

Finding Folds: On the Appearance and Identification of Occlusion

Patrick S. Huggins

Hansen F. Chen

Peter N. Belhumeur*

Steven W. Zucker[†]

Center for Computational Vision and Control

Yale University

New Haven, CT 06520

Abstract

A natural sequel to edge detection is the interpretation of edges. This interpretation can provide useful information to various computer vision processes, including recognition, reconstruction, and tracking. In this paper we consider the problem of identifying occlusion edges in a single image. We examine the appearance of occlusion edges under variable illumination, both analytically and empirically, and find that the pattern of shading in the neighborhood of occlusion edges is a stable feature. Finally, we derive a filter for detecting occlusion and present the results of its application.

1. Introduction

Implicit in the extensive work done on edge detection is the belief that discontinuities in image intensity reveal scene structure [13][31]. However, edge detection alone confounds scene structure: reflectance discontinuities, shadows, surface normal discontinuities, and occlusion all give rise to edges in the image (see Figure 1). Edges resulting from occlusion are of special interest since they correspond to object boundaries. Identifying the occlusion edges is an important step in computer vision: among other things, it facilitates image segmentation, establishes depth-order relations [35], and is a valuable feature for recognition [38] and shape reconstruction [18][30][37], even in domains where multiple images are available [6][36][20]. Surprisingly, however, little has been done to characterize the appearance of occluding edges [40][1][9]. As a result, many algorithms rely on generic edge measurements, which include shadows and so on, when occlusion edge measurements alone would be more appropriate.

In this paper we focus on occlusion edges. Specifically, we consider a subset of occlusion edges, those that correspond to the boundaries of smooth objects; here edges arise as the surface of an object bends out of view, while the surface behind it is obscured. We call the side of the edge corresponding to the foreground a *fold* and the other side a *cut*.

We observe that folds have distinctive shading patterns, due to the projection, while cuts, in general, do not exhibit such a pattern, which suggests that folds and cuts may be distinguished. This leads us naturally to a classification of edges, in which oc-




Figure 1: Edge types and their appearance. (a) Self-occlusion by a smooth object. At point p the projective mapping is singular, while at q it is regular. (b) A surface normal discontinuity. Here the projective mapping is regular on both sides. (c) A reflectance discontinuity. (d) A lighting discontinuity results in a shadow.

clusion edges can be separated from non-occlusion edges and foreground-background relations can be established (Section 2).

As the shading pattern results from projection, we postulate that it is stable under variable illumination [3]. In Section 3 we develop an analytic model of folds to establish this. Then in Section 4 we compute image statistics under variable lighting for verification.

In Section 5, we derive a simple filter for detecting occlusion edges which arise from folds and show the results of its application, demonstrating our proposed categorization.

2. Folds

Occlusion edges arise from discontinuities in depth. In the case of smooth objects these discontinuities arise from singularities in the projective mapping from object surfaces to the image plane called folds [23].

Let Σ be a smooth (C^2) surface, i.e., locally $\Sigma : X \subset \mathbb{R}^2 \rightarrow$

*Supported by a Presidential Early Career Award IIS-9703134, NSF KDI-9980058, NIH R01-EY, and an NSF-ITR.

[†]Supported by AFOSR.

$Y \subset \mathbb{R}^3$, where X is the space of parameters of the surface and Y is a portion of Euclidean space in which the surface is embedded, and let $\hat{c} \in \mathbb{S}^2$ (the unit sphere) be the viewing direction. If $\Pi_{\hat{c}}$ is the orthographic projection ($\Pi_{\hat{c}} : Y \rightarrow Z \subset \mathbb{R}^2$, where Z is the image domain), then the surface to image mapping is $\Pi_{\hat{c}} \circ \Sigma$, which is from $\mathbb{R}^2$ to $\mathbb{R}^2$ (note that for what follows, the choice of orthographic projection can easily be generalized). This implies, via a theorem of Whitney [39][23], that the only stable singularities (i.e., points where the Jacobian is not of full rank) are folds and cusps; in this paper we consider only folds. We define fold points in the image domain as follows

Definition 1 *The fold is the set of points in the image where the mapping $\Pi_{\hat{c}} \circ \Sigma$ is singular and Σ is smooth.*

Complementary to folds are *cuts*. We define them to be the set of image points where image intensity is discontinuous and the projective mapping is regular (non-singular). This includes image edges arising from surface boundaries (e.g., the edge of a sheet of paper), surface normal discontinuities (e.g., polyhedral edges - see Figure 1 (b)), reflectance edges (Figure 1 (c)), and shadows (Figure 1 (d)).

