LIDS-P-1921

Curvatures of Surfaces and their Shadows

W. C. Karl *   G. C. Verghese †

July 24, 1989

Abstract

The relationship between the n-dimensional surfaces of smooth, strictly convex objects and the m-dimensional surfaces of their orthogonal projections, or shadows, is investigated. Our main results concern the relationships between the local properties of the surface at a point and those of its shadow. Specifically, the curvature Hessian of the projection at a boundary point is shown to be simply the projection of the curvature at the point's pre-image. Further, necessary and sufficient conditions are presented for solution of the inverse problem of determining the surface curvature at a point, given the curvatures of a series of projections involving the point. These conditions imply, for example, that knowledge of the projection of a surface point onto two hyperplanes and an additional two-dimensional subspace is necessary and sufficient to determine the local surface curvature there. These local results are then combined with a curvature based object representation on the Gaussian sphere to both construct the shadows of objects and to elucidate the inverse problem of reconstructing object shape from shadows. These results serve to illuminate and extend the work of Van Hove [1] for obtaining the two-dimensional shadow of an object in three-dimensional space.

1 Introduction

The relationship between objects and their shadows is of interest in a variety of disciplines and applications. These applications may be divided into two groups. First, there are those concerned with the forward problem of determining the shadow of an object on a surface given its shape. This is the goal, for example, in image synthesis for computer graphics [2,3] and automatic drawing generation for a part [4,5]. The second set of applications are concerned with the inverse problem of recovering the shape of an object given a series of its shadows. This recovery can be done either directly or indirectly. We can use the shadows to directly reconstruct a shape approximating the underlying generating one, or we can use

*Partially supported by fellowships from the Kodak Company, and the Center for Intelligent Control Systems (CICS) under the U.S. Army Research Office Grant DAAL03-86-K-0171.

†Partially supported by the National Science Foundation under Grant ECS-8700903, and by CICS. Address for correspondence: Room 10-069, MIT, Cambridge, MA 02139

our observations in a recognition framework to choose one from among a set of candidate objects, thus indirectly recovering the shape [6,7,1].

The generation of these shadow observations themselves arises in a variety of settings. Aside from directly obtained shadow images, projection data often contains little more information than the shadow of the object of interest [1]. This is true in certain instances of non-destructive testing, in low-dose X-ray images [8], in images of backlit objects, and in noisy range images [9,10,7,6], to name a few. In all these cases the structure of the projection image data is effectively reduced to a shadow. Mathematically, these shadow observations can be obtained through a number of projection geometries, including orthogonal projection, perspective projection, and spherical projection [11,12,13]. In this work we will only consider the orthogonal projection case.

Once obtained, multiple shadow observations can be utilized in different ways to constrain object shape. Sometimes they are used in a simple and direct manner, as in [14,15], where each shadow is used to find a bounding volume for the object and the combination of all such bounds forms an approximation to the object. Often, however, the shadows are used in conjunction with a representation of the object in some transformed space [13,16,17], such as the Gaussian sphere. These representations are chosen for their convenience when working with shadow information, and because they provide simple relationships between the object in question and its shadows. It is this latter approach that we take, using the shadow observations to update a curvature based representation of the object. Our approach is based on [1], which was the direct stimulus for this work, but is more general and more simple than [1].

To our knowledge, the existing shadow based approaches confine themselves to objects in three dimensions (or less) and shadows in two dimensions (or less). We can imagine, however, situations where such constraints might be restrictive, and the present formulation, with its dimensional generality, could be of use. Any reconstruction problem involving a three-dimensional object that evolves in time is inherently four-dimensional, such as imaging of the beating heart [18, pp. 275]. Indeed, much work related to the Radon transform is already done in such a general dimensional setting [19,20]. Another context in which higher dimensional results may be of interest is suggested by the problem of finding a joint probability distribution function for a set of n random variables given the distributions of linear combinations of the variables. The original distribution is an n-dimensional density (object) and the observations are integral projections onto lower dimensional subspaces, thus fitting the tomography mold [21, pp. 365]. In fact, a result of probability theory essentially states that the collection of *all* such line integral projections suffices to uniquely determine a general density [22, pp. 335]. Shadows may arise either directly in these examples or else, as described earlier, because noisy data or approximation leads to observations that are effectively shadows.

Finally we mention that our results in Section 2.2 on recovering a symmetric (curvature) matrix H from projections essentially provide a solution to the problem of recovering an n-dimensional ellipsoid from its lower dimensional shadows. If H represents an ellipsoid via

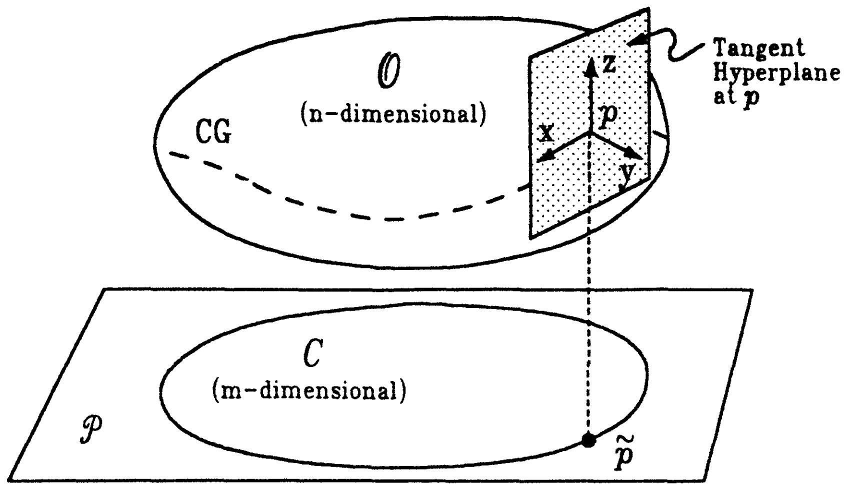

Figure 1: Problem definition

the equation $x^T H x = 1$, then our Result 1 gives the relationship between the ellipsoid and its projections [23] and Result 2 gives the conditions for its reconstruction. Such ellipsoid reconstruction problems have appeared in the literature, both directly and as bounding approximations. In [24] an ellipse is used as a simple parameterized model for objects in an attempt to recover their eccentricity and orientation from low signal to noise tomographic data. In [25,23] the state of a system is assumed confined to an unknown n-dimensional ellipsoid. The goal here is essentially to reconstruct this ellipsoid from observations of its lower dimensional projections. Finally, in [26] a group of closely spaced targets in space is observed through a number of passive sensors. The cluster of targets is modeled as an ellipsoid and the observations as projections of it over time. The desire is to find the evolution of the 3-D shape of the ellipsoid.

Mathematical Preliminaries

We shall be concerned with the orthogonal projection of the n-dimensional surface $\mathcal{O}$ of an object in $(n+1)$-dimensional space onto an $(m+1)$-dimensional projection subspace $\mathcal{P}$, to obtain a shadow with an m-dimensional surface, see Figure 1. A guide to notation is presented in Appendix A. In general, our attention shall be restricted to smooth, strictly convex objects (termed "rotund" [27]). We will represent the curvature of the surface $\mathcal{O}$ at a point p by the symmetric Hessian matrix at the point (notation $H_{\mathcal{O}}(p)$), i.e., by the matrix of second partial derivatives of the surface in some local tangent-based coordinate system [28], as illustrated below.

The boundary of the shadow in the subspace $\mathcal{P}$ will be a "curve" (actually a surface in $\mathcal{P}$), which we label $\mathcal{C}$. This lower dimensional surface, in turn, will be the shadow of a "curve," termed the contour generator (CG) [13, pp. 106], on the surface $\mathcal{O}$ of the object.

