

The Local Projective Shape of Smooth Surfaces and their Outlines

Svetlana Lazebnik
slazebni@uiuc.edu

Jean Ponce
ponce@cs.uiuc.edu

Beckman Institute, University of Illinois
Urbana, Illinois, 61801, USA

Abstract

This paper examines projectively invariant local properties of smooth curves and surfaces. Oriented projective differential geometry is proposed as a theoretical framework for establishing such invariants and describing the local shape of surfaces and their outlines. This framework is applied to two problems: a projective proof of Koenderink's famous characterization of convexities, concavities, and inflections of apparent contours; and the determination of the relative orientation of rim tangents at frontier points.

1 Introduction

Projective differential geometry [12] deals with local properties of smooth curves and surfaces that remain invariant under projective transformations of the ambient space. Previous applications of projective differential geometry to computer vision were concerned with *quantitative* invariants of plane and surface curves [16, 17, 18, 20]. The practical usefulness of such invariants has been limited by the numerical difficulty of computing high-order image derivatives. Our goal here is more modest: we are interested in *qualitative* local invariants of curves, surfaces, and their outlines that require only low-order derivatives of their parametrizations. For this purpose, it is sufficient that only the *sign* of the functions involved remain invariant. A fundamental qualitative invariant is the characterization of the *local shape* of a smooth surface as elliptic, hyperbolic, or parabolic. As shown in Section 2.3, even though familiar constructs from Euclidean differential geometry such as surface normal and Gaussian curvature are missing in the projective setting, the local shape of a surface is determined by the sign of a simple expression in the derivatives of order $k \leq 2$. This is in stark contrast with quantitative invariants such as the projective curvature of a plane curve, which involves derivatives of up to seventh order [5].

The rest of the paper is organized as follows. Section 2 gives a brief treatment of local projective shape, culminating in the definition of a symmetric bilinear form that is the projective precursor of the *second fundamental form* of Euclidean differential geometry. Section 3 transplants projective differential notions into the setting of *oriented projective geometry* [14, 19], and introduces oriented invariants corresponding to notions of convexity and concavity

for surfaces and plane curves. Section 4 describes oriented projective relationships between smooth surfaces and their outlines. Section 5.1 contains a major contribution of this paper: a projective proof of Koenderink's characterization of convexities, concavities, and inflections of the apparent contours of solids bounded by smooth surfaces [8]. Section 5.2 invokes the (oriented) epipolar geometry between pairs of cameras and discusses two results concerning the relative orientation of rim tangents at frontier points and the determination of this orientation from image information.

2 Projective Differential Geometry

2.1 Basic Definitions

In this section, we recall elementary notions of ordinary projective geometry that will be used throughout the paper. A *projective space* \mathbb{P}^n is formed by identifying all nonzero vectors of the space \mathbb{R}^{n+1} that are scalar multiples of each other (in this paper, we are interested only in \mathbb{P}^2 and \mathbb{P}^3). The basic geometric objects in projective space are *flats*. A d -dimensional flat is the subspace of \mathbb{P}^n associated with some $(d+1)$ -dimensional subspace of \mathbb{R}^{n+1} . The flats of \mathbb{P}^3 are points, lines, planes, and the *universe*. A d -dimensional flat is uniquely defined by a set of $d+1$ independent points (a *proper simplex*): a line is defined by two distinct points, a plane is defined by three non-collinear points, etc. A *projective transformation* is a mapping between two projective spaces induced by some linear transformation between the corresponding vector spaces.

A basic operation on two disjoint flats X and Y is the *join*. The result of this operation, denoted $X \vee Y$, is the flat spanned by the simplex formed by concatenating two simplices that span X and Y . For example, the join of two points is a line; the join of a point and a line is a plane (or the universe in \mathbb{P}^2); the join of two skew lines in \mathbb{P}^3 is the universe. Another important operation is the *meet* $X \wedge Y$, which intuitively corresponds to intersection of two flats. We will make only cursory use of meet in this paper.

To represent flats analytically, we must first select a *coordinate system* in the projective space of interest. This subject is beyond the scope of our paper (for a good reference, see e.g. [2]). It is sufficient to note that once we have selected a fixed (but arbitrary) coordinate system, we can represent points in \mathbb{P}^n by $(n+1)$ -dimensional *homo-*

geneous coordinate vectors defined up to non-zero scale. From now on, we will write $X, Y, X \vee Y$ to denote both the corresponding flats and their coordinate vectors. Note that general flats are typically represented using *Plücker coordinates*, but the details of this machinery are not necessary for understanding our notation. We will use lowercase (resp. uppercase) letters to denote flats in \mathbb{P}^2 (resp. \mathbb{P}^3).

Throughout the paper, we will often want to know whether a simplex formed by three points x, y, z in 2D or by four points X, Y, Z, W in 3D is proper (spans the universe). To decide this, we compute the determinants $|x, y, z|$ and $|X, Y, Z, W|$ of the respective 3×3 and 4×4 matrices whose columns are given by homogeneous coordinate vectors of the points. The simplices are proper if and only if the respective determinants are nonzero.

