

On Image Contours of Projective Shapes

Jean Ponce^{1*} and Martial Hebert²

¹ Department of Computer Science
Ecole Normale Supérieure

² Robotics Institute
Carnegie-Mellon University

Abstract. This paper revisits classical properties of the outlines of solid shapes bounded by smooth surfaces, and shows that they can be established in a purely projective setting, without appealing to Euclidean measurements such as normals or curvatures. In particular, we give new synthetic proofs of Koenderink’s famous theorem on convexities and concavities of the image contour, and of the fact that the rim turns in the same direction as the viewpoint in the tangent plane at a convex point, and in the opposite direction at a hyperbolic point. This suggests that projective geometry should not be viewed merely as an analytical device for linearizing calculations (its main role in structure from motion), but as the proper framework for studying the relation between solid shape and its perspective projections. Unlike previous work in this area, the proposed approach does not require an oriented setting, nor does it rely on any choice of coordinate system or analytical considerations.

1 Introduction

Under perspective projection, the image (occluding) contour of a solid shape is the intersection of the retina with the boundary of a cone tangent to the shape’s surface, with apex at the pinhole. It is the projection of the rim curve where the cone and the surface meet tangentially.

What does the occluding contour tell us about solid shape? This is the question asked, and largely answered by Jan Koenderink in his landmark 1984 paper [12]. Specifically, contradicting an earlier claim by David Marr [18] stating that inflections of image contours do not (in general) convey any information about three-dimensional shape, Koenderink proved a remarkable result: The inflections of the image contour of a solid bounded by a smooth surface are the projections of parabolic points, where the Gaussian curvature of the surface vanishes. The convex points of the contour are projections of convex points of the surface, and its concave points are the images of saddle-shaped, hyperbolic points. The concave parts of the surface themselves never show up on the contour, for they are hidden from view by the solid itself.

Koenderink’s proof holds for both orthographic and perspective projection, and it is simple and elegant (see [1, 5, 6, 8, 22] for variants). It is also firmly anchored in Euclidean (differential) geometry, since it largely relies on concepts such as curve and surface normals and curvatures. This paper revisits Koenderink’s question in the more general setting of projective geometry, where Euclidean measurements such as distances,

* Willow project team. DI/ENS, ENS/CNRS/Inria UMR 8548.

angles, and curvatures are forbidden, but more primitive incidence, tangency and more generally contact properties are still available.

Lazebnik and Ponce addressed the same problem in [17], and showed that Koenderink's results are in fact valid in *oriented* projective geometry [21]. In this setting, points, lines, and planes are all oriented, which allows for deciding, for example, on which side of a plane a point lies, but requires somewhat awkward constructions such as maintaining two oriented copies of each point. We go one step further in this presentation, and prove that Koenderink's results holds in classical projective geometry, without the need for such constructions (Theorems 1 and 2). Likewise, we show that the rim turns in the same direction as the pinhole at a convex point, and in the opposite direction at a hyperbolic one (Theorem 3), a well known property in Euclidean geometry, extended to oriented projective geometry in [17].

While the paper does not introduce new theorems, it introduces new ways of manipulating concepts previously restricted to the Euclidean or oriented projective realms. Stated plainly, basic properties of the visual world that were known to be true in Euclidean or oriented projective settings are shown to be true in the much more general setting of plain projective geometry and, unlike most proofs in geometric computer vision, that require global or local coordinate systems and analytical parameterizations [1, 5, 6, 8, 12, 17, 22], ours are purely synthetic and do not require such an apparatus. Thus plain projective geometry is the appropriate framework for studying these properties.

We believe that the type of inquiry pursued in this paper is important because it identifies projective geometry as the natural setting for the qualitative study of the visual world: Koenderink's result (Theorems 1 and 2) tells us about the appearance of a solid shape in one image. Theorem 3 tells us about the rim moves with the viewpoint, and is therefore a first step toward understanding how the appearance of a solid changes in multiple pictures. We will come back to the latter point in Section 6.

The rest of this presentation is organized as follows: We recall basic facts about the local shape of curves and surfaces in classical (differential) projective geometry in Section 2. We prove in Section 3 that inflections of the contour are the projections of parabolic points (Theorem 1), which can be seen as a weak version of Koenderink's theorem. The strong version (Theorem 2) is proven in Section 4, and the rim motion is characterized by Theorem 3 in Section 5. We conclude by a brief discussion of future work in Section 6.

