi^2 + 2)\xi^2 \cos(4\phi),$$

$$b = 4\xi(\xi^2 + 4)(4\xi^2 + 1)\cos(2\phi),$$

$$c = \sqrt{2}[3\xi^2 + 2 - (\xi^2 + 4)\xi \cos(2\phi)].$$

If the surface is isotropic (i.e., $\xi=0$), then $\mu\{T_I\} = \mu\{G_I\} = \phi \bmod 180^\circ$, the correct estimate of the light orientation. For the anisotropic case, a nonzero deviation will occur:

$$\text{dev}_G(\phi) = (\mu\{G_I\} - \phi) \bmod 180^\circ$$

with an equivalent definition for the Hessian-based estimate of the image $\text{dev}_H(\phi)$.

We have so far assumed that $\mu\{G_h\}=0$ (i.e., that the orientation of surface anisotropy is along the x axis). Generalizing to anisotropy of different orientations is easily achieved by shifting the deviations, i.e., $\text{dev}_G(\phi - \mu\{G_h\})$. We will continue to assume $\mu\{G_h\}=0$, knowing that the results are valid for all possible $\mu\{G_h\}$.

The deviations dev_G and dev_H are plotted in Fig. 2 for $\xi=0.1$, a low amount of anisotropy, and yet deviations are quite noticeable. This is indeed expected and shows the impact of anisotropy. The expressions for the deviations are very unwieldy, which is why we choose to approximate them with a truncated Taylor series around the point $\xi=0$:

$$\text{dev}_G(\phi) = 3\xi \sin 2\phi + O(\phi^2), \quad (23)$$

$$\text{dev}_H(\phi) = 4\xi \sin 2\phi + O(\phi^2). \quad (24)$$

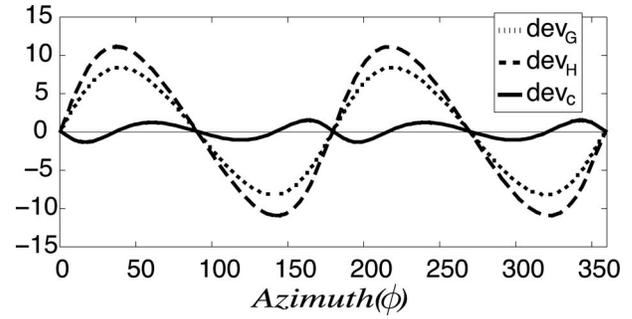


Fig. 2. Theoretical predictions for deviations of the three estimates, dev_H (dashed), dev_G (dotted), and dev_C (solid), as a function of true azimuthal direction of illumination (ϕ), all three for $\xi\{G_h\}=0.1$.

The aim is to have an estimate with no deviation, so Eqs. (23) and (24) are good arguments for a measure that combines both the Hessian-based and the gradient-based tensors of the image. A very simple example of such combination would be

$$\mu_C = 4\mu\{G_I\} - 3\mu\{H_I\}. \quad (25)$$

Expressed in terms of its lower-order Taylor approximation [Eqs. (23) and (24)] μ_C will have zero deviation. When the expressions based on Eqs. (21) and (22) are used, the resulting function $\text{dev}_C = (\mu_C - \phi) \bmod 180^\circ$ is plotted in Fig. 2

While simple and straightforward to derive (given the theory we have presented), the μ_C measure suffers from the drawback of introducing another nonlinearity. The multiplicative factors of -3 and 4 can be expected to amplify any errors that are already present in $\mu\{G_I\}$ and $\mu\{H_I\}$. We stress that this paper focuses on the theory we have presented, i.e., the physics based modeling of the imaging process of anisotropic surface textures. With the combined tensor estimate of Eq. (25) we wish to exemplify the use of this theory.

4. CONNECTION TO OTHER ILLUMINANT ESTIMATORS

To our knowledge, no other local and unsupervised illuminant estimator has been suggested that accounts for surface anisotropy. We will mention a few estimators that are based on the Kube and Pentland modeling, namely those proposed in [18–22]. The premise is a surface texture normally viewed, where a local estimate of ϕ , the illuminant azimuthal angle (tilt) is to be estimated. Iterative nonlocal algorithms, such as [23], that simultaneously estimate the surface texture height profile are not considered; neither are supervised algorithms that deal with a finite set of textures, such as [17].

The early approaches of [20,18] are based on means of directional derivatives, and not variances. For an object that is globally a plane, this means that they will not work (both Knill [19] and Chantler *et al.* [21] discusses this). Many modern approaches still explore the same premise of a convex or spherical surface (for example, [22]).