The points of the CG are thus precisely those that map to the boundary $\mathcal{C}$ under projection. If p is such a point on the CG, we shall label its image in $\mathcal{C}$ by $\tilde{p}$. In addition, we assume the dimension of $\mathcal{P}$ is at least 2, so that the curvature of $\mathcal{C}$ is well defined

In our main result we shall relate $H_{\mathcal{O}}(p)$ to $H_{\mathcal{C}}(\tilde{p})$. Specifically, in Section 2 we show that $H_{\mathcal{C}}(\tilde{p}) = (S^T\, H_{\mathcal{O}}^{-1}(p)\, S)^{-1}$, where the columns of S form an orthonormal basis for a subspace defined by $\mathcal{P}$ and the tangent space of $\mathcal{O}$ at p. This result serves to generalize and unify existing work, such as that in [1]. Following this solution of the forward or projection problem, we treat the inverse problem of *determining* the curvature of a surface at a point from a series of projections involving the point. Necessary and sufficient conditions are presented for recovering $H_{\mathcal{O}}^{-1}(p)$ from a series of observations $H_{\mathcal{C}i}^{-1}(\tilde{p})$, $1 \leq i \leq q$.

In Section 3 we combine these local results with a curvature based representation to make a series of global statements. A formalism based on the work in [1] will be given, allowing us to conveniently find the projection of an object onto an arbitrary subspace. These relationships will serve to generalize and clarify the work in [1,16] on silhouette determination in two-dimensions. We subsequently address issues arising in the inverse problem of overall recovery of object shape from shadows.

In Section 4 we pose some questions and raise some issues for future research. We give a brief discussion there of an alternative curvature representation scheme based on the Gaussian curvature rather than the Hessian. Concluding remarks are presented in Section 5.

2 Local Results

We shall show here that the curvature of the projection at a boundary point $\tilde{p}$ of the shadow is precisely the projection of the curvature at the corresponding point p on the contour generator of the object. We follow this local projection result with necessary and sufficient conditions for the solution of the inverse problem of determining the curvature of a surface at a point from a series of projections.

2.1 Projection of a Surface

The situation under consideration is depicted in Figure 2. Pick a coordinate frame in the tangent hyperplane $\mathcal{T}$ at p and label the n coordinate directions by $\hat{t}_i$. Complete this local frame by appending the local outward normal direction $N(p)$ as the $(n+1)$-st coordinate direction, $\hat{y}$. Let $(t_1, t_2, \ldots, t_n, y)$ be the scalar coordinates of points with respect to this frame. Locally the surface is then representable as (the Monge parameterization [29]) $y = F(t_1, \ldots, t_n)$. The curvature is taken to be the symmetric Hessian matrix $H_{\mathcal{O}}(p)$ of second partial derivatives of $F(t)$ with respect to these coordinates, i.e. the matrix whose ij-th entry is $\partial y/\partial t_i \partial t_j$.

For simplicity, we shall translate the projection subspace $\mathcal{P}$ parallel to itself to the point of interest p, as illustrated. Denote by $\mathcal{S}$ the m-dimensional intersection of $\mathcal{T}$ and the

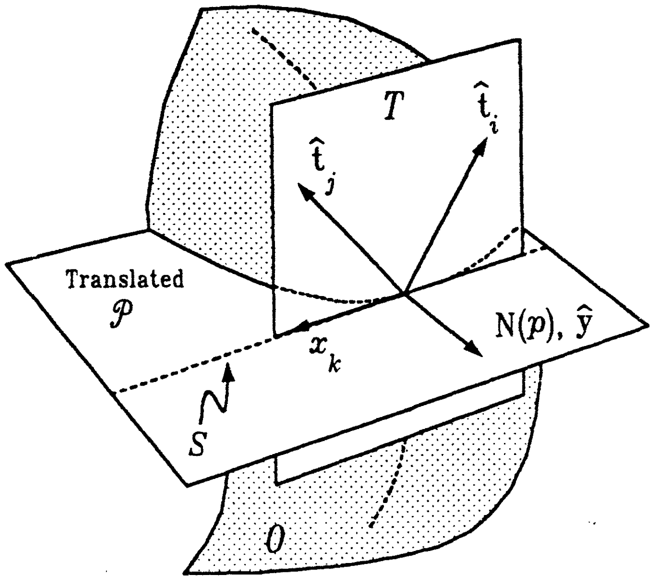

Figure 2: Local situation

translated $\mathcal{P}$, and let the m unit vectors $\hat{S}_i$ define a coordinate frame for this subspace of $\mathcal{T}$. Now, the $m+1$ vectors $\{\hat{S}_i, N(p)\}$ define a local coordinate frame in the *projection* of $\mathcal{O}$ at $\tilde{p}$. As above, let $(x_1, x_2, \ldots, x_m, y)$ be the scalar coordinates of points in the projection with respect to this frame. In the projection then, the curve $\mathcal{C}$ is locally representable as $y = f(x_1, \ldots, x_m)$. The curvature of the shadow $H_{\mathcal{C}}(\tilde{p})$ is the Hessian of $f(x)$ with respect to these coordinates. Finally, denote by S the $n \times m$ matrix whose columns are the representations of the vectors $\{\hat{S}_i\}$ with respect to the coordinate frame $\{\hat{t}_1, \ldots, \hat{t}_n\}$ we have defined in $\mathcal{T}$ at p. The columns of S are thus orthonormal. With this notation the main projection result is:

Result 1 (Curvature of Projection = Projection of Curvature)
The curvature Hessian of the orthogonal projection of a rotund surface at a boundary point is precisely the orthogonal projection of the curvature Hessian of the surface at the pre-image of this point: $H_{\mathcal{C}}(\tilde{p}) = (S^T\, H_{\mathcal{O}}^{-1}(p)\, S)^{-1}$.

The proof of Result 1 is given as Appendix B. Our formulation and proof are done in a general setting and are thus valid for any dimension, unlike the work in [1]. They are also considerably more transparent.

This result is a local one, involving only the surface properties at a point. It establishes a relationship between the second order term of the Taylor series approximation of the surface at p and the corresponding term for the projection. This simple relationship will be of use in Section 3 to find the complete shadow of an object.

2.2 The Inverse Problem – Curvature from Projections

The inverse problem of finding the curvature of a surface at a point from orthogonal projections of the surface at the point will now be examined. From Result 1, we see that the inverse problem may be viewed as the determination of the $n \times n$ symmetric matrix $H_O^{-1}(p)$ from a series of observations of curvature of the form $H_{Ci}^{-1}(\tilde{p}) = S_i^T H_O^{-1}(p) S_i$, where $H_{Ci}^{-1}(\tilde{p})$ is taken as the i-th observation and S_i is a known matrix whose orthonormal columns define the i-th subspace of projection, as described above. We shall cast this task as a standard linear estimation problem and present a simple necessary and sufficient condition for its solvability in terms of a rank test on a matrix derived from the set of S_i. Further, we shall bring out certain implications of this condition for the projection subspaces $\mathcal{P}_i$.

For convenience in what follows, we shall take our observations to be $G_i \equiv S_i\, H_{Ci}^{-1}(\tilde{p})\, S_i^T$ rather than $H_{Ci}^{-1}(\tilde{p})$. The latter can always be recovered from the former, since each S_i has full column rank. Again for simplicity, we shall let G denote $H_O^{-1}(p)$, the inverse of the Hessian of interest. Note that for strictly convex objects, the above inverses are certain to exist. With this notation the inverse problem may be phrased as follows.

Problem 1 (Local Curvature from Projections) *Determine the definite, symmetric, $n \times n$ matrix G, given q quadratic-form observations of the type:*

$$G_i \equiv P_i\, G\, P_i \qquad\qquad 1 \le i \le q \tag{1}$$

where the P_i are orthogonal projection matrices, i.e. are symmetric and satisfy $P_i^2 = P_i$.

In terms of the matrices $\{S_i\}$, we have $P_i \equiv S_i\, S_i^T$. Since the columns of the known $n \times m_i$ matrix S_i form an orthonormal basis for the subspace $\mathcal{S}_i$, P_i is an orthogonal projector onto this subspace and thus defines projection i. Our observations G_i are invariant to the particular bases S_i used for the subspaces $\mathcal{S}_i$.