2 The Local Projective Shape of Curves and Surfaces

2.1 Basic Setting

We assume that objects of interest live in the real projective space \mathbb{P}^m , $m = 2$ or 3 , which can be defined as the quotient of $\mathbb{R}^m \setminus \{0\}$ by the equivalence relation identifying vectors that are nonzero multiples of each other [4]. In our context, it is perhaps better viewed as a manifold, locally affine (in fact the complement of any hyperplane in \mathbb{P}^m has an affine structure) but globally exempt of affine exceptions such as non-intersecting (parallel) coplanar lines or planes. Put more plainly, a projective plane (for example), looks locally just like an affine plane. On the other hand, one has to be mindful of the

fact that, for example, a line does not split a projective plane into two distinct components, nor does a plane split space into two components, and it is in general meaningless to talk about the “two sides” of such a line or plane. The whole arsenal of Euclidean measurements such as length or angle is also missing, along with the notions of normal, curvature, etc. Projective transformations (or collineations) are isomorphisms between projective spaces, and they preserve relations such as incidence or contact, as well as certain numerical invariants such as cross-ratios and homogeneous coordinates.

A perspective camera is defined in \mathbb{P}^3 by its pinhole o and retinal plane Π . It maps any point in $\mathbb{P}^3 \setminus \{o\}$ onto the point y of the retina Π where the line joining o to x intersects that plane [20]. We are interested in this paper in qualitative relations between objects of \mathbb{P}^3 and their perspective projections that are invariant under projective transformations.

The objects considered in this paper are solid, opaque, rigid bodies, bounded by smooth surfaces. These bodies will be called solids for short from now on. Certain visual features such as points, lines or curves drawn on these surfaces, are observed by cameras in the form of their projections onto some retinal plane. The images of the solids themselves, their shadows, so to speak, also form solid regions on the retinas, and they are in general bounded by piecewise-smooth curves.

Our setting excludes “hanging” threads for example, but so be it, and, as argued in [13], it will simplify our arguments. In particular, this will allow us to easily generalize the familiar notions of (local) convexity and concavity to the projective setting. More importantly, the definitions below are intuitive and similar to those used in Euclidean geometry. Yet, they do not rely in any way on the Euclidean machinery of curve and surface normals and curvatures.

2.2 Flatland

Let us start in Flatland—that is, restrict the world to a (projective) plane, equipped with its natural topology. A curve of the projective plane is said to be *smooth* when it admits a unique tangent at every point. Its inflections and cusps of both kinds can be defined as usual. A *piecewise-smooth* curve is continuous and smooth everywhere except at isolated points where it may admit multiple tangents.

From a topological viewpoint, there are two kinds of simple (no crossing) connected curves in \mathbb{P}^2 , one-sided curves (like straight lines), and two-sided ones, called *ovals*, that split the plane into two components, one, called the interior of the oval, homeomorphic to a disc, and the other, called its exterior, homeomorphic to a Möbius strip. We limit our attention in this section to *solid regions* ω of the plane, defined as the closure of the interior of an oval $\partial\omega$.

We shall say that a piecewise-smooth oval $\partial\omega$ is (locally) *convex* at some point x when there exists a line passing through x and (locally) contained in the (closure of the) exterior of ω , and that it is (locally) *concave* in x when there exists a line passing through that point and (locally) contained in the (closure of the) interior of ω . Convex and concave subsets of $\partial\omega$ form connected arcs of this curve, and they are separated by inflections and cusps of the second kind.

A subset ω of the affine plane is said to be convex when the segment joining any two of its points is itself a subset of ω . This definition is problematic in the projective

case because there are always two “antipodal” segments joining two points. This has led previous studies of local shape in computer vision [17] to rely on notions from oriented projective geometry [21]. Limiting our attention to curves bounding solid regions avoids this difficulty.

2.3 In the Round

Let us now move to the slightly more complicated world of surfaces bounding solid bodies defined “in the round”—that is, in the projective model we will adopt from now on for the three-dimensional world surrounding us.

The tangent plane π at a point x on a smooth surface σ (whether σ is orientable or not) is the plane that contains all the tangents to the curves drawn on σ and passing through x . The point x can be classified as elliptic, hyperbolic, or parabolic according to the way σ intersects its tangent plane π there (Figure 1): At an elliptic point, the intersection of σ and π consists of x itself, and σ locally lies on one side of π there (let us emphasize that this statement is only valid *locally*—that is, in some small neighborhood of x since, like a line in a projective plane, π does not split \mathbb{P}^3 into two distinct components). At a hyperbolic point, on the other hand, σ and π meet along two *asymptotic curves* intersecting in x , and the surface traverses its tangent plane. Finally, at a parabolic point, σ and π intersect transversally along a curve that cusps there. Elliptic and hyperbolic subsets of smooth surfaces form connected regions, and they are separated by parabolic curves. Note that all the notions defined in this paragraph are purely projective, and depend in no way on Euclidean devices since they are entirely based on incidence and tangency relations.




