P. Huggins and S.W. Zucker, "How folds cut a scene," Proc. Fourth Int. Workshop on Visual Form, Capri, May 2001.
in C. Arcelli, L. Cordella, and G. Sanniti di Baja (eds), Visual Form 2001, LNCS 2059, Springer, New York, 323 - 332.

How Folds Cut a Scene

Patrick S. Huggins and Steven W. Zucker*

Yale University, New Haven CT 06520, USA
huggins@cs.yale.edu
zucker@cs.yale.edu

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

On which side of an edge is figure; and on which ground? This classical Gestalt question is thought to be locally undecidable, and ambiguous globally (Fig. 1(a)). Even perfect line drawing interpretation is combinatorially difficult (NP-complete for the simple blocks world) [14], and various heuristics, such as closure or convexity, have been suggested [8]. Nevertheless, an examination of natural images suggests that the intensity distribution in the neighborhood of edges does contain relevant information, and our goal in this paper is to show one basic way to exploit it.

The intuition is provided in Fig. 1(b). From a viewer's perspective, 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). But more importantly, the fold side of the edge "cuts" the background scene, which implies that the background cannot exhibit this regularity in general; see Fig. 1(c).

Our main contribution in this paper is to develop the difference between *folds* and *cuts* in a technical sense. We employ the techniques of differential topology to capture qualitative aspects of shape, 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].

* Supported by AFOSR

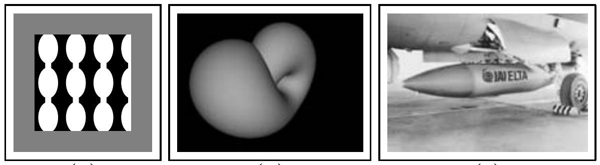

(a) (b) (c)

Fig. 1. (a) An ambiguous image. The edges lack the information present in (b), a Klein bottle. The shading illustrates the difference between the "fold", where the normal varies smoothly to the edge until it is orthogonal to the viewer, and the "cut". (c) An image with pronounced folds and cuts.

2 Folds and Cuts

Figure-ground relationships are determined by the positions of surfaces in the image relative to the viewer, so we are specifically interested in edges resulting from surface geometry and viewing.

Consider an image ($I : Z \subset \mathbb{R}^2 \to \mathbb{R}^+$) of a smooth ($C^2$) surface $\Sigma : X \subset \mathbb{R}^2 \to Y \subset \mathbb{R}^3$; here X is the surface parameter space and Y is 'the world'. For a given viewing direction $\mathbf{V} \in \mathbb{S}^2$ (the unit sphere), the surface is projected onto the image plane by $\Pi_{\mathbf{V}} : Y \to Z \subset \mathbb{R}^2$. 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 $\mathbb{R}^2$ to $\mathbb{R}^2$. See Fig. 2(a).

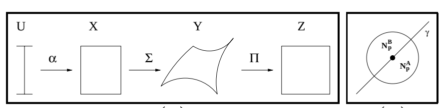

(a) (b)

Fig. 2. (a) The mappings referred to in the paper, from the coordinates of a surface (X), to Euclidean space (Y), to the image domain (Z) we will refer to. The map α is used later to describe surface curves. We omit the intensity mapping as we are interested in geometric discontinuities. (b) At an edge point, the edge defines two semi-open neighborhoods, the characteristics of which can be used to determine figure-ground.

Points in the resulting image are either regular or singular, depending on whether the Jacobian of the surface to image mapping, $d(\Pi_{\mathbf{V}} \circ \Sigma)$ is full rank or not. An important result in differential topology is the Whitney Theorem for mappings from $\mathbb{R}^2$ to $\mathbb{R}^2$ [6], 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 end points.

Note that due to occlusion and opacity, not all of the singularities present in a given image mapping will give rise to edges in the image. The edge in an image corresponding to a fold also corresponds to two curves on the surface: the fold generator and another curve, the locus of points occluded by the fold. We call this the *fold shadow*,

$$\Gamma_{fold-shadow} = \{y_p \in Y | \; \exists t \in \mathbb{R}^+, \; y_p = y_q + t\mathbf{V}, \; y_q \in \Gamma_{fold}\}$$

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,

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

and points where the surface normal is discontinuous,

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

Fig. 3 summarizes the points we've defined.