As discussed in the introduction, Problem 1 is equivalent to recovering an ellipsoid from a series of its shadows (orthogonal projections) on the subspaces defined by the $\{S_i\}$. In the present notation, the positive definite symmetric matrix G represents an ellipsoid via the (support function) definition $\{x : x^T\eta \le \sqrt{\eta^T G\, \eta}\,,\ \forall\, \|\eta\| = 1\}$, and its projection yields the observations given in (1), as shown in [23]. In this framework, G_i can be seen to represent the (degenerate) ellipsoid obtained by projecting the ellipsoid represented by G.

Before we proceed, note that the problem is linear in the elements of the target matrix G. This fact can be demonstrated by writing (1) as:

$$\mathrm{vec}(G_i) = (P_i \otimes P_i)\,\mathrm{vec}(G) \tag{2}$$

where vec($\cdot$) denotes the vector formed by stacking the columns of the argument and $\otimes$ denotes the Kronecker product. The difficulty with this formulation is that the symmetry of the observations G_i and the target G is lost. While the vector vec(G) has n^2 elements, only $n(n+1)/2$ of them are independent. We shall overcome this difficulty by imbedding the problem in a natural way in the space of symmetric matrices. In this space, the target matrix G is an unknown vector and an observation becomes the projection of this vector onto a certain subspace.

Representation in Space of Symmetric Matrices

The set of $n \times n$ symmetric matrices together with the inner product $\langle A, B\rangle \equiv \mathrm{tr}(A^T B)$ defines an $n(n+1)/2$-dimensional Euclidean space which we denote by Ψ (note that this inner product induces the Frobenius norm on a matrix $\langle A, A\rangle^{1/2} = \|A\|_F$). Suppose $\{M_\ell \,|\, 1 \le \ell \le n(n+1)/2\}$ is an orthonormal basis for this space. The symmetric matrices G, $P_i M_j P_i$ (which we shall use shortly), and G_i may then be uniquely represented with respect to the given basis by the following vectors of coefficients:

$$\begin{aligned} G &\xmapsto{\{M_\ell\}} \gamma &&\equiv \left[\gamma_1, \gamma_2, \cdots, \gamma_{\frac{n(n+1)}{2}}\right]^T \\ P_i M_j P_i &\xmapsto{\{M_\ell\}} \wp(i)_j &&\equiv \left[\wp(i)_{j1}, \wp(i)_{j2}, \cdots, \wp(i)_{j\frac{n(n+1)}{2}}\right]^T \\ G_i &\xmapsto{\{M_\ell\}} g(i) &&\equiv \left[g(i)_1, g(i)_2, \cdots, g(i)_{\frac{n(n+1)}{2}}\right]^T \end{aligned} \tag{3}$$

where $\gamma_\ell = \langle G, M_\ell\rangle$, $\wp(i)_{j\ell} = \langle P_i M_j P_i, M_\ell\rangle$, and $g(i)_\ell = \langle G_i, M_\ell\rangle$ for $1 \le i \le q, 1 \le j, \ell \le n(n+1)/2$. With these definitions, it is straightforward to show that:

$$g(i) = \mathbf{P}_i \gamma \tag{4}$$

where

$$\mathbf{P}_i \equiv \left[\wp(i)_1 | \wp(i)_2 | \cdots | \wp(i)_{\frac{n(n+1)}{2}}\right]$$

It is also straightforward, though tedious, to show that $\mathbf{P}_i \mathbf{P}_i = \mathbf{P}_i$ and $\mathbf{P}_i = \mathbf{P}_i^T$, so that $\mathbf{P}_i$ is an orthogonal projector [30] onto the subspace of Ψ spanned by the matrices $P_i M_j P_i$. Thus in the space Ψ, the i-th observation is the vector $g(i)$ obtained as the projection of the (unknown) vector γ onto a subspace specified by $\mathbf{P}_i$ (or equivalently S_i).

We may now stack up all the observations (4) into a single vector to obtain the following overall relation:

$$\begin{bmatrix} g(1) \\ \hline g(2) \\ \hline \vdots \\ \hline g(q) \end{bmatrix} = \begin{bmatrix} \mathbf{P}_1 \\ \hline \mathbf{P}_2 \\ \hline \vdots \\ \hline \mathbf{P}_q \end{bmatrix} \gamma \tag{5}$$

so

$$\mathbf{g} = \mathbf{P}\,\gamma \tag{6}$$

where g and $\mathbf{P}$ are defined in the obvious way. With this formulation, Problem 1 becomes: find the unknown vector γ given the observations in g and the observation geometry specified by $\mathbf{P}$. We have thus phrased the problem as one in standard linear estimation.

Necessary and Sufficient Conditions for Solvability

The formulation in (6) immediately allows us to characterize the solutions of Problem 1. Specifically, a unique solution to the problem exists if and only if the null space of the mapping in (6) is empty. This, in turn, is true if and only if the matrix $\mathbf{P}$ has full column rank ($= n(n+1)/2$). Note that an overdetermined set will be consistent since we have assumed no noise in our measurements. We phrase this condition formally:

Result 2 (General Condition for Solvability) *Problem 1 has a unique solution if and only if the matrix* $\mathbf{P}$ *defined in Equation 6 has rank equal to* $n(n+1)/2$ *(i.e. full column rank). This solution, if it exists, is given by*

$$\gamma = \mathbf{P}^L \mathbf{g} \tag{7}$$

where $\mathbf{P}^L$ *is any left inverse of* $\mathbf{P}$, g *is the vector of observations, and* γ *is the representation of* G *with respect to the basis* $\{M_\ell\}$.

This result provides us with a way to test if a given series of projections, defined by P_i (and thus S_i), is sufficient to determine G, or if others are needed. Later we shall present some conditions phrased *directly* in terms of the subspaces $\mathcal{P}_i$. Note that in the noise free case, the properties of the matrices $\mathbf{P}_i$ imply that $\gamma = \left(\sum_{i=1}^q \mathbf{P}_i\right)^{-1}\left(\sum_{i=1}^q g_i\right)$.

In the noisy case, we may easily find the linear least squared error (LLSE) solution to an inconsistent set of equations given in the form (6). If we choose $\mathbf{P}^L = \mathbf{P}^+$, the Moore-Penrose inverse of $\mathbf{P}$, the equation (7) will yield the LLSE solution to Problem 1 (without the definiteness constraint). Note that we may also implement this least squares solution recursively (via recursive linear least squares), updating our current estimate of the vector γ (and thus of the matrix G) as more observations become available. In this vein, we might imagine using a recursive formulation to track the changing curvature of a *dynamically evolving* object. Such results will be described in a separate paper.

Projection Space Conditions

In this section, some conditions for the reconstruction of G will be stated *directly* in terms of the projection subspaces $\mathcal{P}_i$. These conditions are corollaries of Result 2 but, being phrased in terms of the subspaces of projection, do not require computation of the matrix $\mathbf{P}$. We shall give conditions for projections onto combinations of hyperplanes and two-dimensional subspaces (true planes, which are the smallest projection spaces that yield shadows with well defined curvature).

In [1, pp. 185] the result is given that for objects in three-dimensions ($n = 2$), a minimum of three projections onto planes is required to recover the curvature at a point. We show here that the generalization of this result to higher dimensions is that projections onto two *hyper*planes and onto a single additional two-dimensional plane are required to find the curvature at a point. The result is:

Corollary 1 (Hyperplanes + 1) *Projection onto at least two hyperplanes and a single two-dimensional subspace (plane) is necessary to uniquely recover G. These projections will be sufficient provided the two-dimensional subspace is not contained in either hyperplane.*

This group constitutes a minimal set of observations to recover G, in the sense that any other non-trivial set will increase either the number of observations or the dimension of the observations. We outline the proof in Appendix C.

Now consider the case where the projection subspaces P_i are restricted to be hyperplanes. Applying Corollary 1 to this case, where the two-dimensional subspace referred to there is contained in a third hyperplane, yields the following:

Corollary 2 (Hyperplanes) *If the projection subspaces P_i are restricted to be hyperplanes, then at least three such projections are necessary to uniquely recover G, and three will be sufficient provided the hyperplanes are distinct.*

Other statements of this kind, combining Result 2 with different, specific combinations of projection subspace dimensions, are of course possible.

