P.S. Huggins and S.W. Zucker (2001), "Folds and Cuts: How shading flows into edges," Proc. Eigth International Conf. on Computer Vision, vol. II, pp. 153 - 158.

Folds and Cuts: How Shading Flows Into Edges

Patrick S. Huggins Steven W. Zucker*
Department of Computer Science
Yale University
New Haven, CT 06520

Abstract

We consider the interactions between edges and intensity distributions in semi-open image neighborhoods surrounding them. Locally this amounts to a kind of figure-ground problem, and we analyze the case of smooth figures occluding arbitrary backgrounds. Techniques from differential topology permit a classification into what we call folds (the side of an edge from a smooth object) and cuts (the arbitrary background). Intuitively, cuts arise when an arbitrary scene is "cut" from view by an occluder. The condition takes the form of transversality between an edge tangent map and a shading flow field, and examples are included.

1 Introduction

Edges lie at a representational level between images and models, and can bridge them efficently for many indexing and recognition applications. But this bridge is incomplete when the surface structure in the neighborhood of edges is questionable. Certain questions are thought to be locally undecidable; e.g., the classical Gestalt ownership question: does the edge belong to the figure, or to the ground? Others are known to be combinatorially difficult: perfect line drawing interpretation is NP-complete for the simple blocks world [16]. Various heuristics, such as closure or convexity, have been suggested [9], but these do not clarify the connection between edges and surfaces. Are there circumstances in which the inverse image of an edge onto a surface can be characterized? An examination of natural images suggests that the intensity distribution in the neighborhood of edges contains relevant information, and our goal in this paper is to show one basic way to exploit it.

The intuition is provided in two steps. First, edges often signal occlusion, and one of the principle cues available from edges relates to local occlusions, such as the ”T”-junctions illustrated in Fig. 1. We focus on the observation that, if edge orientation is taken as an explicit ”dimension”, then edge detector responses lifted into (position, orientation)-space can be revealing (fig. 1(c)). In particular, piecewise C^1 curves lift to continuous curves except at orientation discontinuities.

Our second step relates to the shading distribution around an edge. This, of course, connects us to the ways surfaces can approach an edge, and the next example (fig. 2) shows how shading can be lifted into (orientation, position)-space. From a viewer's perspective for this and the previous example, edges arise when the tangent plane to the object ”folds” out of sight; this naturally suggests a type of ”figure”, which we show is both natural and commonplace. In particular, it enjoys a stable pattern of shading (with respect to the edge). Notice in particular that the lift again shows a jump in (position, orientation)-space, now between the cylinder and the background. Furthermore, notice that this lift overlaps the edge lift on the fold side, thereby answering the Gestalt question of ownership for the edge. But most importantly, the fold side of the edge ”cuts” the background scene, which implies that the background cannot exhibit this regularity in general.

Our main contribution in this paper is to develop this difference between *folds* and *cuts* in a technical sense. We employ the techniques of differential topology to capture qualitative aspects of shape (cf. Koenderink [11]), and propose a specific mechanism for classifying folds and cuts based on the interaction between edges and the shading flow field. The result is further applicable to formalizing an earlier classification of shadow edges [1].

2 Folds and Cuts

Consider an image ($I : Z \subset \mathbf{R}^2 \to \mathbf{R}^+$) of a smooth ($C^2$) surface $\Sigma : X \subset \mathbf{R}^2 \to Y \subset \mathbf{R}^3$; here X is the surface parameter space and Y is 'the world'. For a given viewing direction $\mathbf{V} \in \mathbf{S}^2$ (the unit sphere), the surface is projected onto the image plane by $\Pi_{\mathbf{V}} : Y \to Z \subset \mathbf{R}^2$. For simplicity, we assume that Π is orthographic projection, although this particular choice is not crucial to our reasoning. Thus the mapping from the surface domain to the image domain takes $\mathbf{R}^2$ to $\mathbf{R}^2$. See Fig. 3.

Points in the resulting image are either regular or sin-

* Supported by AFOSR

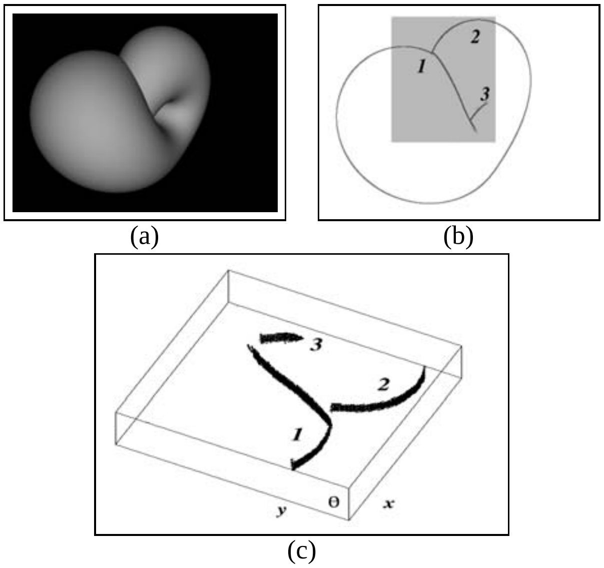

Figure 1: (a) A Klein bottle. (b) Its edge map. (c) The edge map lifted tp x-y-θ space. Notice the jump of the lift at the points of edge orientation discontinuity. Only the portion of the edge map highlighted is shown to simplify the display, and numbers indicate corresponding curve segments in (b) and (c).