Fig. 1. The local shape of a surface. From left to right: an elliptic point, a hyperbolic point, and a parabolic point. (Reproduced with permission from the authors of [17].)

For simplicity, we define a solid as the closure of the interior of an oriented, connected and smooth surface, admitting a single tangent plane at each point. We shall say that the surface $\sigma = \partial\omega$ bounding some solid ω is (locally) *convex* at some elliptic point x when the tangent plane in x (locally) contained in the (closure of the) exterior of ω , and that it is (locally) *concave* in x when it is contained instead in the (closure of the) interior of ω .

As noted earlier, the image of ω is the intersection of the retinal plane II with the solid cone tangent to σ with the pinhole o as its apex. The image contour formed by its boundary is an oval with piecewise-smooth boundary having at most two tangents at any point.³ The intersection curve between the surface of the cone and σ is the rim curve γ . When x is a concave point, it cannot be visible because the line between o and x is tangent to σ , and thus locally contained in the interior of ω . We will use this fact in the proof of Theorem 2 in Section 4.

2.4 The Gauss Map

In Euclidean geometry, the Gauss map associates with every point on an oriented smooth curve γ drawn in a plane the point of the unit circle centered at the origin corresponding to the tip of its normal. Its folds correspond to inflections and cusps of the second kind. Likewise, the Gauss map associates with every point on the smooth oriented surface the point of the unit sphere centered at the origin corresponding to the tip of its normal. Interestingly, under orthographic projection, the Gaussian image of the rim is a great circle of the sphere, orthogonal to the projection direction. It is identical to the Gaussian image of the contour itself. This has been used in [10] to prove the weak version of Koenderink's theorem, namely that inflections of the contour are the projections of parabolic points, in the Euclidean case.

In projective geometry, the Gauss map associates instead with each point on a curve its tangent line (Figure 2[a,b]). Its image is the dual of the curve. This time, the cusps of the dual correspond to inflections and cusps of the second kind. Double points of the Gaussian image correspond to bitangents of the curve. The Gauss map also associates with every point on a surface its tangent plane (Figure 2[c,d]). Note that the definition of the projective Gauss map does not require the curve or surface of interest to be orientable. When it is, however, the image of the Gauss map, that is, the dual, inherits an orientation from the corresponding curve or surface, and we can talk *locally* of the orientation of the corresponding tangent line or surface. This will prove extremely important in the rest of our presentation. Note also that some of the properties of the Gauss map and its cusps established in the Euclidean setting [2] have been generalized to the projective one in [19]. In particular, cusps of the Gauss map correspond to the (tangential) intersections of the parabolic line with the flecnodal one (the locus of inflections of the asymptotic curves). Interestingly, the Gaussian image of the rim in this case is the intersection of the dual surface with the dual plane consisting of all the (primal) planes that contain the viewpoint. We will see in the next section that this can be used to prove the weak version of Koenderink's theorem in the projective case.

Note: The proofs of the classical results used in our presentation all assume that certain genericity conditions are satisfied [7, 19], that is (informally), that the surfaces considered are "sufficiently general", thus excluding quadrics, toruses, etc. We assume throughout the rest of this presentation that these genericity conditions are indeed satisfied.

³ Points with multiple tangents may occur for exceptional viewpoints.