3 Global Statements

The local curvature relationships of Result 1 and Result 2, presented in the previous section, will now be combined with a curvature based representation of objects on an enhanced Gaussian sphere to make global statements about shadow determination and shape reconstruction. From the description in the introduction, it can be seen that finding the contour generator, and hence the projection of an object is, in general, a difficult task. In what follows, a formalism will be given that allows the projection of an object onto an arbitrary subspace to be found in a convenient fashion. This formalism serves to generalize and clarify the work in [1,16] on two-dimensional shadow determination. Following this solution to the forward problem, we will use the enhanced Gauss representation and Result 2 to make statements about the inverse problem of object reconstruction from shadows.

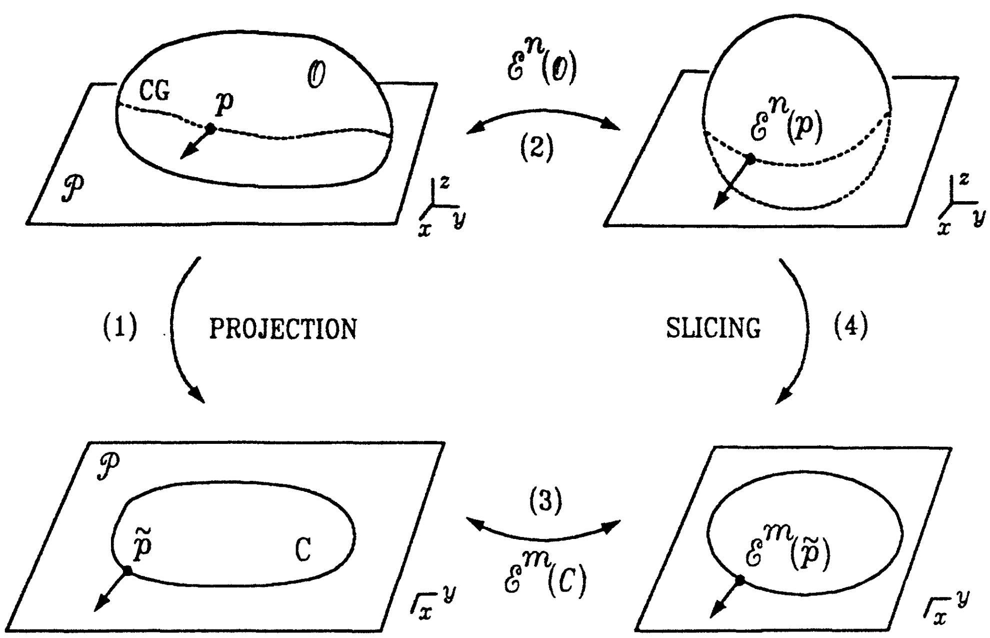

Figure 3: Projection and mapping relations

3.1 Projection of $\mathcal{O}$ onto $\mathcal{P}$

Consider the projection of the object surface $\mathcal{O}$ onto the subspace $\mathcal{P}$ to obtain the region bounded by the curve $\mathcal{C}$. This situation is shown as link (1) in Figure 3. The points p of $\mathcal{O}$ can be separated into two types with respect to the projection onto $\mathcal{P}$: those that project to interior points and those that project to boundary points $\mathcal{C}$. The first type correspond to points within the object's shadow while the second type correspond to points on the object's contour generator (CG). These latter points are precisely those points p of $\mathcal{O}$ where the normal $N(p)$ is parallel to the subspace of projection $\mathcal{P}$. The curve $\mathcal{C}$ may thus be found by first identifying the CG through the normal condition, then projecting this curve onto $\mathcal{P}$.

As can be seen, the CG and thus the curve $\mathcal{C}$ are not simple to find. In spite of this, we may show that the projection $\mathcal{C}$ inherits certain properties from the object surface $\mathcal{O}$. Specifically if the object surface $\mathcal{O}$ is smooth and strictly convex, it is not difficult to show that $\mathcal{C}$ will also be. This provides us an easier, indirect way to find the projection.

3.2 The Enhanced Gauss Images $\mathcal{E}^n(\mathcal{O})$ and $\mathcal{E}^m(\mathcal{C})$

With the insights above, we shall consider a representation of the object better suited to the task of projection than the standard point set or functional one [16]. Such a representation is found in an enhanced Gauss map of the object to the n-dimensional Gaussian sphere.

The Gaussian sphere is a unit sphere with each point on the sphere corresponding to points of an object with the same surface normal orientation. The points corresponding to the contour generator are thus easily found on the Gaussian sphere: they lie on the great circle set obtained by the *intersection* of the Gaussian sphere and the subspace $\mathcal{P}$ [1,31, pp. 536]. This enhanced Gauss mapping operation is shown as link (2) in Figure 3. We define the enhanced Gauss map formally as follows:

Definition 1 (Enhanced Gauss Map – $\mathcal{E}^n(\bullet)$) *The n-dimensional enhanced Gauss map of an object $\mathcal{O}$ (notation $\mathcal{E}^n(\mathcal{O})$) is the composition of the standard Gauss map [32] with a map of each object surface point p to the pair $\{H_{\mathcal{O}}(p), L(p)\}$, where $H_{\mathcal{O}}(p)$ is the Hessian of the surface at p in local coordinates, and $L(p)$ is the transformation from the global coordinate system to this local coordinate system in the tangent hyperplane at p.*

It can be shown that for a rotund hypersurface $\mathcal{O}$ in $\mathbf{R}^{n+1}$ with the associated enhanced Gauss image $\mathcal{E}^n(\mathcal{O})$, the image $\mathcal{E}^n(\mathcal{O})$ (via the curvature) determines $\mathcal{O}$ up to a translation [33,29,34]. Thus we may equivalently work with the enhanced Gauss image of an object as with the object itself. We note that while the object is uniquely defined by its enhanced Gauss image, inverting the image to recover the object is not trivial. In principle, this inversion can always be performed, given suitable initial conditions, but it is only recently that both iterative and closed-form algorithms have been proposed for this problem [35,36,1].

Since the boundary $\mathcal{C}$ inherits its rotundness from $\mathcal{O}$, we may also consider an equivalent representation for $\mathcal{C}$ in terms of its m-dimensional enhanced Gauss image $\mathcal{E}^m(\mathcal{C})$. As in the case of the object surface $\mathcal{O}$, this representation determines $\mathcal{C}$ up to a translation. This tie is shown as link (3) in Figure 3.

3.3 Obtaining $\mathcal{E}^m(\mathcal{C})$ from $\mathcal{E}^n(\mathcal{O})$

In this section, a direct tie is made between $\mathcal{E}^n(\mathcal{O})$ and $\mathcal{E}^m(\mathcal{C})$, the enhanced Gauss representations for $\mathcal{O}$ and $\mathcal{C}$ respectively. The image $\mathcal{E}^m(\mathcal{C})$ depends on the *local* curvature information of $\mathcal{C}$ at each point, which in turn depends only on the local surface shape of $\mathcal{O}$ at points along the contour generator. Since our representations are curvature based, we may thus go directly from $\mathcal{E}^n(\mathcal{O})$ to $\mathcal{E}^m(\mathcal{C})$ without explicitly finding the contour generator or the curve $\mathcal{C}$. The required local relationship between the object surface curvature and the projected surface curvature at a point is provided precisely in our Result 1. By isolating this observation we have focused on the essential element of the generalization.

