Chapter 5

GEOMETRIC IMAGE FEATURES

What is the image of a given geometric figure, for example a line, a plane, a sphere, or a general smooth surface? Answering this question is of course important if we are to interpret pictures of our three-dimensional world. As shown analytically in the Chapter 6, answering it is also easy for linear features such as points and lines: indeed, under perspective projection, points map onto points, and lines project (in general) onto lines. Geometrically, this is also obvious for points, and the image of a line is simply the intersection of the retina with the plane that contains this line and the pinhole (Figure 5.1(a)). In truth, things degenerate a bit for *exceptional views*: when a line passes through the pinhole, its image is reduced to a point (Figure 5.1(b)). Likewise, the image of a surface patch covers in general a finite area of the image plane, but the projection of a plane passing through the pinhole is reduced to a line. Still, almost any small perturbation of the viewpoint restores the usual projection pattern.

By the same argument, general viewpoint assumptions allow us to infer three-dimensional scene information from a single image: in particular, Figure 5.1(b) shows that two points that appear to coincide in the image usually also coincide in space since the set of pinhole positions aligned with two distinct points has a zero volume in the three-dimensional set of possible viewpoints.

More generally, the edges and vertices of a polyhedral solid project onto line segments and points of the image plane. A subset of these segments forms a polygon that bounds the image of the solid, corresponding to surface edges adjacent to faces whose normals point into opposite directions relative to the observer (Figure 5.2(a)). To assert stronger image properties requires additional hypotheses: for example, assuming a general viewpoint and a *trihedral world*, where all polyhedral vertices are formed by the intersection of exactly three planar faces, it is possible to show that the image edges joining at a vertex may only be arranged in four qualitatively distinct configurations, the so-called *arrow*, *fork*, *L-* and *T-junctions* (Figure 5.2(b)).

Using these junctions to construct three-dimensional interpretations of line-drawings is a time-honored artificial intelligence exercise that had its hayday in

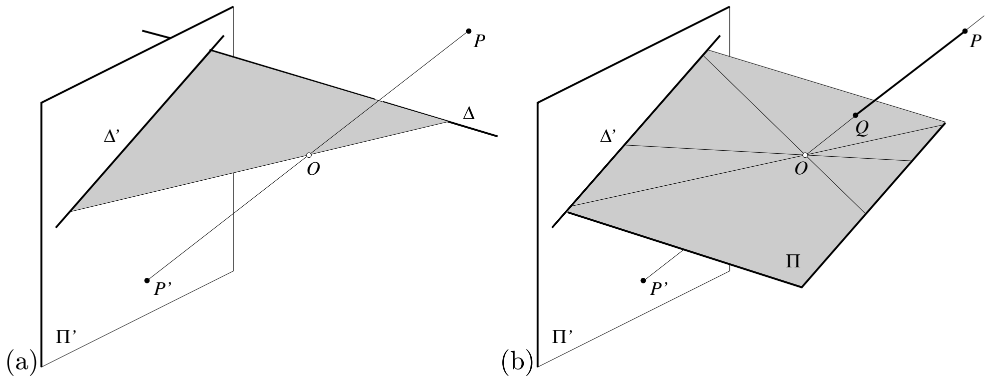

Figure 5.1. Point and line projections: (a) general case: the grey triangle represents the plane defined by the line Δ and the pinhole O; (b) degenerate case: the line supporting the segment PQ and passing through O projects onto the point P', and the plane Π passing through O projects onto the line Δ'.

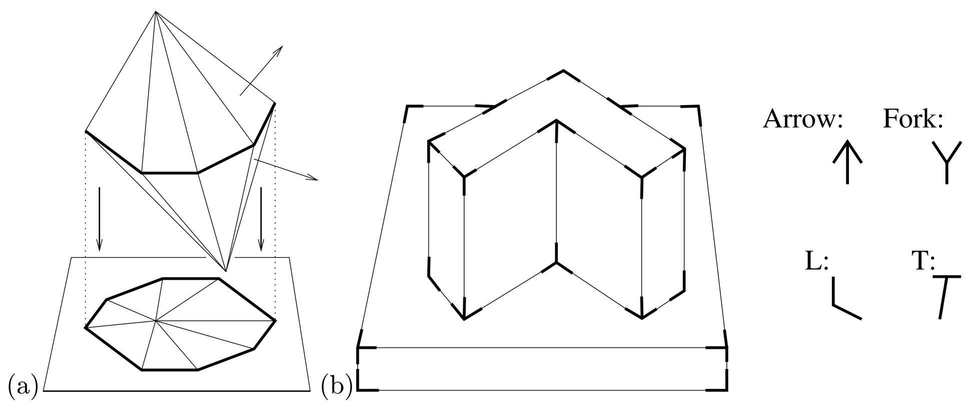

Figure 5.2. The images of polyhedral solids: (a) polygonal edges of the image of a polyhedron; (b) the image of a trihedral object and the corresponding junctions. Note that the solid shown in (a) is not trihedral, and the image junction corresponding to its top vertex is neither an arrow nor a fork, an L- or a T-junction.

the early seventies but has fallen out of favor because of the great difficulty in constructing a program that reliably extracts flawless line-drawings from images. This merely suggests that a purely bottom-up approach to image understanding may fail, but should not detract from the fundamental importance of understanding how lines, points and other simple geometric objects project into images.

What about more complex figures? It is easy to show that any (planar or twisted) curve that can be defined by polynomial equations (e.g., a conic section, which is defined by a quadratic equation) also projects onto a curve defined by polynomial equations. The case of curved surfaces is more interesting: consider for example a solid bounded by a smooth surface with a smooth reflectance function, observed by a perspective camera under a smooth illumination pattern. Ignoring shadows (e.g., considering a single point light source and a camera co-located far from the scene), the image is also smooth, except perhaps along *occlusion boundaries*, where the object surface turns away from the camera, and two points that project on opposite sides of these boundaries belong to spatially separate parts of the surface.

Intuitively, it is clear that occlusion boundaries form a set of image curves, called the *outline*, *silhouette*, or *contour* of the object (Figure 5.3(a)). As in the case of spheres discussed in Chapter 1, these curves are formed by intersecting the retina with viewing cones (or cylinders in the case of orthographic projection) whose apex coincides with the pinhole and whose surface grazes the object along a second set of curves, called the *occluding contour*, or *rim*.

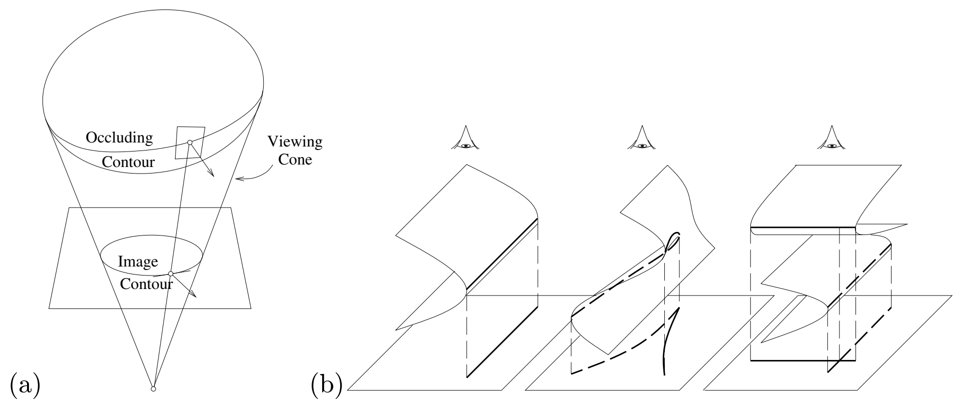

Figure 5.3. The images of solids bounded by smooth surfaces: (a) the occluding boundaries of a smooth surface; (b) contour components: folds, cusps, and T-junctions. Reprinted from [Petitjean *et al.*, 1992], Figure 3.

This intuition can be made mathematically precise. In particular, it can be shown that the occluding contour is in general a smooth curve, formed by *fold* points, where the viewing ray is tangent to the surface, and a discrete set of *cusp* points where the ray is tangent to the occluding contour as well. The image contour is piecewise smooth, and its only singularities are a discrete set of *cusps*, formed by the projection of cusp points, and *T-junctions*, formed by the transversal superposition of pairs of fold points (Figure 5.3(b)). The intuitive meaning of these exotic terms should be pretty clear: a fold is a point where the surface folds away from its viewer, and a contour cusps at a point where it suddenly decides to turn back, following a different path along the same tangent (this is for transparent objects only: contours of opaque objects terminate at cusps, see Figure 5.3(b)). Likewise, two smooth pieces of contour cross at a T-junction (unless the object is opaque and one of the branches terminates at the junction).