Two algorithms that deal explicitly with the problem of unsupervised estimation of illumination direction from a

textured plane are Knill's [19] and Chantler's [21]. Both of these schemes are equivalent to the gradient-based structure tensors, which emphasize the general nature of the theory we have presented.

Chantler's method [21] consists of a circular Fourier series expansion of the power spectrum. This requires the power spectrum to be collapsed (averaged along the radial frequency direction) into a 1D function in the angular frequency variable (a polar frequency plot). Depending on how this collapsing is done, Chantler's method can be made equivalent to either the gradient or the Hessian tensors (given that equivalent selections of derivative filters and scales are selected for the tensor calculations). In Chantler's method, weighting higher frequencies more gives the Hessian tensor estimate, and less the gradient. From Chantler's perspective, the combined measure that we suggest could be implemented by simultaneously making two polar plots of the same power spectrum, one weighting the higher frequencies more, and then combining the two estimates.

The relationship with Chantler's method is most easily seen through the use of so-called complex moments of the local image power spectrum, especially the properties presented in [24]. Structure tensors can be viewed as a circular Fourier series expansion of the image power spectrum (just as in Chantler's method [21]). From the complex moments it is also easy to show that the eigenvector angle $\mu\{G_I\}$ is equivalently calculated as Knill's estimator [19]. Knill, in fact, arrives at the complex moment description explicitly, which was proposed years earlier by Bigün and Granlund [15], thus connecting it for the first time to illumination direction.

5. EXPERIMENTS

An experiment on rendered surfaces was performed to test the theoretical predictions and to analyze the behavior of the theory for violations of the assumptions in terms of higher-relief, nonoblique lighting and larger amounts of anisotropy.

A. Rendering

The rendering was implemented without any 3D graphics libraries, owing to the simple nature of the scene and because we are not aware of any libraries that render surfaces defined with torus topology (described in Subsection 5.B). Images (400×400 pixels) were generated as the orthographic projection of the radiance map of a Gaussian surface illuminated by a collimated beam of light. This was rendered as the product of two entities:

$$I_{rend}(\mathbf{r}) = I(h_{01}(\mathbf{r}), h_{10}(\mathbf{r}), \theta, \phi)S(\mathbf{r}),$$

where $I(h_{01}(\mathbf{r}), h_{10}(\mathbf{r}), \theta, \phi)$ is the formulation of the surface radiance for the low-relief case [Eq. (5)] and $S(\mathbf{r}) \in \{0, 1\}$ is the binary shadow function that was calculated in full (0 indicates self-shadowing for both attached and cast shadows).

Interreflections were not rendered (discussed in Section 8).

B. Surface Generation

The surfaces were generated randomly in the Fourier domain by using the specified power spectrum and pairing it with a randomly generated phase spectrum of independently generated values from a uniform distribution $\in [0, 2\pi)$. This way of generating the surfaces ensures constancy of the estimated autocorrelation function $\rho_h(\mathbf{r})$ and of the height probability distribution function.

As is the case for all signals generated smoothly in the discrete Fourier domain, the surface will be defined on a topological torus, i.e., it can be viewed as having infinite area by repeating itself, without losing its stationarity or continuity at any point. This topology is desirable for our purposes, as there will be no concern about the surface boundaries (there are none!).

Every aspect of the rendering, surface generation, and estimation of statistics was done without violating the torus topology, i.e., $(x, y) \rightarrow \{x \bmod width, y \bmod height\}$ when the surface is represented as a finite 2D array.

C. Setup

Parameters in the experiments were

- σ , surface height: scales the height of the surface linearly.
- θ , elevation angle of the light source
- ξ , the anisotropy, i.e., the value assumed for $\xi\{G_h\}$

For every configuration of these variables a sequence of 360 images is created by varying $\phi \in [0^\circ, 360^\circ)$ in integer steps (two such images of the same surface are shown in Fig. 3, for $\phi=90^\circ$ and $\phi=0^\circ$ (left and center of figure, respectively)). The Hessian-based and the gradient-based structure tensors are applied to each image, and the two different estimates of Eq. (25), $\mu\{G_I\}$ and $\mu\{H_I\}$, are collected from the eigenanalysis outlined in Section 3. This allows relatively smooth curves of the deviations from the true estimates to be compared with the theoretical prediction of Fig. 2.