An illustrative example is provided in Figure 1 (a). Here point p appears as a fold in the image, while point q appears as a cut, since it is occluded by p ; by definition they both project to the same point in the image. The edge point is thus both fold and cut simultaneously. What is key to the classification that we propose here is that in the neighborhood of p which appears projected in the image, the projective mapping becomes singular, while in the neighborhood of q which appears projected in the image, the projective mapping is regular. These neighborhoods form the two sides of the edge.

The geometric difference between folds and cuts results in an observable photometric difference [33][24][22]. If we define the *shading flow field* to be the vector field of tangents to isoluminance contours in the image [4], we can differentiate between the fold and cut neighborhoods of edges as follows. On the fold side of an edge, the shading flow field is tangent to the occluding contour, while on the cut side of an edge the field is transverse (non-tangent). See Figure 2. Furthermore it has been proven that this is the generic case [8][14]. This leads us to propose the following labeling of edges [15].

Classification 1 *At an edge point p of occluding contour γ in an image we can define two semi-open neighborhoods, N_p^A and N_p^B , where the surface to image mapping is continuous in each neighborhood. We can then classify p as follows:*

1. fold-cut: *The shading flow is tangent to γ at p in N_p^A and the shading flow is transverse to γ at p in N_p^B , with exception at isolated points.*
2. cut-cut: *The shading flow is transverse to γ at p in N_p^A and in N_p^B , with exception at isolated points.*
3. fold-fold: *The shading flow is tangent to γ in N_p^A and in N_p^B , with exception at isolated points.*




Figure 2: The shading flow field at a fold. (a) A fold occurs when the view direction ($\hat{c}$) lies in the tangent plane (T) to the surface. (b) In the image, the shading flow field is tangent to the occluding contour as a result of the singularity in the projection mapping.

In general, occlusion edges are the only edges that will appear as *fold-cut*, where the *fold* side corresponds to the foreground and the *cut* side to the background. Other edge types will appear as *cut-cut*, including edges such as shadows and reflectance discontinuities; these may be further classified [40][34][4][42]. We note that the third class, *fold-fold*, is accidental, or non-generic [10], in that it requires a precise alignment of two folds.

This categorization, while useful, is only qualitative. Our goal here is to quantitatively analyze the appearance of folds, so as to identify them and apply our classification.

3. Modeling the Appearance of Folds

Here we construct a model that describes the local appearance of a fold in an image. Our model makes simple (though easily generalized) assumptions about surface geometry, reflectance, and lighting; this enables us to analytically compute the relation of the shading flow field to the occluding contour. We are especially concerned with how this appearance varies as a function of lighting. Our results show that the shading flow is not only tangent to the occluding contour, but is also on average parallel to the occluding contour in a neighborhood around it. This parallelism is stable with respect to illumination variation. Furthermore we show that the magnitude of the image gradient also serves to indicate the presence of a fold.

Consider a surface patch (Σ) viewed such that it exhibits a fold. Let the surface patch be approximated by its second-order Taylor expansion about the origin, where the z -axis direction is the surface normal direction at the origin, the x -axis direction ($\hat{x}$) and the y -axis direction ($\hat{y}$) are the principal directions of the surface at the origin, and a and b are the corresponding principal curvatures with $a > b$ (Σ is a Monge patch)

$$z = \frac{1}{2}ax^2 + \frac{1}{2}by^2 + O((x, y)^3). \quad (1)$$

Let the origin appear as a fold in the image; i.e., we assume that the view direction ($\hat{c}$) lies in the x - y plane and subtends an angle of θ with $\hat{x}$.

Let the surface be Lambertian [27] with constant albedo (without loss of generality we set the albedo to unity), and let the light source be a point light source at infinity in the direction

1 (these assumptions are easily generalized). The radiance of a point on the surface is then given by

$$L = \hat{\mathbf{n}} \cdot \mathbf{l} \quad (2)$$

where $\hat{\mathbf{n}}$ is the unit surface normal at a point, defined as $\hat{\mathbf{n}} = (-\frac{\partial z}{\partial x}, -\frac{\partial z}{\partial y}, 1) / (1 + \frac{\partial z^2}{\partial x^2} + \frac{\partial z^2}{\partial y^2})^{1/2}$.