Interestingly, attached shadows are delineated by the occluding contours associated with the light sources, and cast shadows are bounded by the corresponding object outlines (Figure 5.4). Thus we also know what they look like. (Caveat: the objects onto which shadows are cast may themselves have curved surfaces, which complicates things a great deal. However, the boundaries of attached shadows are really just occluding contours. Of course, light sources are rarely punctual, and this adds further difficulties..)

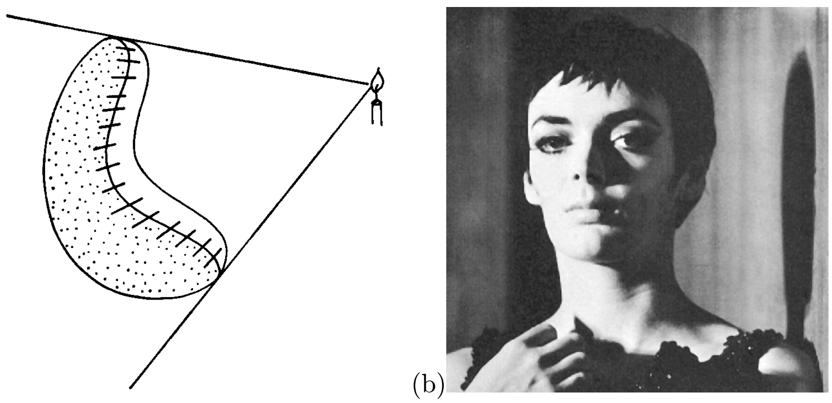

(a) (b)

Figure 5.4. (a) Shadow boundaries and occluding contours. Reprinted from [Koenderink, 1990], Figure 157. (b) The actress Barbara Steele and the cast and attached shadows of her face. Reprinted from [MMF, 1967], p. 33.

It should not come as a surprise that the notion of occluding boundaries carries over to the polyhedral world: the image contour of a polyhedron is exactly the polygon bounding its image, and its occluding contour is formed by the edges separating faces that point in opposite directions relative to the observer (these are drawn as thicker line segments in Figure 5.2(a)). The image contour of a solid shape constrains it to lie within the associated viewing cone but does not reveal the depth of its occluding contour. In the case of solids bounded by smooth surfaces, the contour provides additional information: in particular, the plane defined by the eye and the tangent to the image contour is itself tangent to the surface. Thus the contour orientation determines the surface orientation along the occluding contour.

More generally, the rest of this chapter focuses on the geometric relationship between curved surfaces and their projections and on the type of information about surface geometry that can be infered from contour geometry (we will come back later to the case of polyhedral solids): for example, we will show that the contour curvature also reveals information about the surface curvature. In the mean time, let us start by introducing elementary notions of differential geometry that will provide a natural mathematical setting for our study. Differential geometry will prove useful again later in this book, when we study the changes in object appearance that stem

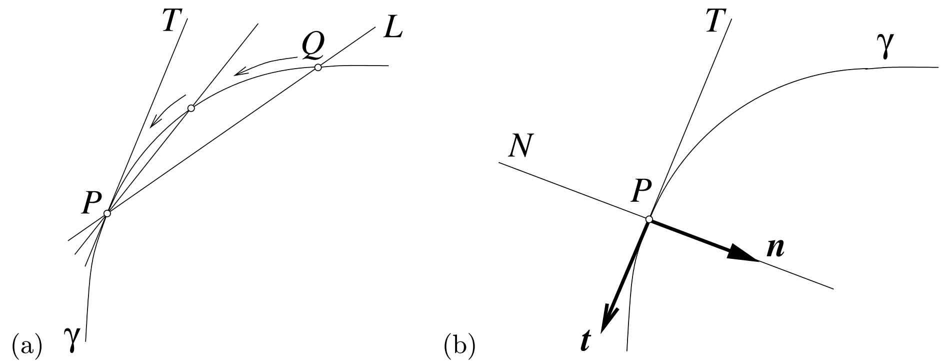

Figure 5.5. Tangents and normals: (a) definition of the tangent as the limit of secants; (b) the coordinate system defined by the (oriented) tangent and normal.

from viewpoint changes.

5.1 Elements of Differential Geometry

In this section we will present the rudiments of differential geometry necessary to understand the local relationship between light rays and solid objects. The topic of our discussion is of course technical, but we will attempt to stay at a fairly informal level, emphasizing descriptive over analytical geometry. We will also illustrate some of the concepts introduced with a simple study of surface specularities and what they reveal about surface shape. Following [Koenderink, 1990] we will limit our study of surfaces to those bounding compact solids in Euclidean space.

5.1.1 Curves

We start with the study of curves that lie in a plane. We will examine a curve γ in the immediate vicinity of some point P, and will assume that γ does not intersect itself, or, for that matter, terminate in P. If we draw a straight line L through P, it will (in general) intersect γ in some other point Q, defining a *secant* of this curve (Figure 5.5(a)). As Q moves closer to P, the secant L will rotate about P and approach a limit position T, called the *tangent line* to γ in P.

By construction, the tangent T has more intimate contact with γ than any other line passing through P. Let us now draw a second line N through P and perpendicular to L, and call it the *normal* to γ in P. Given an (arbitrary) choice for a unit *tangent vector* $\boldsymbol{t}$ along L, we can construct a right-handed coordinate frame whose origin is P and whose axes are $\boldsymbol{t}$ and a unit *normal vector* $\boldsymbol{n}$ along N.

This coordinate system is particularly well adapted to the study of the curve in the neighborhood of P: its axes divide the plane into four quadrants that can be numbered in counterclockwise order as shown in Figure 5.6, the first quadrant being chosen so it contains a particle traveling along the curve toward (and close to) the origin. In which quadrant will this particle end up just after passing P? As shown by Figure 5.6(a)-(d), there are four possible answers to this question, and they characterize the shape of the curve near P: we say that P is *regular* when the moving point ends up in the second quadrant, and *singular* otherwise. When the particle traverses the tangent and ends up in the third quadrant, P is called an *inflection* of the curve, and we say that P is a *cusp* of the first or second kind in the two remaining cases respectively. This classification is in fact independent of the orientation chosen for γ, and it turns out that almost all points of a general curve are regular, while singularities only occur at isolated points.

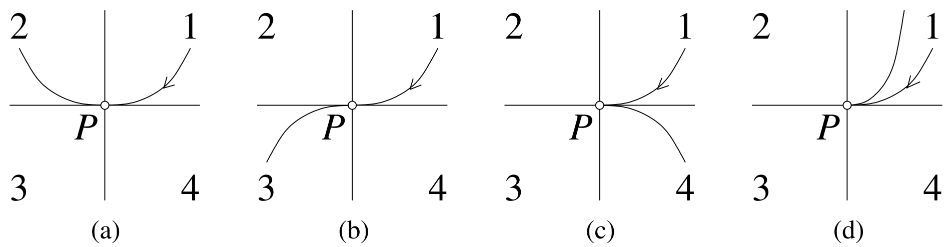

Figure 5.6. A classification of curve points: (a) a regular point; (b) an inflection; (c) a cusp of the first kind; (d) a cusp of the second kind. Note that the curve stays on the same side of the tangent at regular points.

As noted before, the tangent to a curve γ in P is the closest linear approximation of γ passing through this point. In turn, constructing the closest *circular* approximation will now allow us to define the *curvature* in P, another fundamental characteristic of the curve shape: consider a point P' as it approaches P along the curve, and let M denote the intersection of the normal lines N and N' in P and P' (Figure 5.7). As P' moves closer to P, M approaches a limit position C along the normal N, called the *center of curvature* of γ in P.

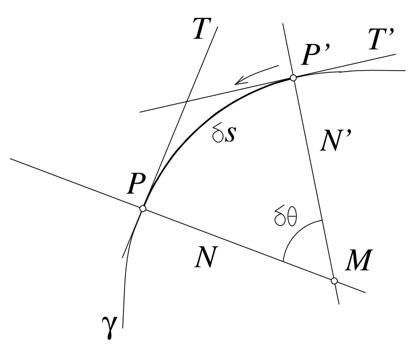

Figure 5.7. Definition of the center of curvature as the limit of the intersection of normal lines through neighbors of P.

At the same time, if $\delta\theta$ denotes the (small) angle between the normals N and N', and δs denotes the length of the (short) curve arc joining P and P', the ratio $\delta\theta/\delta s$ also approaches a definite limit κ, called the *curvature* of the curve in P, as δs nears zero. It turns out that κ is just the inverse of the distance r between C and P (this follows easily from the fact that $\sin u \approx u$ for small angles, see exercises). The circle centered in C with radius r is called the *circle of curvature* in P, and r is the *radius of curvature*.