gular, depending on whether the Jacobian of the surface to image mapping, $d(\Pi_{\mathbf{V}} \circ \Sigma)$ is of full rank or not. An important result in differential topology is the Whitney Theorem for mappings from $\mathbf{R}^2$ to $\mathbf{R}^2$ [6][12], which states that such mappings generically have only two types of singularities, folds and cusps. (By generic we mean that the singularities persist under perturbations of the mapping.)

Let $T_x[A]$ denote the tangent space of the manifold A at the point x.

Definition 1 *The* FOLD *is the singularity locus of the surface to image mapping, $\Pi_{\mathbf{V}} \circ \Sigma$, where Σ is smooth. In the case of orthographic projection the fold is the image of those points on the surface whose tangent plane contains the view direction.*

$$\gamma_{fold} = \{z_p \in Z |\ \mathbf{V} \in T_{y_p}[\Sigma(X)],\ y_p = \Sigma(x_p), \\ z_p = \Pi_{\mathbf{V}}(y_p)\}$$

We denote the *fold generator*, i.e. the pre-image of γ_{fold} on Σ, by

$$\Gamma_{fold} = \{y_p \in Y |\ x_p \in X, \mathbf{V} \in T_{y_p}[\Sigma(X)],\ y_p = \Sigma(x_p)\}$$

Since the singularities of $\Pi_{\mathbf{V}} \circ \Sigma$ lead to discontinuities if we take Z as the domain, they naturally translate into edges in the image corresponding to the occluding contour and its

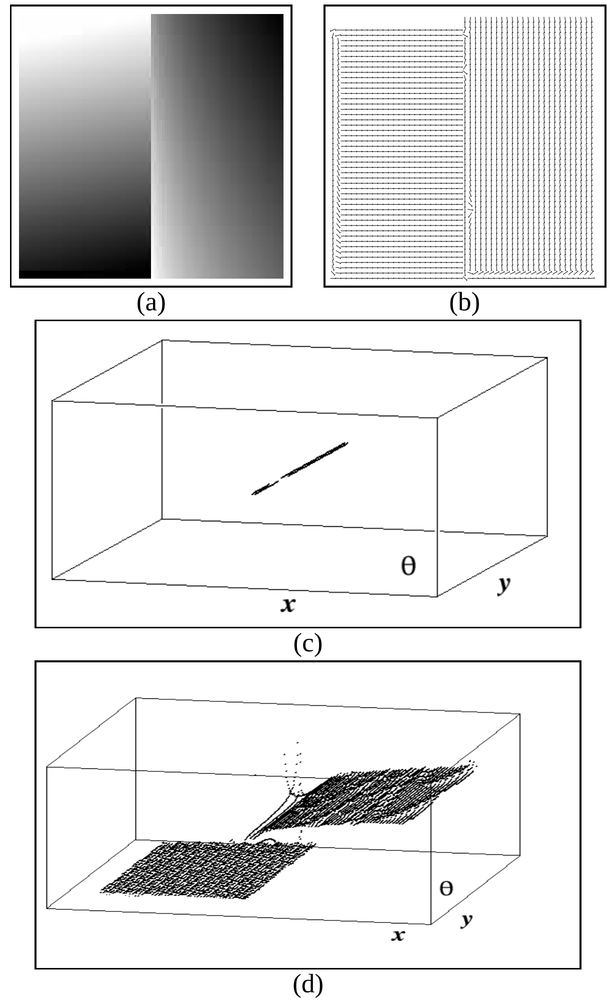

Figure 2: (a) An image of a typical edge due to occlusion. Does the right half of the image appear as a vertically oriented cylinder? (b) The shading flow field of the image; note its orientation relative to that of the edge. (c) The edge map lifted to x-y-θ space. (d) The shading flow field as it appears in x-y-θ space. Here the shading in the righthand side of the image is clearly tangent to the edge, while the shading in the lefthand side is transverse. We denote these configurations as FOLD and CUT, respectively.

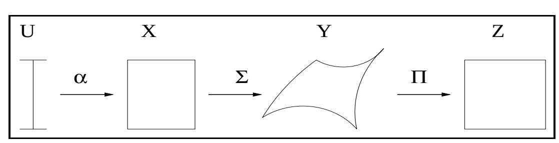

Figure 3: The mappings referred to in the paper, from the coordinates of a surface (X), to Euclidean space (Y), to the image domain (Z). The map α is used later to describe surface curves.