To obtain $\mathcal{E}^m(\mathcal{C})$ from $\mathcal{E}^n(\mathcal{O})$ we thus need only to project the Hessians of points on the contour generator onto $\mathcal{P}$. These points, as mentioned earlier, are easily found from $\mathcal{E}^n(\mathcal{O})$ as the intersection of $\mathcal{E}^n(\mathcal{O})$ and $\mathcal{P}$. All that remains is to define the new transformation to local coordinates in the projection $\tilde{L}(\tilde{p})$, which is straightforward. Pick a global coordinate frame for the *projection* in the subspace $\mathcal{P}$. Let the columns of Π be the representation of these global projection axes with respect to the original global coordinate system, so the columns of Π form an orthonormal basis for $\mathcal{P}$ in global coordinates. The transformation from this set of axes in the projection to the local set at $\tilde{p}$, given by $\{\hat{S}_i, N(p)\}$ and defined in Section 2.1, is then given by $\tilde{L}(\tilde{p})$:

$$\tilde{L}(\tilde{p}) = \left[\begin{array}{c|c} 1 & \leftarrow 0 \rightarrow \\ \hline \uparrow & \\ 0 & S^T(p) \\ \downarrow & \end{array}\right] L(p)\ \Pi \tag{8}$$

where S is defined in Section 2.1.

We thus have the following two-step procedure to go directly from $\mathcal{E}^n(\mathcal{O})$ to $\mathcal{E}^m(\mathcal{C})$; this procedure is a generalization of the method in [1], in that it holds for any combination of object and projection dimension.

Procedure 1 ($\mathcal{E}^n(\mathcal{O})$ **to** $\mathcal{E}^m(\mathcal{C})$)

Step 1) *Identify the points of* $\mathcal{O}$ *on the contour generator by intersecting* $\mathcal{E}^n(\mathcal{O})$ *with the subspace* $\mathcal{P}$.

Step 2) *Use Result 1 to project onto* $\mathcal{P}$ *the curvature information* $H_{\mathcal{O}}(p)$ *at points p along the great circle set obtained in 1). Find the new transformation* $\tilde{L}(\tilde{p})$ *from global to local coordinates in the projection using (8).*

This is a slicing operation followed by a series of local projections of curvature. Note that the method just uses knowledge of $\mathcal{E}^n(\mathcal{O})$ and $\mathcal{P}$ to obtain $\mathcal{E}^m(\mathcal{C})$, so we may work entirely in the domain of the enhanced Gauss representation. This step is shown as link (4) in Figure 3. This final tie completes all the links in Figure 3 relating the various objects and representations. In the spirit of [16] we have the relationships shown in Figure 4.

3.4 The Inverse Problem and the Gauss Map

We will now examine the object reconstruction problem in terms of the enhanced Gauss representation for $\mathcal{O}$. In this framework, to determine $\mathcal{O}$ (within translation) we need to determine the curvature at each surface point, as discussed in Section 2.2. Each projection or shadow provides information about the surface along an m-dimensional great circle set of the Gaussian sphere representation, defined by the intersection of the Gaussian sphere with the projection subspace $\mathcal{P}$. Multiple projections of a point p correspond to places on the Gaussian sphere where these great circle sets intersect. For example, if the point p of an object in $\mathbf{R}^3$ were on the contour generator of two different projections, then it would lie in the intersection of two distinct great circles on the Gauss map of the object, as shown in Figure 5.

With this insight, we see that recovering the shape of an object may be viewed as a series of local curvature reconstructions at each point of the object. We view each projection as its

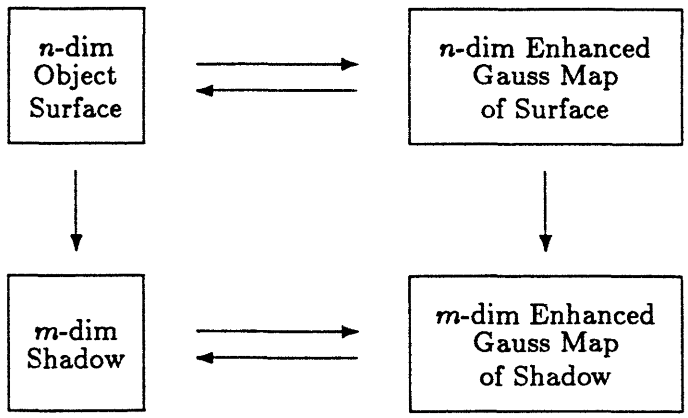

Figure 4: Relations between objects and representations

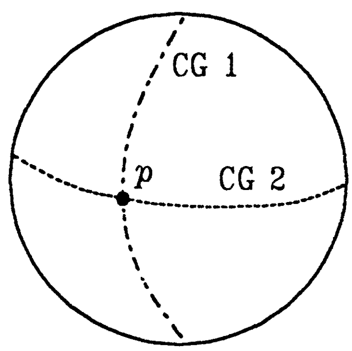

Figure 5: A point on two contour generators

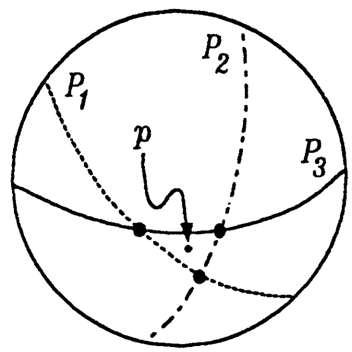

Figure 6: A point near the intersection of a number of projections

great circle set on the Gaussian sphere. The conditions in Section 2.2, specifically Result 2, imply a minimum number of projections involving a given point that are required to find the local surface curvature there. This, in turn, is equivalent to a condition on the number of intersections of great circle sets at the point. Specifically, for uniform projection dimension, to determine the Hessian at p we must have enough distinct great circles intersect at the image of p on the Gaussian sphere so that $\mathbf{P}$ in (7) is invertible. For example, we saw that three distinct hyperplane projections at a point were necessary and sufficient for curvature recovery.

Consider this case of projection onto hyperplanes. This is the largest, non-trivial, projection dimension possible, and is thus a best case in that it yields the most information for each projection. From Corollary 2 and generic intersection counting arguments, it is straightforward to show that three such projections will yield two $(n-3)$-dimensional sets of points on the n-dimensional surface $\mathcal{O}$ of the object, with enough information to recover the curvature. In general, there will also be six $(n-2)$-dimensional sets of points with two hyperplanes worth of information and three $(n-1)$-dimensional sets of points with one hyperplane worth of information. For example, in $\mathbf{R}^3$ where $n = 2$, this implies that projections onto three planes generically yields no points on the Gaussian sphere with three great circle intersections, six points on the Gaussian sphere where two great circles intersect, and three 1-dimensional sets of the great circles themselves. This generic situation is depicted in Figure 6.

For a finite number of projections, it can thus be seen that those points with the required number of intersections will be relatively sparse. In fact, we can see that each new projection only constrains an additional finite number of $(n-3)$-dimensional sets of points on the n-dimensional surface. In general, we will have to use nearby or interpolated projection data, as for the point p in Figure 6. In this case, we might use the inconsistent information from nearby projections together with a least squares solution to the local problem in (6) to estimate a value for the point. This raises the issue of the first order relationship between the orientation of the great circle sets and the corresponding conditioning of the matrix $\mathbf{P}$. It appears desirable to have the great circle sets intersect in angles as large as possible.

4 Questions

We pose some questions and raise some issues for future research. One issue is raised by the insight in Section 3.4 above. Rather than randomly taking projections of an object, we can use the Gaussian sphere to plan a series of views that will lead to the maximum number of intersections and coverage of the sphere. What is the best way to spread the information over the whole sphere, given only a finite number of views? And how can the information from near-intersections best be used, since in general there will be many points with strictly less than the necessary number of intersections. As discussed in [1], using global restrictions on object shape, or continuity, to derive the object from a *finite* set of surface curvatures gives rise to sampling questions on the Gaussian sphere and some type of surface Nyquist criterion.

In the above sections, we have represented the surface curvature by the Hessian matrix in some local coordinate system. This matrix is not the only way we may represent surface curvature. It can be shown that a rotund surface is actually determined (to within a translation) by just the determinant of the Hessian given for all normal orientations (see, for example [33, Vol. 5, pp. 304–305]). This determinant is termed the *Gaussian curvature* (notation $\mathcal{K}$) and is invariant with respect to changes in the system of local coordinates. The *scalar* function $\mathcal{K}(n)$ thus determines $\mathcal{O}$ (to a translation), and the function remains scalar regardless of the dimension of the space. This function defined on the Gaussian sphere (and hence a function of normal orientation) is called the extended Gaussian sphere and has been studied by Horn and others [17].

