

The Local Projective Shape of Smooth Surfaces and Their Outlines

Svetlana Lazebnik (slazebni@uiuc.edu)

Jean Ponce (jponce@uiuc.edu)

*Department of Computer Science and Beckman Institute
University Of Illinois, Urbana, IL 61801, USA*

Abstract. This article examines projectively-invariant local geometric properties of smooth curves and surfaces. *Oriented projective differential geometry* is proposed as a general framework for establishing such invariants and characterizing the local projective shape of surfaces and their outlines. It is applied to two problems: (1) the projective generalization of Koenderink’s famous characterization of convexities, concavities, and inflections of the apparent contours of solids bounded by smooth surfaces, and (2) the image-based construction of *rim meshes*, which provide a combinatorial description of the arrangement induced on the surface of an object by the contour generators associated with multiple cameras observing it.

Keywords: Projective differential geometry, oriented projective geometry, differential invariants, local shape, frontier points.

1. Introduction

As firmly established in the multi-view geometry literature (Faugeras et al., 2001; Hartley and Zisserman, 2000), an essential part of the relationship between “flat” geometric elements—that is, planes, lines, and points—and their perspective images is intrinsically projective, and can be modeled without appealing to the affine and/or Euclidean structure of the ambient space. The present article extends this line of reasoning to smooth curves, surfaces, and their outlines. A natural theoretical framework for this study is offered by *projective differential geometry* (Lane, 1932; Bol, 1950), which is concerned with local properties of smooth curves and surfaces that remain invariant under projective transformations of the ambient space. These properties are typically characterized in terms of scalar functions of (local) curve and surface parametrizations and their derivatives. Our presentation focuses on *qualitative* local invariants, or functions whose

signs, as opposed to actual *values*, are preserved under projective transformations and reparametrizations. The local elliptic, hyperbolic, or parabolic shape of a smooth surface at one of its points is such an invariant. In Euclidean geometry, it is characterized by the sign of the determinant of the second fundamental form, or equivalently, by the sign of the Gaussian curvature. The notion of Gaussian curvature is meaningless in projective geometry; yet, as will be shown in Section 2, the local shape of a surface can still be determined by the sign of a simple expression in the derivatives of order $k \leq 2$ of its parametrization.

Despite its elegance and simplicity, ordinary projective geometry (whether differential or not) lacks a concept of *orientation*, and thus fails to capture such familiar and useful notions as the left and right sides of an oriented line in the plane, or the distinction between the convex and concave elliptic points of an oriented smooth surface. Fortunately, these can be regained at very little additional expense in the *oriented projective geometry* framework recently popularized by Stolfi (1991). This article extends Stolfi's original work, which was limited to the study of *flats* (points, lines, planes, etc.), to encompass smooth curves and surfaces. Specifically, we propose *oriented projective differential geometry* as a language for reasoning about the qualitative invariant relationships between smooth surfaces and their outlines. The study of such relationships was addressed by Koenderink (1984) twenty years ago when he showed that, in *Euclidean spaces*, 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.¹ Koenderink's proof is remarkably simple and elegant, but it deeply relies on Euclidean concepts such as the Gaussian curvature. As shown in Section 4, the oriented projective differential geometry framework reveals the *intrinsically projective* nature of Koenderink's theorem, and it provides a totally different (but once again quite simple) proof that never appeals to any affine or Euclidean concept.

¹ The *rim* (or *contour generator*) is the surface curve that projects onto the apparent contour. See Section 4.2 for more details.

More generally, we contend that many combinatorial structures of interest in computer vision—for example, aspect graphs (Koenderink and Van Doorn, 1979), visual hulls (Baumgart, 1974; Laurentini, 1994), and visibility complexes (Durand et al., 2002)—are oriented projective objects, that are best studied in the oriented projective differential geometry framework. This is illustrated in Section 5 with an image-based algorithm for computing the *rim mesh* of a solid observed by multiple cameras, i.e., a combinatorial description of the arrangement induced on its surface by the corresponding contour generators.

1.1. PREVIOUS WORK

So far, the applications of projective differential geometry to computer vision have been limited to the construction of *quantitative* invariants of plane and surface curves and their image projections. This includes the seminal paper by Weiss (1988) that caused much of the initial interest for invariants in the early 1990s (Mundy and Zisserman, 1992; Mundy et al., 1994), as well as (among others) the work on semi-differential invariants by Moons *et al.* (1995), and the characterization of plane curves in terms of their projective curvature and its evolution as a function of scale and/or projective arc length by Salden *et al.* (1993), Faugeras (1994), and Calabi *et al.* (1998). This line of research can, in principle, be extended to the study of surfaces and their invariants using Cartan’s general *moving frame method* (Cartan, 1992). Unfortunately, the practical usefulness of quantitative invariants is limited by the numerical difficulty of computing high-order image derivatives: For example, the definition of the projective curvature of a plane curve involves derivatives of up to seventh order (Faugeras, 1994). This is our motivation for pursuing instead *qualitative* invariants in this article.

Oriented projective geometry was introduced in computer vision by Laveau and Faugeras (1996). A few researchers have used this framework to derive constraints on valid projective reconstructions from multiple images (Hartley, 1998; Werner et al., 1998; Werner and Pajdla, 2001a; Werner and Pajdla, 2001b;

Chum et al., 2003). Our work differs from these in two key respects: (1) We focus on properties of smooth curves, surfaces and their outlines, instead of flats and their projections; and (2) we do not rely on the “cheirality” assumption (Hartley, 1998) that the true underlying structure of the scene and the cameras is *affine*, so that all points in the scene are in front of the plane at infinity, and all cameras only observe points lying in front of their focal planes.

Parts of the research presented in this article have appeared in (Lazebnik, 2002; Lazebnik and Ponce, 2003). Rim meshes were originally introduced in (Lazebnik et al., 2001) in a setting where all cameras are strongly calibrated, with a Euclidean construction for ordering the rim arcs incident to each vertex and assembling the faces of the mesh. This article presents the first algorithm for computing the rim mesh in the much more difficult setting where the ambient space only has an oriented projective structure, and the cameras’ projection matrices are only known up to an oriented projective ambiguity.

1.2. OBJECTIVES, MAIN CONTRIBUTIONS, AND OUTLINE OF THE PRESENTATION

The main objective of this article is to introduce an oriented version of projective differential geometry as a language for stating *and* proving qualitative local invariant properties of curves, surfaces, and their outlines. Its key contributions are:

- Establishing in Section 3 the existence of *oriented* projective differential invariants of curves and surfaces (Propositions 2 and 5).
- Characterizing in Section 4 the oriented relationship between the points on the rim of a surface and their projections onto the apparent contour, and presenting a new, elementary proof of Koenderink’s theorem (Proposition 9) that reveals its intrinsically projective nature.
- Applying in Section 5 the proposed mathematical language to the task of constructing *rim meshes*. This includes theoretical results on the relative

orientation of rim tangents and its image-based characterization (Propositions 12 and 13), as well as an implemented algorithm taking advantage of these results.

The rest of this article is organized as follows: We start in Section 2 with a brief tutorial introduction to projective differential geometry. Our motivation for including this material is that it is little known to many computer vision researchers, and difficult to find in accessible literature. Our presentation is informal, and—unlike classical treatments of the subject (Lane, 1932)—it emphasizes the intuitive connection between projective and Euclidean differential geometries. In particular, we introduce projective precursors of the second fundamental form and Gaussian curvature in order to extend to the projective setting the familiar Euclidean characterization of the local shape of a surface in terms of elliptic, hyperbolic, and parabolic points (Proposition 1).

Next, Section 3 presents an *oriented* extension of projective differential geometry that, although fairly straightforward, has not (to the best of our knowledge) appeared before in the computer vision literature. This allows us to introduce two qualitative invariants missing from ordinary projective geometry—namely, the local convexity/concavity of points on a plane curve (Proposition 2), and the local convexity/concavity of elliptic points on a surface (Proposition 5).

Section 4 completes the development of our theoretical framework by describing an oriented camera model related to that of Laveau and Faugeras (1996). We present a novel result (Proposition 7) relating the orientation of the tetrahedron formed by the camera center and three scene points to that of the corresponding image triangle. The section concludes with a purely projective proof of Koenrindk’s theorem (Proposition 9). When seen as an oriented projective statement, this theorem elegantly ties together the qualitative invariants of curves and surfaces defined in the previous two sections.

Finally, Section 5 presents a sample application of the proposed framework: We exploit the epipolar constraints associated with a pair of cameras to charac-

terize the relative orientation of the corresponding rims at the *frontier points* where they intersect on the surface (Proposition 12), and to determine this orientation from image information (Proposition 13). This result is used in an implemented algorithm for computing rim meshes from image contours. Other potential applications are listed in Section 6.

Disclaimer. This is *not* a mathematics article, and a mathematician looking for novel theoretical developments will be disappointed by our presentation: Most of our results intuitively make sense, and they are easily shown to hold in *Euclidean* spaces using the proof techniques of classical differential geometry. However, this article is *not* intended to uncover new mathematics. Its main goal is to draw attention to the fact that certain important properties of the visual world are *intrinsically projective*, and to introduce the elementary machinery that is both necessary to precisely state these properties, and sufficient to prove them. We believe that oriented projective differential geometry provides such a machinery, and hope to demonstrate with the elementary results proven in this article that it offers a useful alternative language for computer vision practitioners to reason about the outlines of solids bounded by smooth surfaces, irrespective of their possible embedding in a Euclidean space.