Given the results of the previous section, we would like to compute the shading flow field ($\mathbf{s}$); recall from Section 2 that the shading flow field is defined as the vector field of tangents to the isoluminance contours in the image, i.e., if I is image irradiance, then $\mathbf{s}$ is orthogonal to ∇I . To do this we first compute the gradient of scene radiance, expressed in terms of a local coordinate frame on the surface. We define $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$ to be the principal directions (if $\mathbf{v}$ is a vector, then $\hat{\mathbf{v}} = \mathbf{v} / \|\mathbf{v}\|$)

$$\begin{aligned} \hat{\mathbf{u}} &= \left(1 - \frac{a^2}{2}x^2\right)\hat{\mathbf{x}} - \left(\frac{a^2b}{|a-b|}xy\right)\hat{\mathbf{y}} + (ax)\hat{\mathbf{z}} \\ \hat{\mathbf{v}} &= \left(-\frac{b^2a}{|b-a|}yx\right)\hat{\mathbf{x}} - \left(1 - \frac{b^2}{2}y^2\right)\hat{\mathbf{y}} - (by)\hat{\mathbf{z}}. \end{aligned} \quad (3)$$

The corresponding principal curvatures, κ_u and κ_v , are

$$\begin{aligned} \kappa_u &= a - \frac{3}{2}a^3x^2 - \frac{1}{2}ab^2y^2, \\ \kappa_v &= b - \frac{3}{2}b^3y^2 - \frac{1}{2}ba^2x^2. \end{aligned} \quad (4)$$

The radiance gradient in the $\hat{\mathbf{u}}\text{-}\hat{\mathbf{v}}\text{-}\hat{\mathbf{n}}$ coordinate frame is then

$$\nabla L = l_u \kappa_u \hat{\mathbf{u}} + l_v \kappa_v \hat{\mathbf{v}}, \quad (5)$$

where $\mathbf{l} = (l_u, l_v, l_n)$ in the $\hat{\mathbf{u}}\text{-}\hat{\mathbf{v}}\text{-}\hat{\mathbf{n}}$ frame.

In all subsequent equations we assume a view centered coordinate frame such that $\hat{\mathbf{x}} = \hat{\mathbf{c}}$. The transformation from the Monge patch coordinate system to the view centered coordinate system is simply a rotation by $-\theta$ about the z -axis.

Assuming radiometric calibration of the lens, the image irradiance at a point p in the image can be found as the projection of the surface radiance onto the y - z plane

$$I(p) = L(\Pi_{\hat{\mathbf{x}}}^{-1}(p)). \quad (6)$$

The shading flow field in the image is the projection of the vector field that is orthogonal to the radiance gradient onto the y - z plane

$$\begin{aligned} \mathbf{s} &= \Pi_{\hat{\mathbf{x}}}((\nabla L)^\perp) = \\ &(-l_v \kappa_v u_y + l_u \kappa_u v_y)\hat{\mathbf{y}} + (-l_v \kappa_v u_z + l_u \kappa_u v_z)\hat{\mathbf{z}}, \end{aligned} \quad (7)$$

where $(\nabla L)^\perp$ is the vector orthogonal to ∇L on the surface. This expression gives the shading flow field in terms of our model.

To compare the shading flow field orientation to the edge orientation, we compute the slope of the shading flow field in the image (computing slope as opposed to orientation simplifies the analysis). The slope of the image shading flow field is




Figure 3: The statistics of a fold. A plot of the mean orientation of the shading flow field (shown as vectors) and its variance (indicated by intensity), for one instance of our model. Observe how the shading flow aligns with the fold (located at the top of the image) as the fold is approached and how the variance of this alignment decreases as well, indicating that this alignment is a stable image feature.

$$\omega = \frac{-l_v \kappa_v u_z + l_u \kappa_u v_z}{-l_v \kappa_v u_y + l_u \kappa_u v_y}. \quad (8)$$

Since the view direction is in the x - y plane, the slope of the occluding contour at the origin is zero.

In later computations it will be convenient to express l_v in terms of ω and l_u

$$l_v(\omega, l_u) = l_u \left(\frac{\kappa_u}{\kappa_v} \right) \left(\frac{v_z - v_y \omega}{u_z - u_y \omega} \right). \quad (9)$$

Now that we have computed a description of the shading flow field, we consider its appearance under variable illumination. We are interested in the orientation of the shading flow field relative to the orientation of the occluding contour, specifically in whether the tangency described in Section 2 is observable, and, if so, whether or not it is stable.

A straightforward numerical computation yields the mean and variance of the shading flow field at every point on a surface patch, per our model, assuming the light source direction $\mathbf{l}$ is a random point on the hemisphere centered on $\hat{\mathbf{z}}$. See Figure 3. The computation shows that the mean shading flow field is not only tangent to the occluding contour, but also parallel to it in a large neighborhood. Furthermore, the variance of the shading flow field is small close to the contour and gradually increases away from it. Thus the observed parallelism is stable in the neighborhood of the contour with respect to variable illumination.