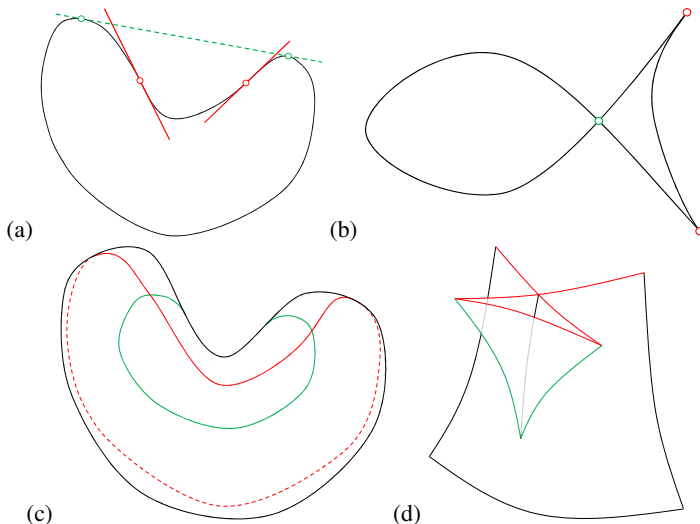


Fig. 2. Dual curves and surfaces: (a) a planar curve, with its inflections and bitangent; (b) its dual, with the cusps and double point associated with the inflections and bitangent; (c) a surface, with its parabolic line (in green) and a rim curve (in red); and (d) (part of) its dual, with the cuspidal edge associated with the parabolic line and the planar section corresponding to the rim. Best seen in color.

3 Koenderink's Theorem: Weak Version

We first characterize the type of rim points corresponding to inflections of the image contour:

Theorem 1. *Under perspective projection, the inflections of the image contour of a smooth surface are the images of parabolic points.*

The Euclidean version of this theorem was proven by Koenderink in [12]. We now show that it also holds under a pure projective setting. We do so by first clarifying the relationship between the Gaussian image of the image contour and that of the surface. We then use this result to connect Theorem 1 to a known (projective) property of the Gaussian images of curves of smooth surfaces.

Lemma 1. *Under perspective projection, the rim of a surface and its image contour have the same Gaussian image, which is a planar section of the Gaussian image of the surface itself by the dual plane associated with the viewpoint.*

Proof. By definition, the rim of a surface and the visual cone tangent to it with apex at the viewpoint share the same Gaussian image. In addition, the dual of a surface's rim is a planar section of its dual. More precisely, the Gaussian image of the rim is the intersection of the Gaussian image of the surface with the dual plane formed by all (primal) planes passing through the viewpoint.

The image contour is a planar section of the visual cone. Its Gaussian image is a curve in the projective plane (line field, or dual of plane π) $l(\pi)$ formed by the lines of the retinal plane II . Now for any π (and in particular for the retinal plane of course), the mapping φ_x that associates with each element of $l(\pi)$ the plane spanned by that line and some point x outside of π is a collineation between $l(\pi)$ and the dual plane x^* formed by all planes passing through x .

To prove that φ_x is a collineation, it is sufficient to show that if three points p_1, p_2, p_3 in $l(\pi)$ are aligned, then $\varphi_x(p_1), \varphi_x(p_2), \varphi_x(p_3)$ are aligned in x^* . Indeed, the lines l_1, l_2, l_3 in π corresponding to p_1, p_2, p_3 intersect at one point q since the points are aligned in $l(\pi)$. Therefore, the planes corresponding to $\varphi_x(p_1), \varphi_x(p_2), \varphi_x(p_3)$ must meet at a common line, namely the line containing x and q . The points in the dual corresponding to any three planes intersecting at a common line are aligned. Therefore, $\varphi_x(p_1), \varphi_x(p_2), \varphi_x(p_3)$ are aligned and φ_x is a collineation. When x is the viewpoint, φ_x induces a bijection between tangent lines to the image contour and tangent planes to the rim. Since this bijection is a collineation, the two Gaussian images can thus be identified with the corresponding planar section of the surface's dual. \square

Lemma 2 ([7, 19]). *The Gaussian image of the parabolic curve of a smooth surface consists of cuspidal edges with isolated swallowtail points.*

Theorem 1 immediately follows from Lemmas 1 and 2 by noting that a planar section of a cuspidal edge is in general a cusp (Figure 2[d]).

4 Koenderink's Theorem: Strong Version

The weak version of Koenderink's theorem does not require the observed surface σ to be oriented. Its strong version, stated below, does. In fact, we assume from now on that σ bounds an opaque solid, with a well defined interior and exterior. It enables us to characterize the concave and convex points of the image contours in addition to the inflection points already addressed in Theorem 1.

Notation. Let us define the perspective camera observing σ by its pinhole o and retina II . Together, the surface and the camera define a conical surface κ with apex in o , tangent to σ along the rim γ . We denote by x' the perspective projection of a rim point x with tangent plane τ onto II , and by γ' the image contour passing through x' with tangent τ' . The contour is obtained (equivalently) as the perspective projection of γ onto II or the intersection of κ with that plane.