In the interest of brevity, our presentation omits the proofs of all classical results (Propositions 1, 6, and 8), as well as others that are straightforward exercises in algebraic manipulation (Propositions 2–5 and 10–11). We refer the interested reader to (Lazebnik, 2002) for details.

2. Projective Differential Geometry

This section informally presents elementary notions of projective differential geometry that will be used throughout the article. The results found in Section 2.3 date back to the beginning of the 20th century (Bol, 1950; Lane, 1932), but a coherent and accessible presentation is not to be found in recent literature. It

is this lack of suitable introductory material, and not the dearth of applications, that accounts for the relatively obscure status of projective differential geometry in the field of computer vision (Koenderink, 1990). We attempt here to fill this gap by presenting an elementary self-contained introduction to projective differential geometry for vision researchers.

2.1. BASIC CONCEPTS

The *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. A d -dimensional *flat* in \mathbb{P}^n is the subspace of \mathbb{P}^n associated with some $(d + 1)$ -dimensional subspace of \mathbb{R}^{n+1} . Note that \mathbb{P}^n contains a unique n -dimensional flat, known as the *universe*. A d -dimensional flat is *spanned* by $d + 1$ *independent* points — that is, points associated with $d + 1$ independent vectors in \mathbb{R}^{n+1} — defining a *proper simplex*. A *projective transformation* is the mapping between two projective spaces induced by some bijective linear transformation between the corresponding vector spaces.

In this presentation, the projective spaces of interest are \mathbb{P}^2 and \mathbb{P}^3 , and the flats of interest are points, lines, and planes. A line is spanned by two distinct points, and a plane is spanned by three non-collinear points. Note that the only plane in \mathbb{P}^2 is the universe. The universe of \mathbb{P}^3 is a three-dimensional flat spanned by any non-coplanar quadruple of points.

Two operators—the *join* and the *meet*—play a fundamental role in projective geometry. The join $F \vee G$ of two disjoint flats F and G is the flat defined by the simplex formed by concatenating two simplices that span F and G . 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); and the join of two skew lines in \mathbb{P}^3 is the universe. The meet $F \wedge G$ of two flats F and G intuitively corresponds to the intersection of these flats. We will only make cursory use of the meet in this presentation.

To represent flats analytically, it is necessary to select a *coordinate system* for \mathbb{P}^n (see Berger [1987] for details), thereby associating each point with a unique

(up to scale) *homogeneous coordinate vector* in \mathbb{R}^{n+1} . More generally, any d -dimensional flat can be associated with a homogeneous vector of $\binom{n+1}{d+1}$ *Plücker coordinates* (note that the Plücker coordinate of the universe is simply a non-zero scalar). We will assume from now on that some fixed but arbitrary coordinate system has been chosen, and will identify flats with their coordinate vectors. Lowercase letters will denote flats in \mathbb{P}^2 , while uppercase letters will denote flats in \mathbb{P}^3 . Likewise, we will identify the geometric join and meet operators with their algebraic counterparts. The latter can be described using simple bilinear formulas (Stolfi, 1991) whose precise form is generally not important for the rest of our presentation. We will, however, explicitly use the fact that the join of $n + 1$ points X_1, \dots, X_{n+1} in \mathbb{P}^n is the determinant $|X_1, \dots, X_{n+1}|$ of their homogeneous coordinate vectors: In particular, the join of three points x, y, z in \mathbb{P}^2 is the 3×3 determinant $|x, y, z|$, and the join of four points X, Y, Z, W in \mathbb{P}^3 is the 4×4 determinant $|X, Y, Z, W|$. These expressions have a nonzero value if and only if the triangle (resp. tetrahedron) in question is *proper*—that is, spans the universe of \mathbb{P}^2 (resp. \mathbb{P}^3).

2.2. CURVES AND SURFACES IN PROJECTIVE SPACE

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 called *regular* when the derivative $X'(s)$ is not identically zero nor 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 $X, X', \dots, X^{(k)}$. Note that in homogeneous coordinates, the successive derivatives $X^{(i)}$ represent points in \mathbb{P}^3 (by contrast, when a non-homogeneous coordinate representation is used, as in the Euclidean case, derivatives are thought of as vectors instead of points).

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




Figure 1. A tangent line of a curve and a tangent plane of a surface. Any linear combination of the partial derivatives X_u and X_v defines a point X' —or *tangent direction*—on the line joining these two points and the corresponding tangent $X \vee X'$.

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 singular. Thus, it is important to realize that 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 . Regular points with a degenerate osculating plane (that is, points for which X'' is contained in the tangent line $X \vee X'$) are known as *inflections*.

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 is a smooth coordinate function. Let Γ be some curve defined on the surface Σ by $X(s) = X(u(s), v(s))$. The derivative in X is $X' = u'X_u + v'X_v$, and it follows that the tangent to Γ lies in the plane $\Pi = X \vee X_u \vee X_v$ (Figure 1). Since Γ and the derivatives u' and v' are arbitrary, Π is geometrically defined as the plane spanned by the tangent lines to all surface curves passing through X , and is known as the *tangent plane* to Σ at X . Although the derivatives X_u and X_v depend on the chosen parametrization, it is easy to show that the tangent plane is invariant under reparametrizations and rescalings.

2.3. LOCAL SHAPE

Projective differential geometry is a study of *projective invariants* of curves and surfaces. A geometric property is said to be a projective invariant when it continues to hold under projective transformations. We also say that a scalar function of a curve or surface parametrization and its derivatives is a projective invariant when it is invariant under any transformation that changes the analytical representation of the curve or surface but not its geometry—that is, (1) a *rescaling* of homogeneous coordinates by a nonzero scalar function of the parameters; (2) a *reparametrization*; or (3) a *projective transformation*.

The most important projective invariant of a smooth surface in the neighborhood of one of its points is its *local shape*, leading, as in the Euclidean setting, to the partition of the surface into elliptic, hyperbolic, and parabolic points (Figure 2). Our goal in the remainder of this section is to arrive at an analytic characterization of local shape (Proposition 1). To this end, let us first introduce the *local shape matrix* \mathcal{S} defined at a surface point X by

$$\mathcal{S} = \begin{pmatrix} l & m \\ m & n \end{pmatrix}, \quad \text{where} \quad \begin{cases} l &= |X, X_u, X_v, X_{uu}| \\ m &= |X, X_u, X_v, X_{uv}| \\ n &= |X, X_u, X_v, X_{vv}| \end{cases}.$$

A *tangent direction* is a point on the line spanned by the first-order partial derivatives X_u and X_v (Figure 1). Given two tangent directions $U_1 = \alpha_1 X_u + \beta_1 X_v$ and $U_2 = \alpha_2 X_u + \beta_2 X_v$, the symmetric bilinear form associated with \mathcal{S} is

$$\text{II}(U_1, U_2) = U_1^T \mathcal{S} U_2 = l\alpha_1\alpha_2 + m(\alpha_1\beta_2 + \alpha_2\beta_1) + n\beta_1\beta_2. \quad (1)$$

Our choice of notation is justified by the fact that $\text{II}(U_1, U_2)$ is the (negated) *second fundamental form* in the Euclidean case where the coordinates of X have the form $[X_1, X_2, X_3, 1]^T$. We say that two tangent directions U_1 and U_2 (resp. tangent lines $X \vee U_1$ and $X \vee U_2$) are *conjugate* when $\text{II}(U_1, U_2) = 0$. Although this is not obvious from our purely algebraic definition of conjugacy, it is a projectively invariant relationship that can be defined in purely geometric terms using

the language of *conjugate nets*: Namely, a net² of curves on a surface is called *conjugate* if the tangents of the curves of one family of the net constructed at the points of each fixed curve of the other family form a developable surface (Lane, 1932). Then the tangents to the two curves passing through any point on the surface are said to be conjugate. Conversely, these geometric properties can be used to define conjugate tangents and prove that conjugate directions U_1 and U_2 verify $\text{II}(U_1, U_2) = 0$ (Lazebnik, 2002).

In particular, *self-conjugate* tangent directions $U = \alpha X_u + \beta X_v$ satisfying $\text{II}(U, U) = l\alpha^2 + 2m\alpha\beta + n\beta^2 = 0$ are called *asymptotic directions*. Geometrically, they are characterized by the fact that the corresponding tangent $X \vee U$ has (at least) second-order contact with the surface: Informally, a line and a surface have contact of order (at least) k at a common point if they intersect at $k + 1$ “consecutive” points. For example, a stabbing line has contact of order zero with the surface; any ordinary (generic) tangent has contact of order one; and a surface point may (in general) have zero, one, or two asymptotic tangents with contact of order two or higher. Order of contact is a projective invariant, and so is the number of asymptotic tangents at a given point. As shown by the following proposition, the sign of the determinant $K = ln - m^2$ of the shape matrix provides an analytic characterization of the latter.

Proposition 1. *The sign of K is a projective invariant. The local shape of a surface at some point is (Figure 2):*

elliptic: *when $K > 0$, with 0 asymptotic tangent;*

parabolic: *when $K = 0$, with 1 asymptotic tangent;*

hyperbolic: *when $K < 0$, with 2 asymptotic tangents.*

In summary, the projectively invariant *geometric* characterization of local shape is given by the number of asymptotic tangents to the surface at the

² Two one-parameter families of curves on the surface are said to form a *net* if exactly one curve of each family passes through each point of the surface, and the tangents of the two curves passing through the same point are always distinct.