6. RESULTS

Typical results are illustrated in Fig. 4 (each point of the graphs was calculated on images like those in Fig. 6). The output is bounded in $[-90^\circ, 90^\circ)$, as it is an orientation estimation and not a direction estimation. Comparing the top panel of Fig. 4 with Fig. 2 yields, as expected, a strong resemblance. In this range of the parameters the variation in curves, from one surface to the other, is very low. For the case where the surface is rough, as in the second

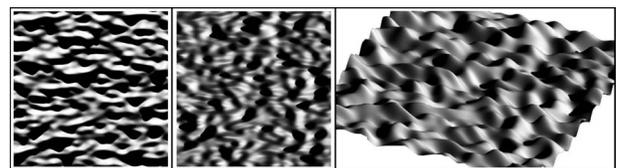


Fig. 3. Rendered Gaussian surface with $\xi=0.6$, $\sigma=100$, and $\theta=30$. Left, $\phi=90^\circ$. Middle, $\phi=0^\circ$. Right, $\phi=0^\circ$ and viewed from another direction.

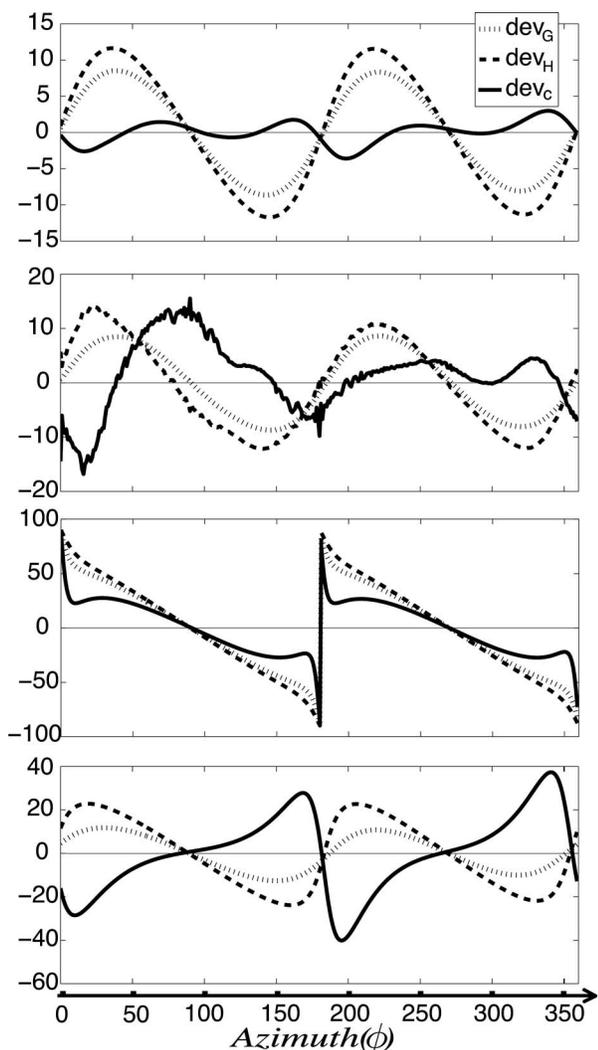


Fig. 4. Typical outputs from the renderings for dev_H (dashed), dev_G (dotted), and dev_C (solid). X and Y axes, same as Fig. 2 (note that the ranges of the Y axes differ in the figure). First graph (Top), close to the assumptions of the theory ($\sigma=30, \theta=30^\circ, \xi=0.1$). Second, higher relief ($\sigma=60, \theta=30^\circ, \xi=0.1$). Third, strongly anisotropic surface ($\sigma=30, \theta=30^\circ, \xi=0.6$). Fourth (bottom), high light elevation (nonoblique lighting, $\sigma=30, \theta=70^\circ, \xi=0.1$).

panel of Fig. 4 and the lighting angle is still relatively low, shadows will make a big impact, as is evident from the noisy curves.

In the experiment of the third panel of Fig. 4 a large amount of anisotropy is present. All three estimates deviate toward the axis of anisotropy, as expected. Consider, as an extreme case, a completely anisotropic surface ($\xi=1$), say, a sinusoidal wave surface. It is impossible to determine the light orientation from the information given by the irradiance of the wave surface by analogy with the aperture problem in optical flow. At some point, while anisotropy increases, any algorithm must falter. It is worth noting, however, that the combined estimate still performs noticeably better than separate estimates in the high-anisotropy case. Finally, the fourth (bottom) panel of Fig. 4 shows what happens when the light elevation is high. The deviation of the combined estimate worsens, but the overall shape of the three curves remains quite consistent with the theoretically predicted output.