Given the observed parallelism of the shading flow to the occluding contour, we restrict our analysis to image points normal to the edge point under consideration. In what follows we take the origin to be our occluding edge and we consider the statistics

of the shading flow field at image points normal to the edge, i.e., surface points in our model satisfying $y = 0$.

If in a given image we observe an edge and the shading flow field normal to the edge, we would like to be able to use the relative orientation of the shading flow field to the edge as evidence of the edge's type: e.g., if the shading flow field on one side of the edge is roughly parallel to the edge we might label that side of the edge as a fold.

Towards this, we compute the probability density of shading flow field slope under variable illumination. We assume a given surface viewed from a given direction, with a distribution on light sources that is uniform over light source directions and falls off in magnitude. Choosing simply a Gaussian, the probability density of light sources is

$$\rho_l(l) = \frac{1}{(\sqrt{2\pi}\sigma)^3} e^{-\frac{1}{2\sigma^2}(l_x^2 + l_y^2 + l_z^2)}. \quad (10)$$

We acknowledge that there are many choices for $\rho_l(l)$, some of which may arguably be more appropriate. In the context of examining appearance under variable illumination, we choose our distribution so as not to favor any particular light source direction.

The probability density of the slope of the image shading flow field can be obtained by combining this result with Equation 9

$$\rho_\omega(\omega; z, a, b, \theta) = \frac{1}{2\pi\sigma^2} \int_{-\infty}^{\infty} e^{-\frac{1}{2\sigma^2}(l_u^2 + l_v(\omega, l_u)^2)} \left| \frac{\partial(l_u, l_v)}{\partial(\omega, l_u)} \right| dl_u, \quad (11)$$

where z is the distance to the occlusion edge in the image, a and b are the principal curvatures of the Monge patch, and θ is the angle between the viewing direction and the first principal direction.

The resulting family of probability density functions is plotted in Figure 4(a). The most prominent feature of ρ_ω is the ridge close to $\omega = 0$, which indicates that the shading flow is on average parallel to the occluding contour. The sharpness of the ridge indicates the stability of this parallelism with respect to illumination variation. As expected from the tangency of the shading flow to the contour, the function ρ_ω approaches a Dirac delta function $\delta(\omega)$ when $z \rightarrow 0$. The observed ridge is always present in ρ_ω , though its spread and position are modulated by the view direction and the curvatures.

Since we do not know the viewing angle in advance, it is also important to look at the marginal density of the slope $\rho_\omega(\omega; z, a, b, \theta)$ integrated over the viewing angle θ

$$\rho_\omega^\theta(\omega; z, a, b) = \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \rho_\omega(\omega; z, a, b, \theta) d\theta. \quad (12)$$

We present numerical results in Figure 4(b). ρ_ω^θ has the same qualities as ρ_ω ; the observed ridge is centered on $\omega = 0$ and remains quite sharp, especially near the occluding contour.

A further observation about folds is that the image irradiance gradient magnitude depends on foreshortening: a radiance gradient on the surface will be amplified by the foreshortening that

occurs as the fold is approached. In the limit, the foreshortening will give rise to an infinite image irradiance gradient. This foreshortening is maximal when the radiance gradient is in the direction of the tilt of the surface near the fold, i.e., when the shading flow field is parallel to the occluding contour. Thus we expect the image irradiance gradient magnitude as a function of shading flow field slope to correlate with the probability density function of slope.

The image irradiance gradient is

$$\nabla I = \frac{\|\nabla L\|}{\left\| \Pi_{\hat{\mathbf{x}}} \left(\frac{\nabla L}{\|\nabla L\|} \right) \times \hat{\mathbf{s}} \right\|} \hat{\mathbf{x}} \times \hat{\mathbf{s}}. \quad (13)$$

Using Equations 13 and 10, we calculate the expectation of the magnitude of the image gradient

$$\mathbb{E}[\|\nabla I\|] = \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{\infty} \|\nabla I\| e^{-\frac{1}{2\sigma^2}l_u^2} dl_u \quad (14)$$

where l_v is given by Equation 9.

Figure 4(c) shows the mean magnitude of the image irradiance gradient. It is clearly correlated with the observed parallelism of shading flow field. In this sense, the gradient magnitude can be seen as an indicator of where in the distribution (Figure 4(a)) a particular instance of shading flow field orientation lies.