Theorem 2. *Under perspective projection, the convexities of the image contour of an oriented smooth surface are the images of convex points, its concavities are the projections of hyperbolic points, and its inflections are the images of parabolic points.*

The following lemma, given without proof, summarizes well known properties of a conical surface κ [11]. These will prove to be a key to the proof of Theorem 2. These properties are true whether κ is a visual cone or not.

Lemma 3. *A conical surface κ is made up entirely of parabolic points. Contrary to the generic case depicted in Figure 1[right], its intersection with its tangent plane τ at*




Fig. 3. Koenderink’s theorem: (a) convex points project onto convex points of the image contour; (b) hyperbolic points project onto concave ones; (c) parabolic points project onto inflections.

some point x is a straight line joining its apex to x . A necessary and sufficient condition for κ to (locally) lie entirely on one side of τ at x is that there exists an open segment of a curve γ passing through x and (locally) lying entirely on that side of τ .

Proof (of Theorem 2). The conical surface κ , the rim γ , the tangent plane τ to σ in x , the contour γ' and its tangent τ' all inherit their orientation from σ (in particular, γ' is also an orientable, and in fact oriented, curve). At a convex point x (Figure 3[a]), τ is (locally) outside σ and therefore κ , and thus τ' is also outside γ' , which in turn implies that the contour is convex in x' . When x is a concave point, on the other hand, τ is locally inside σ , and thus inside the opaque solid it bounds, which in turn implies that x is hidden from the pinhole o .

Consider now a hyperbolic point x (Figure 3[b]). The orientation of σ in x induces as before an orientation of the tangent plane τ , which locally splits the surface into an inner part “below” τ —that is, on the same side as the interior of the solid, and an outer part “above” this plane. Any ray passing through a point below the plane τ in a sufficiently small neighborhood of x intersects the solid’s interior. Thus the rim γ must (locally) be drawn on the part of σ emerging above τ , i.e., there exists an open segment of γ containing x and entirely lying on the outer side of τ . According to Lemma 3, it follows that τ is (locally) inside κ . Likewise, τ' is locally inside κ and therefore inside γ' , which in turn implies that the contour is concave in x' . The parabolic case, already covered by Theorem 1, is illustrated by Figure 3(c). \square

Consider a hyperbolic point and a viewpoint moving in its tangent plane. From Theorem 2, the contour is concave everywhere it is visible—that is, within the sector that lies outside the intersection curve γ of the observed solid and its tangent plane (Figure 4). This sector is bounded by the two *asymptotic lines* tangent to the two branches of γ . As illustrated by Figure 4, the contour cusps when the viewpoint belongs to one of the asymptotic lines.




Fig. 4. The rim at a hyperbolic point. Left: Two generic visual rays (in red and green) and the corresponding rim curves. For a surface bounding a solid, one of the two rays and the corresponding rim would be (locally) in the interior of the object. Right: An asymptotic ray and the corresponding rim curve (in blue). The ray is tangent to the rim, giving birth to a cusp in the image. For a surface bounding an opaque solid, one branch of the rim and the corresponding branch of the contour at the cusp would be invisible.

5 Relative Rim and Viewpoint Motions in the Tangent Plane

Having established the relationship between the configuration of the local surface at a point on the rim and that of the local contour at its projection, we now turn our attention to the relationship between a change of viewpoint and the motion of the rim on the surface which it induces. Informally, given a pair of viewpoints in the tangent plane of a point which is on the rims corresponding to the two viewpoints, can we say something about the direction in which the rims locally turn around the point? We show in this section that the following theorem holds in the general projective case:

Theorem 3. *The rim turns in the same direction as the viewpoint in the tangent plane of a smooth oriented surface at a convex point, and in opposite ones at a hyperbolic point (Figure 5).*

In fact, this result can be shown in the Euclidean setting by showing explicitly that the Gauss map preserves orientation at elliptic points and reverses orientation at hyperbolic points [9]. The same result was shown in the *oriented* projective geometry setting by tracing oriented rays from the pinhole to the point [17]. Here, we go one step further by showing that the result holds without requiring the machinery of oriented projective geometry.

Of course, since we do not have any ready concept of signed angle or orientation at our disposal in the pure projective setting, what we mean by “turning in the same in the same direction or in the opposite direction” is not so obvious. Our first order of business is to introduce a clear definition. For that, we first need a notion of separability for pairs of lines in the same flat pencil (Figure 6).