We generated a large amount of data by varying $\xi \in \{0, 0.1, \dots, 0.8\}$, $\theta \in \{30^\circ, 35^\circ, \dots, 80^\circ\}$, and $\sigma \in \{30, 35, \dots, 100\}$. Furthermore, we averaged the results over five independent surfaces. The absolute value of each curve of the type in Fig. 4 was averaged, yielding one scalar error measure (average error) for every configuration of θ, ξ , and σ . The results are 3D, and thus cannot be plotted in paper form. Figure 5 shows contour plots over pairs of dimensions, while the third dimension is kept fixed, close to the assumptions of the theory.

The average error should be bounded $\in [0^\circ, 45^\circ]$; if it is any higher, the output is worse than random guesses (e.g., an average error of 90 indicates an output that is always perpendicular to the true orientation).

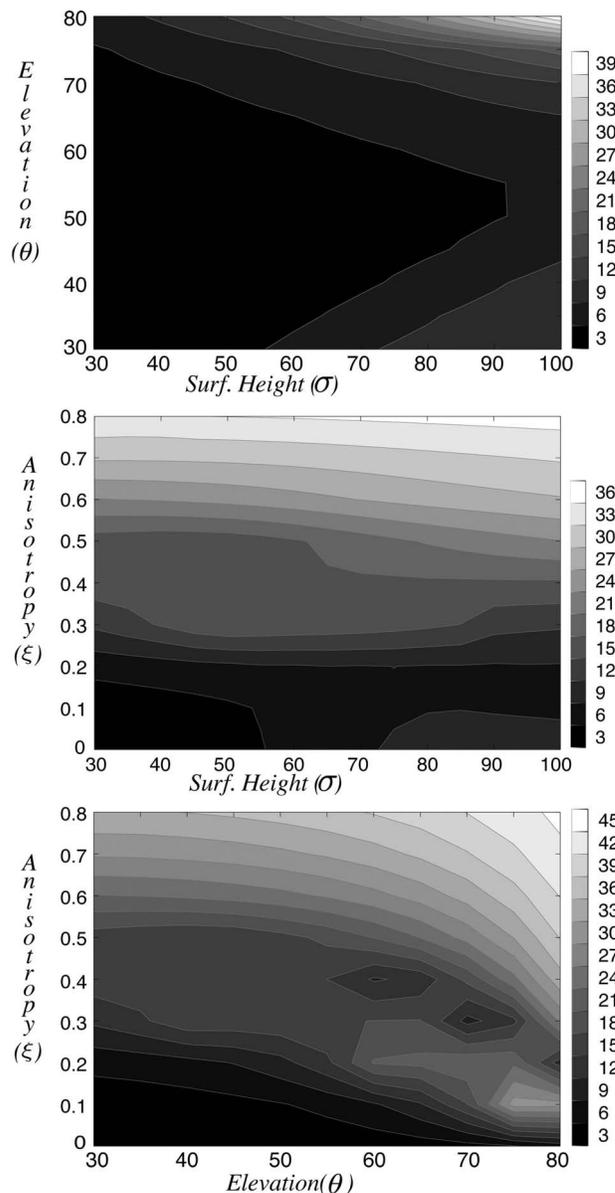


Fig. 5. Contour plots of the average magnitude deviations. The absolute value of each curve of the type in Fig. 4 is averaged (one nonnegative scalar for each curve). The gray value of each region indicates an upper bound on the error, specified by the bars on the right of each figure. Top, θ versus σ with $\xi=0.1$. Middle, ξ versus σ with $\theta=30^\circ$. Bottom, ξ versus θ with $\sigma=30$.

The main observation to make from Fig. 5 is low error and stable monotonic behavior around the point corresponding to the assumptions ($\xi=0$, $\sigma=0$, and $\theta\approx 45^\circ$). As we go further away from the assumptions in parameter space, the output gradually becomes worse. We also note that ξ completely dominates as source of error in the plots, with the exception of the elevation angle θ when it comes close to 90° .

For an accepted error of say 6° , one should stay within the bounds of $\xi < 0.2$, $\sigma < 65$ and expect even better behavior for oblique lighting of around $\theta=45^\circ$.

As expected, for large amounts of anisotropy, the measure becomes inaccurate, especially when the illumination orientation is in the direction of the anisotropy axis of the surface (at $\phi=0^\circ$ and 180° in the graphs of Fig. 4).