The use of the Gaussian curvature $\mathcal{K}$ as a representation of curvature in our enhanced Gauss mapping thus appears attractive. The key lies in the relationship between $\det[H_{\mathcal{O}}(p)]$ and $\det[H_{\mathcal{C}}(\tilde{p})]$, or more generally between the eigenvalues of a symmetric, definite H and those of $(S^T H^{-1} S)^{-1}$, where S has orthonormal columns (since $H_{\mathcal{C}}(\tilde{p}) = [S^T H_{\mathcal{O}}^{-1}(p) S]^{-1}$). Let $U \equiv [S|S^{\perp}]$, so U is orthogonal. A determinantal and matrix identity then yields the following simple relationship between the Gaussian curvature of the surface and that of its projection:

$$
\begin{aligned}
\text{Gauss curvature of surface} &= \det(H) = \det(U^T H U) \\
&= \det\left(S^T H^{-1} S\right)^{-1} \bullet \det\left(S^{\perp^T} H S^{\perp}\right) \\
&= (\text{Gauss curvature of projection}) \bullet \\
&\qquad (\text{Gauss curvature of } \textit{slice perpendicular} \text{ to projection})
\end{aligned}
$$

The relationship between the two Gaussian curvatures thus involves the extra term $\det(S^{\perp^T} H S^{\perp})$, arising from the Gaussian curvature of a curve obtained as a *slice* of the surface *perpendicular* to the projection defined by S. The problem is how to obtain this extra term given only the Gaussian curvature over the Gaussian sphere and the projection geometry? If we restrict our attention to a point p on the object, all we have is the single number $\mathcal{K}(p)$. This single number in isolation gives no directional information, and thus no way of distinguishing between projections onto different subspaces. It seems impossible to produce a *local* result (as in (1)) given only knowledge of $\mathcal{K}(p)$ and $\mathcal{P}$. We must somehow take into account surrounding information from the surface, perhaps looking at how the Gaussian curvature changes along the contour generator, or perpendicular to it. By looking at the Gaussian curvature of points perpendicular to the CG perhaps we could 'divide out' the curvature information in that direction. We leave these as questions for future work.

5 Conclusions

In summary, we have generalized and clarified results on the relationship between local surface curvature of an object and that of its orthogonal projection. In addition to the forward or projection problem, we have given necessary and sufficient conditions for solution of the inverse or curvature reconstruction problem. These conditions, in turn, allowed the formulation of statements directly in terms of the projection subspaces. The above local considerations were then used to generalize to arbitrary dimensions the representation and projection scheme found in [1,16]. Our setting and solution, in addition to clarifying existing work, also suggest directions for fruitful future investigation on problems such as reconstruction from finite and misaligned projections.

A Notation

n	Dimension of object surface.
m	Dimension of shadow surface.
$\mathcal{O}$	n-dimensional object surface.
$\mathcal{P}$ ($\mathcal{P}_i$)	$(m+1)$-dimensional ((m_i+1)-dimensional) projection subspace.
$\mathcal{H}_{\mathcal{O}}(p)$	The Hessian of $\mathcal{O}$ at p in local coordinates.
$\mathcal{C}$	Boundary of projection of $\mathcal{O}$ in $\mathcal{P}$.
CG	The contour generator. The curve on $\mathcal{O}$ whose image is $\mathcal{C}$.
p	Point on the CG of $\mathcal{O}$.
$\tilde{p}$	Image of p in $\mathcal{C}$.
$\mathcal{H}_{\mathcal{C}}(\tilde{p})$	The Hessian of $\mathcal{C}$ at $\tilde{p}$ in local coordinates.
$\mathcal{T}$	Tangent hyperplane to $\mathcal{O}$ at p.
$\{\hat{t}_1,..,\hat{t}_n\}$	Local coordinate directions at p in $\mathcal{T}$.
$N(p), \hat{y}$	Unit outward normal to $\mathcal{O}$ at p.
$(t_1,..,t_n,y)$	Coordinates of points with respect to the frame $\{\hat{t}_1,..,\hat{t}_n,\hat{y}\}$.
$F(t_1,..,t_n)$	Local representation of $\mathcal{O}$ as $y = F(t_1,..,t_n)$.
$\mathcal{S}$ ($\mathcal{S}_i$)	Subspace of intersection between $\mathcal{T}$ and $\mathcal{P}$ ($\mathcal{P}_i$).
$\{\hat{S}_i\}$	Set of coordinate vectors for $\mathcal{S}$.
$\{\hat{S}_i, N(p)\}$	Set of vectors defining a coordinate frame in $\mathcal{P}$ at p.
$(x_1,..,x_m,y)$	Coordinates of points with respect to the frame $\{\hat{S}_1,..,\hat{S}_m,N(p)\}$.
$f(x_1,..,x_m)$	Local representation of $\mathcal{C}$ as $y = f(x_1,..,x_m)$.
S (S_i)	$n \times m$ (m_i) matrix with columns representations of $\{\hat{S}_i\}$ in frame $\{\hat{t}_1,..,\hat{t}_n\}$.
G	The $n \times n$ target matrix $H_{\mathcal{O}}(p)$.
$G(i)$	The i-th observation, $P_i\, G\, P_i$.
P_i	Projector onto projection subspace i, $P_i \equiv SS^T$.
$\langle\cdot,\cdot\rangle$	The inner product defined on the space Ψ. $\langle A, B\rangle = \mathrm{tr}(A^T B)$.
Ψ	The space of symmetric $n \times n$ matrices with inner product $\langle\cdot,\cdot\rangle$.
$\{M_\ell\}$	Orthonormal symmetric basis for Ψ.
γ	Representation of G in basis $\{M_\ell\}$: $[\gamma_1,..,\gamma_{n(n+1)/2}]^T$; $\gamma_k = \langle G, M_k\rangle$.
$\wp(i)_j$	Representation of $P_i M_j P_i$ in basis $\{M_\ell\}$: $[\wp(i)_{j1},..,\wp(i)_{j\,n(n+1)/2}]^T$.
$\wp(i)_{j\ell}$	$\langle P_i M_j P_i, M_\ell\rangle$.
$\mathbf{P}_i$	$[\wp(i)_1, \wp(i)_2,..,\wp(i)_{n(n+1)/2}]^T$.
$g(i)$	Representation of $G(i)$ in basis $\{M_\ell\}$: $[g(i)_1,..,g(i)_{n(n+1)/2}]^T$; $g(i)_k = \langle G_i, M_k\rangle$.
$\mathbf{g}$	The overall observation vector: $[g(1),..,g(q)]^T$.
$\mathbf{P}$	The overall projection matrix: $[\mathbf{P}_1 \vert \mathbf{P}_2 \vert .. \vert \mathbf{P}_q]^T$.
$L(p)$	Transformation to local coordinates $(t_1,..,t_n,y)$ at p.
$\tilde{L}(\tilde{p})$	Transformation to local coordinates $(x_1,..,x_m,y)$ at $\tilde{p}$ in the projection.

B Proof of Result 1

In this appendix, we shall prove Result 1 relating the curvature of a projection to the curvature of the original surface: $H_C(\tilde{p}) = (S^T\, H_O^{-1}(p)\, S)^{-1}$.

The curvature we are discussing here is the Hessian of the surface or curve at a point in some local coordinate system. Recall, that the Hessian is defined as the matrix of second order partial derivatives of the surface, and may be viewed as generating a second order approximation to the surface at the point. In what follows, we shall assume we are working at a point p on the surface of an object and henceforth suppress any mention of that fact to simplify notation (e.g. all partial derivatives are assumed to be evaluated at p). We also assume all surfaces are oriented by their outward normal.