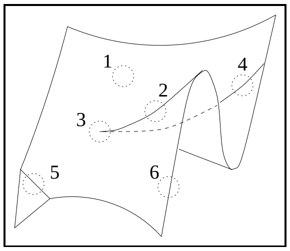

Figure 4: The categories of points of a mapping from $\mathbf{R}^2$ to $\mathbf{R}^2$: (1) a regular point, (2) a fold point, (3) a cusp, (4) a Γ-shadow point, (5) a crease point, (6) a boundary point. The viewpoint is taken to be at the upper left. From this position the fold (solid line) and the Γ-shadow (dashed line) appear aligned.

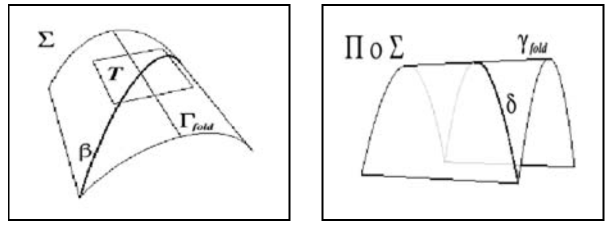

(a) (b)

Figure 5: (a) A curve, $\beta = \Sigma \circ \alpha$, passing through a point on the fold generator, Γ_{fold}; as defined by the viewpoint at the lower left. The tangent to the curve $T[\beta]$ at the point of intersection lies in the tangent plane to the surface at that point, $\mathbf{T}$, as does the tangent to the fold generator, $T[\Gamma_{fold}]$. (b) In the image, the tangent plane to the surface at the fold projects to a line, and so the curve, $\delta = \Pi(\beta)$,is tangent to the fold, γ_{fold}.

end points (although due to occlusion and opacity not all of these points, as defined, are visible).

Now suppose Σ is piecewise smooth, i.e. we permit discontinuities of all orders in Σ. We now have two additional sources of discontinuity in the image mapping: points where the surface itself is discontinuous,

$$\Gamma_{boundary} = \{y_p \in Y|\ \exists\delta \in \mathbf{S}^1,\ \lim_{\varepsilon\to 0} \Sigma(x_p + \varepsilon\delta) \neq \Sigma(x_p),\\ y_p = \Sigma(x_p)\}$$

and points where the surface normal is discontinuous,

$$\Gamma_{crease} = \{y_p \in Y|\ \exists\delta \in \mathbf{S}^1,\ \lim_{\varepsilon\to 0} N(x_p + \varepsilon\delta) \neq N(x_p),\\ y_p = \Sigma(x_p)\}$$

As a result of occlusion, the occluding edges present in the image have two pre-images in the scene: the edge of the occluder, and the curve that this projects to, along the view direction, on the occluded background. We denote this second locus of points as Γ-*shadow*,

$$\Gamma_{\Gamma-shadow} = \{y_p \in Y|\ \exists t \in \mathbf{R}^+,\ y_p = y_q + t\mathbf{V},\\ y_q \in \Gamma_{fold} \cup \Gamma_{boundary} \cup \Gamma_{crease}\}$$

Fig. 4 summarizes the points defined so far.

Definition 2 *The* CUT *is the set of points in the image where the image is discontinuous due to occlusion, surface discontinuities, or surface normal discontinuities.*

$$\gamma_{cut} = \{z_p \in Z|\ z_p \in \Pi(\Gamma_{\Gamma-shadow} \cup \Gamma_{boundary} \cup \Gamma_{crease})\}$$

Note that $\gamma_{fold} \cap \gamma_{cut}$ may be non-empty, while their respective pre-images are disjoint, except at special points such as corners and cusps.

If a surface has a pattern on it, such as shading, the geometry of folds and cuts give rise to distinct patterns in the image. Identifying folds and cuts is useful as a prerequisite for geometrical analysis [8][11][14]. We show in the next section how the geometry of folds and cuts naturally leads to a basis for distinguishing between γ_{fold} and γ_{cut}.

2.1 Curves and Flows at Folds and Cuts

Consider a surface viewed such that its image has a fold, with a curve on the surface which runs through the fold. In general, the curve in the image osculates the fold (Fig. 5).

Let α be a smooth (C^2) curve on the surface parameterized by X; $\alpha : U \subset R \to X$. If $\Sigma \circ \alpha$ passes through point y_p on the surface then $T_{y_p}[\Sigma \circ \alpha(U)] \subset T_{y_p}[\Sigma(X)]$. An immediate consequence of this for images is that, if we choose $\mathbf{V}$ such that $\Pi_{\mathbf{V}}(y_p) \in \gamma_{fold}$, then the image of α is tangent to the fold, i.e. $T_{z_p}[\Pi \circ \Sigma \circ \alpha(U)] = T_{z_p}[\gamma_{fold}(Y)]$, where $z_p = \Pi_{\mathbf{V}}(y_p)$.