It can also be shown that a circle drawn through P and two close-by points P' and P'' approaches the circle of curvature as P' and P'' move closer to P. This circle is indeed the closest circular approximation to γ passing through P. The curvature is zero at inflections, and the circle of curvature degenerates to a straight line (the tangent) there: inflections are the "flattest" points along a curve.

We will now introduce a device that will prove to be extremely important in the study of both curves and surfaces, the *Gauss map*. Let us pick an orientation for the curve γ and associate with every point P on γ the point Q on the unit circle S^1 where the tip of the associated normal vector meets the circle (Figure 5.8). The corresponding mapping from γ to S^1 is the Gauss map associated with γ.[1]

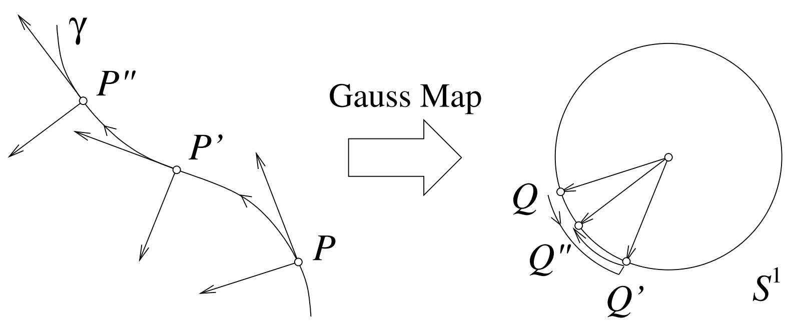

Figure 5.8. The Gaussian image of a plane curve. Observe how the direction of traversal of the Gaussian image reverses at the inflection P' of the curve. Also note that there are close-by points with parallel tangents/normals on either side of P'. The Gaussian image folds at the corresponding point Q'.

Let us have another look at the limiting process used to define the curvature. As P' approaches P on the curve, so does the Gaussian image Q' of P' approach the image Q of P. The (small) angle between N and N' is equal to the length of the arc joining Q and Q' on the unit circle. The curvature is therefore the limit of the ratio between the lengths of corresponding arcs of the Gaussian image and of the curve as both approach zero.

The Gauss map also provides an interpretation of the classification of curve points introduced earlier: consider a particle traveling along a curve and the cor-

[1] The Gauss map could have been defined just as well by associating with each curve point the tip of its unit tangent on S^1. The two representations are equivalent in the case of planar curves. The situation will be different when we generalize the Gauss map to twisted curves and surfaces.

responding motion of its Gaussian image. The direction of traversal of γ stays the same at regular points and inflections, but reverses for both types of cusps (Figure 5.6). On the other hand, the direction of traversal of the Gaussian image stays the same at regular points and cusps of the first kind, but it reverses at inflections and cusps of the second kind (Figure 5.8). This indicates a double covering of the unit circle near these singularities: we say that the Gaussian image *folds* at these points.

A conventional sign can be chosen for the curvature at every point of a plane curve γ by picking some orientation for this curve, and deciding, say, that the curvature will be positive when the center of curvature lies on the same side of γ as the tip of the oriented normal vector, and negative when these two points lie on opposite sides of γ. Thus the curvature changes sign at inflections, and reversing the orientation of a curve also reverses the sign of its curvature.

Twisted space curves are more complicated animals that their planar counterparts. Although the tangent can be defined as before as a limit of secants, there is now an infinity of lines perpendicular to the tangent at a point P, forming a *normal plane* to the curve at this point (Figure 5.9).

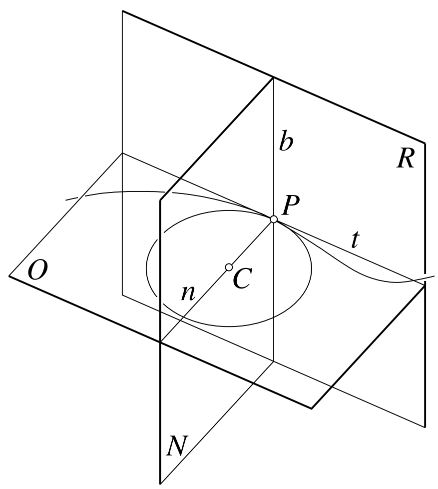

Figure 5.9. The local geometry of a space curve: N, O and R are respectively the normal, osculating, and rectifying plane; t, n and b are respectively the tangent, (principal) normal and binormal lines, and C is the center of curvature.

In general, a twisted curve does not lie in a plane in the vicinity of one of its points, but there exists a unique plane that lies closest to it. This is the *osculating plane*, defined as the limit of the plane containing the tangent line in P and some close-by curve point Q as the latter approaches P. We finish the construction of a local coordinate frame in P by drawing a *rectifying plane* through P, perpendicular to both the normal and osculating planes. The axes of this coordinate system, called *moving trihedron*, or *Frénet frame*, are the tangent, the *principal normal* formed by the intersection of the normal and osculating planes, and the *binormal* defined by the intersection of the normal and rectifying planes (Figure 5.9).

As in the planar case, the curvature of a twisted curve can be defined in a number of ways: as the inverse of the radius of the limit circle defined by three curve points as they approach each other (this circle of curvature lies in the osculating plane), as the limit ratio of the angle between the tangents at two close-by points and the distance separating these points as it approaches zero, etc. Likewise, the concept of Gaussian image can be extended to space curves, but this time the tips of the tangents, principal normals and binormals draw curves on a unit *sphere*. Note that it is not possible to give a meaningful sign to the curvature of a twisted curve: in general, such a curve does not have inflections, and its curvature is positive everywhere.

The curvature can be thought of as a measure of the rate of change of the tangent direction along a curve. It is also possible to define the rate of change of the osculating plane direction along a twisted curve: consider two close-by points P and P' on the curve; we can measure the angle between their osculating planes, or equivalently between the associated binormals, and divide this angle by the distance between the two points. The limit of this ratio as P' approaches P is called the *torsion* of the curve in P. Not surprisingly, its inverse is the limit of the ratio between the lengths of corresponding arcs on the curve and the spherical image of the binormals.

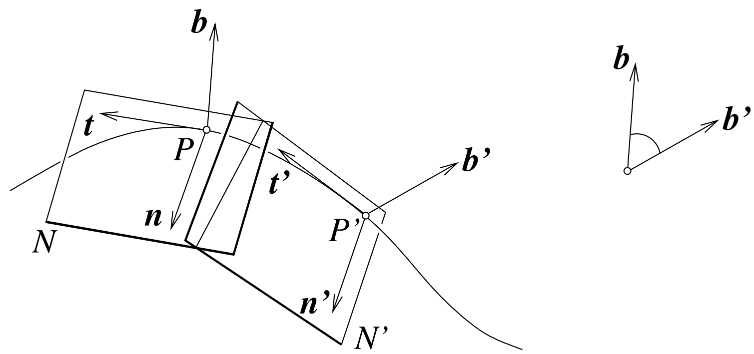

Figure 5.10. Geometric definition of the torsion as the limit, as both quantities approach zero, of the ratio obtained by dividing the angle between the binormals by the distance between the associated surface points.

A space curve can be oriented by considering it as the trajectory of a moving particle and picking a direction of travel for this particle. Furthermore, we can pick an arbitrary reference point P_0 on the curve and define the *arc length* s associated with any other point P as the length of the curve arc separating P_0 and P. Although the arc length depends on the choice of P_0, its differential does not (moving P_0 along the curve amounts to adding a constant to the arc length), and it is often convenient to parameterize a curve by its arc length, where some unspecified choice of origin P_0 is assumed. In particular, the tangent vector at the point P is the unit velocity

$\boldsymbol{t} = dP/ds$. Reversing s also reverses $\boldsymbol{t}$. It can be shown that the acceleration d^2P/ds^2, the curvature κ, and the (principal) normal $\boldsymbol{n}$ are related by

$$\frac{d^2P}{ds^2} = \frac{d\boldsymbol{t}}{ds} = \kappa\boldsymbol{n}.$$

Note that κ and $\boldsymbol{n}$ are both independent of the curve orientation (the negative signs introduced by reversing the direction of traversal of the curve cancel during differentiation), and the curvature is the magnitude of the acceleration. The binormal vector can be defined as $\boldsymbol{b} = \boldsymbol{t} \times \boldsymbol{n}$, and, like $\boldsymbol{t}$, it depends on the orientation chosen for the curve. In general, it is easy to show that

$$\frac{d}{ds}\begin{pmatrix} \boldsymbol{t} \\ \boldsymbol{n} \\ \boldsymbol{b} \end{pmatrix} = \begin{pmatrix} 0 & \kappa & 0 \\ -\kappa & 0 & \tau \\ 0 & -\tau & 0 \end{pmatrix} \begin{pmatrix} \boldsymbol{t} \\ \boldsymbol{n} \\ \boldsymbol{b} \end{pmatrix},$$

where τ denotes the torsion in P. Unlike the curvature, the torsion may be positive, negative or zero for a general space curve. Its sign depends on the direction of traversal chosen for the curve, and it has a geometric meaning: in general, a curve crosses its osculating plane at every point with non-zero torsion, and it emerges on the positive side of that plane (i.e., the same side as the binormal tangent) when the torsion is positive, and on the negative side otherwise. The torsion is of course identically zero for planar curves.