2.2 Curves and Surfaces

Curves. A curve in \mathbb{P}^3 is defined in the neighborhood of one of its points by a smooth mapping $s \mapsto X(s) = [X_1(s), \dots, X_4(s)]^T$ from an open interval of \mathbb{R} into \mathbb{P}^3 . A point $X(s)$ on the curve is said to be *regular* when the derivative $X'(s)$ is not identically zero and not a scalar multiple of $X(s)$. At a regular point X of a curve Γ , the *osculating flat* of order k is the flat spanned by $X, X', X'', \dots, X^{(k)}$. In the projective setting, the successive derivatives $X^{(i)}$ are *points* and not *directions* as in the Euclidean case. The order-1 osculating flat is the *tangent line* $X \vee X'$, defined as the limit of a line through X and another point on Γ that approaches X . Consider what happens when X is multiplied by a scalar function $\mu(s)$, a transformation that does not change the image of the curve in \mathbb{P}^3 . Since $(\mu X)' = \mu' X + \mu X'$, it follows that we can scale X to place the derivative point $(\mu X)'$ anywhere on the tangent to Γ at X , though it cannot coincide with X unless Γ is degenerate. In short, even though the derivative point X' is not invariant under algebraic transformations such as homogeneous rescaling, the tangent line $X \vee X'$ is. The order-2 osculating flat $X \vee X' \vee X''$, or *osculating plane*, is the limit of the plane through X and two points that independently approach X . In particular, an *inflection point* has a degenerate osculating plane, which means that the point X'' is in the span of X and X' .

Surfaces. A surface in \mathbb{P}^3 is defined in the neighborhood of one of its points by a smooth mapping $(u, v) \mapsto X(u, v) = [X_1(u, v), \dots, X_4(u, v)]^T$ from an open set of \mathbb{R}^2 into \mathbb{P}^3 , where each $X_i(u, v)$ is a smooth coordinate function. The *tangent plane* at a point of a surface is the plane containing the tangent lines of all curves on the surface through the point. Note that this definition makes no reference to the concept of a normal to the surface, which does not exist in projective space. We can easily verify that the tangent plane is $X \vee X_u \vee X_v$, where subscripts denote partial differentiation with respect to u and v . Let Γ be a curve on the

surface Σ passing through the point X . The equation of Γ is $X(s) = X(u(s), v(s))$. The tangent to Γ at X is spanned by X and $X' = u'X_u + v'X_v$. Thus, we can see that the tangent to Γ at X lies in the subspace spanned by X, X_u , and X_v (Fig. 1). We will use the word *direction* to refer to any point on the line spanned by X_u and X_v .




Figure 1: The tangent plane of a surface.

Order of Contact. An important invariant in projective differential geometry is the *order of contact* between curves and surfaces. Two curves have contact of order $\geq k$ at a common point X (informally speaking, they intersect at $k+1$ “consecutive” points) if they can be parametrized at X so that the first k derivatives match. For example, a tangent has order of contact ≥ 1 with a curve, and an *inflectional* tangent has order of contact ≥ 2 with the curve. A curve Γ and a surface Σ have contact of order $\geq k$ at a common point X if there exists a curve $\Gamma' \subset \Sigma$ passing through X such that Γ and Γ' have contact of order $\geq k$. For example, a tangent line has order ≥ 1 contact with the surface, and a curve has order ≥ 2 contact with its osculating plane.

2.3 The Local Shape of Surfaces

Projective Invariants. A geometric property is a *projective invariant* when it continues to hold under projective transformations. To make sure that a scalar function of a surface parametrization $X(u, v)$ and its derivatives is a projective invariant, we must show that it is invariant under any transformations that change the analytic form of the representation of the surface but not its geometry. These transformations are (1) *rescaling* of homogeneous coordinates by a nonzero scalar function of the parameters; (2) *reparametrization*; and (3) *projective transformation* (multiplication by a nonsingular 4×4 matrix).

Asymptotic Tangents and Conjugate Directions. An *asymptotic tangent* is a tangent line having at least order 2 contact with a surface. The following classical result [12] concerning asymptotic tangents is of fundamental importance to the rest of our presentation.

Proposition 1. *If $T = X \vee (\alpha X_u + \beta X_v)$ is an asymptotic tangent to Σ at X , then*

$$l\alpha^2 + 2m\alpha\beta + n\beta^2 = 0, \quad (1)$$

where $l = |X, X_u, X_v, X_{uu}|$, $m = |X, X_u, X_v, X_{uv}|$, and $n = |X, X_u, X_v, X_{vv}|$.