There is one specific choice of $\mathbf{V}$ for which this does not hold: $\mathbf{V} \in T_{y_p}[\Sigma \circ \alpha(U)]$. At such a point $\Pi \circ \Sigma \circ \alpha(U)$ has a cusp and is transverse (non-tangent) to γ_{fold}.

Intuitively, it seems that the image of α should be tangent to γ_{fold} "most of the time". Situations in which the image of α is not tangent to γ_{fold} result from the "accidental" alignment of the viewer with the curve. The notion of "generic viewpoint" is often used in computer vision to discount such accidents. We use the concept of *general position*, or *transversality*, from differential topology, to distinguish between typical and atypical situations.

Definition 3 *[7]: Let M be a manifold. Two submanifolds $A, B \subset M$ are* IN GENERAL POSITION, *or* TRANSVERSAL, *if $\forall p \in A \cap B$, $T_p[A] + T_p[B] = T_p[M]$.*

We call a situation typical if the conditions under which it occurs are transversal, atypical (accidental) otherwise.

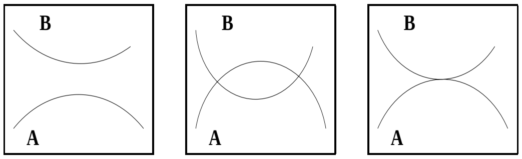

Figure 6: Transversality. (a) A and B do not intersect. Thus they are transversal. (b) A and B intersect transversally. A small motion of either curve leaves the intersection intact. (c) A non-transverse intersection. A small motion of either curve transforms (c) into (a) or (b).

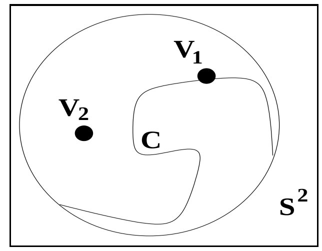

Figure 7: The tangent field of a curve, α, on a surface is $C = T[\Sigma \circ \alpha(U)]$, which traces a curve in $\mathbf{S}^2$. When the viewpoint $\mathbf{V}$ intersects C, the curve α is tangent to the fold in the image. This situation ($\mathbf{V}_1$) is not transversal, and thus only occurs accidentally. The typical situation is $\mathbf{V}_2$, and so in the image α is typically tangent to the fold where they intersect.

See Fig. 6. Other attempts to characterize such differences are probabilistic [5].

We show that if we view an arbitrary smooth curve, on an arbitrary smooth surface, from an arbitrary viewpoint, then typically at the point where the curve crosses the fold in the image, the curve is tangent to the fold. We do so by showing that in the space of variations, the set of configurations for which this holds is transversal, while the non-tangent configurations are not transversal.

Theorem 1 *If, in an image of a surface with a curve lying on the surface, the curve on the surface crosses the fold generator, then the curve in the image will typically appear tangent to the fold at the corresponding point in the image.*

Proof: For the image of α to appear transverse to the fold, we need $T_{y_p}[\Sigma \circ \alpha(U)] = \mathbf{V}$ at some point $y_p \in \Gamma_{fold}$. $T[\Sigma \circ \alpha(U)]$ traces a curve in $\mathbf{S}^2$, possibly with self intersections. $\mathbf{V}$ however is a single point in $\mathbf{S}^2$. At $T[\Sigma \circ \alpha(U)] = \mathbf{V}$ we note that $T_{\mathbf{V}}[T[\Sigma \circ \Sigma \circ \alpha(U)]] \cup T_{\mathbf{V}}[\mathbf{V}] = T_{\mathbf{V}}[T[\Sigma \circ \Sigma \circ \alpha(U)]] \cup \emptyset \neq T_{\mathbf{V}}[\mathbf{S}^2]$, thus this situation is not transversal. If $T[\Sigma \circ \alpha(U)] \neq \mathbf{V}$ then $T[\Sigma \circ \alpha(U)] \cap \mathbf{V} = \emptyset$. See Fig. 2.1.

This result appears in different forms in several domains, e.g. work on line drawing interpretation [15], shadows [3][10], and silhouettes [19], the main difference being the physical nature of the curve α.

For a family of curves on a surface, the situation is similar: along a fold, the curves are typically tangent to the fold. However, along the fold we expect the tangents to the curves vary, and perhaps at some point coincide with the view direction.

Theorem 2 *In an image of a surface with a family of smooth curves lying on the surface, the curves crossing the fold generator typically are everywhere tangent to the fold in the image, except at isolated points.*

Proof: Let $A : (U, V) \subset \mathbf{R}^2 \to X$ define a family of curves on a surface. As before, a curve appears transverse to the fold if its tangent is the same as the view direction: $T_{y_p}[\Sigma \circ A(U, V)] = \mathbf{V}$, and $\mathbf{V}$ is a point in $\mathbf{S}^2$. Now $T_U[\Sigma \circ A(U, V)]$ is a surface in $\mathbf{S}^2$. The singularities of such a field are generically folds and cusps (again applying the Whitney Theorem), and so $\mathbf{V}$ does not intersect the singular points transversally. However, $\mathbf{V}$ will intersect the regular portion of $T_U[\Sigma \circ A(U, V)]$, and such an intersection is transversal: $T_{\mathbf{V}}[T_U[\Sigma \circ A(U, V)]] = T_{\mathbf{V}}[\mathbf{S}^2]$. The dimensionality of this intersection is zero: thus non-tangency occurs at isolated points along γ_{fold}. The number of such points depends on the singular stucture of the vector field [18].