5.1.2 Surfaces

Most of the discussion of the local characteristics of plane and space curves can be generalized in a simple manner to surfaces. Consider a point P on the surface S and all the curves passing through P and lying on S. It can be shown that the tangents to these curves lie in the same plane Π, appropriately called the *tangent plane* in P (Figure 5.11(a)). The line N passing through P and perpendicular to Π is called the *normal line* to P in S, and the surface can be oriented (locally) by picking a sense for a unit *normal vector* along N (unlike curves, surfaces admit a single normal but an infinity of tangents at every point). The surface bounding a solid admits a canonical orientation defined by letting the normal vectors locally point toward the outside of the solid.[2]

Intersecting a surface with the planes that contain the normal in P yields a one-parameter family of planar curves, called *normal sections* (Figure 5.11(b)). These curves are in general regular in P, or they may exhibit an inflection there. The curvature of a normal section is called the *normal curvature* of the surface in the

[2]Of course the reverse orientation, where, as Koenderink [1990, p. 137] puts it, "the normal vector points into the 'material' of the blob like the arrows in General Custer's hat", is just as valid. The main point is that either choice yields a coherent global orientation of the surface. Certain pathological surfaces (e.g., Möbius strips) do not admit a global orientation, but they do not bound solids.

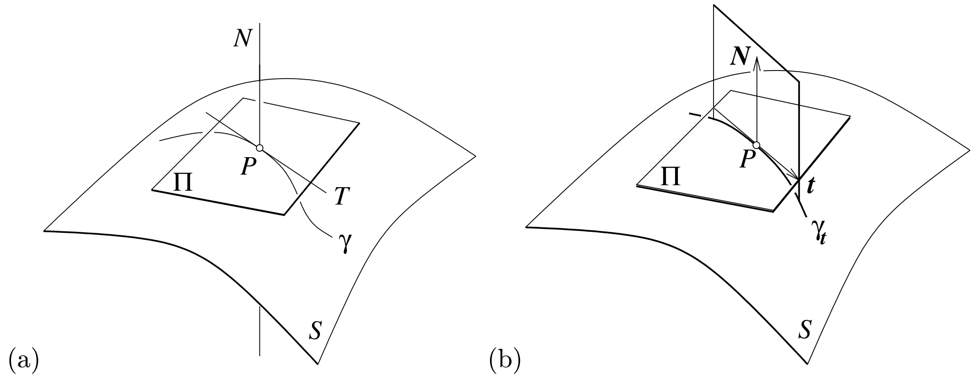

Figure 5.11. Tangent plane and normal sections: (a) the tangent plane Π and the associated normal line N at a point P of a surface; γ is a surface curve passing through P, and its tangent line T lies in Π; (b) the intersection of the surface S with the plane spanned by the normal vector $\boldsymbol{N}$ and the tangent vector $\boldsymbol{t}$ forms a normal section $\gamma_{\boldsymbol{t}}$ of S.

associated tangent direction. By convention, we will choose a positive sign for the normal curvature when the normal section lies (locally) on the same side of the tangent plane as the inward-pointing surface normal, and a negative sign when it lies on the other side. The normal curvature is of course zero when P is an inflection of the corresponding normal section.

With this convention, we can record the normal curvature as the sectioning plane rotates about the surface normal. It will (in general) assume its maximum value κ_1 in a definite direction of the tangent plane, and reach its minimum value κ_2 in a second definite direction. These two directions are called the *principal directions* in P and it can be shown that, unless the normal curvature is constant over all possible orientations, they are orthogonal to each other (see exercises). The *principal curvatures* κ_1 and κ_2 and the associated directions define the best local quadratic approximation of the surface: in particular, we can set up a coordinate system in P with x and y axes along the principal directions and z axis along the outward-pointing normal; the surface can be described (up to second order) in this frame by the paraboloid $z = -1/2(\kappa_1 x^2 + \kappa_2 y^2))$.

The neighborhood of a surface point can locally take three different shapes, depending on the sign of the principal curvatures (Figure 5.12). A point P where both curvatures have the same sign is said to be *elliptic*, and the surface in its vicinity is cup-shaped (Figure 5.12(a)): it does not cross its tangent plane and looks like the surface of an egg (positive curvatures) or the inside surface of its broken shell (negative curvatures). We say that P is convex in the former case and concave in the latter one. When the principal curvatures have opposite signs, we have a *hyperbolic* point. The surface is locally saddle-shaped and crosses its tangent plane along two curves (Figure 5.12(b)). The corresponding normal sections have an inflection in P and their tangents are called the *asymptotic directions* of the surface in P. They are

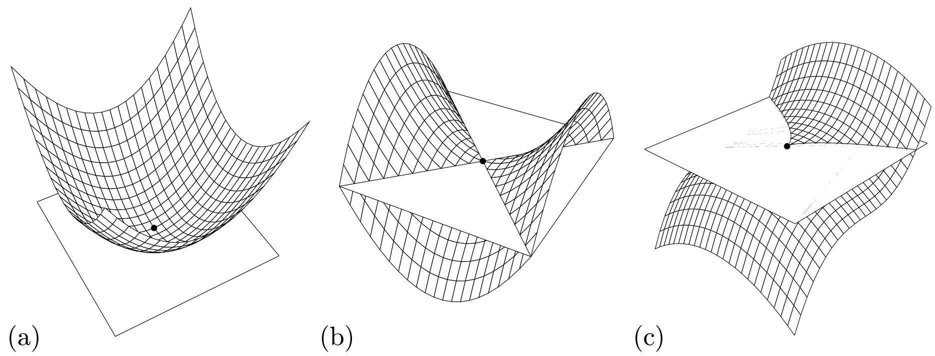

Figure 5.12. Local shape of a surface: (a) an elliptic point, (b) a hyperbolic point, and (c) a parabolic point (there are actually two distinct kinds of parabolic points [Koenderink, 1990]; we will come back to those in Chapter 23).

bisected by the principal directions. The elliptic and hyperbolic points form patches on a surface. These areas are in general separated by curves formed by *parabolic* points where one of the principal curvature vanishes. The corresponding principal direction is also an asymptotic direction, and the intersection of the surface and its tangent plane has (in general) a cusp in that direction (Figure 5.12(c))

Naturally, we can define the Gaussian image of a surface by mapping every point onto the place where the associated unit normal pierces the unit sphere. In the case of plane curves, the Gauss map is one-to-one in the neighborhood of regular points, but the direction of traversal of the Gaussian image reverses in the vicinity of certain singularities. Likewise, it can be shown that the Gauss map is one-to-one in the neighborhood of elliptic or hyperbolic points. The orientation of a small closed curve centered at an elliptic point is preserved by the Gauss map, but the orientation of a curve centered at a hyperbolic point is reversed (Figure 5.13).

The situation is a bit more complex at a parabolic point: in this case, any small neighborhood will contain points with parallel normals, indicating a double covering of the sphere near the parabolic point (Figure 5.13): we say that the Gaussian image *folds* along the parabolic curve. Note the similarity with inflections of planar curves.

Let us now consider a surface curve γ passing through P, parameterized by its arc length s in the neighborhood of P. Since the restriction of the surface normal to γ has constant (unit) length, its derivative with respect to s lies in the tangent plane in P. It is easy to show that the value of this derivative only depends on the unit tangent $\boldsymbol{t}$ to γ and not on γ itself. Thus we can define a mapping $d\boldsymbol{N}$ that associates with each unit vector $\boldsymbol{t}$ in the tangent plane in P the corresponding derivative of the surface normal (Figure 5.14). Using the convention $d\boldsymbol{N}(\lambda\boldsymbol{t}) \stackrel{\text{def}}{=} \lambda d\boldsymbol{N}(\boldsymbol{t})$ when $\lambda \neq 1$, we can extend $d\boldsymbol{N}$ to a linear mapping defined over the whole tangent plane and called the *differential of the Gauss map* in P.