The above analysis suggests that the alignment of the shading flow with the contour of the fold is stable in the neighborhood of the occluding contour. It is insensitive to the lighting conditions, the geometry of the surface, and the view direction. We now proceed to verify the conclusions of our model with real images.

4. Measuring the Appearance of Folds

We need to evaluate the results of the previous section in the context of real images. The model discussed in the previous section does not capture all of the variability present in real images: albedos vary, surfaces are often non-Lambertian, small scale surface structure can significantly affect image measurements, and the imaging process is noisy. Given these factors, how robust is the parallelism between the shading flow field and the occluding contour, which we observed in the model, in real images?

To answer this question, we considered the statistics of the shading flow field and the gradient in images taken under variable illumination.

We measured the statistics of eight different objects under variable illumination. For each object, 64 images were taken, each illuminated by a single distinct flash. The flashes were uniformly distributed over the hemisphere centered over the object. This array of point light sources is effectively a sampling of the theoretical uniform distribution we considered in the previous section. Due to space limitations we present here the results from a single object. Figure 5 shows the images captured for this object.

We computed the mean and variance of the shading flow field and the irradiance gradient in the image. Some care needs to be taken in these computations. Consider the shading flow field




Figure 4: (a) The family of probability density functions for the slope of the shading flow field ($\rho_\omega(\omega)$). (b) $\rho_\omega^\theta(\omega)$, (a) integrated over viewing directions. (c) The expectation of the image gradient magnitude ($\mathbb{E}[\|\nabla I\|]$). In all plots *Distance* is measured in the image plane normal to the contour. The surface parameters are $a = 1$ and $b = 0.3$. In (a) and (c) $\theta = -\pi/6$. The plots are clipped for display purposes.




Figure 5: (left) An experimental subject: a toy rubber duck. (right) Images of the duck under illumination from 64 different point sources.

measured at a point in image i illuminated by light source direction $\mathbf{l}_i = (l_u, l_v, l_n)$; if the light source illuminating image j is $\mathbf{l}_j = (-l_u, -l_v, l_n)$, then the shading flow vectors in images i and j will cancel one another at that point, even though in some sense their orientation is the same. Thus we define orientation to be π -periodic in our computations.

The expected shading flow field is then

$$\mathbb{E}[\hat{\mathbf{s}}] = \angle_{\times \frac{1}{2}} \left(\int_{\mathbf{l} \in \mathbb{S}^2} \angle_{\times 2}(\hat{\mathbf{s}}) \|\mathbf{l} \times \hat{\mathbf{z}}\| d\mathbf{l} \right), \quad (15)$$

where $\angle_{\times 2}$ and $\angle_{\times \frac{1}{2}}$ map vectors to double- and half-angled vectors respectively. We define the variance of the shading flow field as

$$\mathbb{V}[\hat{\mathbf{s}}] = \int_{\mathbf{l} \in \mathbb{S}^2} \|\angle_{\times 2}(\hat{\mathbf{s}}) - \angle_{\times 2}(\mathbb{E}[\hat{\mathbf{s}}])\|^2 \|\mathbf{l} \times \hat{\mathbf{z}}\| d\mathbf{l}. \quad (16)$$

Figure 6 shows the expected shading flow field for the duck image set (Figure 5), in a region of the image centered on the back of the duck's neck. The alignment of the shading flow field with the occluding contour on the fold side of the edge is clear. In contrast, the shading flow field on the cut side of the edge exhibits no such alignment. Furthermore, on the fold side of the edge the variance is low, thus the observed parallelism is robust with respect to illumination variation, while the variance is high on the cut side, and so even if there exist lighting conditions where the shading flow field here appears parallel to the occluding contour, those situations occur with low probability. These statistics agree nicely with those of our model (Figure 3).

Figure 7(a) shows $\mathbb{V}[\hat{\mathbf{s}}]$ for the rubber duck. The regions of low variance are where the shading flow field is stable; these are

most notably the folds and the reflectance edges (the eye of the duck). In regions where the fold is smooth (i.e., the curvature of the surface in the view direction is low), such as the wing of the duck, the shading flow is stable over a greater extent.

As we saw in the previous section, the magnitude of the image gradient is usually largest at the occluding contour and, on the fold side of the contour, decreases as the distance from the contour increases. We define the expected gradient similarly to the expectation of the shading flow field (Equation 15)

$$\mathbb{E}[\nabla I] = \angle_{\times \frac{1}{2}} \left(\int_{\mathbf{l} \in \mathbb{S}^2} \angle_{\times 2}(\nabla I) \|\mathbf{l} \times \hat{\mathbf{z}}\| d\mathbf{l} \right), \quad (17)$$

In Figure 7(b) we show $\|\mathbb{E}[\nabla I]\|$. The folds appear accentuated as the gradient magnitude gives more weight to aligned shading flow vectors, as discussed in Section 3. The results suggest that the gradient magnitude is a useful measure of the significance of a particular orientation measurement.