Dufour [4] classifies the possible diffeomorphic forms families of curves under smooth mappings from $\mathbf{R}^2$ to $\mathbf{R}^2$ can take. At folds a feature of this classification is the tangency condition we've just shown, as noted by Rieger[17].

For a discontinuity in the image not due to a fold, the situation is reversed: for a curve to be tangent to the edge locus, it must have the exact same tangent as the edge (Fig. 8), and similarly for a family of curves. This is stated in the next two theorems.

Theorem 3 *If, in an image of a surface with a curve lying on the surface, the curve on the surface crosses the cut generator, then the curve in the image will typically appear transverse to the cut at the corresponding point in the image.*

Proof: For $\Pi_{\mathbf{V}} \circ \Sigma \circ \alpha$ to be tangent to γ_{cut}, we need $T_{z_p}[\Pi \circ \Sigma \circ \alpha(U)] = T_{z_p}[\gamma_{cut}]$, which only occurs when $T_{y_p}[\Sigma \circ \alpha(U)] = T_{x_p}[\Gamma_{cut}]$, or equivalently $T_{x_p}[\alpha(U)] = T_{x_p}[\Sigma^{-1} \circ \Gamma_{cut}]$. Consider the space $\mathbf{R}^2 \times \mathbf{S}^1$. $\alpha \times T[\alpha]$ traces a curve in this space, as does $\Sigma^{-1} \circ \Gamma_{cut} \times T[\Sigma^{-1} \circ \Gamma_{cut}]$. We would not expect these two curves to intersect transversally in this space, and indeed: $p \in \alpha \times T[\alpha] \cap \Sigma^{-1} \circ \Gamma_{cut} \times T[\Sigma^{-1} \circ \Gamma_{cut}] \neq T_p[\mathbf{R}^2 \times \mathbf{S}^1]$.

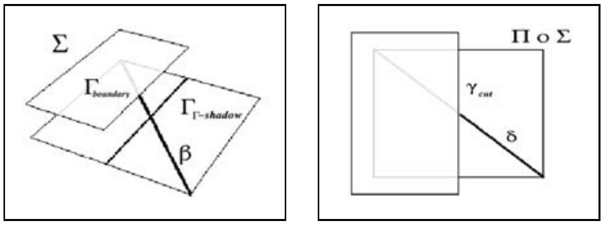

(a) (b)

Figure 8: (a) A curve, $\beta = \Sigma \circ \alpha$, passing through a point on the cut generator, $\Gamma_{\Gamma-shadow}$; their tangents are unequal. Here the viewpoint is above and the cut is caused by occlusion from another cut, $\Gamma_{boundary}$. (b) In the image, the curve, $\delta = \Pi(\beta)$, and the cut, γ_{cut}, are transverse at their intersection, as there is no degeneracy in the projection.

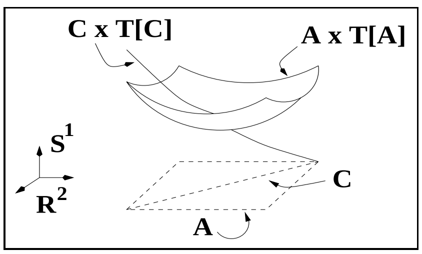

Figure 9: If A is a family of smooth curves lying on a surface, then $A \times T[A(U,V)]$, traces a surface in $\mathbf{R}^2 \times \mathbf{S}^1$, while, letting $C = \Sigma^{-1} \circ \Gamma_{cut}$, i.e. the cut locus in surface coordinates, $C \times T[C]$ traces a curve. When the two intersect, the curves of A are tangent to the cut in the image. This situation is transversal, but the intersection itself has dimension zero.

Theorem 4 *In an image of a surface with a family of smooth curves lying on the surface, the curves crossing the cut generator typically are everywhere transverse to the cut in the image, except at isolated points.*

Proof: For $\Pi_{\mathbf{V}} \circ \Sigma \circ A(U,V)$ to be tangent to γ_{cut}, we need $T_{z_p}[\Pi \circ \Sigma \circ A(U,V)] = T_{z_p}[\gamma_{cut}]$, which only occurs when $T_{y_p}[\Sigma \circ A(U,V)] = T_{y_p}[\Gamma_{cut}]$. In $\mathbf{R}^2 \times \mathbf{S}^1$, $A \times T[A]$ is a surface, and $\Sigma^{-1} \circ \Gamma_{cut} \times T[\Sigma^{-1} \circ \Gamma]$ is a curve. The intersection of these two objects is transverse: $p \in A \times T[A] \cap \Sigma^{-1} \circ \Gamma_{cut} \times T[\Sigma^{-1} \circ \Gamma_{cut}] = T_p[\mathbf{R}^2 \times \mathbf{S}^1]$. See Fig. 9.