The *second fundamental form* in P is the bilinear form that associates with any

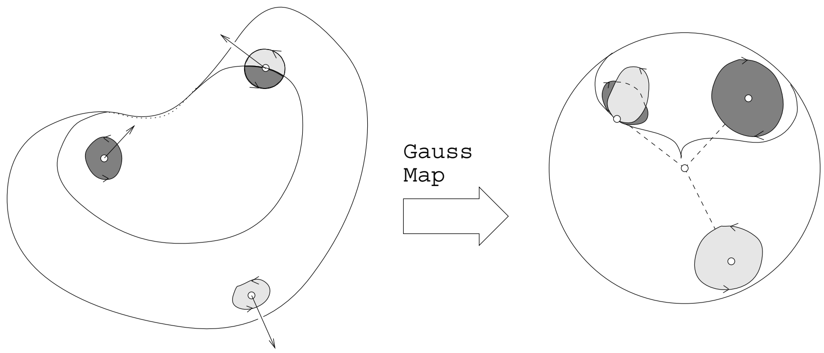

Figure 5.13. Left: a surface in the shape of a kidney bean. It is formed of a convex area, a hyperbolic region, and the parabolic curve separating them. Right: the corresponding Gaussian image. Darkly shaded areas indicate hyperbolic areas, lightly shaded ones indicate elliptic ones. Note that the bean is not convex but does not have any concavity.

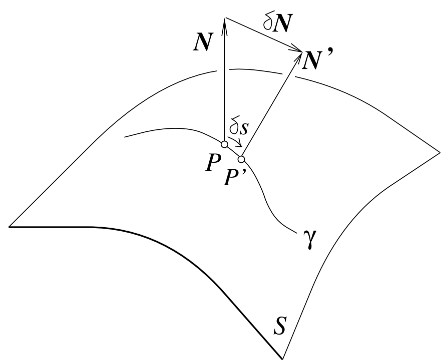

Figure 5.14. The directional derivative of the surface normal: if P and P' are nearby points on the curve γ, and $\boldsymbol{N}$ and $\boldsymbol{N}'$ denote the associated surface normals, with $\delta\boldsymbol{N} = \boldsymbol{N}' - \boldsymbol{N}$, the derivative is defined as the limit of $\delta\boldsymbol{N}/\delta s$ as the length δs of the curve arc separating P and P' tends toward zero.

two vectors $\boldsymbol{u}$ and $\boldsymbol{v}$ lying in the tangent plane the quantity[3]

$$\mathrm{II}(\boldsymbol{u}, \boldsymbol{v}) \stackrel{\text{def}}{=} \boldsymbol{u} \cdot d\boldsymbol{N}(\boldsymbol{v}).$$

[3] The second fundamental form is sometimes defined as $\mathrm{II}(\boldsymbol{u}, \boldsymbol{v}) = -\boldsymbol{u} \cdot d\boldsymbol{N}(\boldsymbol{v})$ (see, for example [do Carmo, 1976; Struik, 1988]). Our definition allows us to assign positive normal curvatures to the surfaces bounding convex solids with outward-pointing normals.

Application: The shape of specularities

Specularities offer hints about the colour of the illuminant, which we shall exploit in Chapter **??**. They also offer cues to the local geometry of a surface. Understanding these cues is a simple exercise in differential geometry that will serve to illustrate the concepts introduced in this chapter. We consider a smooth specular surface and assume that the radiance reflected in the direction $\boldsymbol{V}$ is a function of $\boldsymbol{V} \cdot \boldsymbol{P}$, where $\boldsymbol{P}$ is the specular direction. We expect the specularity to be small and isolated, so we can assume that the source direction $\boldsymbol{S}$ and the viewing direction $\boldsymbol{V}$ are constant over its extent. Let us further assume that the specularity can be defined by a threshold on the specular energy, i.e., $\boldsymbol{V} \cdot \boldsymbol{P} \geq 1 - \varepsilon$ for some constant ε, denote by $\boldsymbol{N}$ the unit surface normal, and define the **half-angle direction** as $\boldsymbol{H} = (\boldsymbol{S} + \boldsymbol{V})/2$ (Figure 5.15(left)). Using the fact that the vectors $\boldsymbol{S}$, $\boldsymbol{V}$ and $\boldsymbol{P}$ have unit length and a whit of plane geometry, it can easily be shown that the boundary of the specularity is defined by (see exercises):

$$1 - \varepsilon = \boldsymbol{V} \cdot \boldsymbol{P} = 2\frac{(\boldsymbol{H} \cdot \boldsymbol{N})^2}{(\boldsymbol{H} \cdot \boldsymbol{H})} - 1. \tag{5.1.1}$$

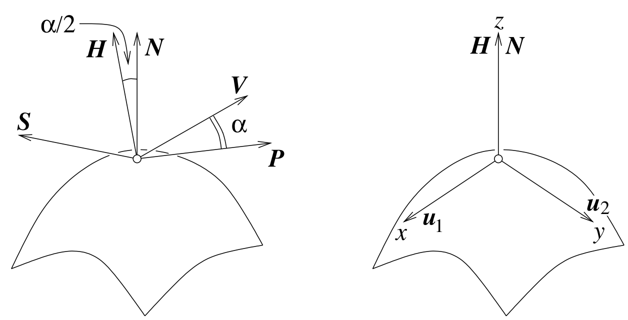

Figure 5.15. A specular surface viewed by a distant observer. We establish a coordinate system at the brightest point of the specularity (where the half-angle direction is equal to the normal) and orient the system using the normal and principal directions.

Because the specularity is small, the second-order structure of the surface will allow us to characterize the shape of its boundary as follows: there is some point on the surface inside the specularity (in fact, the brightest point) where $\boldsymbol{H}$ is parallel to $\boldsymbol{N}$. We set up a coordinate system at this point, oriented so that the z-axis lies along $\boldsymbol{N}$ and the x- and y-axes lie along the principal directions $\boldsymbol{u}_1$ and $\boldsymbol{u}_2$ (Figure 5.15(right)). As noted earlier, the surface can be represented up to second order as $z = -1/2(\kappa_1 x^2 + \kappa_2 y^2)$ in this frame, where κ_1 and κ_2 are the principal curvatures. Now, let us define a *parametric surface* as a differentiable mapping $\boldsymbol{x} : U \subset \mathbb{R}^2 \to \mathbb{R}^3$ associating with any couple $(u, v) \in U$ the coordinate vector $(x, y, z)^T$ of a point in some fixed coordinate system. It is easily shown (see exercises) that the normal to a parametric surface is along the vector $\partial\boldsymbol{x}/\partial u \times \partial\boldsymbol{x}/\partial v$. Our second-order surface model is a parametric surface parameterized by x and y, thus its unit surface normal is defined in the corresponding frame by

$$\boldsymbol{N}(x, y) = \frac{1}{\sqrt{1 + \kappa_1^2 x^2 + \kappa_2^2 y^2}} \begin{pmatrix} \kappa_1 x \\ \kappa_2 y \\ 1 \end{pmatrix},$$

and $\boldsymbol{H} = (0,0,1)^T$. Since $\boldsymbol{H}$ is a constant, we can rewrite (5.1.1) as $\kappa_1^2 x^2 + \kappa_2^2 y^2 = \zeta$ where ζ is a constant depending on ε. In particular, the shape of the specularity on the surface contains information about the second fundamental form. The specularity is an ellipse, with major and minor axes along the principal directions, and eccentricity equal to the ratio of the principal curvatures. Unfortunately, the shape of the specularity on the surface is not, in general, directly observable, so this property can only be exploited when a fair amount about the viewing and illumination setup is known [Healey and Binford, 1986]. While we cannot get much out of the shape of the specularity in the image, it is possible to tell a convex surface from a concave one by watching how a specularity moves as the view changes (you can convince yourself of this with the aid of a spoon). Let us consider a point source at infinity and assume that the specular lobe is very narrow so the viewing direction and the specular direction coincide. Initially, the specular direction is $\boldsymbol{V}$ and the specularity is at the surface point P; after a small eye motion, $\boldsymbol{V}$ changes to $\boldsymbol{V}'$ while the specularity moves to the close-by point P' (Figure 5.16).

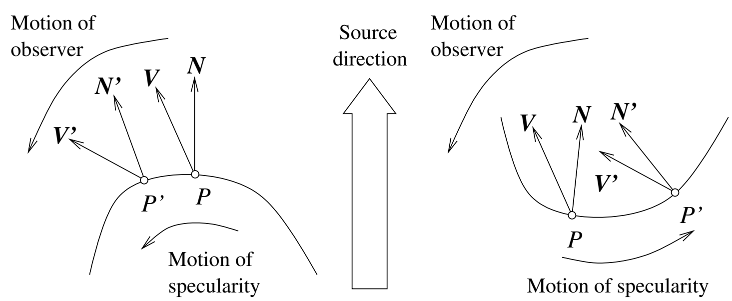

Figure 5.16. Specularities on convex and concave surfaces behave differently when the view changes.