Proof (sketch). Asymptotic curves are the integral curves of the asymptotic tangent field on the surface, and they are characterized by the fact that their osculating plane coincides with the tangent plane [12]. Let the asymptotic curve be defined by $X(u(s), v(s))$. Then its osculating plane is given by $X \vee X' \vee X''$. The points X and $X' = u'X_u + v'X_v$ that span the asymptotic tangent at X are clearly already contained in the tangent plane $X \vee X_u \vee X_v$. Thus, to express the geometric fact that the osculating plane coincides with the tangent plane we need only write $|X, X_u, X_v, X''| = 0$. Expanding X'' and rewriting this determinant yields the expression $l\alpha^2 + 2m\alpha\beta + n\beta^2 = 0$, where $\alpha = u'(s)$ and $\beta = v'(s)$. \square

Let us denote by $\text{II}(U_1, U_2)$ the symmetric bilinear form that associates with every pair of tangent directions $U_1 = \alpha_1 X_u + \beta_1 X_v$ and $U_2 = \alpha_2 X_u + \beta_2 X_v$ the quantity

$$\begin{pmatrix} \alpha_1 & \beta_1 \end{pmatrix} \begin{pmatrix} l & m \\ m & n \end{pmatrix} \begin{pmatrix} \alpha_2 \\ \beta_2 \end{pmatrix}.$$

Our choice of notation is justified by the fact that the matrix $\begin{pmatrix} l & m \\ m & n \end{pmatrix}$ defines the *second fundamental form* II in the Euclidean case where the coordinates of X have the form $[X_1, X_2, X_3, 1]^T$. We will say that two tangent directions U_1 and U_2 are *conjugate* when $\text{II}(U_1, U_2) = 0$. If $U = \alpha X_u + \beta X_v$ is an asymptotic direction, clearly (1) can be written as $\text{II}(U, U) = 0$, recalling the well-known characterization of asymptotic directions as *self-conjugate*.

The discriminant of II is $-4(ln - m^2)$, and it is easy to confirm that a surface point will admit two distinct asymptotic tangents when $ln - m^2 < 0$, a double one when $ln - m^2 = 0$, and none at all when $ln - m^2 > 0$. Since the existence of asymptotic tangents is a geometric property of the surface, and order of contact is preserved under projective transformations, it is not surprising that the sign of the discriminant is a projective invariant.

Proposition 2. *The sign of $ln - m^2$ is a projective invariant. The local shape of a surface at some point is (Fig. 2):*

Elliptic: $ln - m^2 > 0$, 0 asymptotic tangents.
 Parabolic: $ln - m^2 = 0$, 1 asymptotic tangent.
 Hyperbolic: $ln - m^2 < 0$, 2 asymptotic tangents.




Figure 2: The local shape of a surface at a point.

The proof involves routine computations and is not shown here. For the rest of the paper, we will omit many proofs for the sake of brevity, but most of the omitted proofs may be found in [15].

3 Oriented Framework

3.1 Basic Definitions

Following the framework of Stolfi [19], we briefly introduce *oriented* analogues of projective notions listed in Section 2. An *oriented projective space* \mathbb{T}^n is formed by identifying all nonzero vectors of \mathbb{R}^{n+1} that are *positive* multiples of each other. Thus, homogeneous coordinate vectors of points in oriented space are defined up to positive scale.

Two proper simplices spanning the same flat are said to have the same orientation when they are related by an *orientation-preserving* projective transformation represented by a matrix with a positive determinant. All the simplices of a given flat form two classes under this equivalence relation. An *oriented flat* is a flat to which an orientation has been assigned by choosing as “positive” one of the two classes of simplices that span it. There are two oriented flats of dimension n , the *positive* and the *negative universe*.

To form an oriented simplex spanning the join of two oriented flats X and Y , we concatenate the simplices spanning the respective flats *in that order*. The *relative orientation* of two disjoint flats X and Y such that $\dim(X) + \dim(Y) + 1 = n$ is positive (resp. negative) if $X \vee Y$ spans the positive (resp. negative) universe. The two cases will be denoted by $X \vee Y > 0$ and $X \vee Y < 0$, respectively. In \mathbb{T}^3 , relative orientation is defined for pairs of lines and for points and planes. We will say that a point X lies on the positive (resp. negative) side of a plane Π when $\Pi \vee X > 0$ (resp. $\Pi \vee X < 0$). In \mathbb{T}^2 , relative orientation is defined for points and lines, and we will say that a point x lies on the positive (resp. negative) side of a line l when $l \vee x > 0$ (resp. $l \vee x < 0$). A proper simplex (x, y, z) in \mathbb{T}^2 is *positive* (i.e., spans the positive universe) if and only if $|x, y, z| > 0$. Similarly, the determinant test can be used in 3D to find out whether a 4-simplex is positive.

3.2 Orienting Curves and Surfaces

In this section, we will show how to orient curves in \mathbb{T}^2 and surfaces in \mathbb{T}^3 , and characterize their invariants.