B.1 Curvature of the Projection

Assume we are at a point p on the surface that will map to a boundary point of the projection (so p is a point of the contour generator). Also assume that there is a local coordinate system at p oriented in the following way (see Figure 7). As in Section 2.1, $\hat{y}$ points along the local normal $N(p)$ and the set of m vectors $\{\hat{x}_i\}$ lies in a subspace parallel to the subspace of projection. We complete the set with $n - m$ vectors $\{\hat{z}_j\}$, which lie in a perpendicular subspace. Let y, x, and z, respectively, be vectors of coordinates along these axes. With this choice, we will be projecting along the z direction onto the x-y plane (the $\{\hat{x}_i\}$ and $\hat{y}$ set of axes together spanning the space of projection). We term these coordinates 'projection coordinates'. In this coordinate system we may represent the surface by the function $y = \tilde{F}(x, z)$. Now the boundary of the projection in these coordinates is defined by the set of equations:

$$y = \tilde{F}(x,z) \tag{9}$$
$$\tilde{F}_z = 0 \tag{10}$$

where letter subscripts denote vectors or matrices of partial derivatives (for example $\tilde{F}_z = [\tilde{F}_{z_1}, \tilde{F}_{z_2}, \cdots, \tilde{F}_{z_{n-m}}]$ - a row vector). The matrix of second partial derivatives y_{xx} derived subject to the constraints (9) and (10) is $H_C(\tilde{p})$, which it is our goal to find.

Start by taking the first partial derivative of (9) with respect to x, Applying the chain rule, this gives:

$$y_x = \tilde{F}_x + \tilde{F}_z z_x \tag{11}$$

where z_x is the (Jacobian) matrix of partial derivatives. Next, find the unconstrained second partial derivative matrix y_{xx} using (11) and the chain rule.

$$y_{xx} = \tilde{F}_{xx} + \tilde{F}_{xz} z_x + \left(\tilde{F}_{zx} + \tilde{F}_{zz} z_x\right) z_x + \tilde{F}_z \left(z_{xx} + z_{xz} z_x\right) \tag{12}$$

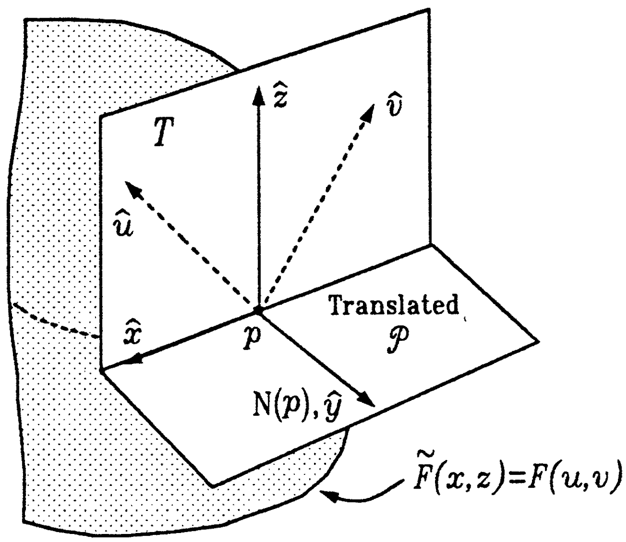

Figure 7: Local coordinate system configuration

Now from (10), $\tilde{F}_z = 0$, and taking the partial of this equation with respect to x we obtain:

$$\tilde{F}_{zx} + \tilde{F}_{zz} z_x = 0$$

which gives:

$$z_x = -\tilde{F}_{zz}^{-1} \tilde{F}_{zx} \tag{13}$$

where $\tilde{F}_{zz}^{-1}$ exists since the surface is strictly convex.

Applying (10) and (13) to (12), we get for the curvature of the projection:

$$y_{xx} = \tilde{F}_{xx} - \tilde{F}_{xz}\tilde{F}_{zz}^{-1}\tilde{F}_{zx} - \left(\tilde{F}_{xz} - \tilde{F}_{zz}\tilde{F}_{zz}^{-1}\tilde{F}_{zx}\right)\tilde{F}_{zz}^{-1}\tilde{F}_{zx}$$

or

$$y_{xx} = \tilde{F}_{xx} - \tilde{F}_{xz}\tilde{F}_{zz}^{-1}\tilde{F}_{zx} \equiv H_C(\tilde{p}) \tag{14}$$

B.2 Projection of the Curvature

We shall now relate $H_C(\tilde{p})$ to $H_O(p)$. Start by partitioning the set $(t_1, \ldots, t_n)$ of general coordinates in $\mathcal{T}$, defined in Section 2.1, so that $\mu = (t_1, \ldots, t_m)$ and $\nu = (t_{m+1}, \ldots, t_n)$. Schematically, we have the relationship between (x, z) and (μ, ν) shown in Figure 7. Given these definitions, we may always make the following association between these two sets of coordinates, by proper labeling:

$$\begin{bmatrix} \mu \\ \nu \end{bmatrix} = U \begin{bmatrix} x \\ z \end{bmatrix} \tag{15}$$

where $U = [S|S^{\perp}]$ is orthogonal and the columns of S are the representations of the $\{\hat{x}_i\}$ axes in (μ, ν) coordinates. The columns of S thus span the space of the projection in these coordinates. Now given the definition of F in Section 2.1, we have the following relationship between F and $\tilde{F}$:

$$F(\mu,\nu) = \tilde{F}(x,z).$$

It is straightforward to show that the Hessians of F and $\tilde{F}$ in the two coordinate systems are related by:

$$\begin{bmatrix} \tilde{F}_{xx} & \tilde{F}_{xz} \\ \tilde{F}_{zx} & \tilde{F}_{zz} \end{bmatrix} = U^T \begin{bmatrix} F_{\mu\mu} & F_{\mu\nu} \\ F_{\nu\mu} & F_{\nu\nu} \end{bmatrix} U \tag{16}$$

Now note that (see, for example, [23, pp. 539]):

$$\tilde{F}_{xx} - \tilde{F}_{xz}\tilde{F}_{zz}^{-1}\tilde{F}_{zx} = \left([\, I \,|\, 0 \,] \begin{bmatrix} \tilde{F}_{xx} & \tilde{F}_{xz} \\ \tilde{F}_{zx} & \tilde{F}_{zz} \end{bmatrix}^{-1} \left[\frac{I}{0} \right] \right)^{-1} \tag{17}$$

Substituting (16) into the right hand side yields:

$$\tilde{F}_{xx} - \tilde{F}_{xz}\tilde{F}_{zz}^{-1}\tilde{F}_{zx} = \left(S^T \begin{bmatrix} F_{\mu\mu} & F_{\mu\nu} \\ F_{\nu\mu} & F_{\nu\nu} \end{bmatrix}^{-1} S \right)^{-1} \tag{18}$$

By definition, the Hessian of $\mathcal{O}$ at p in the coordinate system (μ, ν) is:

$$H_{\mathcal{O}}(p) = \begin{bmatrix} F_{\mu\mu} & F_{\mu\nu} \\ F_{\nu\mu} & F_{\nu\nu} \end{bmatrix}. \tag{19}$$

Combining (14), (18), and (19) we obtain:

$$H_{\mathcal{C}}(\tilde{p}) = \left(S^T\, H_{\mathcal{O}}^{-1}(p)\, S \right)^{-1} \tag{20}$$

We have thus related $H_{\mathcal{C}}(\tilde{p})$ to $H_{\mathcal{O}}(p)$ through the projection type operation given in (20) and hence the result is proved. ∎

C Outline of Proof of Corollary 1

The proof of Corollary 1 proceeds as follows. First, we consider just the two hyperplanes defined by P_1 and P_2 and show that $\text{rank}[\mathbf{P}_1^T, \mathbf{P}_2^T]^T = n(n+1)/2 - 1$, just one independent row short of the required number. We then show that a single additional *planar* shadow, defined by P_3, adds this independent observation.