Figure 2. The local shape of a surface. Elliptic point: the surface does not cross its tangent plane, no asymptotic tangent. Hyperbolic point: the surface intersects the tangent plane along two curves whose tangents are the two asymptotic directions. Parabolic point: the surface intersects the tangent plane along a curve that cusps in the single asymptotic direction.

point, while the equivalent *analytic* characterization is in terms of the sign of the determinant K of the local shape matrix, which may be thought of as a projective precursor of Gaussian curvature.

3. Oriented Projective Differential Geometry

Next, we present an extension of projective differential geometry in the *oriented* setting due to Stolfi (1991). In particular, we derive in Section 3.2 qualitative invariants corresponding to local convexity/concavity of points on a plane curve and local convexity/concavity of elliptic points on a surface.

3.1. BASIC CONCEPTS

First, let us briefly review the basics of oriented projective geometry (the interested reader is referred to [Stolfi, 1991] for details). 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 flats in oriented projective 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




Figure 3. (a) Relative orientation of a line and a point in 2D. The counterclockwise arrow indicates the positive orientation of the 2D universe. (b) Relative orientation of a plane and a point in \mathbb{T}^3 . The counterclockwise arrow indicates the orientation of the plane Π , and the straight arrow points to the positive side of Π , which is located using the right-hand rule.

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. To form an oriented simplex spanning the join of two oriented flats F and G , we concatenate the simplices spanning the respective flats *in that order*.

There are two oriented flats of dimension n , the *positive* and the *negative universe*, whose Plücker coordinates are given by a positive and a negative scalar, respectively. Assuming from now on that right-handed triangles in \mathbb{T}^2 and tetrahedra in \mathbb{T}^3 have (arbitrarily) been chosen as the positive ones, a necessary and sufficient condition for the triangle (x, y, z) in \mathbb{T}^2 (resp. tetrahedron (X, Y, Z, W) in \mathbb{T}^3) to span the positive universe is that the determinant $|x, y, z|$ (resp. $|X, Y, Z, W|$) be positive.

The *relative orientation* of two disjoint flats F and G such that $\dim(F) + \dim(G) + 1 = n$ is positive (resp. negative) if $F \vee G$ spans the positive (resp. negative) universe. The two cases will be denoted by $F \vee G > 0$ and $F \vee G < 0$. 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. $l \vee x < 0$). In Figure 3 (a), we have $l \vee x = |y, z, x| > 0$. In \mathbb{T}^3 , relative orientation is defined for points and planes, as well as pairs of (skew) lines. In this article, we are mostly interested in the former, and we say that a point X lies on the positive (resp. negative) side of a plane Π when $\Pi \vee X > 0$ (resp. $\Pi \vee X < 0$).




Figure 4. The local relationship between a curve and its tangent line: γ lies on the positive side of the tangent line at a convex point (left), lies on its negative side at a concave point (middle), and crosses the tangent at an inflection (right).

In Figure 3 (b), we have $\Pi \vee X = |Y, Z, W, X| > 0$. Finally, note that the join of a line and a point in 2D—or of a plane and a point in 3D—is given by the dot product of the respective coordinate vectors, $l \vee x = l^T x$ and $\Pi \vee X = \Pi^T X$ (Stolfi, 1991).

3.2. ORIENTING CURVES AND SURFACES

We show in this section how to orient curves in \mathbb{T}^2 and surfaces in \mathbb{T}^3 , and characterize their invariants. A geometric property is an *oriented projective invariant* when it continues to hold under orientation-preserving projective transformations. We also say that 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 any (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 naturally oriented in the direction of increasing values of s (Figure 4). 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. In the former case, we say that γ is locally *convex*, and in the latter case, it is locally *concave*.

When the curve γ is the boundary $\partial\omega$ of a solid region ω of the plane, it is possible to orient it so that ω is on the *positive* side of the tangent at every point of the curve (Figure 5). Intuitively, this orientation can be thought of as




Figure 5. Orientation convention for a curve γ bounding a solid region ω in the plane. Note that at the concave point x , the tangent is locally inside ω , and at the convex point y , the tangent is locally outside ω .

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).

Consider the determinant $\kappa = |x, x', x''|$, a simple quantity that can be seen as a precursor of Euclidean plane curvature. Geometrically, the sign of κ 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 shows that the sign of κ can be used to determine the local shape of γ at x .

Proposition 2. *The sign of $\kappa = |x, x', x''|$ is an oriented projective invariant. The point x is convex (resp. concave, an inflection) when κ is positive (resp. negative, zero).*

The proof is based on considering the determinant $|x, x', x(s + \delta s)|$, where $x(s + \delta s)$ is a point on γ infinitesimally close to x . With the help of a second-order Taylor expansion of $x(s + \delta s)$, it is easy to show that the sign of this determinant is the same as the sign of κ as δs approaches zero.

Orienting surfaces. The parametrization of a smooth surface in \mathbb{T}^3 induces a natural orientation of its tangent plane $\Pi = X \vee X_u \vee X_v$. According to Proposition 1, the sign of the determinant K of the shape matrix is an ordinary projective invariant, and thus an oriented one as well. As shown by the next proposition, the sign of $\text{II}(U, U)$ is also an oriented projective invariant.

Proposition 3. *Let U be a tangent direction at the point X of the surface Σ . The sign of $\text{II}(U, U)$ is an oriented projective invariant. At elliptic points, it is*

independent of the direction U . If U_1 and U_2 are two conjugate tangent directions at a hyperbolic point, then $\text{II}(U_1, U_1)$ and $\text{II}(U_2, U_2)$ have opposite signs.

When the surface Σ is the boundary $\partial\Omega$ of some solid Ω , it is possible to orient it so that Ω is always located on the positive side of the tangent plane. 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 the tangent plane.³ The sign of II can be used to characterize tangents that are locally inside or outside Ω , as shown by the next proposition.

Proposition 4. *Let Σ denote the surface bounding a solid Ω , oriented so that Ω lies on the positive side of its tangent planes. Given a point X on Σ and a direction U in the tangent plane at X , the tangent line $X \vee U$ is locally outside the solid when $\text{II}(U, U) > 0$, and locally inside it when $\text{II}(U, U) < 0$.*

The proof of this important result is somewhat technical, and is omitted here for the sake of conciseness. See (Lazebnik, 2002) for details.

We say that an elliptic point on the surface of a solid is *convex* when the solid is locally on the positive side of its tangent plane, and *concave* when it is locally on the negative side. Combining Propositions 3 and 4, we finally obtain a characterization of the oriented local shape of elliptic points.

Proposition 5. *Let Σ be as in Proposition 4. An elliptic point on Σ is convex when $\text{II}(U, U) > 0$ for some tangent direction U , and concave otherwise.*

4. Surfaces and Their Outlines

This section concludes the development of our theoretical framework by showing how oriented projective differential geometry can be used for reasoning about

³ Because the notion of an inward-pointing tangent is not readily 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 omitted here (Lazebnik, 2002).




Figure 6. An oriented projective camera. The circular arrow indicates the intrinsic orientation of the image plane, so that $O \vee R > 0$. The three projection planes are $U = O \vee B \vee C$, $V = O \vee C \vee A$, and $W = O \vee A \vee B$.

surfaces and their projections. Specifically, it presents an oriented camera model and gives consistent conventions for orienting the rim on a surface and the outline in the picture. With all the mathematical ingredients in place, we end the section by presenting a novel proof of Koenderink’s famous characterization of the convexities, concavities and inflections of the apparent contours of solids bounded by smooth surfaces (Koenderink, 1984).

4.1. CAMERA MODEL

Geometrically, a pinhole camera is defined by its *optical center* O and its *image plane* R . By convention, we orient O and R such that O lies on the negative side of R , i.e., $R \vee O < 0$ or $O \vee R > 0$ (note that relative orientation for points and planes is anti-commutative). The corresponding perspective projection maps every point X distinct from O onto the point $(O \vee X) \wedge R$, i.e., the (oriented) intersection of the ray joining O to X with the image plane. The convention that $O \vee R$ be positive ensures that the projection maps any point X lying in the image plane onto itself (as opposed to the *antipodal* point $-X$) (Stolfi, 1991).

The perspective projection mapping is a generalized projective transformation from \mathbb{T}^3 to R (Stolfi, 1991). In order to model this transformation algebraically, we choose an arbitrary but fixed coordinate system for R and identify R with \mathbb{T}^2 .

Thus, if X denotes a point in \mathbb{T}^3 and x is its projection, we can write $x \simeq \mathcal{P}X$, where “ \simeq ” denotes equality up to positive scale, and \mathcal{P} is the 3×4 *projection matrix* associated with the camera.

Note that selecting a projective coordinate system for R amounts to taking three points A, B, C in R that form a *positive* triangle, i.e., $R \simeq A \vee B \vee C$.⁴ Using the methodology of (Faugeras et al., 2001), it is easy to show that

$$\mathcal{P} = \begin{pmatrix} U^T \\ V^T \\ W^T \end{pmatrix},$$

where $U = O \vee B \vee C$, $V = O \vee C \vee A$, and $W = O \vee A \vee B$ are the *projection planes* of the camera (Figure 6).