Oriented Projective Invariants. A geometric property is an oriented projective invariant when it continues to hold under orientation-preserving projective transformations. A scalar function of a curve parametrization in \mathbb{T}^2 or surface parametrization in \mathbb{T}^3 and its derivatives is an oriented projective invariant when it is invariant under (1) *positive rescaling*, i.e., multiplication by a positive scalar function of the parameters, (2) *orientation-preserving reparametrization*, i.e., a change of parameters with positive Jacobian, and (3) *orientation-preserving projective transformation*.

Orienting Curves. A curve γ in \mathbb{T}^2 is locally defined by a smooth mapping $s \mapsto x(s) = [x_1(s), x_2(s), x_3(s)]^T$ and is naturally oriented in the increasing direction of s . The oriented tangent $x(s) \vee x'(s)$ is the limit of the line joining $x(s)$ to the point $x(s + \delta s)$. If a point x of γ is not an

inflection, then γ locally lies either on the positive or on the negative side of the oriented tangent line (Fig. 3). In the former case, we can say that γ is locally *convex*, and in the latter case, it is locally *concave*. When the curve γ is the boundary of a solid region ω of the plane, it is possible to orient γ it so that ω is on the *positive* side of the tangent at every point of the curve. Intuitively, this orientation can be thought of as the direction in which we must traverse γ so as to always see ω to our left (we conventionally designate the left side of the line as positive).




Figure 3: γ is on the positive side of the tangent line (left), on the negative side (middle), crossing (right).

Consider the determinant $|x, x', x''|$, a simple quantity that can be seen as a precursor of Euclidean plane curvature. Geometrically, the sign of $|x, x', x''|$ tells us about the relative orientation of the tangent $x \vee x'$ and the second derivative point x'' at a point x of an oriented curve γ . The next proposition states that this sign can be used to determine the local shape of γ at x .

Proposition 3. *The sign of $|x, x', x''|$ is an oriented projective invariant. The point x is convex (resp. concave, inflection) when $|x, x', x''|$ is positive (resp. negative, zero).*

Orienting Surfaces. The parametrization of a smooth surface in \mathbb{T}^3 induces a natural orientation of its tangent plane $\Pi = X \vee X_u \vee X_v$. Giving the surface a well-defined orientation makes the sign of Π an invariant.

Proposition 4. *Let U be a tangent direction at the point X of Σ . The sign of $\Pi(U, U)$ is an oriented projective invariant. If X is elliptic, $\Pi(U, U)$ has the same sign for all U . If X is hyperbolic and U_1, U_2 are conjugate directions, then $\Pi(U_1, U_1)$ and $\Pi(U_2, U_2)$ have opposite signs.*

When Σ is the boundary $\partial\Omega$ of some solid Ω , we can select its orientation so that Ω is always on the positive side of the tangent plane. Because the notion of an inward-pointing normal is not available in projective differential geometry, the precise statement of what it means for a solid to lie on the positive side of the tangent plane is a bit involved, and is omitted here. Intuitively, we want to orient Σ so that any line stabbing the surface at X penetrates Ω along some interval beginning at X and lying on the positive side of Π . As shown by the next proposition, the sign of Π characterizes tangents that are locally inside or outside the solid.

Proposition 5. *Consider a surface $\Sigma = \partial\Omega$ oriented so that Ω is on the positive side of its tangent plane, and let $U = \alpha X_u + \beta X_v$ be a tangent direction in X . Then the tangent $X \vee U$ is locally outside Ω when $\Pi(U, U) > 0$ and locally inside it when $\Pi(U, U) < 0$.*

We say that a point on the surface of a solid is *convex* (resp. *concave*) when the solid is locally on the positive (resp. negative) side of its tangent plane. Combining Propositions 4 and 5, we obtain a characterization of the oriented local shape of surfaces bounding solids.

Proposition 6. *An elliptic point on the surface a solid oriented so this solid lies on the positive side of its tangent plane is convex when $\Pi(U, U) > 0$ for some tangent direction U , and concave otherwise.*

4 Surfaces and their Outlines

4.1 Camera Model

Let us model the 3D world and the image plane by the oriented spaces \mathbb{T}^3 and \mathbb{T}^2 , respectively. A nonsingular 3×4 camera projection matrix P describes a transformation from the scene to the image plane. If X is a 3D point and x is its projection, then we write $x \simeq PX$, where \simeq denotes equality up to positive scale. The rows P_1, P_2, P_3 of the matrix P can be interpreted as coefficients of three projection planes [6]. The null space of P is the point O such that $PO = 0$. In the standard projective camera model, O determines, up to an arbitrary scale factor, the camera center. It is possible to assign a unique sign to O by defining it as the oriented meet of the three projection planes: $O \simeq P_1 \wedge P_2 \wedge P_3$ (see [6, 14, 19]).