Also expected was the increase in deviation due to higher relief. When the surface is rough (σ is high), and as θ decreases, the surface tends to cast shadows. Shadows affect the outcome of the tensor estimates, because they give rise to new edges in the image. Shadow edges are nonlocal and more noisy with respect to ϕ changes than are regular shading patterns. On the other hand, as θ increases it comes closer to 90° where the problem is undefined.

For high values of θ , the modeling of the imaging process as a directional derivative is less accurate; other structure becomes evident as is seen in Fig. 6. The images on the left of Fig. 6 are images of the same surface, but only the top one, with low θ , resembles a directional derivative. The structure of the top image resembles more the square magnitude of the gradient and is often referred to as second-order shading. Second-order shading is independent of the orientation of illumination and will inherit anisotropy from the surface (if the surface contains anisotropy, then so will the square magnitude of its gradient).

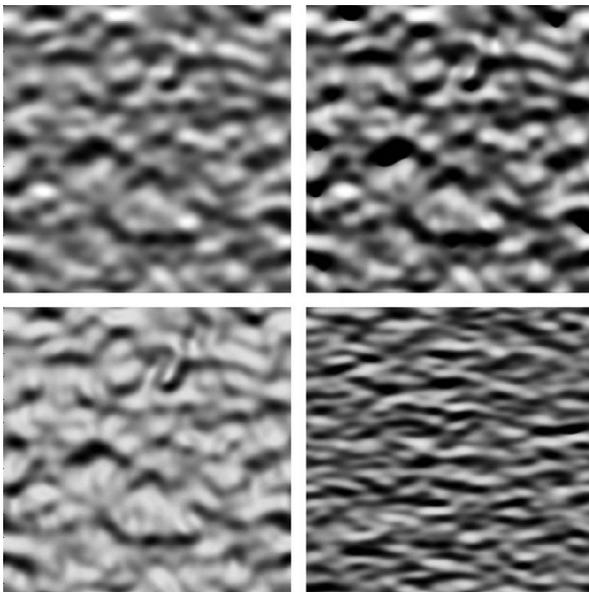


Fig. 6. Example images corresponding to $\phi=90^\circ$ (top illumination), each one corresponding to one out of 360 images used to draw the graphs in Fig. 4. The images have been histogram normalized for visualizing details. Top left, $\sigma=30$, $\theta=30^\circ$, $\xi=0.1$. Top right, $\sigma=60$, $\theta=30^\circ$, $\xi=0.1$. Bottom left, $\sigma=30$, $\theta=70^\circ$, $\xi=0.1$. Bottom right, $\sigma=30$, $\theta=30^\circ$, $\xi=0.6$.

For anisotropic surfaces with a higher relief, we expect to see higher average errors when θ approaches 0° and 90° , because of shadows and second-order shading, respectively. This explains the concave shape of the top graph of Fig. 5 as σ increases, varying θ ; oblique lighting at around $\theta=45^\circ$ will allow greater violations of the assumptions.

We mentioned in Section 2 that we assumed (for convenience) a Gaussian-shaped power spectrum for the surface. All the renderings so far have been on such Gaussian surfaces. Preliminary experiments were performed on Brownian surfaces [10], i.e., the power spectrum of the shape $\hat{\rho}_h(\mathbf{k})=\mathbf{k}^T\mathbf{k}^{-\beta/2}$ where we used $\beta=2$. Anisotropy was introduced in the same way as for the Gaussian surfaces, i.e., as a binomial form with the structure tensor: $\hat{\rho}_h(\mathbf{k})=(\mathbf{k}^TG_h^{-1}\mathbf{k})^{-1}$. The Brownian surfaces needed to be smoothed to be implemented (the roll-off frequencies were suppressed by multiplying the power spectrum with a Gaussian with identical G_h). Preliminary results on the anisotropic fractal surfaces were identical to that of the regular anisotropic Gaussian ones.

7. REAL-WORLD TEXTURES

The question naturally arises of how the combined measure will fare on real-world textures. We tested it on a selection of the Photex database [16,25]. From this database, we selected 26 textures, namely, aaa, aab, aaf, aai, aaj, aam, aan, aao, aap, aar, aas, aba, abj, abk, acc, acd, acf, acg, ack, acl, acm, afa, afb, afd, afe, and afg. We picked these because they appear to have only slight variations in albedo, with diffuse reflectance and varying degrees of anisotropy. We did not tweak any of the many hidden parameters (such as implementations of derivatives) for the benefit of the combined estimate, nor did we preprocess the images in any way. We ran the algorithm with the same parameters as for the rendered textures, except that we implemented smooth local averaging windows (2×2 as seen in Fig. 7). Figure 7 illustrate typical outputs for the texture aab (one of the textures where the combined estimate performed noticeably better). It is noteworthy that ground truth for ϕ is slightly inaccurate for local estimates. The black arrows in Fig. 7 illustrates the direction of illumination as indicated by the Photex database, but the light source was likely positioned quite close to the texture, resulting in slight variation in the illuminance flow field over the surface. Each texture has a great number of variations in light direction (both azimuthal and elevation), and we average over the entire collection.