In standard projective geometry, the *unoriented* center O of a camera is determined by the null space of the camera projection matrix, since $\mathcal{P}O = 0$. In the oriented setting, it is also possible to recover the pinhole given the projection matrix \mathcal{P} . Indeed, given our convention that $O \vee R > 0$, it is easy to show that O is given by the oriented meet $U \wedge V \wedge W$ of the three projection planes. We can recover the three *projection rays* $V \wedge W \simeq O \vee A$, $W \wedge U \simeq O \vee B$, and $U \wedge V \simeq O \vee C$. However, it is impossible to recover the basis points A , B , and C , or even the image plane itself. In fact, it is easy to see that any two cameras with the same pinhole but different image planes can be represented by the same (up to positive scale) projection matrix provided that we choose the same projection rays. On the other hand, a different choice of projection rays will result in a different projection matrix for the same geometric camera. These facts compel us to seek a more fundamental invariance relationship between cameras. The following proposition is an oriented statement of the well-known result that cameras with the same center are projectively equivalent (Faugeras et al., 2001).

Proposition 6. *Any two cameras with the same center O and two properly oriented image planes R and R' (i.e., $O \vee R > 0$ and $O \vee R' > 0$) are in oriented*

⁴ A fourth *unit* point is actually required to complete the specification of a projective frame and the corresponding projective coordinates (see Berger [1987] for details). The unit point will not, however, play an explicit role in the rest of this presentation.




Figure 7. Illustration of Proposition 7. The circular arrows indicate the orientation of the planes, and the straight arrows point toward the positive sides of the planes. Here the tetrahedron (O, X, Y, Z) and the triangle (x, y, z) are both positive.

projective equivalence: if \mathcal{P} and \mathcal{P}' are the projection matrices corresponding to the two cameras for any projective bases of R and R' , respectively, then for any point X in \mathbb{T}^3 , its two images $x \simeq \mathcal{P}X$ and $x' \simeq \mathcal{P}'X$ are related by the same orientation-preserving transformation: $x \simeq \mathcal{M}x'$, where \mathcal{M} is a 3×3 matrix with positive determinant.

The next proposition, which to our knowledge has not previously appeared in the computer vision literature, relates the orientation of the tetrahedron formed by the camera center and three points in \mathbb{T}^3 to that of the corresponding image triangle (Figure 7). This result is central to the development of our theoretical framework: It alone will allow us to connect the qualitative invariants of 3D surfaces to those of their outlines, as will be seen, for example, in Propositions 9 and 13.

Proposition 7. *Let $X, Y,$ and Z be three points in \mathbb{T}^3 with projections $x, y,$ and z in the image plane. Then the tetrahedron (O, X, Y, Z) is positive (resp. negative, improper) if and only if the triangle (x, y, z) is positive (resp. negative, improper) in the image plane:*

$$|x, y, z| \simeq |O, X, Y, Z|.$$

Proof. Let us first suppose that the tetrahedron (O, X, Y, Z) is positive, and let $R' = X \vee Y \vee Z$. Because $O \vee R' > 0$, the optical center O and the image plane R' form a properly oriented camera. Let x', y' , and z' denote the images of X, Y , and Z under some perspective projection matrix associated with this camera.⁵ According to Proposition 6, the points x, y, z and x', y', z' are related by the same orientation-preserving transformation, so the corresponding triangles have the same orientation, and, in particular, the triangle (x, y, z) is positive. If the tetrahedron is negative, we define $R' = -X \vee -Y \vee -Z$ and apply the same reasoning to the points $-X, -Y$, and $-Z$. It follows that the triangles $(-x, -y, -z)$ and $(-x', -y', -z')$ have the same orientation, thus (x, y, z) is negative. When the tetrahedron (O, X, Y, Z) is not proper, it is easy to see that the points x, y, z are collinear. \square

Note that the above proof assumes a particular choice of coordinate systems for the image planes. However, as a consequence of Proposition 6, Proposition 7 holds for any such choice.

Before concluding this section, let us briefly mention a distinction between the oriented camera model presented above and the ones described in previous literature (Hartley, 1998; Laveau and Faugeras, 1996; Werner et al., 1998; Werner and Pajdla, 2001a; Werner and Pajdla, 2001b; Chum et al., 2003). In the (usual) setting where \mathbb{T}^3 is considered as the projective completion of the “physical” affine space (Faugeras et al., 2001), the third projection plane, W can be interpreted as a *focal plane* parallel to the image plane and passing through the pinhole O . Since the third coordinate of an image point $x = \mathcal{P}X$ is $x_3 = W^T X = W \vee X$, the projections of points in the focal plane verify $x_3 = 0$ and lie on the line at infinity of the image plane. For all other points, x_3 is positive (resp. negative) when X is on the positive (resp. negative) side of W . In this way, the oriented camera model can be used to distinguish between scene points that lie in front

⁵ Geometrically speaking, in this case the points X, Y , and Z project to themselves since they lie on the image plane R' . However, the algebraic formalism of projection matrices requires us to make a distinction between the points and their images.

of the camera focal plane from the ones that lie behind (Laveau and Faugeras, 1996). Clearly, only points located in front of the focal plane can be visible in the image. Of course, when one insists on a purely projective setting, as we do here, such a simple interpretation is not possible, and it is only *by convention* that “physical” image points verify $x_3 > 0$. Thus, we do not use this constraint in the article and content ourselves with ensuring consistent orientation between points, curves, surfaces, and their projections.

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 $|X, X_u, X_v, O| = 0$ (Figure 8). The corresponding perspective projection of Γ is an image plane curve γ called the *apparent contour*, or *outline* of Σ . The following well-known and fundamental result (Koenderink, 1990) relates the viewing direction from the camera center to a rim point and the tangent to the rim at that point:

Proposition 8. *Let X be a point on the rim Γ , O the center of a perspective projection camera, and T the (unoriented) tangent line to Γ at X . Then $O \vee X$ and T lie in conjugate directions.*

This fact will play a key role in our proofs of Propositions 9 and 12. It makes intuitive sense if one recalls the geometric interpretation of conjugate directions given in Section 2.2: It is clear that the viewing rays $O \vee X$ associated with all points X belonging to the rim sweep out a developable surface, namely, a cone with apex O .⁶

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

⁶ More generally, one can specify how to obtain a conjugate net on the surface by moving the camera center along a line. This result is known as Blaschke’s Theorem (Blaschke, 1967; Koenderink, 1990): The curves along which osculating cones with their apices on a straight line touch a surface, together with the curves of intersection of that surface with the planes through that line, define a net of conjugate directions on the surface.




Figure 8. The rim and the apparent contour of a smooth surface.

prior to grazing it at X . In general, visibility is not a purely local phenomenon, since 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 4)—hence the well-known fact that concavities never show up on the silhouette of an opaque object. We call a point X on the rim *locally visible* if the viewing ray $O \vee X$ is locally outside Ω .




Figure 9. Local visibility: X is locally visible (left), locally invisible (middle), locally visible but globally invisible (right).




Figure 10. 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 .

Orienting the rim. Let us assume that Σ is oriented so that Ω is everywhere on the positive side of the tangent plane. The projection process induces a relationship between the rim tangent $T = X \vee X'$ at the point X and the contour tangent t at the point $x \simeq \mathcal{P}X$, namely $t = x \vee x' \simeq (\mathcal{P}X) \vee (\mathcal{P}X')$. Clearly, all points in the tangent plane $\Pi = X \vee X_u \vee X_v$ (except O) 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 $y \simeq \mathcal{P}Y$ must satisfy $t \vee y > 0$. This can be written as $|X, X_u, X_v, Y| \simeq |x, x', y|$, and by Proposition 7 we have $|x, x', y| \simeq |O, X, X', Y|$. Therefore, we must have $X \vee X_u \vee X_v \simeq O \vee X \vee X'$. This is satisfied if we orient the rim tangent T such that $O \vee T \simeq \Pi$ (Figure 10). From now on, we will assume that the rim is oriented according to this convention, and that the orientation of the apparent contour is induced by the orientation of the rim.

4.3. KOENDERINK'S THEOREM

The following result, referred to as *Koenderink's theorem* in this presentation, was first proved in (Koenderink, 1984) for smooth surfaces embedded in Euclidean space and observed by orthographic and spherical perspective cameras.

Proposition 9. *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 (Figure 11).*




Figure 11. 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.

Koenderink's theorem is in fact also valid under planar perspective projection, and variants of the original proof can be found in (Brady et al., 1985; Koenderink, 1990; Arbogast and Mohr, 1991; Cipolla and Blake, 1992; Vaillant and Faugeras, 1992; Boyer, 1996). It is usually presented as a corollary of the simple formulas that relate the curvature of the apparent contour to the Gaussian curvature of a surface under various projection models in a three-dimensional Euclidean space. Such a formula does not exist in the projective case, but the proposition itself still holds. The proof presented below is purely projective, and, as such, encompasses both the orthographic and perspective imaging situations, which are both subsumed by our oriented camera model. More importantly, it does not rely on the ambient three-dimensional space being endowed with a Euclidean structure, revealing the intrinsically projective nature of Koenderink's theorem.