Next, we introduce a useful lemma stating that the orientation of the 3-simplex formed in the image plane by the projections of three 3D points matches the orientation of the 4-simplex formed by the camera center and the three points.

Lemma 1. *Let X, Y , and Z be points in \mathbb{T}^3 with images $x \simeq PX, y \simeq PY$, and $z \simeq PZ$. Then $|x, y, z| \simeq |O, X, Y, Z|$.*

4.2 Rims and Apparent Contours

The rim or contour generator of a surface Σ associated with a camera center O is the curve Γ formed by the points X on Σ whose tangent plane contains O , or $|X, X_u, X_v, O| = 0$. The perspective projection of Γ is a curve γ in \mathbb{T}^2 called the *apparent contour*, or *outline* of Σ (Fig. 5). It well known that the tangent to the rim and the viewing ray $O \vee X$ are in conjugate directions [9].

Visibility. When the surface Σ is the boundary of an opaque solid Ω , a rim point X will be hidden from view if the ray $L = O \vee X$ enters the object Ω prior to grazing it at X . In general, visibility is not a local phenomenon: the infinitesimal properties of Σ in the neighborhood of X do not tell us whether the viewing ray L has already passed through the object elsewhere. However, there is one *necessary* local condition for visibility: L must be locally outside Ω . Note that this condition is never satisfied for concave points (recall Proposition 5)—hence the well-known fact that concavities never show up on the silhouette of an opaque ob-

ject. We will call a point X on the rim *locally visible* if the viewing ray $O \vee X$ is locally outside Ω .

Orienting the Rim. Let us assume that Σ is oriented so that Ω is everywhere on the positive side of the tangent plane. The contour tangent t at a point $x \simeq PX$ is related to the rim tangent $T = X \vee X'$ by the camera projection: $t = x \vee x' \simeq (PX) \vee (PX')$. Clearly, all points (except O) lying on the tangent plane $\Pi = X \vee X_u \vee X_v$ project onto t in the image. We want the orientation of t to be *consistent* with the orientation of Π in the following way: if Y is a point such that $\Pi \vee Y > 0$, then $y \simeq PY$ must satisfy $t \vee y > 0$. This can be written as $|X, X_u, X_v, Y| \simeq |x, x', y|$, and by Lemma 1 we have $|x, x', y| \simeq |O, X, X', Y|$. Thus, we must have $X \vee X_u \vee X_v \simeq O \vee X \vee X'$, and we must orient the rim tangent T to satisfy

$$O \vee T \simeq \Pi. \quad (2)$$




Figure 4: Orienting the rim. Left: the intrinsic orientation of the surface tangent plane, $X \vee X_u \vee X_v$ (the orientation is indicated with a counterclockwise arrow). Right: orienting the rim tangent such that $O \vee X \vee X'$ matches the intrinsic orientation of Π .

From now on, we will assume that the rim is oriented according to (2) (see also Fig. 4), and that the rim induces the orientation of the apparent contour.

5 Applications

5.1 Koenderink's Theorem

The following theorem was proven by Koenderink [8] using Euclidean concepts such as the curvature of plane curves and the Gaussian curvature of surfaces.

Proposition 7. *A convex (resp. concave, inflection) point on the apparent contour of a smooth solid is the projection of a convex (resp. hyperbolic, parabolic) point on the rim of its surface (Fig. 5).*

Our proof, presented below, is based on the oriented projective camera model described in Section 4.1 and, like Koenderink's original proof, is general enough to encompass both orthographic and perspective projection. More importantly, however, our proof does not rely on the metric structure of the ambient space, and shows that Koenderink's theorem is purely projective, in the sense that it can be established using only (oriented) projective invariants of curves and surfaces.

Proof. Let X denote the point on Σ , $U = \alpha_1 X_u + \beta_1 X_v$ the direction such that $O \vee X \simeq X \vee U$, $X' = \alpha_2 X_u + \beta_2 X_v$ the




Figure 5: A smooth solid and its perspective projection. The dashed curve on the surface is the rim and the dotted curve is the *parabolic curve* or the locus of parabolic points. The rim points X , Y , and Z are respectively convex, hyperbolic, and parabolic, and their images x , y , and z are respectively convex, concave, and inflection points of the apparent contour.

derivative point of the properly oriented rim tangent, and x the projection of X . Invoking Lemma 1, we get

$$\begin{aligned} |x, x', x''| &\simeq |O, X, X', X''| \simeq |X, U, X', X''| \\ &= (\alpha_1 \beta_2 - \beta_1 \alpha_2) |X, X_u, X_v, X''| \\ &= (\alpha_1 \beta_2 - \beta_1 \alpha_2) \Pi(X', X''). \end{aligned}$$