5. Deriving Fold-Filters

Having established the statistical stability of the appearance of folds under variable illumination, we consider the problem of their identification. The process we suggest here consists of edge detection, followed by the application of a fold-filter, which we develop here, at each side of every edge point, followed by a decision process based on the fold-filter output.

Edge detection is still an active field of research in computer vision [12], and the quality of the edge measurements we use as a substrate for our fold filter will have a significant impact upon our results. We use the logical-linear edge detector [19] to measure edge position and orientation. Ideally we would like to measure the curvature of the edge so as to better fit our model; however we do not explore this possibility here.

Our fold-filters are constructed using the Karhunen-Loève transform (principal component analysis) [29] (see also [16][7][17][2]) based on the model developed in Section 3. The essential idea is to compare an image region to the space of images generated by our fold model to assess the probability that the image region is a fold. The filters serve to project the image region under consideration onto the space of fold-images; from




Figure 6: The mean shading flow field, $\mathbb{E}[\hat{s}]$, in the neck region of the duck (highlighted region of Figure 5), shown as vectors, superimposed on its variance, $\mathbb{V}[\hat{s}]$. Observe the parallelism between the shading flow field on the fold side of the occluding contour and the contour itself, and the accompanying low variance. Compare with Figure 3.

this projection a measure of “fold-ness” can be computed. Note that these filters are derived from observations and thus learnable; previous work [15] has .

In order to generate the space of fold images, we need to determine the distributions of the model parameters, in this case, a , b , θ , and l . Reasonable assumptions might be that a and b are uniformly distributed over some range, θ is uniformly distributed over $[0, 2\pi)$, and l is, as before, uniformly distributed over the hemisphere. Sampling the parameter space, the images are defined by Equation 6.

The probability density of these images is then

$$\rho(I) \propto \exp\left(-\frac{1}{2} \iint_{p, p' \in Z} (I(p) - \mathbb{E}[I(p)]) C^{-1}(p, p') (I(p') - \mathbb{E}[I(p')]) dp dp'\right) \quad (18)$$

where C is the covariance of I , i.e. $C(p, p') = \mathbb{E}[(I(p) - \mathbb{E}[I(p)])(I(p') - \mathbb{E}[I(p')])]$ where the expectation is taken over the set of images generated by the parameters a , b , θ , and l under our assumed distributions, Z is the image domain, and p and p' are points in the image. By Mercer’s Theorem,

$$C(p, p') = \sum_{n=1}^{\infty} \sigma_n^2 \psi_n(p) \psi_n(p'), \quad (19)$$

where the σ_n^2 ’s and ψ_n ’s are the eigenvalues and corresponding eigenfunctions of C ; the ψ_n ’s will become the image filters.




Figure 7: Statistics for the rubber duck images. (a) Shading flow field variance, $\mathbb{V}[\hat{s}]$, where dark/bright indicates low/high variance. The shading flow field is clearly stable at folds. The gradient magnitude amplifies the appearance of the folds, as can be seen in (b) the magnitude of the mean image irradiance gradient, $\|\mathbb{E}[\nabla I]\|$.

Substituting Equation 19 into Equation 18, and taking an N -dimensional approximation to the result, we obtain

$$\rho(I) \propto \exp\left(-\frac{1}{2} \left[\sum_{n=1}^N \left(\frac{(I - \mathbb{E}[I]) \cdot \psi_n}{\sigma_n} \right)^2 + \frac{\|I - \mathbb{E}[I]\|^2 - \sum_{n=1}^N ((I - \mathbb{E}[I]) \cdot \psi_n)^2}{\sigma_{N+1}^2} \right] \right). \quad (20)$$

The resulting filters are $\psi_1, \dots, \psi_N$, which are applied directly to the image, the outputs of which are combined as in Equation 20 to give $\rho(I)$. We now illustrate the application of this method.

6. Results

Using the described construction, we designed a fold filter and applied it to several images. We assumed that $\theta = 0$, a is uniformly distributed over $[0, 1]$, and $b = 0$. The filter used is a window of 6×12 square pixels. With three free parameters, we chose $N = 3$. The three eigenfunctions are shown in Figure 6. We make the fold/cut distinction simply by thresholding $\rho(I)$. This particular choice of parameters is equivalent to locally approximating the occluding contour with a cylinder of variable radius under variable lighting.