Proof. Let X be the point of interest on the surface Σ , Π the tangent plane at X , X' and X'' the corresponding derivatives along the properly oriented rim, x the projection of X , and x' and x'' the corresponding derivatives along the

apparent contour. We define the *projection direction* $U_1 = \alpha_1 X_u + \beta_1 X_v$ such $O \vee X \simeq X \vee U_1$, and the *rim tangent direction* $U_2 = X' = \alpha_2 X_u + \beta_2 X_v$. Invoking Proposition 7 yields

$$\kappa = |x, x', x''| \simeq |O, X, X', X''| \simeq |X, U_1, U_2, X''| = (\alpha_1 \beta_2 - \beta_1 \alpha_2) |X, X_u, X_v, X''|.$$

Now, it is easy to show that, for any surface curve with tangent direction $X' = U_2$, the determinant $|X, X_u, X_v, X''|$ is equal to $\text{II}(U_2, U_2)$ (Lane, 1932), and it follows that $\kappa \simeq (\alpha_1 \beta_2 - \beta_1 \alpha_2) \text{II}(U_2, U_2)$.

Note that $O \vee X \vee X' = X \vee U_1 \vee U_2 = (\alpha_1 \beta_2 - \alpha_2 \beta_1) (X \vee X_u \vee X_v)$. Since the rim tangent is oriented by convention such that $O \vee X \vee X' \simeq \Pi$, we must have $\alpha_1 \beta_2 - \alpha_2 \beta_1 > 0$, and it follows that the signs of κ and $\text{II}(U_2, U_2)$ are the same. For the point X to be visible, $\text{II}(U_1, U_1)$ must be positive. When κ is positive, $\text{II}(U_2, U_2)$ must be positive as well, and since the viewing ray and the rim tangent are conjugate, the point X must be elliptic according to Proposition 3, and therefore convex according to Proposition 5. By the same token, X must be hyperbolic when κ is negative. Finally, the rim tangent must be an asymptotic direction when $\kappa = 0$, and thus self-conjugate. But since the viewing ray is conjugate to the rim tangent as well, this means that any tangent direction is conjugate to U_2 , and the point X must be parabolic. \square

5. Application: Computing Rim Meshes

As stated in the Introduction, one of the main contentions of this paper is that oriented projective differential geometry forms a useful language for establishing intrinsically projective properties of smooth surfaces and their outlines. In the previous section, we have demonstrated this idea using the example of Koenderink's theorem, perhaps the simplest and most fundamental of the properties in question. In the final part of our paper, we provide an additional, comprehensive demonstration of the proposed mathematical framework in action by applying it

to the task of computing the *rim mesh* of a solid viewed by multiple cameras—that is, the combinatorial arrangement of faces (surface patches), edges (rim arcs), and vertices (rim intersections, or *frontier points*) induced on its surface by the corresponding contour generators. Section 5.1 uses the language of oriented projective differential geometry to define a new qualitative invariant, the *relative orientation* of a pair of rims in the vicinity of a frontier point. We also give analytic characterizations of this invariant on the surface (Proposition 12) and in the image (Proposition 13). These theoretical results form the basis for the first algorithm for computing the rim mesh of a surface from a set of image contours and (oriented) projective cameras (Section 5.2). Finally, Section 5.3 presents experiments with several real-world datasets.

It is important to emphasize that the present section is not meant as a culmination or the end result of the theoretical developments presented in Sections 2 to 4. Rather, its goal is to demonstrate how the mathematical language proposed in this article can be applied to a computational task involving curves and surfaces. We seek to give a concrete illustration of the key theoretical and practical components of any such task, from identifying an oriented projective data structure that describes the geometry and/or topology of a surface seen from different viewpoints, to stating and proving its local properties (qualitative invariants), and finally to using these theoretical results in the specification of an algorithm for computing the data structure from real-world input data. In Section 6, we will go beyond the relatively limited application considered here and indicate additional data structures that possess an intrinsically projective nature and may be treated in the same framework.

5.1. RELATIVE ORIENTATION OF RIM TANGENTS AT A FRONTIER POINT

This section briefly reviews the geometric properties of *frontier points* (Rieger, 1986; Porrill and Pollard, 1991; Cipolla et al., 1995) where pairs of rims intersect, and defines the *relative orientation* of the rims in the vicinity of a frontier point.

The two key theoretical results are Proposition 12, which shows that the relative orientation of rims depends on the relative orientation of the viewing rays in 3D and on the local shape of the surface at the frontier point, and Proposition 13, which states that the relative orientation can be determined from image information alone. The latter result will be an essential ingredient of the image-based algorithm for computing rim meshes discussed in the next section.

Frontier Points. Suppose that X is a point where the rims Γ_i and Γ_j meet (Figure 12). Then the visual rays $O_i \vee X$ and $O_j \vee X$ both lie in the tangent plane to Σ at X , and the tangent plane coincides with the (unoriented) epipolar plane defined by O_i , O_j , and X . Since O_j belongs to the plane $O_i \vee T_i = O_i \vee X \vee X'_i$, we have $|O_i, X, X'_i, O_j| = 0$. In the first image, let $x_i \simeq \mathcal{P}_i X$ be the projection of X , and $x'_i \simeq \mathcal{P}_i X'_i$ be the derivative point of the outline γ_i at x_i . By Proposition 7, we conclude that

$$|x_i, x'_i, e_{ij}| = 0, \quad (2)$$

where $e_{ij} \simeq \mathcal{P}_i O_j$ is the *epipole* in the first view. This expression can be rewritten as $t_i \vee e_{ij} = 0$, where $t_i = x_i \vee x'_i$ is the tangent to γ_i at x_i . Thus, x_i is distinguished by the geometric property that its tangent line t_i passes through the epipole e_{ij} . Equivalently, we can say that the derivative point x'_i lies on the epipolar line $l_{ij} \simeq e_{ij} \vee x_i$. In this view, an analogous relationship holds:

$$|x_j, x'_j, e_{ji}| = 0. \quad (3)$$

We will refer to x_i and x_j as *matching frontier points*. Note that x_i and x_j satisfy the epipolar constraint $x_j^T \mathcal{F}_{ij} x_i = 0$ where \mathcal{F}_{ij} is the *fundamental matrix* associated with views i and j (Luong and Faugeras, 1996).

Relative Orientation. We say that the *relative orientation* of two directions U_1 and U_2 in the tangent plane of Σ in X is *positive* when $X \vee U_1 \vee U_2 \simeq X \vee X_u \vee X_v$. In this case, we also say that the relative orientation of the tangent lines $X \vee U_1$ and $X \vee U_2$ is positive. As shown by the following proposition, whose proof is




Figure 12. Frontier points: see text for details.

elementary and omitted here for the sake of brevity, it is a simple matter to characterize relative orientation analytically.

Proposition 10. *A necessary and sufficient condition for the relative orientation of two tangent directions $U_1 = \alpha_1 X_u + \beta_1 X_v$ and $U_2 = \alpha_2 X_u + \beta_2 X_v$ to be positive is that $\alpha_1 \beta_2 - \alpha_2 \beta_1 > 0$.*

In the sequel, we will be particularly interested in pairs of *conjugate* directions (viewing rays and rim tangents). The next proposition shows how to obtain a positively oriented pair of conjugate directions.

Proposition 11. *Two tangent directions $U_1 = \alpha_1 X_u + \beta_1 X_v$ and $U_2 = \alpha_2 X_u + \beta_2 X_v$ such that*

$$\begin{pmatrix} \alpha_2 \\ \beta_2 \end{pmatrix} \simeq \mathcal{C} \begin{pmatrix} \alpha_1 \\ \beta_1 \end{pmatrix}, \quad \text{where} \quad \mathcal{C} = \begin{pmatrix} -m & -n \\ l & m \end{pmatrix},$$

are conjugate. A necessary and sufficient condition for their relative orientation to be positive is that $\Pi(U_1, U_1) > 0$ —that is, the tangent $X \vee U_1$ is locally outside the surface.

Since we are concerned with frontier points where pairs of rims meet, there are two conjugate pairs to worry about. The next proposition relates the relative orientation of the rim tangents $T_i = X \vee X'_i$ and $T_j = X \vee X'_j$ to that of the viewing rays $O_i \vee X$ and $O_j \vee X$.

Proposition 12. *The relative orientation of the rim tangents T_i and T_j is the same as (resp. opposite of) the relative orientation of the viewing rays $O_i \vee X$ and $O_j \vee X$ when X is a convex (resp. hyperbolic) point (Figure 13).*




Figure 13. Relative orientation of rims and camera centers for a convex point (left), and a hyperbolic point (right).

Note. The relative orientation of two rim tangents at a parabolic point is not defined since, as noted in the proof of Koenderink's theorem, these tangents must run along the unique asymptotic direction there. This only occurs for exceptional pairs of viewpoints since the rims associated with two views will not, in general, intersect at parabolic points.

Proposition 12 can be understood intuitively as follows. Consider putting a tiny flashlight on the tangent plane and shining the light beam 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 (Figure 13, bottom). In the Euclidean setting, this qualitative result is a simple consequence of the fact that the Gauss map pre-

serves orientation at elliptic points and reverses it at hyperbolic ones (do Carmo, 1976). The proof below is purely projective.