Definition 1. *Given four distinct, coplanar and concurrent lines ξ_1 , δ_1 , ξ_2 and δ_2 , we say that (ξ_1, ξ_2) and (δ_1, δ_2) form separable pairs when both δ_1 and δ_2 lie in the interior*

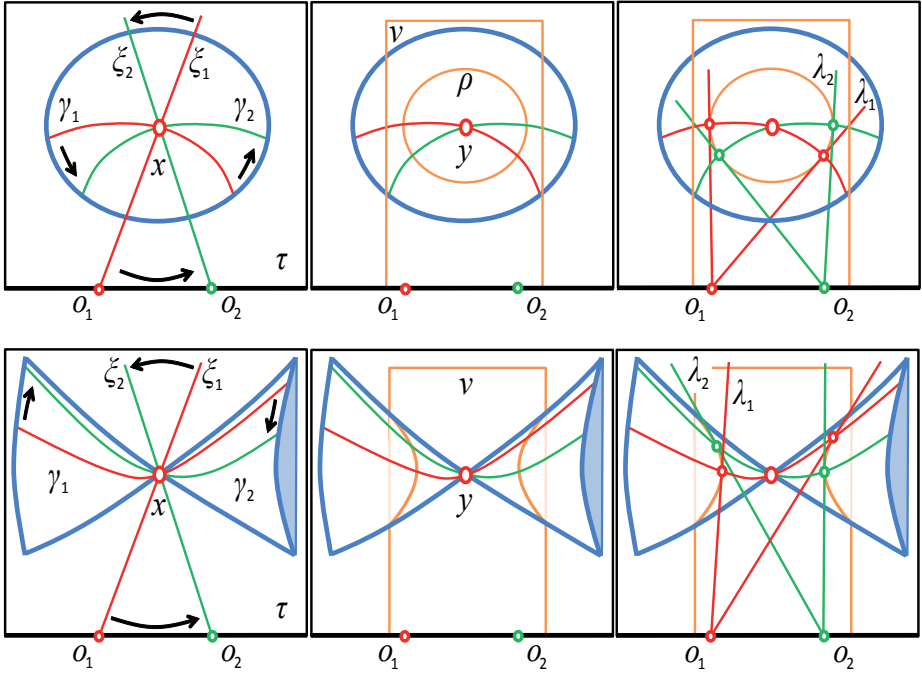


Fig. 5. Visual rays and rim curves turn in the same direction at an elliptic point (top) and in opposite directions at a hyperbolic one (bottom).

of one of the sectors delimited by ξ_1 and ξ_2 in their common plane or, equivalently, when both ξ_1 and ξ_2 lie in the interior of one of the sectors delimited by δ_1 and δ_2 .

When δ_1 and δ_2 (resp. ξ_1 and ξ_2) are fixed, a sufficient condition for separability is (for example) that ξ_2 (resp. δ_2) be close enough to ξ_1 (resp. δ_1) in the natural topology of the projective line formed by a pencil of lines. Whether the elements of two separable pairs of lines turn in the same direction or not can now be defined, again in terms of separability (Figure 6).

Definition 2. Given two separable pairs of coplanar and concurrent lines (ξ_1, ξ_2) and (δ_1, δ_2) , we say that δ_1 and δ_2 turn in the same direction as ξ_1 and ξ_2 when (ξ_1, δ_2) and (ξ_2, δ_1) are separable, and that they turn in opposite directions otherwise.

This is indeed a projective notion, and a strictly analogous line of reasoning leads by duality to a notion of relative order for two pairs of separable points on a line.

Proof (of Theorem 3). Let us consider a point x on the surface σ with tangent plane τ , two nearby pinholes o_1 and o_2 in τ , the corresponding visual rays ξ_1 and ξ_2 through x , and the corresponding rim curves γ_1 and γ_2 with tangents δ_1 and δ_2 in the neighborhood of x (Figure 5). We assume that o_1 is fixed, and that ξ_1 and δ_1 are distinct—that is, either x is convex, or it is hyperbolic but ξ_1 is not an asymptotic direction in x . We also assume




Fig. 6. Comparing the way separable pairs of lines in the same pencil turn in their plane. The lines ξ_1 and ξ_2 are separable from the lines δ_1 and δ_2 in each one of the four cases depicted. Top: The lines ξ_1 and δ_2 define two bidirectional sectors (here one is drawn in light blue, the other one is white). When ξ_2 and δ_1 lie in the same sector, (ξ_1, δ_2) and (ξ_2, δ_1) are separable, and δ_1 and δ_2 turn in the same direction as ξ_1 and ξ_2 . Bottom: When (ξ_1, δ_2) and (ξ_2, δ_1) are not separable, (ξ_1, δ_1) and (ξ_2, δ_2) are separable instead, and the pairs (δ_1, δ_2) and (ξ_1, ξ_2) turn in opposite directions.