Figure 8: The top three eigenfunctions of the fold-image space, which we use as fold-filters.

We applied these filters to the synthetic image of a Klein bottle, an image of the rubber duck, and an image of a human face. Figure 9 shows the results, where the right column shows our classification of edges.




Figure 9: Edge classification. Edge detection is run on an image (left); at each detected edge (middle) our fold filters are applied, then $\rho(I)$ is computed and thresholded (at 0.4) to give our classification (right). Black indicates that the side of the edge is a *cut*, white indicates a *fold*.

The results demonstrate the potential of our filter. The occluding contour of the Klein bottle is successfully identified as *fold-cut* everywhere, except right at the cusp points, where the shading flow swings around too sharply for our filter to detect its non-tangency to the contour on the cut side of the edge. Similarly, the occluding contour of the duck is picked out correctly everywhere, the exception being the duck's beak, which is too sharp for our particular parameter settings. Finally, on the face image we successfully categorize the chin, the cheeks, the lips, the nose, and parts of the ears.

7. Conclusions

We have examined the appearance of occlusion edges, both analytically and empirically, and we have shown that folds exhibit a stable shading flow pattern with respect to variable illumination. This pattern does not depend on the assumptions made by our model: it is the result of degeneracy in the projection mapping, which suggests that our results generalize to a variety of reflectance and illumination functions [32][28]. Based on this observed regularity, we derived a filter for fold detection and demonstrated its viability. The filter was designed using our model to generate a set of images upon which the filter was based.

While our filter demonstrates both the feasibility of local edge

classification and the potential for learning filters from variable illumination data, several issues remain unexplored. The parameters we use in our experiments represent the simplest possible choice; clearly the parameter space needs to be explored. For example we could incorporate contour curvature measurements [25][41] to steer our model and thus the filters. We could also apply the filter at various scales to detect folds of varying curvature; this could provide an estimate of the surface curvature in the viewing direction [21]. Furthermore, our computations are purely local; we expect the use of global constraints would improve our detection. For simplicity our filters are derived from an imposed analytic model; similar, perhaps more accurate, filters could be derived directly from a database of images where the occluding contour is identified (e.g., [26]). Finally, our filter is based on the probability of a fold generating the encountered image - we have not modeled other image formation processes, which would then lead to a more informed decision process. All of these issues are subjects for future research.

Finally, we note that though edges vary dramatically under variable illumination, occlusion edges are stable. This suggests that edge detection can be successfully applied in domains where illumination has proven problematic, such as recognition, provided it includes edge classification, as presented here.