Proof. Let us define constants $\alpha_i, \beta_i, \alpha_j, \beta_j$, and $\alpha'_i, \beta'_i, \alpha'_j, \beta'_j$ such that

$$\begin{cases} O_i \vee X \simeq X \vee (\alpha_i X_u + \beta_i X_v), \\ O_j \vee X \simeq X \vee (\alpha_j X_u + \beta_j X_v), \end{cases} \quad \text{and} \quad \begin{cases} X \vee X'_i \simeq X \vee (\alpha'_i X_u + \beta'_i X_v), \\ X \vee X'_j \simeq X \vee (\alpha'_j X_u + \beta'_j X_v). \end{cases}$$

According to Proposition 10, the relative orientation of the viewing rays $O_i \vee X$ and $O_j \vee X$ is determined by the sign of $\alpha_i \beta_j - \alpha_j \beta_i$. Likewise, the relative orientation of $X \vee X'_i$ and $X \vee X'_j$ is given by the sign of $\alpha'_i \beta'_j - \alpha'_j \beta'_i$. We can use the *conjugate mapping* \mathcal{C} as defined in Proposition 11 to express the derivative points as functions of the viewing directions: $\begin{pmatrix} \alpha'_i \\ \beta'_i \end{pmatrix} \simeq \mathcal{C} \begin{pmatrix} \alpha_i \\ \beta_i \end{pmatrix}$ and $\begin{pmatrix} \alpha'_j \\ \beta'_j \end{pmatrix} \simeq \mathcal{C} \begin{pmatrix} \alpha_j \\ \beta_j \end{pmatrix}$. Note that \mathcal{C} obeys the orientation convention of Figure 10: This is guaranteed by Proposition 11 and the fact that the rays $O_1 \vee X$ and $O_2 \vee X$ both lie locally outside the surface since X is locally visible by both cameras. Thus, \mathcal{C} expresses the properly oriented relationship between the determinants $\alpha_i \beta_j - \alpha_j \beta_i$ and $\alpha'_i \beta'_j - \alpha'_j \beta'_i$:

$$\begin{vmatrix} \alpha'_i & \alpha'_j \\ \beta'_i & \beta'_j \end{vmatrix} \simeq |\mathcal{C}| \begin{vmatrix} \alpha_i & \alpha_j \\ \beta_i & \beta_j \end{vmatrix} = K (\alpha_i \beta_j - \alpha_j \beta_i).$$

Whenever X is convex (resp. hyperbolic), we have $K > 0$ (resp. < 0) and the signs of $\alpha_i \beta_j - \alpha_j \beta_i$ and $\alpha'_i \beta'_j - \alpha'_j \beta'_i$ are the same (resp. opposite). \square

From now on, we will simplify our terminology and refer to the relative orientation of the rims, implying the relative orientation of their tangents at a frontier point. Next, we give a proposition showing that the relative orientation of the rims can be decided from image information alone.

Proposition 13. *Given a frontier point x_i in the first image, the corresponding tangent t_i to the apparent contour, and the epipolar line l_{ij} , the relative orientation of the rims Γ_i and Γ_j is positive if and only if (a) x_i is convex and the lines t_i and l_{ij} have the same orientations (i.e. $t_i \simeq l_{ij}$), or (b) x_i is concave and $t_i \simeq -l_{ij}$ (Figure 14).*




Figure 14. Illustration of Proposition 13. The epipolar lines are oriented from the epipole e_{ij} toward the image frontier points labeled 1 to 4, and the orientations of the epipolar tangents at the frontier points are indicated by the arrows. The four points illustrate the four possible cases. Point 1: x_i is convex and $l_{ij} \simeq t_i$. Point 2: x_i is concave and $l_{ij} \simeq t_i$. Point 3: x_i is concave and $l_{ij} \simeq -t_i$. Point 4: x_i is convex and $l_{ij} \simeq -t_i$. The relative orientation of Γ_i and Γ_j is positive in cases 1 and 3, and negative in cases 2 and 4.

Proof. The epipolar line l_{ij} is the oriented projection of the viewing ray $O_j \vee X$, and the tangent t_i is the projection of the rim tangent $X \vee X'_i$. It follows immediately from Proposition 7 that l_{ij} and t_i have the same orientation exactly when $O_i \vee O_j \vee X \simeq O_i \vee X \vee X'_i$. But $O_i \vee O_j \vee X = X \vee O_i \vee O_j$, and $O_i \vee X \vee X'_i \simeq X \vee X_u \vee X_v$. Therefore, a necessary and sufficient condition for the orientation of l_{ij} and t_i to be the same is that the relative orientation of the viewing rays $O_i \vee X$ and $O_j \vee X$ be positive. Since we know from Proposition 9 that the preimage X is convex when x_i is convex, and hyperbolic otherwise, Proposition 13 follows immediately from Proposition 12. \square

For completeness, let us consider the role of the second image in the statement of the Proposition 13. Since the local shape of X does not depend on the viewpoint and since both rays $O_i \vee X$ and $O_j \vee X$ are locally outside the surface by assumption, x_i and x_j are either both convex or both concave. However, the relative orientation of X'_i and X'_j is the opposite of the relative orientation of X'_j and X'_i . Therefore, whenever $l_{ij} \simeq t_i$ in the first image, we must have $l_{ji} \simeq -t_j$ in the second image.

5.2. THE RIM MESH

Let us suppose now that we observe the object Ω using n oriented projective cameras. The rims $\Gamma_1, \dots, \Gamma_n$ associated with the centers of these cameras form an *arrangement* on the surface Σ , the *rim mesh*. The *vertices* of the rim mesh are frontier points where two rims intersect, its *edges* are rim segments between successive vertices, and its *faces* are maximal connected regions of Σ bounded by closed paths consisting of vertices and edges (Lazebnik et al., 2001).⁷

The rest of this section presents an image-based algorithm for computing the rim mesh. Specifically, we assume that we are given as input the apparent contours $\gamma_1, \dots, \gamma_n$ and the oriented projective camera matrices $\mathcal{P}_1, \dots, \mathcal{P}_n$.⁸ The knowledge of the camera matrices enables us to compute the properly oriented epipolar geometry between each pair of views (Lazebnik, 2002).

Image data alone provides very little information about the geometry of the rim mesh: We can recover the tangent planes along the rims and the positions of the mesh vertices, but this only constrains the edges (resp. faces) of the mesh to lie on the surface of (resp. inside) an outer approximation of the observed solid, its *visual hull* (Baumgart, 1974; Laurentini, 1994). On the other hand, we will show that the same image information is sufficient to recover the *combinatorial* structure of the rim mesh, thus providing an exact *topological representation* of the surface in the form of the incidence relationships of the vertices, edges, and faces of the mesh.

Although the rim mesh is well defined for any smooth surface Σ and any combination of viewpoints, its construction from image data is severely complicated by the loss of information inherent in the projection process. Apart from

⁷ See (Cross and Zisserman, 2000) for the related notion of an *epipolar net*, a term which informally refers to the set of all frontier points on a surface.

⁸ In practice, these camera matrices are usually obtained with the help of some (projective) calibration or structure-and-motion estimation technique. To guarantee the correctness of the computed rim mesh, we must first ensure that the orientations of these matrices are consistent: Namely, we must enforce the constraint that any point x_{ij} observed in the image plane of the i th camera is equal up to a *positive* scale factor to the computed projection $P_i X_j$. These sign consistency constraints can either be built into the estimation process itself or imposed a posteriori using the technique described in Hartley (1998).

-
-
1. *Find Vertices.* For each two views i, j , find all pairs of matching frontier points. Each matching pair in the image corresponds to a single vertex X belonging to the intersection of rims Γ_i and Γ_j .
 2. *Find Edges.* For each contour γ_i , construct a circular list of frontier points where the ordering is induced by the orientation of the apparent contour. For each interval between two successive pair of frontier points corresponding to vertices X and Y , add (X, Y) as the edge of the rim mesh.
 3. *Find Faces.* For each edge $E = (X, Y)$, find the boundary of its positive and negative faces (faces that lie to the left and to the right of the edge, respectively):

If the positive face pointer of E is not already initialized, initialize it to a new face F , and find the boundary of F using the following loop:

```

Set  $X_0 = X$ .
Append  $E$  to the boundary list of  $F$ .
Set the positive face pointer of  $E$  to  $F$ .
While  $X_0 \neq Y$ 
  Find the successor  $E'$  of  $E$  as shown in Figure 16.
  Append  $E'$  to the boundary list of  $F$ .
  If  $E' = (Y, Z)$ 
    Set the positive face pointer of  $E'$  to  $F$ .
  Else If  $E' = (Z, Y)$ 
    Set the negative face pointer of  $E'$  to  $F$ .
  End If
  Set  $Y = Z, E = E'$ .
End While

```

If the negative face pointer of $E = (X, Y)$ is not initialized, create a new face F and find the boundary of F using the same loop as above, but with different initial conditions:

```

Set  $X_0 = Y$ .
Append  $E$  to the boundary list of  $F$ .
Set the negative face pointer of  $E$  to  $F$ .
While  $X_0 \neq X \dots$ 

```

Figure 15. Algorithm for computing the rim mesh.

the obvious problem of self-occlusion, it may not be possible to determine the subdivision induced on Σ by a rim Γ when either Γ or Σ consists of multiple connected components. To sidestep these difficulties, we impose the following conditions on the input of our algorithm: that a visual ray passing through a rim point only intersects Σ at that point (in particular, this precludes contours from having T-junctions); that the edges of the rim mesh cannot form multiple loops, and faces cannot have holes; and finally, that the surface Σ is connected.