Fig. 7. Left: A convex point. Nearby rays λ_1 and λ_2 and the lines μ_1 and μ_2 joining y to the points where they graze the curve ρ . Right: A hyperbolic point. We have also highlighted the inner side of the lines λ_1 and λ_2 with hash marks.

that the second pinhole o_2 is fixed, but sufficiently close to o_1 for the two pairs of lines (ξ_1, ξ_2) and (δ_1, δ_2) to be separable.

Let us now consider a plane ν in the pencil of planes passing through o_1 and o_2 , close enough to τ to intersect the rim. When x is a convex point, ν is “below” (on the inner side of) τ and locally intersects σ along a non-empty, closed and (everywhere locally) convex curve ρ (Figure 5, top). The curve ρ inherits its orientation from that of σ . Given some point y in its interior, this point approaches x as ν approaches τ , and the ray η_1 passing through o_1 and y approaches ξ_1 . Let us further assume from now on, without loss of generality, that ν is close enough to τ that all “inside” and “outside” relations below are meaningful (they indeed are, *locally*, in general). In particular, among the two rays passing through the first viewpoint o_1 and tangent to ρ , let us choose λ_1 to be the one with η_1 on its inner side.

Let us also denote by μ_1 the line joining y to t_1 , and define similarly η_2, t_2, λ_2 and μ_2 for the second viewpoint o_2 . For ν close enough to τ , the pairs (η_1, η_2) and (μ_1, μ_2) are separable since the lines η_1 and η_2 approach the rays ξ_1 and ξ_2 , and μ_1 and μ_2 approach the rim tangents δ_1 and δ_2 . Finally, let u_1 (resp. u_2) denote the intersection of λ_1 and μ_2 (resp. λ_2 and μ_1), and let v_1 denote the intersection of λ_1 and η_2 . We can (again) always pick ν close enough to τ that the points t_1, u_1, t_2, u_2 , and v_1 are all in a small open neighborhood of y where it is meaningful to say that one of these points is on the inner or outer side of λ_1 or λ_2 . In particular, t_1 (resp. t_2) is by construction on the inner side of λ_2 (resp. λ_1), thus u_1 (resp. u_2) is on the outer side of λ_2 (resp. λ_1).

Let us show that the two pairs of points (o_1, u_1) and (v_1, t_1) on λ_1 are separable. By duality, this will show in turn that the pairs (η_1, μ_2) and (η_2, μ_1) are separable. Let us define $[v_1 t_1]$ as the intersection of the line λ_1 with the sector formed by η_2 and μ_1 which does *not* contain o_1 . Since t_1 and v_1 are on the inner side of λ_2 (and for ν close

enough to τ), the entire segment $[v_1 t_1]$ is also on the inner side of λ_2 . Since u_1 is by construction outside λ_2 , it is thus also outside $[v_1 t_1]$. It follows that the two pairs of points (o_1, u_1) and (v_1, t_1) are separable, and thus the lines (η_1, μ_2) and (η_2, μ_1) are separable too. Since the lines η_1 and η_2 approach the rays ξ_1 and ξ_2 , and the lines μ_1 and μ_2 approach the rim tangents δ_1 and δ_2 as ν approaches τ , it finally follows that the rim turns in the same direction as the viewpoint in x .

The same reasoning can be applied at a hyperbolic point, except that the plane ν intersecting the rim must lie “above” (on the outer side of) τ (Theorem 2), and locally intersects the surface along two curves (Figure 5, bottom). These curves inherit the orientation of σ , and we can choose some point y lying outside both these curves and approaching x as ν approaches τ . When ν is close enough to τ , we can always choose a ray λ_1 tangent to one of these curves such that y lies on the outer side of λ_1 (Figure 7[right]). A similar line of reasoning one shows that (v_1, u_2) can be separated from (o_2, t_2) , and thus (ξ_1, μ_1) can be separated from (ξ_2, μ_2) . Since μ_1 and μ_2 converge to δ_1 and δ_2 as ν approaches τ , the rim and the viewpoint turn in opposite directions in x . \square

6 Discussion

We have shown that two fundamental questions on the relationship between a three-dimensional shape and its image contour could be answered in a purely projective setting without resorting to analytical tools or oriented projective geometry. Specifically, we have shown how the inflections, convex, and concave parts of an image contour can be related to local properties of the surface at the corresponding points with a new proof of Koenderink’s theorem, and we have shown how the motion of the rim relates to that of the viewpoint.