The average error for the three estimators are, gradient based, 13.3° ; Hessian based, 11.4° ; and combined estimate, 17.2° . Thus the combined estimate cannot be expected to perform better than the gradient or the Hessian by themselves in the real world. This error amplification is due to the multiplicative factors of -3 and 4 evident in Eq. (25). Any error stemming from non-Lambertian, higher-relief, oblique viewing, etc. will be amplified. Also, errors due to image noise or the camera transfer function (from irradiance to image intensity) will all give rise to deviations that are amplified by the combined measure. Compared with the rendered surfaces, the real-world tex-

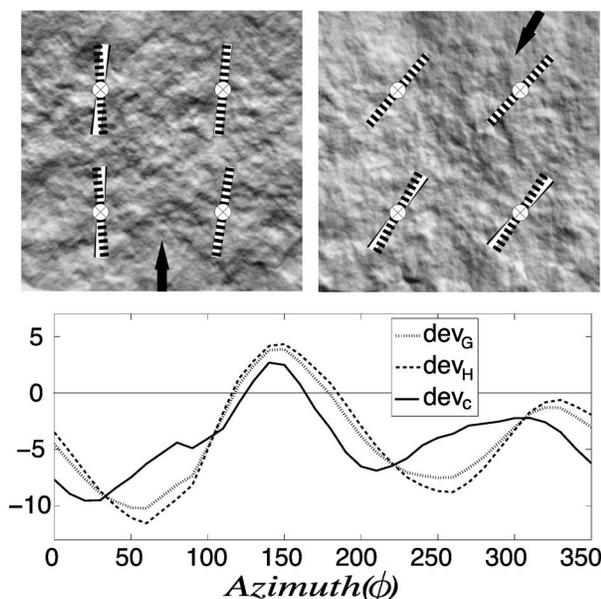


Fig. 7. Example of results from real-world textures. Top row, two images of the texture aab with light elevation (θ) = 45° ; they vary in azimuthal light angle (ϕ). The dotted line is the combined measure; the black line the gradient, and the white line the Hessian. The black arrows indicate global light direction: top left, $\phi = -90^\circ$; top right, $\phi = 60^\circ$. Bottom graph, results for the 36 different images corresponding to equidistant sampling of ϕ .

tures most certainly deviate more from the model, giving rise to more unstable behavior. A hint of this error amplification is also seen by the results on the rendered textures when strong shadows occur. In the second graph of Fig. 4, one can just make out small erratic deviations from the predictions for the gradient and the Hessian, yet these small erratic deviations are made very noticeable when they reemerge in the combined estimate.

8. CONCLUSIONS

The theory of illuminance flow estimation has been successfully generalized to the case of small amounts of surface anisotropy. This has been done by introducing structure tensor analysis of the surface (not just of the image, as was done previously). The surface structure tensors are not observable, yet their output (which constitutes measures of surface geometry) will affect the output of the image tensors; the form of this dependence is described by our theory. The theory is based on a simple approximation of the imaging process of the surface [Eq. (8)], valid under the assumptions of a low-relief isotropic Lambertian surface of uniform albedo, viewed normally in a light field [1] of one dominant directional component (e.g., a collimated beam).

A new measure, μ_C , which combines the image tensor outputs [Eq. (25)], has been devised based on the theory. The predicted output has been compared with results from computer-rendered Gaussian surfaces, where it has been found to be stable for small deviations from the assumptions of low relief and isotropy (as predicted), especially when the light elevation is oblique (say close to 45° or so). Results on real-world textures show the behavior for larger deviations from the model where the combined

estimate performs worse than the gradient or the Hessian by themselves. This is because of the error amplification due to the multiplicative factors of -3 and 4 [as seen in Eq. (25)]. The combined estimate is a nonlinear transform of the original vectors (seen as complex numbers, we take them to the power of -3 and 4 to change their angle) that propagates and amplifies the error. More stable estimates based on our theory can be derived and are a worthy subject of future work.