Thus, in an image of a surface with a family of curves on the surface, there are two situations: (FOLD) the curves are typically tangent to the fold, with isolated exceptional points; (CUT) the curves are typically transverse to the cut, with isolated exceptional points.

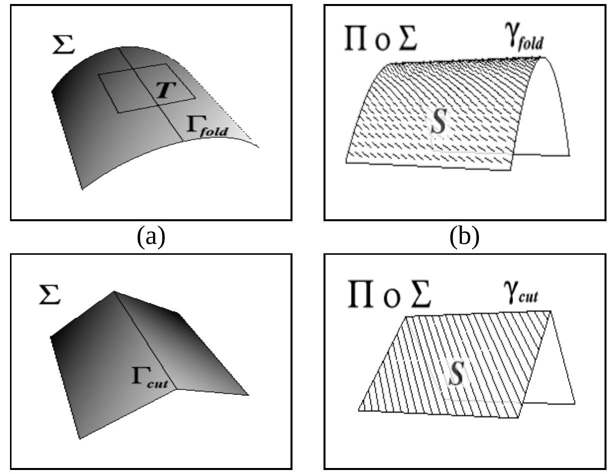

(c) (d)

Figure 10: Shaded surfaces with a fold (a) and a cut (c); the viewpoint is from the lower left. Their respective shading flow fields approach tangency near the fold (b), but remain transverse at the cut (d).

2.2 The Shading Flow Field at an Edge

Now consider a surface Σ under illumination from a point source at infinity in the direction L. If the surface is Lambertian then the shading at a point p is $s(p) = N \cdot L$ where N is the normal to the surface at p; this is the standard model assumed by most shape-from-shading algorithms. We define the *shading flow field* to be the unit vector field tangent to the level sets of the shading field:

$$\mathbf{S} = \frac{1}{\sqrt{(\frac{\partial s}{\partial x})^2 + (\frac{\partial s}{\partial y})^2}}(-\frac{\partial s}{\partial y}, \frac{\partial s}{\partial x})$$

The structure of the shading flow field can be used to distinguish between several types of edges, e.g. cast shadows and albedo changes [1]. Applying the results of the previous section, the shading flow field can also be used to categorize edge neighborhoods as *fold* or *cut*.

Since Σ is smooth (except possibly at Γ_{cut}), N varies smoothly, and as a result so does s. Thus $\mathbf{S}$ is the tangent field to a family of smooth curves. Consider $\mathbf{S}$ at an edge point p. If p is a fold point, then in the image $\mathbf{S}(p) = T_p[\gamma_{fold}]$. If p is a cut point, then $\mathbf{S}(p) \neq T_p[\gamma_{cut}]$. (Fig. 10)

Proposition 1 *At an edge point $p \in \gamma$ 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 (Fig. 2.2). We can then classify p as follows:*

1. FOLD-FOLD*: The shading flow is tangent to γ in N_p^A and in N_p^B, with exception at isolated points.*

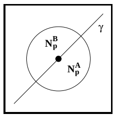

Figure 11: The semi-open neighborhoods of Proposition 1.

2. 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.*

3. 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.*

Figs. 12 thru 2.2 illustrate the applicability of our categorization.

These categorizations are computable locally. Furthermore, the advantage of introducing the differential topological analysis for this problem is that it is readily generalized to more realistic shading distributions (or unrealistic ones, see Fig. 15). For example, shading that results from diffuse lighting can be expressed in terms of an aperture function that smoothly varies over the surface [13], meeting the conditions we described in Section 2, thus enabling us to make the fold-cut distinction. The same analysis could be applied to texture or range data (see Fig. 16).

References

[1] Breton, P. and Zucker, S.W.: Shadows and shading flow fields. CVPR (1996) 782–789

[2] Do Carmo M.P.: Differential Geometry of Curves and Surfaces. (1976) Prentice-Hall

[3] Donati, L., Stolfi, N.: Singularities of illuminated surfaces. International Journal Computer Vision **23** (1997) 207–216

[4] Dufour, J.P.: Familles de courbes planes differentiables. Topology **4** (1983) 449–474

[5] Freeman, W.T.: The generic viewpoint assumption in a framework for visual perception. Nature **368** (1994) 542–545

[6] Golubitsky M., Guillemin, M.: Stable Mappings and Their Singularities. (1973) Springer-Verlag

[7] Hirsch, M.: Differential Topology. (1976) Springer-Verlag

[8] Ikeuchi, K., Horn, B.K.P.: Numerical shape from shading and occluding boundaries, AI **17** (1981) 141–184