The quantity of interest is $\delta a = (\boldsymbol{V}' - \boldsymbol{V}) \cdot \boldsymbol{t}$, where $\boldsymbol{t} = \frac{1}{\delta s}\overrightarrow{PP'}$ is tangent to the surface, and δs is the (small) distance between P and P': if δa is positive, then the specularity moves in the direction of the view (back of the spoon), and if it is negative, the specularity moves against the direction of the view (bowl of the spoon). By construction, we have $\boldsymbol{V} = 2(\boldsymbol{S} \cdot \boldsymbol{N})\boldsymbol{N} - \boldsymbol{S}$, and

$$\begin{aligned}\boldsymbol{V}' &= 2(\boldsymbol{S} \cdot \boldsymbol{N}')\boldsymbol{N}' - \boldsymbol{S} = 2(\boldsymbol{S} \cdot (\boldsymbol{N} + \delta\boldsymbol{N}))(\boldsymbol{N} + \delta\boldsymbol{N}) - \boldsymbol{S} \\ &= \boldsymbol{V} + 2(\boldsymbol{S} \cdot \delta\boldsymbol{N})\boldsymbol{N} + 2(\boldsymbol{S} \cdot \delta\boldsymbol{N})\delta\boldsymbol{N} + 2(\boldsymbol{S} \cdot \boldsymbol{N})\delta\boldsymbol{N},\end{aligned}$$

where $\delta\boldsymbol{N} \stackrel{\text{def}}{=} \boldsymbol{N}' - \boldsymbol{N} = \delta s \, d\boldsymbol{N}(\boldsymbol{t})$. Since $\boldsymbol{t}$ is tangent to the surface in P, ignoring second-order terms yields

$$\delta a = 2\delta s\, (\boldsymbol{S} \cdot \boldsymbol{N})\mathrm{II}(\boldsymbol{t}, \boldsymbol{t}).$$

Thus, for a concave surface, the specularity always moves against the view and for a convex surface it always moves with the view. The effect is clearly visible if you move a spoon back and forth and look at the images of light sources (of course you could also tell you are looking into the bowl of the spoon by just noticing you see yourself upside-down there). Things are more complex with hyperbolic surfaces; the specularity may move with the view, against the view, or perpendicular to the view (when $\boldsymbol{t}$ is an asymptotic direction).

Since II is easily shown to be symmetric, i.e., $\mathrm{II}(\boldsymbol{u},\boldsymbol{v}) = \mathrm{II}(\boldsymbol{v},\boldsymbol{u})$, the mapping that associates with any tangent vector $\boldsymbol{u}$ the quantity $\mathrm{II}(\boldsymbol{u},\boldsymbol{u})$ is a quadratic form. In turn, this quadratic form is intimately related to the curvature of the surface curves passing through P. Indeed, note that the tangent $\boldsymbol{t}$ to a surface curve is everywhere orthogonal to the surface normal $\boldsymbol{N}$. Differentiating the dot product of these two vectors with respect to the curve arc length yields

$$\kappa\boldsymbol{n}\cdot\boldsymbol{N} + \boldsymbol{t}\cdot d\boldsymbol{N}(\boldsymbol{t}) = 0,$$

where $\boldsymbol{n}$ denotes the principal normal to the curve, and κ denotes its curvature. This can be rewritten as

$$\mathrm{II}(\boldsymbol{t},\boldsymbol{t}) = -\kappa\cos\phi. \tag{5.1.2}$$

where ϕ is the angle between the surface and curve normals. For normal sections, we have $\boldsymbol{n} = \mp\boldsymbol{N}$, and it follows that the normal curvature in some direction $\boldsymbol{t}$ is

$$\kappa_{\boldsymbol{t}} = \mathrm{II}(\boldsymbol{t},\boldsymbol{t}),$$

where, as before, we use the convention that the normal curvature is positive when the principal normal to the curve and the surface normal point in opposite directions.

In addition, (5.1.2) shows that the curvature κ of a surface curve whose principal normal makes an angle ϕ with the surface normal is related to the normal curvature $\kappa_{\boldsymbol{t}}$ in the direction of its tangent $\boldsymbol{t}$ by $\kappa\cos\phi = -\kappa_{\boldsymbol{t}}$. This is known as Meusnier's theorem (Figure 5.17).

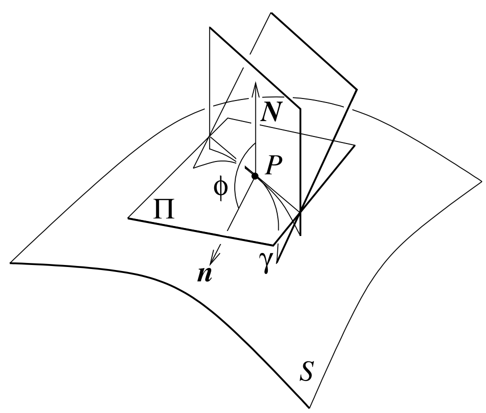

Figure 5.17. Meusnier's theorem.

It turns out that the principal directions are the eigenvectors of the linear map $d\boldsymbol{N}$, and the principal curvatures are the associated eigenvalues. The determinant K of this map is called the Gaussian curvature, and it is equal to the product of the principal curvatures. Thus, the sign of the Gaussian curvature determines the local shape of the surface: a point is elliptic when $K > 0$, hyperbolic when $K < 0$, and parabolic when $K = 0$. If δA is the area of a small patch centered in P on a surface S, and $\delta A'$ is the area of the corresponding patch of the Gaussian image of S, it can also be shown that the Gaussian curvature is the limit of the (signed) ratio $\delta A'/\delta A$ as both areas approach zero (by convention, the ratio is chosen to be positive when the boundaries of both small patches have the same orientation, and negative otherwise, see Figure 5.13). Note again the strong similarity with the corresponding concepts (Gaussian image and plain curvature) in the context of planar curves.

5.2 Contour Geometry

Before studying the geometry of surface outlines, let us pose for a minute and examine the relationship between the local shape of a twisted curve Γ and that of its orthographic projection γ onto some plane Π (Figure 5.18). Let us denote by α the angle between the tangent to Γ and the plane Π, and by β the angle betweem the principal normal to Γ and Π. These two angles completely define the local orientation of the curve relative to the image plane.

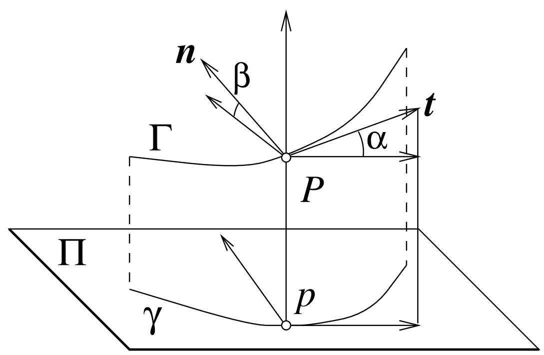

Figure 5.18. A space curve and its projection Note that the tangent to γ is the projection of the tangent to Γ (think for example of the tangent as the velocity of a particle traveling along the curve). The normal to γ is *not*, in general, the projection of the normal to Γ.

If κ denotes the curvature of Γ at some point, and κ_a denotes its *apparent curvature*, i.e., the curvature of γ at the corresponding image point, it is easy to show analytically (see exercises) that

$$\kappa_a = \kappa \frac{\cos \beta}{\cos^2 \alpha}. \tag{5.2.1}$$

In particular, when the viewing direction is in the osculating plane ($\cos \beta = 0$), the apparent curvature κ_a vanishes and the image of the curve acquires an inflection.

When, on the other hand, the viewing direction is tangent to the curve ($\cos\alpha = 0$), κ_a becomes infinite and the projection acquires a cusp.

These two properties of the projections of space curves are well known and mentioned in all differential geometry textbooks. Is it possible to relate in a similar fashion the local shape of the surface bounding a solid object to the shape of its image contour? The answer is a resounding "Yes!", as shown by Koenderink [1984] in his delightful paper "What does the occluding contour tell us about solid shape?," and we present in this section a few elementary properties of image contours, before stating and proving the main theorem of Koenderink's paper, and concluding by discussing some of its implications.

5.2.1 The Occluding Contour and the Image Contour

As noted earlier, the image of a solid bounded by a smooth surface is itself bounded by an image curve, called the contour, silhouette or outline of this solid. This curve is the intersection of the retina with a viewing cone whose apex coincides with the pinhole and whose surface grazes the object along a second curve, called the occluding contour, or rim (Figure 5.3).