For the first part, note that the number of independent columns of $\mathbf{P} = [\mathbf{P}_1^T, \mathbf{P}_2^T]^T$ equals the number of independent matrices in the set $\{P_i M_j P_i \,|\, i = 1,2; j = 1,\ldots,n(n+1)/2\}$, where the $\{M_j\}$ are the elements of an orthonormal basis on Ψ. Thus, showing the first part of the proof may be reduced to a counting argument on the number of independent matrices in the given set, $\{P_i M_j P_i\}$.

We may, without loss of generality, align $n-1$ of the local axes $\{\hat{t}_i\}$ of $\mathcal{T}$ to lie in $\mathcal{S}_1$, say $\{\hat{t}_1,\ldots,\hat{t}_{n-1}\}$. Since $\mathcal{S}_1$ and $\mathcal{S}_2$ will intersect in an $(n-2)$-dimensional subspace in $\mathcal{T}$, we may further align $n-2$ of the above set of axes, say $\{\hat{t}_1,\ldots,\hat{t}_{n-2}\}$, with this intersection space, again without loss of generality. This choice of alignment for the local axes results in the following special form for the projectors $P_1 = S_1 S_1^T$ and $P_2 = S_2 S_2^T$:

$$P_1 = \begin{bmatrix} \mathrm{I}_{n-1} & 0 \\ 0 & 0 \end{bmatrix}, \qquad P_2 = \begin{bmatrix} \mathrm{I}_{n-2} & 0 \\ 0 & X \end{bmatrix}$$

where X is a 2×2, 1-dimensional orthogonal projector with $X_{22} \neq 0$. We couple these observations about the form of P_1 and P_2 with a convenient choice of symmetric basis M_j. Specifically we chose the appropriately normalized set of matrices $\{e_k e_\ell^T + e_\ell e_k^T; k, \ell = 1,\ldots,n\}$ where e_k is the k-th unit vector. The counting argument on the set $\{P_i M_j P_i\}$ is now straightforward, but tedious.

To show the second part, we observe that $P_3 = yy^T$ for some unit n-vector y, with $P_1 y \neq y$ and $P_2 y \neq y$. Thus $P_3 M_j P_3 = \alpha\, yy^T$ for some scalar α and we need only show that yy^T is independent of $\{P_i M_j P_i \,|\, i = 1,2; j = 1,\ldots,n(n+1)/2\}$. ■

References

[1] P. L. Van Hove. *Silhouette-Slice Theorems.* PhD thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, September 1986.

[2] F. C. Crow. Shadow algorithms for computer graphics. *Computer Graphics*, 11(2):242–248, Summer 1977.

[3] L. Williams. Casting curved shadows on curved surfaces. *Computer Graphics*, 12(3):270–274, August 1978.

[4] D. F. Rogers and R. A. Earnshaw, editors. *Techniques for Computer Graphics.* Springer-Verlag, New York, 1987. pp. 152.

[5] N. Magnenat-Thalmann and D. Thalmann, editors. *Computer-Generated Images: The State of the Art.* Springer-Verlag, Tokyo, 1985. pp. 45, 356.

[6] D. Dudgeon. Extraction and recognition of silhouettes from range imagery. November 13, 1987. M.I.T. Stochastic Systems Group Seminar.

[7] P. L. Van Hove. Model-based silhouette recognition. November 18, 1987. M.I.T. Digital Signal Processing Group Seminar.

[8] J. L. Prince. *Geometric Model-based Etimation from Projections.* PhD thesis, Massachusetts Institute of Technology, January 1988.

[9] J. G. Verly, P. L. Van Hove, R. I. Walton, and D. E. Dudgeon. Silhouette understanding system for laser-radar range imagery. In *Proceedings of the Meeting of the IRIS Specialty Group on Active Systems*, Monterey, CA, November 5–7, 1985.

[10] J. G. Verly, P. L. Van Hove, R. I. Walton, and D. E. Dudgeon. Silhouette understanding system. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, Tokyo, Japan, April 7–11, 1986.

[11] B. K. P. Horn and E. J. Weldon, Jr. Filtering closed curves. *I.E.E.E. Journal of Pattern Analysis and Machine Intelligence*, PAMI–8:665–668, 1986.

[12] B. L. Yen and T. S. Huang. Determining 3-d motion and structure of a rigid body using the spherical projection. *Computer Vision, Graphics, and Image Processing*, 21:21–32, 1983.

[13] S. A. Shafer. *Shadows and Silhouettes in Computer Vision.* Kluwer, Boston, 1985.

[14] W. N. Martin and J. K. Aggarwal. Occluding contours in dynamic scenes. In *Proceedings of the I.E.E.E. Conference on Pattern Recognition and Image Processing*, pages 189–192, 1981.

[15] Y. C. Kim and J. K. Aggarwal. Rectangular parallelepiped coding: A volumetric representation of three-dimensional objects. *I.E.E.E. Journal of Robotics and Automation*, RA-2:127–134, 1986.

[16] P. Van Hove and J. Verly. A silhouette slice theorem for opaque 3-D objects. In *International Conference on Acoustic and Speech Signal Processing*, pages 933–936, March 1985.

[17] B. K. P. Horn. Extended gaussian images. *Proceedings of the I.E.E.E.*, 72:1671–1686, December 1984.

[18] E. H. Wood, J. H. Kinsey, B. K. Robb, B. K. Gilbert, L. D. Harris, and E. L. Ritman. Applications of high temporal resolution, computerized tomography to physiology and medicine. In G. T. Herman, editor, *Image Reconstruction from Projections*, chapter 6, pages 247–279, Springer-Verlag, Berlin, 1979.

[19] Fritz Keinert. Inversion of k-plane transforms and applications in computer tomography. *SIAM Review*, 31(2):273–289, June 1989.

[20] G. M. Constantine. On the invertibility of the discrete radon transform. *SIAM Journal on Discrete Mathematics*, 2(3):300–306, August 1989.

[21] D. E. Dudgeon and R. M. Mersereau. *Multidimensional Digital Signal Processing.* Prentice-Hall, New Jersey, 1984.

[22] P. Billingsley. *Probability and Measure.* John Wiley & Sons, New York, 1979.

[23] F. C. Schweppe. *Uncertain Dynamic Systems.* Prentice-Hall, New Jersey, 1973.

[24] D. J. Rossi. *Reconstruction from Projections Based on Detection and Estimation of Objects.* PhD thesis, Massachusetts Institute of Technology, August 1982.

[25] F. C. Schweppe and H. K. Knudsen. The theory of amorphous cloud trajectory prediction. *I.E.E.E. Transactions on Information Theory*, IT-14:415–427, 1968.

[26] O. E. Drummond, S. S. Blackman, and K. C. Hell. Multiple sensor tracking of clusters and extended objects. In *Proceedings of the 1988 Tri-Service Data Fusion Symposium*, Laurel, MD, May 19, 1988.

[27] H. Guggenheimer. *Applicable Geometry: Global and Local Convexity. Applied Mathematics Series*, Krieger, Huntington, New York, 1977.

[28] J. A. Thorp. *Elementary Topics in Differential Geometry.* Springer, New York, 1979.

[29] C. C. Hsiung. *A First Course in Differential Geometry.* John Wiley & Sons, New York, 1981.

[30] P. Lancaster and M. Tismenetsky. *The Theory of Matrices.* Academic, Orlando, 1985.

[31] M. E. Mortenson. *Geometric Modeling.* John Wiley & Sons, New York, 1985.

[32] M.P. do Carmo. *Differential Geometry of Curves and Surfaces.* Prentice-Hall, New Jersey, 1976.

[33] M. Spivak. *A Comprehensive Introduction to Differential Geometry.* Volume V, Publish or Perish, Boston, MA, 1975.

[34] S. S. Chern, editor. *Studies in Global Geometry and Analysis.* Volume 4 of *Studies in Mathematics*, Prentice-Hall, 1967.

[35] B. K. P. Horn. *Robot Vision.* MIT Press, Cambridge, 1986.

[36] Shankar Moni. A closed-form solution for the recovery of a convex polyhedron from its extended gaussian image. Machine Vision Laboratory, University of Hawaii. Preprint.