Our algorithm for computing the rim mesh (Figure 15) is divided into three main steps: (1) find the vertices of the mesh; (2) find its edges; (3) determine the relative orientation of the edges incident to each vertex, and use this information to trace the loops bounding each face. In step (1), the vertices corresponding to pairs of matching frontier points are computed for each pair of images. Recall that pairs of frontier points are identified according to Equations (2) and (3) as simultaneous epipolar tangencies in the two images. Step (2)—identification of the edges—is straightforward: The four edges incident to a particular vertex are simply given by the segments of the two apparent contours that are incident to the respective image frontier points. Note that the first two steps completely determine the graph structure, or the *1-skeleton*, of the rim mesh. All that remains is to find the faces of the rim mesh in step (3). This involves finding the relative ordering of the rim segments incident to each vertex, as specified in Proposition 13. With the help of the orientation information, the loops of edges bounding rim faces are traced as follows. Suppose we start with one edge $E = (X, Y)$ of the rim Γ_i and want to find the face F that lies on its left (the *positive* face). Informally, we traverse E in its forward direction (as determined by the orientation of the corresponding segment of the apparent contour) until we reach its endpoint Y , and take a left turn to get to the *successor edge* E' . This edge is one of the two segments of the second rim Γ_j incident to Y . The choice between the segments depends on the relative orientation of Γ_i and Γ_j (Figure 16): If the relative orientation is positive (resp. negative), E' is the segment of Γ_j originating (resp. terminating) at Y . We move from vertex to vertex in this manner, traversing edges either forward or backward, taking a left turn each time until we complete a cycle. The *negative* face of an edge, or the face that lies to its right, is found in a similar way, as shown in Figure 15.

Note. Let v , e , and f denote the numbers of vertices, edges, and faces of the rim mesh. Under the three assumptions listed at the beginning of this section, the Euler formula $v - e + f = 2$ holds when the surface Σ has genus 0 (which




Figure 16. Choosing the successor E' of the edge E along the boundary of the face F of the rim mesh. The circular arrows indicate the relative orientation of Γ_i and Γ_j (the intrinsic orientation of the tangent plane is assumed to be counterclockwise). Left: The relative orientation is positive, and E' is the segment of Γ_j originating at Y . Right: The relative orientation is negative, and E' is the segment of Γ_j terminating at Y .

is the case in the examples shown in Figure 17). Since each vertex has degree 4, and the sum of the degrees⁹ of all the vertices is equal to twice the number of edges, we have $e = 2v$. Substituting this expression into the Euler formula and solving for the number of faces yields $f = 2 + v$. These two relations can be used to verify the validity of the rim mesh representation computed by the algorithm outlined in this section.

5.3. RESULTS

We have implemented the algorithm described in the previous section and successfully applied it to several real-world image datasets. Figure 17 shows the results for three calibrated turntable image sequences: the vase (six images), the teapot (nine images), and the gourd (nine images). The top row of the figure shows a representative image from each of the three sequences. The rim mesh computation program takes as input a set of outlines (extracted to sub-pixel precision from each image) and fundamental matrices between each pair of views (computed from the calibration data). The 1-skeletons computed by the program are shown in the second and third rows of Figure 17. To visualize the

⁹ By the degree of a vertex we mean the number of edges incident on that vertex, regardless of edge orientation.

vertices (frontier points), we reconstruct their position in 3D and reproject them into each respective view. The edges are drawn as straight lines connecting the corresponding vertices. Because of the turntable setup, the camera remains in the same plane throughout the sequence, and most of the frontier points end up being densely clustered near the top and the bottom of the objects — an “ugly” geometric configuration that makes these sequences a rather difficult case for rim mesh computation. Finally, the bottom row of the figure shows an alternative visualization of the 1-skeletons, where they are “flattened” with the help of a graph-drawing program *Graphviz* (Gansner and North, 1999). The 1-skeletons are planar graphs, despite the fact that Graphviz is unable to compute a planar embedding from them. The figure also does not reflect their directed nature, where each edge inherits an orientation from the corresponding segment of the outline. Because of the scarcity of geometric information associated with the rim mesh, it is difficult to visualize its faces. Nevertheless, note that the numbers of vertices, edges and faces listed at the bottom of the figure verify the relationships $e = 2v$ and $f = 2 + v$ derived earlier from the Euler formula, thus demonstrating the topological consistency of the computed rim meshes.

6. Discussion

Computer vision applications of projective differential geometry have focused—so far—on the use of high-order quantitative differential invariants for object recognition. By contrast, the oriented projective differential framework presented in this article is aimed at deriving *low-order, qualitative* invariants suitable for reconstruction tasks. The past decade has seen intense study of projective reconstruction techniques for points, lines, and planes (Faugeras et al., 2001; Hartley and Zisserman, 2000). In our opinion, the addition of mathematical tools applicable to the reconstruction of more complex geometric entities such as curves

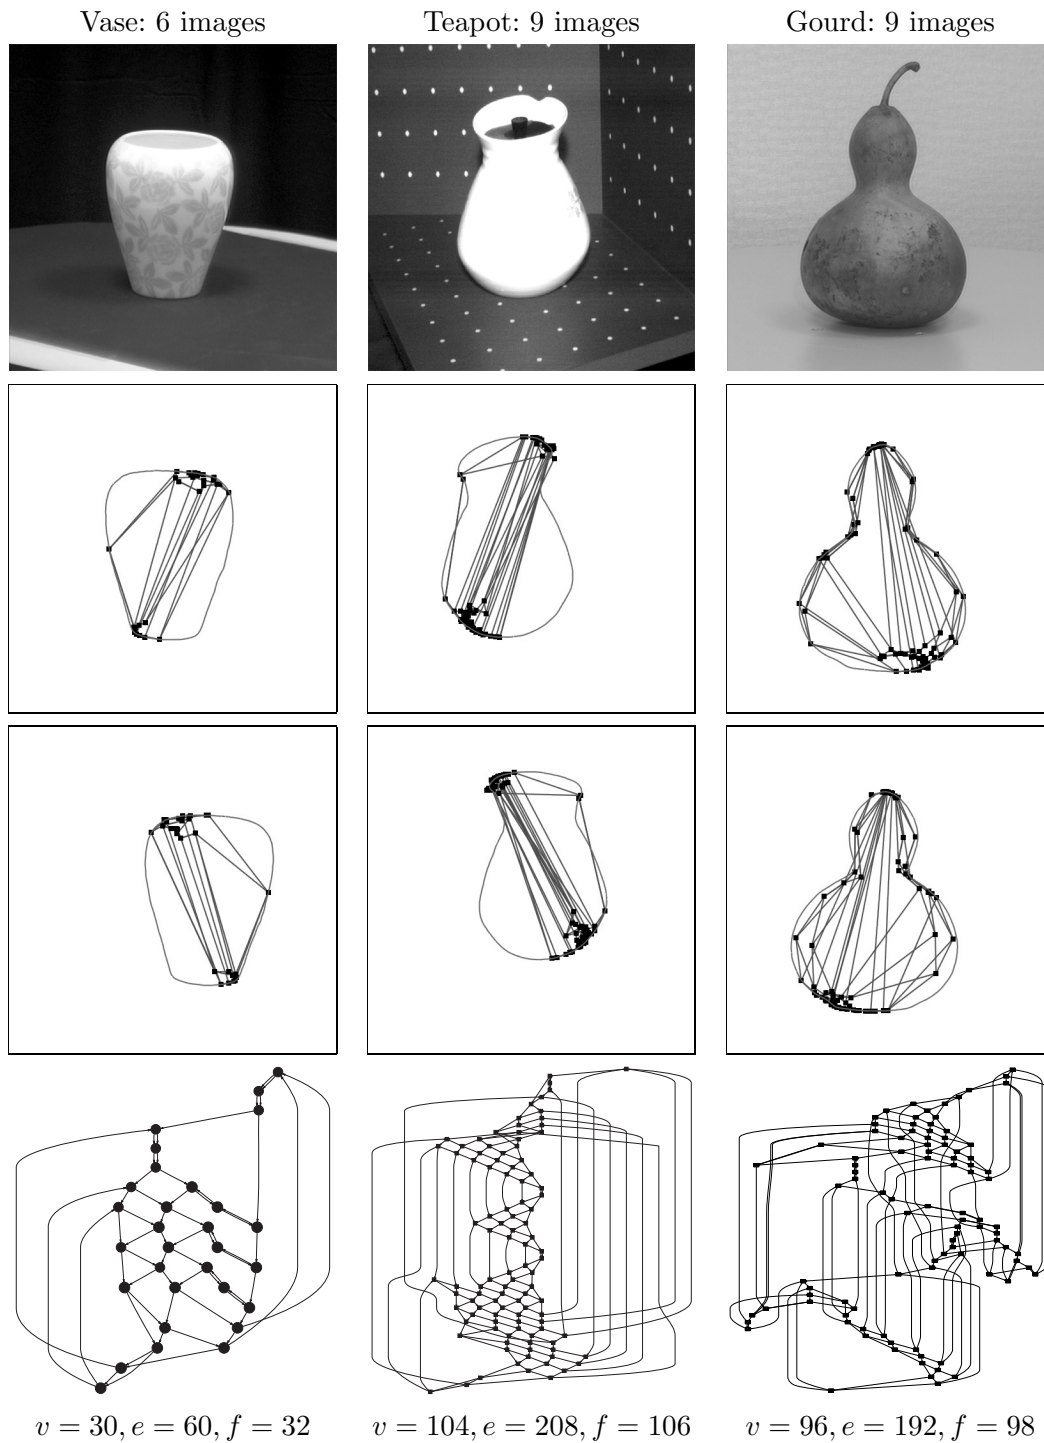


Figure 17. Top: Sample images from the vase, teapot, and gourd image sequences. Second and third row: The 1-skeleton of the rim mesh from two of the original viewpoints. Bottom: An alternative visualization of the 1-skeleton of the rim mesh. Data courtesy of Edmond Boyer.

and smooth surfaces should greatly enrich the subject of multi-view geometry (see Faugeras and Papadopoulos [1993] for related work in the Euclidean case).