We will assume orthographic projection in the rest of this section. In this case, the pinhole moves to infinity and the viewing cone becomes a cylinder whose generators are parallel to the (fixed) viewing direction. The surface normal is constant along each one of these generators, and it is parallel to the image plane (Figure 5.19). The tangent plane at a point on the occluding contour projects onto the tangent to the image contour, and it follows that the normal to this contour is equal to the surface normal at the corresponding point of the occluding contour.

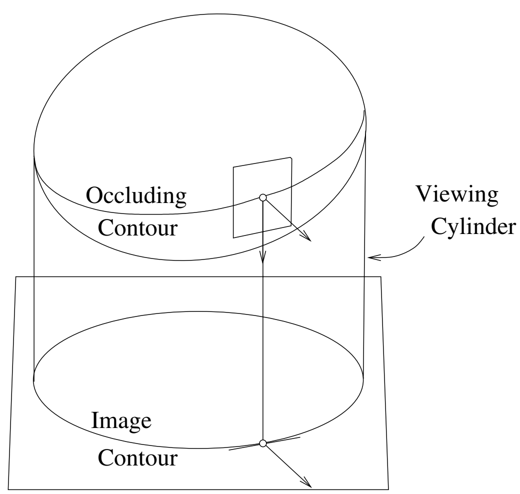

Figure 5.19. Occluding boundaries under orthographic projection.

It is important to note that the viewing direction $\boldsymbol{v}$ is *not*, in general, perpendicular to the occluding contour tangent $\boldsymbol{t}$ (as noted by Nalwa [1988] for example, the occluding contour of a tilted cylinder is parallel to its axis and not to the image plane). Instead, it can be shown that these two directions are *conjugated*, an extremely important property of the occluding contour.

5.2.2 The Cusps and Inflections of the Image Contour

Two directions $\boldsymbol{u}$ and $\boldsymbol{v}$ in the tangent plane are said to be conjugated when $\mathrm{II}(\boldsymbol{u}, \boldsymbol{v}) = 0$. For example, the principal directions are conjugated since they are orthogonal eigenvectors of $d\boldsymbol{N}$, and asymptotic directions are self-conjugated.

It is easy to show that the tangent $\boldsymbol{t}$ to the occluding contour is always conjugated to the corresponding projection direction $\boldsymbol{v}$: indeed, $\boldsymbol{v}$ is tangent to the surface at every point of the occluding contour, and differentiating the identity $\boldsymbol{N} \cdot \boldsymbol{v} = 0$ with respect to the arc length of this curve yields

$$0 = (\frac{d}{ds}\boldsymbol{N}) \cdot \boldsymbol{v} = d\boldsymbol{N}(\boldsymbol{t}) \cdot \boldsymbol{v} = \mathrm{II}(\boldsymbol{t}, \boldsymbol{v}).$$

Let us now consider a hyperbolic point P_0 and project the surface onto a plane perpendicular to one of its asymptotic directions. Since asymptotic directions are self-conjugated, the occluding contour in P_0 must run along this direction. As shown by (5.2.1), the curvature of the contour must be infinite in that case, and the contour acquires a cusp of the first kind.

We will state in a moment a theorem by Koenderink [1984] that provides a quantitative relationship between the curvature of the image contour and the Gaussian curvature of the surface. In the mean time, we will prove (informally) a weaker, but still remarkable result: *under orthographic projection, the inflections of the contour are images of parabolic points* (Figure 5.20).

First note that, under orthographic projection, the surface normal at a point on the occluding contour is the same as the normal at the corresponding point of the image contour. Since the Gaussian image of a surface folds at a parabolic point, it follows that the Gaussian image of the image contour must reverse direction at such a point. As shown earlier, the Gaussian image of a planar curve reverses at its inflections and cusps of the second kind. It is easily shown that the latter singularity does not occur for a general viewpoint, which proves the result.

In summary, the occluding contour is formed by points where the viewing direction $\boldsymbol{v}$ is tangent to the surface (the fold points mentioned in the introduction). Occasionally, it becomes tangent to $\boldsymbol{v}$ or crosses a parabolic line, and cusps (of the first kind) or inflections appear on the contour accordingly. Unlike the curves mentioned so far, the image contour may also cross itself (transversally) when two distinct branches of the occluding contour project onto the same image point, forming a T-junction (Figure 5.3). For general viewpoints, these are the only possibilities: there is no cusp of the second kind, nor any tangential self-intersection for example.

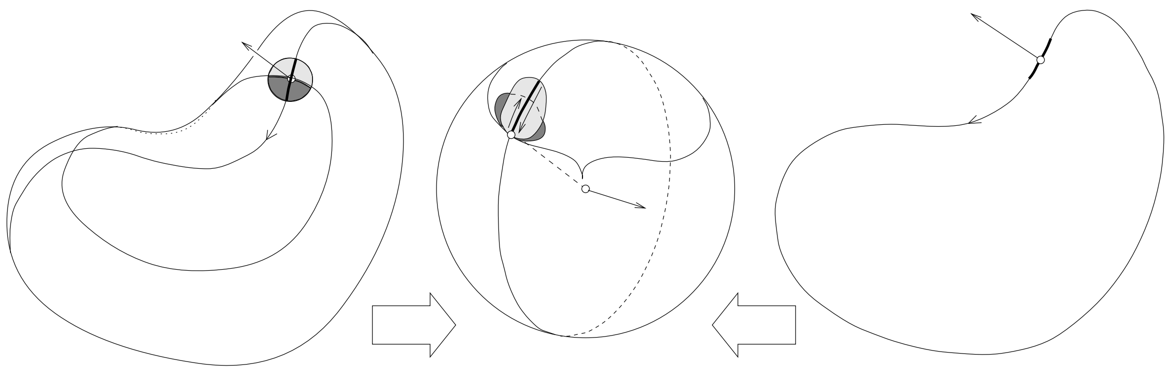

Figure 5.20. The inflections of the contour are images of parabolic points: the left side of this diagram shows the bean-shaped surface with an occluding contour overlaid, and its right side shows the corresponding image contour. The Gaussian image folds at the parabolic point, and so does its restriction to the great circle formed by the image of the occluding and image contours.

We will come back to the study of exceptional viewpoints and the corresponding contour singularities in a latter chapter.

5.2.3 Koenderink's Theorem

Let us now state the theorem by Koenderink [1984] that has already been mentioned several times. We assume as before orthographic projection, consider a point P on the occluding contour of a surface S, and denote by p its image on the contour.

Theorem 1: *The Gaussian curvature K of S in P and the contour curvature κ_c in p are related by*

$$K = \kappa_c \kappa_r,$$

where κ_r denotes the curvature of the radial curve *formed by the intersection of S with the plane defined by the normal to S in P and the projection direction (Figure 5.21).*

This remarkably simple relation has several important corollaries. Note first that κ_r remains positive (or zero) along the occluding contour since the projection ray locally lies inside the imaged object at any point where $\kappa_r < 0$. It follows that κ_c will be positive when the Gaussian curvature is positive, and negative otherwise. In particular, the theorem shows that convexities of the contour corresponds to elliptic points of the surface, while contour concavities correspond to hyperbolic points and contour inflections correspond to parabolic points.

Among elliptic surface points, it is clear that concave points never appear on the occluding contour of an opaque solid since their tangent plane lies (locally) completely inside this solid (Figure 5.21(b)). Thus convexities of the contour also

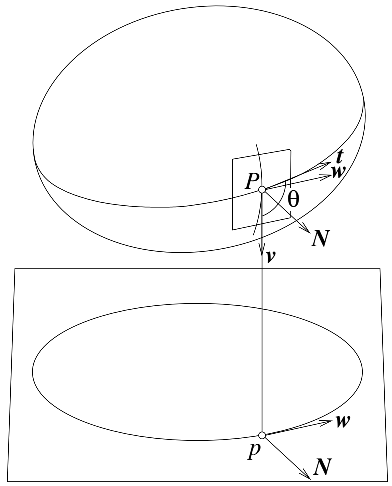

Figure 5.21. The relationship between occluding contour and image contour: the viewing direction $\boldsymbol{v}$ and the occluding contour tangent $\boldsymbol{t}$ are conjugated, and the radial curvature is always nonnegative at a visible point of the contour for opaque solids.

correspond to convexities of the surface. Likewise, we saw earlier that the contour cusps when the viewing direction is an asymptotic direction at a hyperbolic point. In the case of an opaque object, this means that concave arcs of the contour may terminate at such a cusp, where a branch of the contour becomes occluded. Thus we see that Koenderink's theorem strengthens and refines the earlier characterization of the geometric properties of image contours.