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.*

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

Note that $\gamma_{fold} \subset \gamma_{cut}$, while their respective pre-images are disjoint, except at special points such as T-junctions.

If a surface has a pattern on it, such as shading, the geometry of folds gives rise to a distinct pattern in the image. Identifying the fold structure is naturally useful as a prerequisite for geometrical analysis [10][15]. It is the contrast of this structure with that of cuts which is intriguing in the context of figure-ground. Our contribution develops this as a basis for distinguishing between γ_{fold} and γ_{cut}.

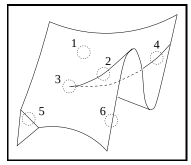

Fig. 3. Categories of points of a mapping from $\mathbb{R}^2$ to $\mathbb{R}^2$: (1) a regular point, (2) a fold point, (3) a cusp, (4) a fold-shadow point, (5) a crease point, (6) a boundary point. The viewpoint is taken to be at the upper left. From this position the fold (solid line) and the fold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

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. 4).

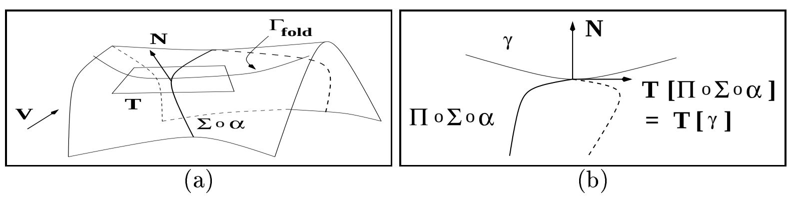

Fig. 4. A curve, $\Sigma \circ \alpha$, passing through a point on the fold generator, Γ_{fold}. (a) The tangent to the curve $T[\Sigma \circ \alpha]$, lies in the tangent plane to the surface, $\mathbf{T}$, as does the tangent to the fold generator, $T\Gamma_{fold}$. (b) In the image, the tangent plane to the surface at the fold projects to a line, and so the curve is tangent to the fold.

Let α be a smooth (C^2) curve on Σ; $\alpha : U \subset R \to X$. If α passes through point $y_p = \Sigma \circ \alpha(u_p)$ 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 $z_p = \Pi_{\mathbf{V}}(y_p) \in \gamma_{fold}$, then the image of α is tangent to the fold, i.e. $T_{z_p}[\Pi \circ \Sigma\alpha(U)] = T_{z_p}[\gamma_{fold}(Y)]$.

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 its configuration is transversal, atypical (accidental) otherwise. See Fig. 5. Other attempts to characterize such differences are probabilistic [5][17].

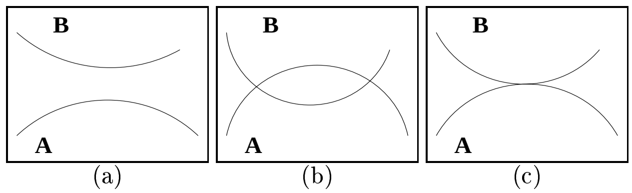

Fig. 5. Transversality. (a) A and B do not intersect. Thus they are transversal. (b) A and B intersect transversally. A small motion of either curve leaves the intersection intact. (c) A non-transverse intersection. A small motion of either curve transforms (c) into (a) or (b).

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.

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 $\mathbb{S}^2$, possibly with self intersections. $\mathbf{V}$ however is a single point in $\mathbb{S}^2$. At $T[\Sigma \circ \alpha(U)] = \mathbf{V}$ we note that $T_{\mathbf{V}}[T[\Sigma \circ \alpha(U)]] \cup T_{\mathbf{V}}[\mathbf{V}] = T_{\mathbf{V}}[T[\Sigma \circ \alpha(U)]] \cup \emptyset \neq T_{\mathbf{V}}[\mathbb{S}^2]$, 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 is our first result:

Result 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.*

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 the tangents to the curves vary, and may at some point coincide with the view direction. The typical situation is that the curves are tangent to the fold, except at isolated points on the fold, where they are transverse.