[9] Kanizsa, G.: Organization in Vision. (1979) Praeger

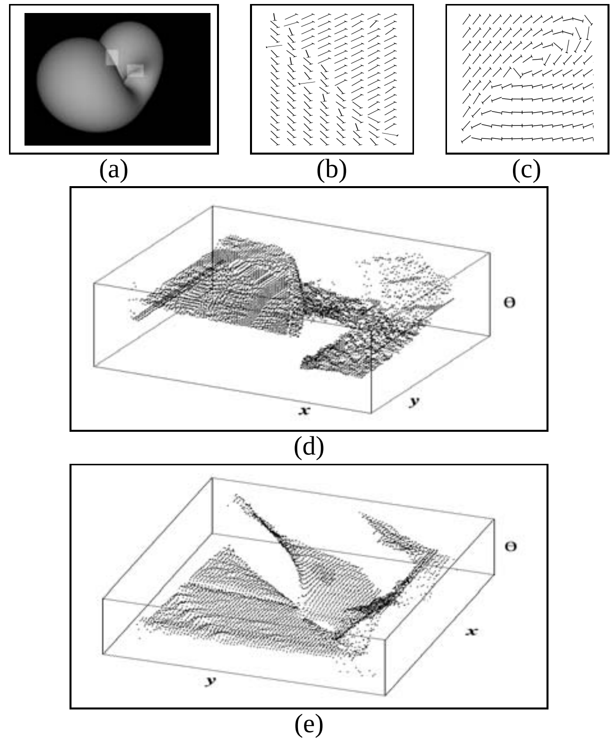

Figure 12: The Klein bottle (a) and its shading flow field at a fold (b) and a cusp (c) (the highlighted regions of (a)). On the fold side of the edge, the shading flow field is tangent to edge, while on the cut side it is transverse. In the vicinity of a cusp, the transition is evident as the shading flow field swings around the cusp point and becomes discontinuous. (d) and (e) depict these flow fields in x-y-θ space.

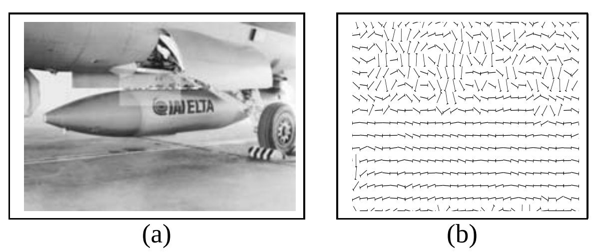

Figure 13: A scene with folds and cuts. We highlight the region where the pod obscures the hatch: the shading flow field indicates the edge to be of the FOLD-CUT type, suggesting that the pod is the "figure" side of the edge.

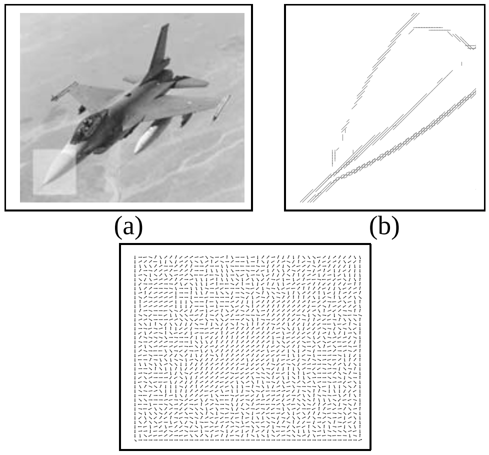

Figure 14: An F-16. We highlight the jet's nose cone region where the edge is clearly of the FOLD-CUT type, suggesting that the airplane is the "figure".

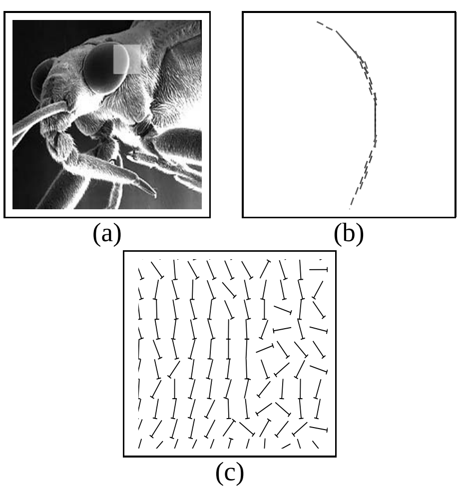

Figure 15: (a) An electron micrograph. The shading in the image is unnatural, yet occlusion relationships are clear. (b) The edge map of the highlighted region. (c) The shading flow field is tangent to the occlusion edge, as expected.