Note that $O \vee X \vee X' = X \vee U \vee X' = (\alpha_1 \beta_2 - \alpha_2 \beta_1) (X \vee X_u \vee X_v)$. Since the rim tangent is oriented to satisfy (2), we must have $\alpha_1 \beta_2 - \alpha_2 \beta_1 > 0$, and it follows that the signs of $|x, x', x''|$ and $\Pi(X', X')$ are the same. For X to be visible, $\Pi(U, U)$ must be positive. When $|x, x', x''|$ is positive, $\Pi(X', X')$ must be positive as well, and since the viewing ray and the rim tangent are conjugate, X must be elliptic according to Proposition 4, and therefore convex according to Proposition 6. By the same token, X must be hyperbolic when $|x, x', x''|$ is negative. Finally, the rim tangent must be an asymptotic direction when $|x, x', x''| = 0$, and thus self-conjugate. But since the viewing ray is conjugate to the rim tangent as well, this means that any direction is conjugate to X' , and X must be parabolic. \square

5.2 Ordering Rim Tangents

This section focuses on surfaces observed by two cameras and gives an image-based characterization of the relative orientation of the tangents to two rims at *frontier points* where they intersect. This result is used in the image-based algorithm for computing visual hulls presented in [15].

Relative Orientation of Tangents. Consider two tangents $T_1 \simeq X \vee U_1$ and $T_2 \simeq X \vee U_2$, with $U_1 = \alpha_1 X_u + \beta_1 X_v$ and $U_2 = \alpha_2 X_u + \beta_2 X_v$. We say that the relative orientation of T_1 and T_2 (also, of U_1 and U_2) is *positive* when $X \vee U_1 \vee U_2 \simeq X \vee X_u \vee X_v$. As shown by the following two lemmas, it is a simple matter to characterize relative orientation analytically.

Lemma 2. *The relative orientation of T_1 and T_2 is positive if and only if $\alpha_1 \beta_2 - \alpha_2 \beta_1 > 0$.*

Lemma 3. *Let $U_1 = \alpha_1 X_u + \beta_1 X_v$ and $U_2 = \alpha_2 X_u + \beta_2 X_v$ be two tangent directions such that*

$$\begin{pmatrix} \alpha_2 \\ \beta_2 \end{pmatrix} \simeq S \begin{pmatrix} \alpha_1 \\ \beta_1 \end{pmatrix}, \text{ where } S = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} l & m \\ m & n \end{pmatrix}.$$

Then U_1 and U_2 are conjugate. Their relative orientation is positive if and only if $\Pi(U_1, U_1) > 0$ (that is, the tangent $X \vee U_1$ is locally outside the surface).

Let us take a point X on the surface, and place two cameras O_1 and O_2 in the tangent plane Π at X such that X is locally visible to both cameras. Let Γ_1 and Γ_2 denote the two corresponding rims. The following proposition relates the relative orientation of the rim tangents to that of the viewing rays $O_1 \vee X$ and $O_2 \vee X$.

Proposition 8. *The relative orientation of the rim tangents $X \vee X'_1$ and $X \vee X'_2$ is the same as (resp. opposite of) the relative orientation of the viewing rays $O_1 \vee X$ and $O_2 \vee X$ when X is a convex (resp. hyperbolic) point (Fig. 6).*




Figure 6: Relative orientation of rims and camera centers for a convex point (left), and a hyperbolic point (right).

To visualize the above statement, imagine putting a tiny flashlight on the tangent plane and shining it toward X . As we rotate the flashlight while keeping it aimed at X , the boundary of the shadow will rotate in the same direction as the light beam if X is elliptic and in the opposite direction if X is hyperbolic (Fig. 6). In the Euclidean setting, this qualitative result follows from the fact that the Gauss map is orientation-preserving at elliptic points and orientation-reversing at hyperbolic points [3]. Below we give a proof that uses only oriented projective differential geometry.

Proof. Define $\alpha_1, \alpha'_1, \beta_1, \beta'_1$ and $\alpha_2, \alpha'_2, \beta_2, \beta'_2$ such that

$$\begin{cases} O_1 \vee X \simeq X \vee (\alpha_1 X_u + \beta_1 X_v), \\ X \vee X'_1 \simeq X \vee (\alpha'_1 X_u + \beta'_1 X_v), \\ O_2 \vee X \simeq X \vee (\alpha_2 X_u + \beta_2 X_v), \\ X \vee X'_2 \simeq X \vee (\alpha'_2 X_u + \beta'_2 X_v). \end{cases}$$

According to Lemma 2, the relative orientation of the viewing rays $O_1 \vee X$ and $O_2 \vee X$ is determined by the sign of $\alpha_1 \beta_2 - \alpha_2 \beta_1$. Likewise, the relative orientation of $X \vee X'_1$ and $X \vee X'_2$ is given by the sign of $\alpha'_1 \beta'_2 - \alpha'_2 \beta'_1$.