These results are first steps toward a complete reformulation of several classical models of shape and appearance in purely projective terms. In particular, the oriented projective version of Theorem 3 is used in [17] as a building block for an algorithm for computing the *rim mesh* [15] defined on the surface of a solid by the rim curves associated with multiple cameras. It would be interesting to see whether the purely projective version of the theorem can be used in a similar fashion. This would have practical implications in the context of visual hulls as well [3, 16]. The geometry of the Euclidean Gauss map has also been shown to dictate (in part) the evolution of the image contour and its topology under changes in viewpoint (see [10, Ch. 20] for an overview). Using the projective Gauss map in a similar way would be very interesting. Results in this direction [19] showing that the cusps of the Gauss map correspond to the (tangential) intersections of the parabolic line with the flecnodal one, in the projective case as well as in the Euclidean one, are encouraging, and may lead to a purely projective formulation of aspect graphs [14] for example.

More generally, we have introduced new concepts and tools such as separability in order to generalize concepts such as ordering or “in betweenness” normally defined in a Euclidean or oriented projective setting to a general projective setting. We believe that they are crucial in formulating many geometric properties of solid shapes and their projections in the projective realm.

Acknowledgments. This work was supported in part by the ERC grant VideoWorld, the Institut Universitaire de France, and ONR MURI N000141010934.

References

1. E. Arbogast and R. Mohr. 3D structure inference from image sequences. *Journal of Pattern Recognition and Artificial Intelligence*, 5(5), 1991.
2. T. Banchoff, T. Gaffrey, and C. McCrory. *Cusps of Gauss mappings*. Pitman, 1982.
3. B.G. Baumgart. Geometric modeling for computer vision. Technical Report AIM-249, Stanford University, 1974. Ph.D. Thesis. Department of Computer Science.
4. M. Berger. *Geometry*. Springer-Verlag, 1987.
5. E. Boyer and M.-O. Berger. 3D surface reconstruction using occluding contours. *IJCV*, 22(3):219–233, 1997.
6. J.M. Brady, J. Ponce, A. Yuille, and H. Asada. Describing surfaces. In H. Hanafusa and H. Inoue, editors, *Proceedings of the 2nd International Symposium on Robotics Research*, pages 5–16. MIT Press, 1985.
7. J.W. Bruce. The duals of generic hypersurfaces. *Math. Scand.*, 48:36–60, 1981.
8. R. Cipolla and A. Blake. Surface shape from the deformation of the apparent contour. *IJCV*, 9(2):83–112, 1992.
9. M.P. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, Englewood Cliffs, NJ, 1976.
10. D.A. Forsyth and J. Ponce. *Computer Vision: A Modern Approach*. Prentice-Hall, 2003.
11. D. Hilbert and S. Cohn-Vossen. *Geometry and the Imagination*. Chelsea, New York, 1952.
12. J.J. Koenderink. What does the occluding contour tell us about solid shape? *Perception*, 13:321–330, 1984.
13. J.J. Koenderink. *Solid Shape*. MIT Press, Cambridge, MA, 1990.
14. J.J. Koenderink and A.J. Van Doorn. The internal representation of solid shape with respect to vision. *Biological Cybernetics*, 32:211–216, 1979.
15. S. Lazebnik, E. Boyer, and J. Ponce. On computing exact visual hulls of solids bounded by smooth surfaces. In *CVPR*, pages 156–161, 2001.
16. S. Lazebnik, Y. Furukawa, and J. Ponce. Projective visual hulls. *IJCV*, 74(2):137–165, 2007.
17. S. Lazebnik and J. Ponce. The local projective shape of smooth surfaces and their outlines. *IJCV*, 63(1):65–83, 2005.
18. D. Marr. Analysis of occluding contour. *Proc. Royal Society, London*, B-197:441–475, 1977.
19. C. McCrory and T. Shifrin. Cusps of the projective Gauss map. *Journal of Differential Geometry*, 19, 1984.
20. J. Ponce. What is a camera? In *CVPR*, 2009.
21. J. Stolfi. *Oriented Projective Geometry: A Framework for Geometric Computations*. Academic Press, 1991.
22. R. Vaillant and O.D. Faugeras. Using extremal boundaries for 3D object modeling. *PAMI*, 14(2):157–173, February 1992.