Let $A : (U, V) \subset \mathbb{R}^2 \to X$ 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 $\mathbb{S}^2$. Now $T_U[\Sigma \circ A(U, V)]$ is a surface in $\mathbb{S}^2$. The singularities of such a field are generically folds and

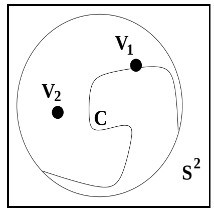

Fig. 6. The tangent field of α, $C = T[\Sigma \circ \alpha(U)]$, traces a curve in $\mathbb{S}^2$. When $\mathbf{V}$ intersects C, the curve α is tangent to the fold in the image. This situation ($\mathbf{V}_1$) is not transversal, and thus only occurs accidentally. The typical situation ($\mathbf{V}_2$) is α tangent to the fold when it crosses.

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: $T_{\mathbf{V}}[T_U[\Sigma \circ A(U,V)]] = T_{\mathbf{V}}[\mathbb{S}^2]$. The dimensionality of this intersection is zero: and so non-tangency occurs at isolated points along γ_{fold}. The number of such points depends on the singular stucture of the vector field [16]. This gives us:

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

Similar arguments can be made for more general projective mappings. Dufour [4] has classified the possible diffeomorphic forms families of curves under mappings from $\mathbb{R}^2$ to $\mathbb{R}^2$ can take.

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. 7).

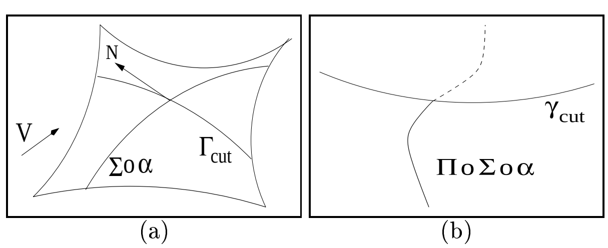

Fig. 7. The appearance of a curve intersecting a cut. (a) At a cut, the tangent plane to the surface does not contain the view direction. As a result there is no degeneracy in the projection, and so the curve will appear transverse to the cut in the image (b).

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 $\mathbb{R}^2 \times \mathbb{S}^1$. $\alpha \times T[\alpha]$ 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: $p \in \alpha \times T[\alpha] \cap \Sigma^{-1} \circ \Gamma_{cut} \times T[\Sigma^{-1} \circ \Gamma_{cut}] \neq T_p[\mathbb{R}^2 \times \mathbb{S}^1]$.

Result 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.*

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 $\mathbb{R}^2 \times \mathbb{S}^1$, $A \times T[A]$ 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: $p \in A \times T[A] \cap \Sigma^{-1} \circ \Gamma_{cut} \times T[\Sigma^{-1} \circ \Gamma_{cut}] = T_p[\mathbb{R}^2 \times \mathbb{S}^1]$. See Fig. 8.

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

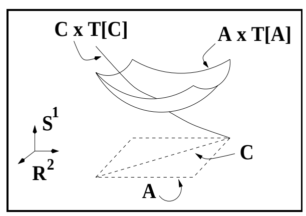

Fig. 8. $A \times T[A(U,V)]$, traces a surface in $\mathbb{R}^2 \times \mathbb{S}^1$, while, letting $C = \Sigma^{-1} \circ \Gamma_{cut}$, $C \times T[C]$ traces a curve. When the two intersect, the curves of A are tangent to the cut in the image. This situation is transversal, but has dimension zero.

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.

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. Applying the results of the previous section, the shading flow field can 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. 9)

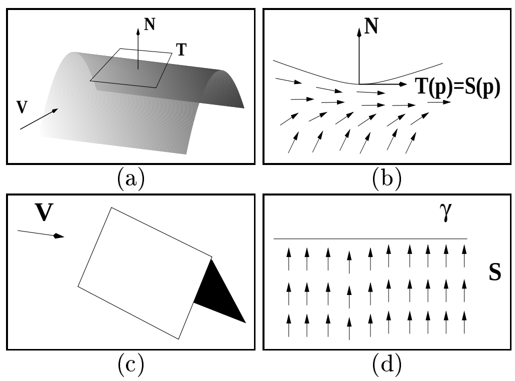

Fig. 9. The shading flow field at an edge. Near a fold (a) the shading flow field becomes tangent to the edge (b). At a cut (c), the flow is transverse (d).

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(b)). 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.*
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. 10,11, and 12 illustrate the applicability of our categorization.