We can use the *conjugate mapping* S to express the derivative points as functions of the viewing directions: $\begin{pmatrix} \alpha'_1 \\ \beta'_1 \end{pmatrix} \simeq S \begin{pmatrix} \alpha_1 \\ \beta_1 \end{pmatrix}$ and $\begin{pmatrix} \alpha'_2 \\ \beta'_2 \end{pmatrix} \simeq S \begin{pmatrix} \alpha_2 \\ \beta_2 \end{pmatrix}$. Note that S obeys the

orientation convention of Eq. (2): this is guaranteed by Lemma 3 and the fact that the $O_1 \vee X$ and $O_2 \vee X$ both lie locally outside the surface since X is locally visible by both cameras. Thus, S expresses the properly oriented relationship between $\alpha_1 \beta_2 - \alpha_2 \beta_1$ and $\alpha'_1 \beta'_2 - \alpha'_2 \beta'_1$:

$$\begin{vmatrix} \alpha'_1 & \alpha'_2 \\ \beta'_1 & \beta'_2 \end{vmatrix} \simeq |S| \begin{vmatrix} \alpha_1 & \alpha_2 \\ \beta_1 & \beta_2 \end{vmatrix} = (ln - m^2) (\alpha_1 \beta_2 - \alpha_2 \beta_1).$$

Whenever X is convex (resp. hyperbolic), we have $ln - m^2 > 0$ (resp. < 0) and the signs of $\alpha_1 \beta_2 - \alpha_2 \beta_1$ and $\alpha'_1 \beta'_2 - \alpha'_2 \beta'_1$ are the same (resp. opposite). \square

Frontier Points in the Image. Since O_2 belongs to the plane $O_1 \vee T_1 = O_1 \vee X \vee X'_1$, we have $|O_1, X, X'_1, O_2| = 0$. In the first image, let $x_1 \simeq P_1 X$ be the projection of X , and $x'_1 \simeq P_1 X'_1$ is the derivative point of the outline γ_1 at x_1 . By Lemma 1, we have $|P_1 X, P_1 X'_1, P_1 O_2| \simeq |x_1, x'_1, e_{12}| = 0$, where $e_{12} \simeq P_1 O_2$ is the *epipole* in the first view. The expression $|x_1, x'_1, e_{12}| = 0$ can be rewritten as $t_1 \vee e_{12} = 0$, where $t_1 = x_1 \vee x'_1$ is the tangent to γ_1 at x_1 . The contour point x_1 is distinguished by the property that its tangent line t_1 passes through the epipole e_{12} . Equivalently, we could say that the derivative point x'_1 lies on the epipolar line $l_{12} \simeq e_{12} \vee x_1$. In the second image, an analogous relationship holds: $|x_2, x'_2, e_{21}| = 0$. In short, x_1 and x_2 are points of *epipolar tangency*, where the epipolar lines l_{12} and l_{21} have order 1 contact with γ_1 and γ_2 , respectively (Fig. 7). We will refer to x_1 and x_2 as *frontier points*.




Figure 7: Frontier points.

Next, we show that the relative orientation of the rim tangents can be decided *from image information alone*.

Proposition 9. *Given a frontier point x_1 in the first image, the tangent t_1 to the apparent contour, and the epipolar line l_{12} , the relative orientation of the tangents to the rims Γ_1 and Γ_2 is positive if and only if (a) x_1 is convex and the lines t_1 and l_{12} have the same orientations (i.e. $t_1 \simeq l_{12}$), or (b) x_1 is concave and $t_1 \simeq -l_{12}$ (Fig. 8).*

Proof. The epipolar line l_{12} is the oriented projection of the viewing ray $O_2 \vee X$, and the tangent t_1 is the projection of the rim tangent $X \vee X'_1$. It follows from Proposition 1 that l_{12} and t_1 have the same orientation exactly when $O_1 \vee O_2 \vee X \simeq O_1 \vee X \vee X'_1$. But $O_1 \vee O_2 \vee X = X \vee O_1 \vee O_2$, and




Figure 8: Illustration of Proposition 9. The dashed epipolar lines are oriented from the epipole e_{12} toward the frontier points 1 – 4, and the orientations of the epipolar tangents are indicated by the arrows. Case 1: x_1 is convex and $l_{12} \simeq t_1$. Case 2: x_1 is concave and $l_{12} \simeq t_1$. Case 3: x_1 is concave and $l_{12} \simeq -t_1$. Case 4: x_1 is convex and $l_{12} \simeq -t_1$. The relative orientation of Γ_1 and Γ_2 is positive in cases 1 and 3, and negative in cases 2 and 4.

$O_1 \vee X \vee X'_1 \simeq X \vee X_u \vee X_v$. Therefore, the orientation of l_{12} and t_1 is the same if and only if the relative orientation of the viewing rays $O_1 \vee X$ and $O_2 \vee X$ is positive. Since we know from Proposition 7 that the preimage X is convex when x_1 is convex, and hyperbolic otherwise, Proposition 9 follows immediately from Lemma 8. \square

For completeness, let us consider the role of the second image in the proposition. Since the local shape of X does not depend on the viewpoint and since both rays $O_1 \vee X$ and $O_2 \vee X$ are locally outside the surface by assumption, x_1 and x_2 are either both convex or concave. However, the relative orientation of X'_1 and X'_2 is the opposite of the relative orientation of X'_2 and X'_1 . Therefore, whenever $l_{12} \simeq t_1$ in the first image, we must have $l_{21} \simeq -t_2$ in the second.