Interreflections are not factored into the approximation due to the low-relief assumption. For the statistics we investigate here, under the Lambertian assumption, the interreflections should have only a minor effect. This is especially true if the surface has low albedo, that is, if some amount of light is absorbed upon each reflection (most natural surfaces have a rather low albedo, with a reported average of 0.2 [26]). Even with unit albedo and high relief, however, there is no indication that incorporating more reflections will change the anisotropy of the resulting image. For isotropic Lambertian surfaces one can show [27] that interreflections will not (in theory) affect the output of the tensors.

In real scenes one expects surfaces to have nonuniform albedo and non-Lambertian reflectances and to be viewed from an arbitrary direction; therefore these cases will be discussed briefly.

If the albedo is nonuniform, we are faced with an additional 2D signal: an albedo map $a(\mathbf{r})$. The results of our findings will still hold if we relax the nonuniform albedo assumption to the case of an $a(\mathbf{r})$ that is uncorrelated with the surface $h(\mathbf{r})$, with the scale of the contents being sufficiently different; i.e., it boils down to a scale selection question with the outer scale defining areas of interest and the inner scale defining the size of derivative filters.

For non-Lambertian surfaces, we consider the bidirectional reflectance distribution function (BRDF) [28] denoted $B(\mathbf{v}, \mathbf{n}, \mathbf{i})$. Here, \mathbf{v} and \mathbf{i} indicate the viewing and light directions and \mathbf{n} the local normal. A Lambertian BRDF is constant, whereas a perfectly reflective surface (for which the present analysis is not intended) has a BRDF equalling a weighted δ -function. The local radiance due to a collimated beam is $I^{(B)}(x, y) = B(\mathbf{v}, \mathbf{n}(x, y), \mathbf{i})I(x, y)$, where I is the irradiance of the surface given by Eq. (5). So, for a nonconstant BRDF the structure of the surface radiance will be similar to that of a Lambertian surface if $\|\nabla_n B(\mathbf{v}, \mathbf{n}, \mathbf{i})\| \ll B(\mathbf{v}, \mathbf{n}, \mathbf{i})$. Thus, the Lambertian assumption can be relaxed to a smooth BRDF, which applies to most natural materials.

The issue of nonnormal views is important in applications and has been investigated by Pont and Koenderink [8]. They showed that, in theory, the illuminance orientation estimates can be corrected for oblique views by using surface attitude estimates, up to a viewing angle of 55° . For viewing angles larger than 55° there are no unique solutions. For the case where the surface is anisotropic and nonobliquely viewed one could model camera deviations from the normal position as an affine transformation of the (low relief) surface, that is, approximating the rotation as foreshortening along an axis. Similarly, surface anisotropy could be introduced as an affine transformation of an originally isotropic surface. Thus, we could model the nonnormal viewing of an anisotropic surface as

the composition of two affine transformations (which is itself affine) applied to an isotropic surface. Therefore, we expect that the combined effects of surface anisotropy and obliqueness will affect the illuminance flow estimates in a very systematic way. However, it cannot be corrected for because the individual effects on the radiance structure are very similar (or even the same in the low-relief approximation). This will be a subject of our future work.

To conclude, we have described the second-order statistics of anisotropic surfaces viewed under directed illumination. We have shown theoretically and experimentally that the anisotropy of a surface can be corrected for, for small amounts of anisotropy, by using a combination of the Hessian-based and gradient-based structure tensors. Finally, we discussed physical arguments for the generalization of our results far beyond our rather restrictive assumptions.

ACKNOWLEDGMENT

This work has been funded by EU project VISIONTRAIN (MRTN-CT-2004-005439).