These categorizations are computable locally, and are intimately related to figure-ground discrimination. Furthermore, the advantage of introducing the differential topological analysis for this problem is that it is readily generalized to more realistic shading distributions. For example, shading that results from diffuse lighting can be expressed in terms of an aperture function that smoothly varies over the surface [12], 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. Examples of single curves on surfaces which can be treated in the same way are occluding contours [13] and cast shadows [9].

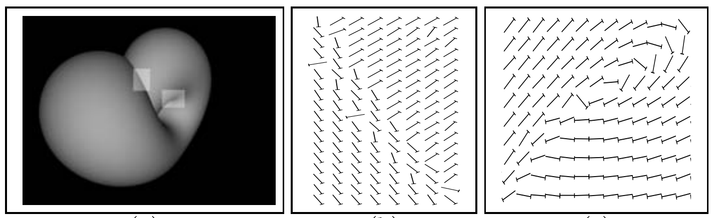

(a) (b) (c)

Fig. 10. The Klein bottle (a) and its shading flow field at a fold (b) and a cusp (c). 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.

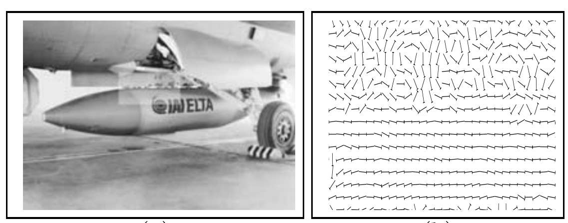

(a) (b)

Fig. 11. A real scene with folds and cuts. We highlight the region where the rocket obscures the hatch: the shading flow field clearly indicates the edge to be of the FOLD-CUT type, suggesting that the rocket is the "figure" side of the edge.

References

1. Breton, P. and Zucker, S.W.: Shadows and Shading Flow Fields. CVPR (1996) 782–789

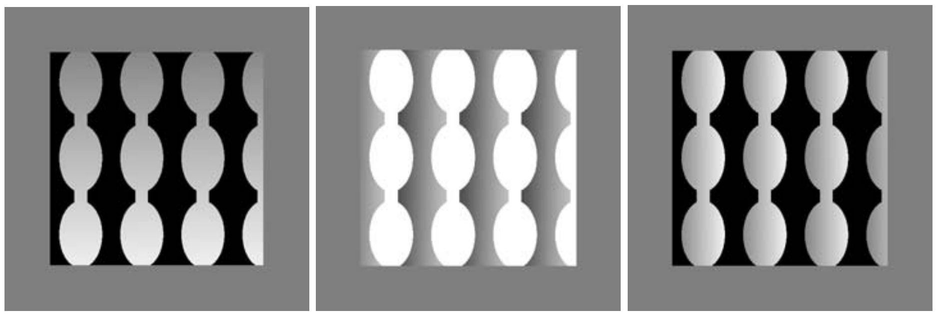

(a) (b) (c)

Fig. 12. Shaded versions of the ambiguous figure from Fig.1, after Kanizsa[8]. In (a) the light shading is transverse to the edges. In (b) the dark shading is approximately tangent to the to the edges. In (c) the light shading is tangent to the edges.

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. Kanizsa, G.: Organization in Vision. (1979) Praeger
9. Knill, D.C., Mamassian, P., Kersten D.: The geometry of shadows. J Opt Soc Am A **14** (1997)
10. Koenderink, J.J.: What does the occluding contour tell us about solid shape? Perception **13** (1976) 321–330
11. J.J. Koenderink and A.J. van Doorn: The singularities of the visual mapping. Biological Cybernetics **24** (1976) 51–59
12. Langer, M.S., Zucker, S.W.: Shape from shading on a cloudy day. J Opt Soc Am A **11** (1994) 467–478
13. Nalwa, V.S.: Line-drawing interpretation: a mathematical framework. International Journal of Computer Vision **2** (1988) 103–124
14. Parodi, P.: The complexity of understanding line-drawings of origami scenes, International Journal of Computer Vision **18** (1996) 139–170
15. Rieger, J.H.: The geometry of view space of opaque objects bounded by smooth surfaces. Artificial Intelligence **44** (1990) 1–40
16. 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
17. Weinshall, D., Werman, M.: On view likelihood and stability. IEEE-PAMI **19** (1997) 97–108