6 Discussion

In the past, most applications of projective differential geometry in computer vision have dealt with high-order quantitative differential invariants in the context of object recognition. By contrast, the oriented projective differential framework presented in this paper is aimed at deriving *low-order qualitative* invariants suitable for reconstruction tasks. Projective differential geometry is the natural framework for the problem of reconstructing smooth curves and surfaces based on projective information alone. The past decade has seen intense study of projective reconstruction techniques for points, lines, and planes [6, 7]. In our opinion, the addition of tools applicable to the reconstruction of more complex geometric entities should greatly enrich the subject of multi-view geometry.

The theoretical framework of this paper can be used for reasoning about geometric constructs such as visual hulls [13], visibility complexes [4], and aspect graphs [11]. These constructs, whose combinatorial structure is determined by local and multilocal events corresponding to special kinds of contact of lines and surfaces, can be seen as qualitative projective invariants of a scene. Projective differential geometry may also be an appropriate setting for identifying the class of projective transformations that leave certain other geometric structures unchanged, such as the shadow

field [1], or even perhaps the set of reconstructions compatible with a fixating stereo system with unknown vergence angles [10].

Acknowledgments. This research was partially funded by the UIUC Campus Research Board and by the National Science Foundation grants IRI-990709 and IIS-0308087.

References

- [1] P. Belhumeur, D. Kriegman, and A. Yuille, “The Bas-Relief Ambiguity,” in *Proc. CVPR*, 1997, pp. 1060–1066.
- [2] M. Berger, *Geometry*, Springer-Verlag, 1987.
- [3] M. do Carmo, *Differential Geometry of Curves and Surfaces*, Prentice-Hall, Englewood Cliffs, NJ, 1976.
- [4] F. Durand, G. Drettakis, and C. Puech, “The 3D Visibility Complex,” *ACM Trans. on Graphics*, 21(2), 2002, pp. 176–206.
- [5] O. Faugeras, “Cartan’s Moving Frame Method and its Application to the Geometry and Evolution of Curves in the Euclidean, Affine and Projective Planes,” in [17], pp. 11–46.
- [6] O. Faugeras, Q.-T. Luong, and T. Papadopoulos, *The Geometry of Multiple Images*, MIT Press, Cambridge, MA, 2001.
- [7] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000.
- [8] J. Koenderink, “What Does the Occluding Contour Tell Us About Solid Shape?” *Perception* 13, 1984, pp. 321–330.
- [9] J. Koenderink, *Solid Shape*, MIT Press, Cambridge, MA, 1990.
- [10] J. Koenderink and A. Van Doorn, “Geometry of Binocular Vision and a Model for Stereopsis,” *Biological Cybernetics* 21, 1976, pp. 29–35.
- [11] J. Koenderink and A. Van Doorn, “The Internal Representation of Solid Shape with Respect to Vision,” *Biological Cybernetics* 32, 1979, pp. 211–216.
- [12] E. Lane, *Projective Differential Geometry of Curves and Surfaces*, The University of Chicago Press, 1932.
- [13] A. Laurentini, “The Visual Hull Concept for Silhouette-Based Image Understanding,” *IEEE Trans. PAMI*, 16(2), 1994, pp. 150–162.
- [14] S. Laveau and O. Faugeras, “Oriented Projective Geometry for Computer Vision,” in *Proc. ECCV*, 1996, pp. 147–156.
- [15] S. Lazebnik, *Projective visual hulls*, M.S. thesis, University of Illinois at Urbana-Champaign, 2002. Also Beckman CVR Technical Report CVR-TR-2002-01, available at <http://www-cvr.ai.uiuc.edu/ponce-grp/>.
- [16] J. Mundy and A. Zisserman, *Geometric Invariance in Computer Vision*, MIT Press, Cambridge, MA, 1992.
- [17] J. Mundy, A. Zisserman, and D. Forsyth, *Applications of Invariance in Computer Vision, Lecture Notes in Computer Science* 825, Springer-Verlag, 1994.
- [18] A. Salden, B. t. Romeny, and M. Vierverger, “Affine and Projective Differential Geometric Invariants of Space Curves,” in *Geometric Methods in Computer Vision II*, B. Vemuri, ed., SPIE, 1993, pp. 60–74.
- [19] J. Stolfi, *Oriented Projective Geometry: A Framework for Geometric Computations*, Academic Press, 1991.
- [20] I. Weiss, “Projective Invariants of Shapes,” in *Proc. CVPR*, 1988, pp. 291–297.