References

- [1] N. Asada, H. Fujiwara, and T. Matsuyama. "Seeing behind the scene: analysis of photometric properties of occluding edges by the reversed projection blurring model." In *Proc. ICCV*, pp. 150-155, 1995.
- [2] S. Baker, S.K. Nayar, and H. Murase. "Parametric feature detection." *IJCV*, 27(1):27-50, 1998.
- [3] P.N. Belhumeur and D.J. Kriegman. "What is the set of images of an object under all possible illumination conditions?" *IJCV*, 28(3):245-260, 1998.
- [4] P. Breton and S.W. Zucker. "Shadows and shading flow fields." In *Proc. CVPR*, pp. 782-789, 1996.
- [5] H.F. Chen, D. Jacobs, and P.N. Belhumeur. "In search of illumination invariants." In *Proc. CVPR*, pp. 254-261, 2000.
- [6] R. Cipolla and A. Blake. "Surface shape from the deformation of apparent contours." *IJCV*, 9(2):83-112, 1992.
- [7] P.A. Devijver and J. Kittler. *Pattern Recognition: A Statistical Approach*. Prentice Hall, 1982.
- [8] J.P. Dufour. "Familles de courbes planes différentiables." *Topology*, 4:449-474, 1983.
- [9] D.A. Forsyth and A. Zisserman. "Reflections on shading." *IEEE PAMI*, 13(7):671-679, 1991.
- [10] W.T. Freeman. "The generic viewpoint assumption." *Nature*, 368:542-545, 1994.
- [11] E.R. Hancock and J.V. Kittler. "Edge-labeling using dictionary-based relaxation." *IEEE PAMI* 12(2):165-181, 1990.
- [12] M.D. Heath, S. Sarkar, T. Sanocki, and K.W. Bowyer. "A robust visual method for assessing the relative performance of edge detection algorithms." *IEEE PAMI*, 19(12):1338-1359, 1997.
- [13] B.K.P. Horn. "Understanding image intensities." *Artificial Intelligence*, 8:201-231, 1977.
- [14] P.S. Huggins and S.W. Zucker. "How folds cut a scene." In *International Workshop on Visual Form 4*, pp. 323-332, 2001.
- [15] P.S. Huggins and S.W. Zucker. "Folds and cuts: how shading flows into edges." In *Proc. ICCV*, Vol. II, pp. 153-158, 2001.
- [16] R.A. Hummel. "Feature detection using basis functions." *CVGIP*, 9(1):40-45, 1979.
- [17] R.A. Hummel and S.W. Zucker. "A three-dimensional edge operator." *IEEE PAMI*, 3(3):324-331, 1981.
- [18] K. Ikeuchi and B.K.P. Horn. "Numerical shape from shading and occluding boundaries." *Artificial Intelligence*, 17:141-184, 1981.
- [19] L.A. Iverson and S.W. Zucker. "Logical/linear operators for image curves." *IEEE PAMI*, Vol. 17(10):982-996, 1995.
- [20] F. Kahl and K. Åström. "Motion estimation in image sequences using the deformation of apparent contours." In *Proc. ICCV*, pp. 934-944, 1998.
- [21] J.J. Koenderink. "What does the occluding contour tell us about solid shape?" *Perception*, 13:321-330, 1976.
- [22] J.J. Koenderink. *Solid Shape*. MIT Press, 1990.
- [23] J.J. Koenderink and A.J. van Doorn. "Singularities of the visual mapping." *Biological Cybernetics*, 24:51-59, 1976.
- [24] J.J. Koenderink and A.J. van Doorn. "Local features of smooth shapes: ridges and courses." In *SPIE Vol. 2031 Geometric Methods in Computer Vision II*, pp. 2-13, 1993.
- [25] J.J. Koenderink and W. Richards. "Two-dimensional curvature operators." *J. Opt. Soc. Amer. A*, 5(7), 1988.
- [26] S. Konishi, A.L. Yuille, J.M. Coughlan, and S.C. Zhu. "Fundamental bounds on edge detection: an information theoretic evaluation of different edge cues." In *Proc. CVPR*, Vol. I, pp. 573-579, 1999.
- [27] J. Lambert. *Photometria sive de mensura et gradibus luminis, colorum et umbrae*. Eberhard Klett, 1760.
- [28] M. Langer and S.W. Zucker. "What is a light source?" In *Proc. CVPR*, pp. 172-178, 1997.
- [29] M.M. Loève. *Probability Theory*. Van Nostrand, 1955.
- [30] J. Malik. "Interpreting line drawings of curved objects." *IJCV*, 1:73-103, 1987.
- [31] D. Marr. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. W.H. Freeman and Co., 1982.
- [32] M. Oren and S.K. Nayar. "Generalization of the Lambertian model and implications for machine vision." *IJCV*, 14(3):227-251, 1995.
- [33] J.H. Rieger. "The geometry of view space of opaque objects bounded by smooth surfaces." *Artificial Intelligence*, 44:1-40, 1990.
- [34] J.M. Rubin and W.A. Richards. "Color vision and image intensities: when are changes material?" *MIT AI Memo*, 1981.
- [35] P.A. Smith, T. Drummond, and R. Cipolla. "Edge tracking for motion segmentation and depth ordering." In *BMVC*, Vol. II, pp. 584-593, 1999.
- [36] R. Szeliski and R. Weiss. "Robust shape recovery from occluding contours using a linear smoother." *IJCV*, 28(1):27-44, 1998.
- [37] R. Vaillant, and O.D. Faugeras. "Using extremal boundaries for 3-D object modeling." *IEEE PAMI*, 14:157-173, 1992.
- [38] B. Vijayakumar, D.J. Kriegman, and J. Ponce. "Invariant-based recognition of complex curved 3D objects from image contours." In *Proc. ICCV*, pp. 508-514, 1995.
- [39] H. Whitney. "On singularities of mappings of Euclidean spaces I, Mappings of the plane into the plane." *Ann. of Math.*, 62:374-410, 1955.
- [40] A.P. Witkin. "Intensity-based edge classification." In *AAAI-82*, pp. 36-41, 1982.
- [41] M. Worring and A.W.M. Smeulders. "Digital curvature estimation." *CVGIP*, 58:366-382, 1993.
- [42] W. Zhang and F. Bergholm. "Multiscale blur estimation and edge type classification for scene analysis." *IJCV*, 24(3):219-250, 1997.