Since the present article is concerned primarily with developing a theoretical framework of qualitative invariants, we have limited our discussion of applications to the relatively self-contained subject of rim meshes. However, we have extensively used the oriented constraints derived in this article in practical algorithms for computing visual hulls from image data (Lazebnik, 2002). Other potential applications include the construction of visibility complexes (Durand et al., 2002) and aspect graphs (Koenderink and Van Doorn, 1979). These objects, whose combinatorial structure is determined by local and multilocal events corresponding to special kinds of contact between 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 (Belhumeur et al., 1997), or even perhaps the set of reconstructions compatible with a fixating stereo system with unknown vergence angles (Helmholtz, 1909; Koenderink and Van Doorn, 1976).

Acknowledgments

This research was partially supported by the Beckman Institute, the UIUC Campus Research Board, and the National Science Foundation under grants IRI-990709, IIS-0308087, and IIS-0312438. Many thanks to Edmond Boyer for providing the gourd, teapot, and vase data sets, and for his participation in our work on Euclidean rim meshes (Lazebnik et al., 2001).

References

- Arbogast, E. and R. Mohr: 1991, '3D structure inference from image sequences'. *Journal of Pattern Recognition and Artificial Intelligence* **5**(5).
- Baumgart, B.: 1974, 'Geometric modeling for computer vision'. Technical Report AIM-249, Stanford University. Ph.D. Thesis. Department of Computer Science.

- Belhumeur, P., D. Kriegman, and A. Yuille: 1997, 'The bas-relief ambiguity'. In: *Proc. IEEE Conf. Comp. Vision Patt. Recog.* San Juan, Puerto Rico, pp. 1060–1066.
- Berger, M.: 1987, *Geometry*. Springer-Verlag.
- Blaschke, W.: 1967, *Differential Geometrie*. New York: Chelsea. 2 vols.
- Bol, G.: 1950, *Projektive Differentialgeometrie*. Göttingen: Vandenhoeck & Ruprecht.
- Boyer, E.: 1996, 'Object Models from Contour Sequences'. In: *Proceedings of Fourth European Conference on Computer Vision, Cambridge, (England)*. pp. 109–118. Lecture Notes in Computer Science, volume 1065.
- Brady, J., J. Ponce, A. Yuille, and H. Asada: 1985, 'Describing Surfaces'. *Computer Vision, Graphics and Image Processing* **32**(1), 1–28.
- Calabi, E., P. Olver, C. Shakiban, A. Tannenbaum, and S. Haker: 1998, 'Differential and numerically invariant signature curves applied to object recognition'. *Int. J. of Comp. Vision* **26**(2), 107–135.
- Cartan, E.: 1992, *La théorie des groupes finis et continus et la géométrie différentielle traitée par la méthode du repère mobile*. Jacques Gabay. Original edition, Gauthier-Villars, 1937.
- Chum, O., T. Werner, and T. Pajdla: 2003, 'Joint Orientation of Epipoles'. In: *Proc. British Machine Vision Conference*. pp. 73–82.
- Cipolla, R., K. Åström, and P. Giblin: 1995, 'Motion from the Frontier of Curved Surfaces'. In: *Proc. Int. Conf. Comp. Vision*. pp. 269–275.
- Cipolla, R. and A. Blake: 1992, 'Surface Shape from the Deformation of the Apparent Contour'. *Int. J. of Comp. Vision* **9**(2), 83–112.
- Cross, G. and A. Zisserman: 2000, 'Surface Reconstruction from Multiple Views Using Apparent Contours and Surface Texture'. In: *NATO Advanced Research Workshop on Confluence of Computer Vision and Computer Graphics*. pp. 25–47.
- do Carmo, M.: 1976, *Differential Geometry of Curves and Surfaces*. Englewood Cliffs, NJ: Prentice-Hall.
- Durand, F., G. Drettakis, and C. Puech: 2002, 'The 3D visibility complex'. *ACM Transactions on Graphics* **21**(2), 176–206.
- Faugeras, O.: 1994, 'Cartan's Moving Frame Method and its Application to the Geometry and Evolution of Curves in the Euclidean, Affine and Projective Planes'. In: J. Mundy, A. Zisserman, and D. Forsyth (eds.): *Applications of Invariance in Computer Vision*, Vol. 825 of *Lecture Notes in Computer Science*. Springer-Verlag, pp. 11–46.
- Faugeras, O., Q.-T. Luong, and T. Papadopoulos: 2001, *The Geometry of Multiple Images*. MIT Press.
- Faugeras, O. and T. Papadopoulos: 1993, 'A Theory of the Motion Fields of Curves'. *Int. J. of Comp. Vision* **10**(2), 125–156.
- Gansner, E. and S. North: 1999, 'An open graph visualization system and its applications to software engineering'. *Software—Practice and Experience* **30**(11), 1203–1233.
- Hartley, R.: 1998, 'Chirality'. *Int. J. of Comp. Vision* **26**(1), 41–61.
- Hartley, R. and A. Zisserman: 2000, *Multiple view geometry in computer vision*. Cambridge University Press.
- Helmholtz, H.: 1909, *Physiological optics*. Dover. 1962 edition of the English translation of the 1909 German original, first published by the Optical Society of America in 1924.
- Koenderink, J.: 1984, 'What does the occluding contour tell us about solid shape?'. *Perception* **13**, 321–330.
- Koenderink, J.: 1990, *Solid Shape*. Cambridge, MA: MIT Press.
- Koenderink, J. and A. Van Doorn: 1976, 'Geometry of binocular vision and a model for stereopsis'. *Biological Cybernetics* **21**, 29–35.
- Koenderink, J. and A. Van Doorn: 1979, 'The internal representation of solid shape with respect to vision'. *Biological Cybernetics* **32**, 211–216.
- Lane, E.: 1932, *Projective Differential Geometry of Curves and Surfaces*. The University of Chicago Press.
- Laurentini, A.: 1994, 'The Visual Hull Concept for Silhouette-based Image Understanding'. *IEEE Trans. Patt. Anal. Mach. Intell.* **16**(2), 150–162.
- Laveau, S. and O. Faugeras: 1996, 'Oriented Projective Geometry for Computer Vision'. In: *Proc. European Conf. Comp. Vision*. pp. 147–156.

- Lazebnik, S.: 2002, 'Projective Visual Hulls'. Master's thesis, University of Illinois at Urbana-Champaign. Also Beckman CVR Technical Report CVR-TR-2002-01, available at http://www-cvr.ai.uiuc.edu/ponce_grp.
- Lazebnik, S., E. Boyer, and J. Ponce: 2001, 'On Computing Exact Visual Hulls of Solids Bounded by Smooth Surfaces'. In: *Proc. IEEE Conf. Comp. Vision Patt. Recog.* pp. 156–161.
- Lazebnik, S. and J. Ponce: 2003, 'The Local Projective Shape of Smooth Surfaces and Their Outlines'. In: *Proc. Int. Conf. Comp. Vision.* To appear.
- Luong, Q.-T. and O. Faugeras: 1996, 'The Fundamental matrix: theory, algorithms, and stability analysis'. *Int. J. of Comp. Vision* **17**(1), 43–76.
- Moons, T., E. Pauwels, L. VanGool, and A. Oosterlinck: 1995, 'Foundations Of Semi-Differential Invariants'. *Int. J. of Comp. Vision* **14**(1), 25–47.
- Mundy, J. and A. Zisserman: 1992, *Geometric Invariance in Computer Vision*. Cambridge, Mass.: MIT Press.
- Mundy, J., A. Zisserman, and D. Forsyth: 1994, *Applications of Invariance in Computer Vision*, Vol. 825 of *Lecture Notes in Computer Science*. Springer-Verlag.
- Porrill, J. and S. Pollard: 1991, 'Curve Matching and Stereo Calibration'. *Image and Vision Computing* **9**(1), 45–50.
- Rieger, J.: 1986, 'Three-dimensional motion from fixed points of a deforming profile curve'. *Optics Letters* **11**, 123–125.
- Salden, A., B. t. Romeny, and M. Vierverger: 1993, 'Affine and projective differential geometric invariants of space curves'. In: B. Vemuri (ed.): *Geometric Methods in Computer Vision II*. pp. 60–74, SPIE.
- Stolfi, J.: 1991, *Oriented Projective Geometry: A Framework for Geometric Computations*. Academic Press.
- Vaillant, R. and O. Faugeras: 1992, 'Using extremal boundaries for 3D object modeling'. *IEEE Trans. Patt. Anal. Mach. Intell.* **14**(2), 157–173.
- Weiss, I.: 1988, 'Projective invariants of shapes'. In: *Proc. IEEE Conf. Comp. Vision Patt. Recog.* Ann Arbor, MI, pp. 291–297.
- Werner, T. and T. Pajdla: 2001a, 'Chirality in Epipolar Geometry'. In: *ICCV*. pp. 548–553.
- Werner, T. and T. Pajdla: 2001b, 'Oriented Matching Constraints'. In: *Proc. British Machine Vision Conference*. pp. 441–450.
- Werner, T., T. Pajdla, and V. Hlavac: 1998, 'Oriented Projective Reconstruction'. In: *Proc. meeting of the Austrian Association for Pattern Recognition*. pp. 245–254.