Let us now prove the theorem. It is related to a general property of conjugated directions: if $\kappa_{\boldsymbol{u}}$ and $\kappa_{\boldsymbol{v}}$ denote the normal curvatures in conjugated directions $\boldsymbol{u}$ and $\boldsymbol{v}$, and K denotes the Gaussian curvature, then

$$K \sin^2 \theta = \kappa_{\boldsymbol{u}} \kappa_{\boldsymbol{v}}, \tag{5.2.2}$$

where θ is the angle between $\boldsymbol{u}$ and $\boldsymbol{v}$. This relation is easy to prove by using the fact that the matrix associated with the second fundamental form is diagonal in the basis of the tangent plane formed by conjugated directions (see exercises). It is obviously satisfied for principal directions ($\theta = \pi/2$) and asymptotic ones ($\theta = 0$).

In the context of Koenderink's theorem, we obtain

$$K \sin^2 \theta = \kappa_r \kappa_{\boldsymbol{t}},$$

where $\kappa_{\boldsymbol{t}}$ denotes the normal curvature of the surface along the occluding contour direction $\boldsymbol{t}$ (which is of course different from the actual curvature of the occluding contour).

To complete the proof of the theorem, we use another general property of surfaces: the apparent curvature of any surface curve with tangent $\boldsymbol{t}$ is

$$\kappa_a = \frac{\kappa_{\boldsymbol{t}}}{\cos^2 \alpha}, \tag{5.2.3}$$

where α denotes as before the angle between $\boldsymbol{t}$ and the image plane. Indeed, we can use Meusnier's theorem to rewrite (5.2.1) as

$$\kappa_a = \kappa \frac{\cos\beta}{\cos^2 \alpha} = \frac{\kappa_{\boldsymbol{t}} \cos\beta}{-\cos^2 \alpha \cos\phi} = \frac{\kappa_{\boldsymbol{t}}}{\cos^2 \alpha}$$

since in this case $\phi = \beta + \pi$. In other words, the apparent cuvature of any surface curve is obtained by dividing the associated normal curvature by the square of the cosine of the angle between its tangent and the image plane. Applying this result to the occluding contour yields

$$\kappa_c = \frac{\kappa_{\boldsymbol{t}}}{\sin^2 \theta} \tag{5.2.4}$$

since $\alpha = \theta - \pi/2$. Substituting (5.2.4) into (5.2.2) concludes the proof of the theorem.

5.3 Notes

There is a rich literature on the three-dimensional interpretation of line-drawings, including the seminal work by Huffman [1971], Clowes [1971] and Waltz [1975], that uses the relatively small set of possible junction labels associated with typical scenes to control the combinatorial explosion of the line-labeling process. More quantitative approaches have also been proposed (see, for example [Draper, 1981; Kanade, 1981]), and the technique based on linear programming proposed by Sugihara [1984] is often considered as the ultimate achievement in this field. The relationship between contour and surface orientation is studied in [Barrow and Tenenbaum, 1981]. See also the work of Malik [1987] for an extension of the junction catalogue to piecewise-smooth surfaces.

There is of course a large number of excellent textbooks on differential geometry, including the very accessible presentations found in the books of do Carmo [1976] and Struik [1988]. Our presentation is closer in spirit (if not in elegance) to the descriptive introduction to differential geometry found in Hilbert's and Cohn-Vossen's wonderful book "Geometry and the Imagination" [Hilbert and Cohn-Vossen, 1952].

It was not always recognized that the image contour carries vital information about surface shape: see [Marr, 1977; Horn, 1986] for arguments to the contrary. The theorem proven in this chapter clarified the situation and appeared first in [Koenderink, 1984]. Our proof is different from the original one, but it is close in spirit to the proof given in [Koenderink, 1990], which is based on Blaschke's dual version of Euler's theorem. Our choice here was motivated by our reluctance to use any formulas that require setting a particular coordinate system. Alternate proofs for various kinds of projection geometries can be found in [Brady *et al.*, 1985; Arbogast and Mohr, 1991; Cipolla and Blake, 1992; Vaillant and Faugeras, 1992; Boyer, 1996].

5.4 Assignments

Exercises

1. Prove that the curvature κ of a planar curve in a point P is the inverse of the radius of curvature r at this point.

   Hint: use the fact that $\sin u \approx u$ for small angles.

2. Prove that, unless the normal curvature is constant over all possible directions, the principal directions are orthogonal to each other.

3. Prove that the second fundamental form is bilinear and symmetric.

4. Consider a fixed coordinate system, and a parametric surface $\boldsymbol{x} : U \subset \mathbb{R}^2 \rightarrow \mathbb{R}^3$. Show that the normal to a parametric surface is parallel to the vector $\partial \boldsymbol{x}/\partial u \times \partial \boldsymbol{x}/\partial v$.

   Hint: consider the two surface curves $u = u_0$ and $v = v_0$ passing through the point $\boldsymbol{x}(u_0, v_0)$.

5. Consider a specular surface observed by a distant observer. Denote by $\boldsymbol{V}$ the viewing direction, $\boldsymbol{P}$ the specular direction, and $\boldsymbol{S}$ the light source direction, all unit vectors. Define the **half-angle direction** $\boldsymbol{H} = (\boldsymbol{S} + \boldsymbol{V})/2$. Show that

   $$\boldsymbol{V} \cdot \boldsymbol{P} = 2\frac{(\boldsymbol{H} \cdot \boldsymbol{N})^2}{(\boldsymbol{H} \cdot \boldsymbol{H})} - 1,$$

   where $\boldsymbol{N}$ is the unit surface normal.

6. Let us denote by α the angle between the tangent to Γ and the plane Π, by β the angle betweem the principal normal to Γ and Π, and by κ the curvature of Γ at some point. Prove that if κ_a denotes its *apparent curvature*, i.e., the curvature of γ at the corresponding point, then

   $$\kappa_a = \kappa \frac{\cos \beta}{\cos^2 \alpha}.$$

   (Note: this is a variant of a result by Koenderink [1990, p. 191], who uses a different notation.)

7. Let $\kappa_{\boldsymbol{u}}$ and $\kappa_{\boldsymbol{v}}$ denote the normal curvatures in conjugated directions $\boldsymbol{u}$ and $\boldsymbol{v}$ at a point P, and let K denote the Gaussian curvature, prove that:

$$K \sin^2 \theta = \kappa_{\boldsymbol{u}} \kappa_{\boldsymbol{v}},$$

where θ is the angle between the $\boldsymbol{u}$ and $\boldsymbol{v}$.

Proof: Note that the second fundamental form can be written in the basis $(\boldsymbol{u}, \boldsymbol{v})$ as

$$\mathrm{II}(x\boldsymbol{u}+y\boldsymbol{v}, x\boldsymbol{u}+y\boldsymbol{v}) = (x, y) \begin{pmatrix} \mathrm{II}(\boldsymbol{u},\boldsymbol{u}) & \mathrm{II}(\boldsymbol{u},\boldsymbol{v}) \\ \mathrm{II}(\boldsymbol{u},\boldsymbol{v}) & \mathrm{II}(\boldsymbol{v},\boldsymbol{v}) \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = (x, y) \begin{pmatrix} \kappa_{\boldsymbol{u}} & 0 \\ 0 & \kappa_{\boldsymbol{v}} \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}.$$

In other words, the second fundamental form has a diagonal matrix in the basis $(\boldsymbol{u}, \boldsymbol{v})$. If $\boldsymbol{e}_1$ and $\boldsymbol{e}_2$ denote the principal directions, and $x\boldsymbol{u} + y\boldsymbol{v} = x'\boldsymbol{e}_1 + y'\boldsymbol{e}_2$, we have

$$(x', y')\,(\boldsymbol{e}_1, \boldsymbol{e}_2)^T\,(\boldsymbol{u}, \boldsymbol{v})^{-T} \begin{pmatrix} \kappa_{\boldsymbol{u}} & 0 \\ 0 & \kappa_{\boldsymbol{v}} \end{pmatrix} (\boldsymbol{u}, \boldsymbol{v})^{-1}\,(\boldsymbol{e}_1, \boldsymbol{e}_2) \begin{pmatrix} x' \\ y' \end{pmatrix}.$$

Since the Gaussian curvature is the determinant of the differential of the Gauss map, we have therefore

$$K = \frac{\kappa_{\boldsymbol{u}} \kappa_{\boldsymbol{v}}}{|\,(\boldsymbol{u}, \boldsymbol{v})\,|^2} = \frac{\kappa_{\boldsymbol{u}} \kappa_{\boldsymbol{v}}}{\sin^2 \theta}.$$

■

8. Show that the occluding is a smooth curve that does not intersect itself.

   Hint: use the Gauss map.