REFERENCES

1. A. Gershun, "The light field," P. Moon and G. Timoshenko, translators, *J. Math. Phys.* (Cambridge, Mass.) **18**, 51–151 (1939).
2. B. M. ter Haar Romeny, L. Florack, J. J. Koenderink, and M. A. Viergever, "Scale-space: its natural operators and differential invariants," in *Information Processing in Medical Imaging*, A. C. F. Colchester and D. J. Hawkes, eds., Vol. 511 of Lecture Notes in Computer Science (Springer, 1991) pp. 239–255.
3. D. G. Lowe, "Object recognition from local scale-invariant features," in *The Proceedings of the Seventh IEEE International Conference on Computer Vision, 1999* (IEEE, 1999) Vol. 2, pp. 1150–1157.
4. J. Bigün, T. Bigün, and K. Nilsson, "Recognition by symmetry derivatives and the generalized structure tensor," *IEEE Trans. Pattern Anal. Mach. Intell.* **26**, 1590–1605 (2004).
5. D. J. Kriegman and P. N. Belhumeur, "What is the set of images of an object under all possible illumination conditions?" *Int. J. Comput. Vis.* **28**, 245–260 (1998).
6. W. Li, C. Wang, D. Xu, B. Luo, and Z. Chen, *A Study on Illumination Invariant Face Recognition Methods Based on Multiple Eigenspaces*, Vol. 3947 of Lecture Notes in Computer Science (Springer, 2005), pp. 131–136.
7. J. J. Koenderink and S. C. Pont, "Irradiation direction from texture," *J. Opt. Soc. Am. A* **20**, 1875–1882 (2003).
8. S. C. Pont and J. J. Koenderink, "Irradiation orientation from obliquely viewed texture," in *Deep Structure, Singularities, and Computer Vision*, O. F. Olsen, L. Florack, and A. Kuijper, eds. (Springer, 2005), pp. 205–210.
9. M. J. Chantler, "Why illuminant direction is fundamental to texture analysis." *IEE Proc. Vision Image Signal Process.* **142**, 199–206 (1995).
10. P. Kube and A. Pentland, "On the imaging of fractal surfaces," *IEEE Trans. Pattern Anal. Mach. Intell.* **10**, 704–707 (1988).
11. P. N. Belhumeur, D. Kriegman, and A. Yuille, "The bas-relief ambiguity," *Int. J. Comput. Vis.* **35**, 33–44 (1999).
12. M. Varma and A. Zisserman, "Estimating illumination direction from textured images," in *IEEE Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004* (IEEE, 2004), pp. I-179–I-186.
13. M. Berry and V. Hannay, "Umbilic points on Gaussian random surfaces," *J. Phys. A* **10**, 1809–1821 (1977).
14. M. S. Longuet-Higgins, "The statistical analysis of a random moving surface," *Philos. Trans. R. Soc. London, Ser. A* **249**, 321–364 (1956).
15. J. Bigün and G. H. Granlund, "Optimal orientation detection of linear symmetry," in *Proceedings of the IEEE First International Conference on Computer Vision* (IEEE, 1987), pp. 433–438.
16. O. Drbohlav and M. Chantler, "Illumination-invariant texture classification using single training images," in *Texture 2005: Proceedings of the 4th International Workshop on Texture Analysis and Synthesis* (IEEE, 2005), pp. 31–36.
17. M. Chantler, M. Petrou, A. Penirsche, M. Schmidt and G. McGunnigle, "Classifying surface texture while simultaneously estimating illumination direction," *Int. J. Comput. Vis.* **62**, 83–96 (2005).
18. A. P. Pentland, "Finding the illuminant direction," *J. Opt. Soc. Am.* **72**, 448–455 (1982).
19. D. Knill, "Estimating illuminant direction and degree of surface relief," *J. Opt. Soc. Am. A* **7**, 759–775 (1990).
20. Q. Zheng and R. Chellappa, "Estimation of illuminant direction, albedo and shape from shading," *IEEE Trans. Pattern Anal. Mach. Intell.* **13**, 680–702 (1991).
21. M. Chantler and G. Delguste, "Illuminant-tilt estimation from images of isotropic texture," *IEE Proc. Vision Image Signal Process.* **144**, 213–219 (1997).
22. Y. Zhang and Y. H. Yang, "Illuminant Direction Determination for Multiple Light Sources," in *IEEE Conference on Computer Vision and Pattern Recognition, 2000. Proceedings* (IEEE, 2000), Vol. 1, pp. 269–276.
23. M. Brooks and B. Horn, "Shape and source from shading," *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, A. K. Joshi, ed. (Morgan Kaufmann, 1985), 932–936.
24. S. Karlsson and J. Bigun, "Multiscale complex moments of the local power spectrum," *J. Opt. Soc. Am. A* **24**, 618–625 (2007).
25. M. Chantler, Photex Photometric Image Database, <http://www.macs.hw.ac.uk/texturelab/resources/databases/Photex/index.htm> (October 2007).
26. A. Pentland, "The visual inference of shape: computation from local features," Ph.D. dissertation (MIT, 1982).
27. J. J. Koenderink and A. J. van Doorn, "Geometrical modes as a general method to treat diffuse interreflections in radiometry," *J. Opt. Soc. Am.* **73**, 843–850 (1983).
28. F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis, "Geometrical considerations and nomenclature for reflectance" (National Bureau of Standards, 1977).