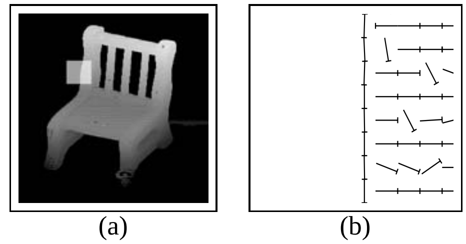

Figure 16: (a) A range image. As the iso-depth contours lie on the surface of the object, our categorization should apply. (b) The tangent field to the iso-depth contours in the highlighted region. Despite sensor noise, the flow remains transverse to the edge almost everywhere and agrees with the interpretation of the edge as a crease.

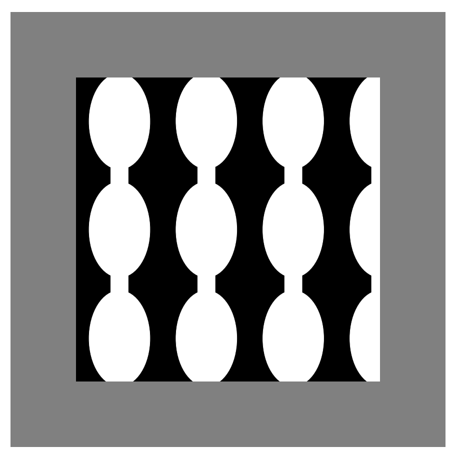

Figure 17: An ambiguous figure-ground image, after Kanizsa [9]. The lack of shading information prevents us from typing the edges.

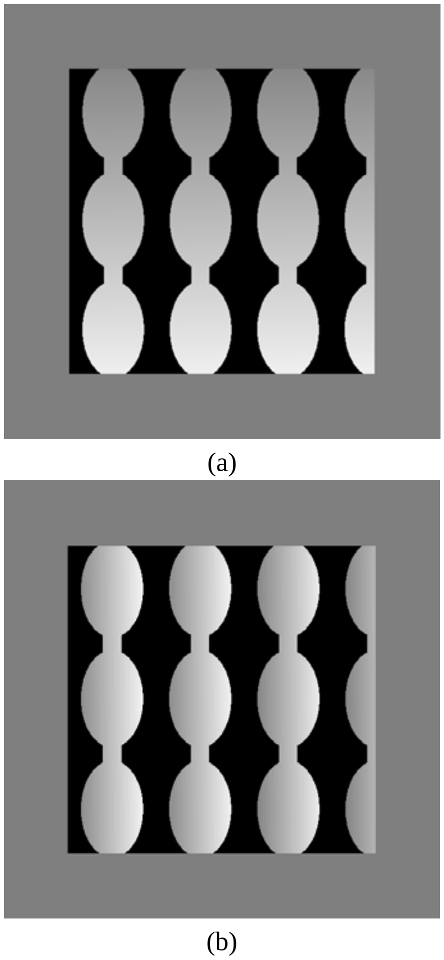

Figure 18: The same as Fig. 17 except the white region is shaded so as to be the (a) cut and (b) fold side of the image edges.

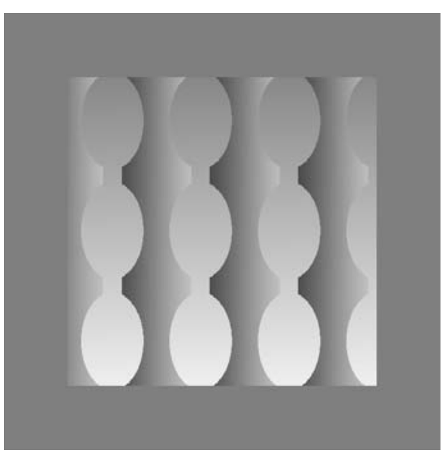

Figure 19: Fig. 18(a) with additional shading so that the dark region appears folded.

[10] Knill, D.C., Mamassian, P., Kersten D.: The geometry of shadows. J Opt Soc Am A **14** (1997)

[11] Koenderink, J.J.: What does the occluding contour tell us about solid shape? Perception **13** (1976) 321–330

[12] J.J. Koenderink and A.J. van Doorn: The singularities of the visual mapping. Biological Cybernetics **24** (1976) 51–59

[13] Langer, M.S., Zucker, S.W.: Shape from shading on a cloudy day. J Opt Soc Am A **11** (1994) 467–478

[14] Malik, J. Interpreting line drawings of curved objects. International Journal of Computer Vision **1** (1987) 73–107

[15] Nalwa, V.S.: Line-drawing interpretation: a mathematical framework. International Journal of Computer Vision **2** (1988) 103–124

[16] Parodi, P.: The complexity of understanding line-drawings of origami scenes, International Journal of Computer Vision **18** (1996) 139–170

[17] Rieger, J.H.: The geometry of view space of opaque objects bounded by smooth surfaces. Artificial Intelligence **44** (1990) 1–40

[18] Thorndike, A.S., Cooley, C.R., Nye, J.F.: The structure and evolution of flow fields and other vector fields. J. Phys. A. **11** (1978) 1455–1490

[19] Tse, P.U., Albert, M.K.: Amodal completion in the absence of image tangent discontinuities