A LINEAR
TIME THEORY FOR
RECOGNIZING SURFACES
IN 3-D

by

G.L. Bilbro
W.E. Snyder
and
J.E. Franke

Center for Communications and Signal Processing
Department of Electrical Computer Engineering
North Carolina State University

April 1986

CCSP-TR-86/8

A Linear Time Theory for Recognizing Surfaces in 3-D

Abstract

We present a new theoretical approach to rigid object recognition in range imagery. It is based on a special re-parametrization of a smooth surface that does not depend on the condition of rotation, translation, or parametrization of the surface, or on occlusion of remote parts of the surface. If a point on the surface appears, then it always transforms to numerically identical coordinates in the special parameter system. This amounts to a continuous correspondence between points on two surfaces and permits the comparison of the surfaces by correlation techniques. We present a new method for the determination of relative pose between two congruent surfaces from the Cartesian coordinates of corresponding pairs of points. The pose is the minimum eigensolution of a rank 4 real symmetric matrix whose entries are correlation coefficients. We also introduce a means for predicting correspondences between primitives in any geometric matching scheme.

1. Introduction

The new technology of range imagery has propelled researchers in image understanding past the problem of determining the shape of a object from ambiguous clues in luminance images such as shading. A range image defines the visible geometry of the surface of an object, except for the more familiar effects of noise and sampling. This has brought us to the next problem of deciding[1] when two shapes are "the same", discounting rotation and translation. It has become important to characterize shapes in such a way that they can be compared.

The classical mathematical study of differential geometry has supplied several techniques for dealing with surfaces. In particular, the curvatures of the surface have been considered. Since the curvatures do not depend on rotation, translation, or parametrization, they have been attractive candidates for this characterization of surfaces. Brady et. al. have approached[9,4] the problem of characterizing surfaces by finding a few special curves in the surface that conveyed decisive information about the surface. Most of the surface does not appear in such a representation except in the very indirect way of not qualifying as special. It appears however that this very sparse characterization may be sufficient to permit the comparison of objects and therefore the recognition of objects.

Faugeras and Hebert[6] have recently reported an important geometric technique for object recognition in range imagery. They partition the surface of the imaged object into a set of plane segments, each defined by its normal **n**, together with a distance **d** from the origin in the image coordinate system. A model object is similarly represented in its coordinate system. A "backtrack tree search" is then used to determine the correspondence that minimizes a sum of mismatches between the parameters of each imaged plane and the parameters of the correspondent model plane

$$\sum_i |\mathbf{n}_i - R\,\mathbf{n}_{\pi(i)}|^2 + W\,|d_i - d_{\pi(i)} - \mathbf{t}\cdot R\,\mathbf{n}_{\pi(i)}|^2, \tag{1.1}$$

where the rotation, R, followed by the translation, $\mathbf{t}$, together define an unknown rigid motion, and π defines an unknown correspondence from imaged planes to model planes. The constant weight, W, is an empirically chosen adjustment for the relative effect of the orientation error and the translation error. Since the total mismatch is computed as a sequence of partial sums, they can use the current partial sum as a "branch and bound" cost to control the search. This exploits the powerful geometric constraints relating plane segments rather than just the adjacency constraints [10,3,11] used previously.

However the planar segmentation of the surface is not "visible-invariant" for a general curved surface and this deficiency compromises both the theory and performance of the algorithm. Jain and Besl define a "visible-invariant" feature[2] as one that does not depend on the pose of the object or the parametrization of its surface or the occlusion of other features, if only it is visible at all. The planar partition fails in two ways on smooth surfaces. First it is sensitive to occlusion of remote parts of the surface: The plane segments neighboring an occlusion edge must conform to the detailed shape of the edge, just as the next-nearest neighbors must in turn conform to the edge presented by the nearest neighbors. This effect propagates into the surface so that the particular planar segment that a pixel falls into depends on the details of occlusion several segments away. Second, each plane segment will contain a directional uncertainty of order $\sigma_1 = k_{rms} L$ radians, where L is a length typical of the size of the planar segment and $k_{rms}^2 = (k_1^2 + k_2^2)/2$ defines the root mean square curvature in terms of the principal curvatures of the surface. Furthermore, the planar segmentations of two similar but not identical views of a large smooth surface will drift relative to each other: The unit normal of a segment near the center of the large surface will have accumulated a variance of order $\sigma_N = \sqrt{N}\ \sigma_1$ radians, where N is the number of planar segments along a typical diameter of the large surface. The tradeoff in the size of the segments could be optimized, but we are interested in a more fundamental reform.

Clearly our objection to planar partitions must be tempered by Faugeras and Hebert's success. Evidently global consistency between corresponding segments results in a minimum of the mismatch that is still sharp enough to identify the correct rigid body motion in spite of the noise introduced by the planar modeling, which must broaden and weaken the minimum. We view the speed that Faugeras and Hebert report as doubly strong support for the use of rigid body geometric constraints.

With the expectation of better performance, the previous authors also introduce a version of their algorithm based on quadric segments rather than planes. We agree, since exact quadric surfaces are common in industrial applications; the partition of a surface composed of exact and distinct quadric segments will be visible-invariant just as a planar partition of a polyhedron will be visible-invariant. Imaged surfaces that are not exactly quadric, however, will be subject to the same errors.

Faugeras and Hebert properly divide the problem of recognizing segmented surfaces into two subproblems: (1) finding the correct correspondence between the set of imaged segments and the set of model segments, and (2) determining the rigid motion that will bring the two surfaces into coincidence. In this paper, we will present improved solutions to both subproblems. In section 2, we introduce a new surface parametrization and show how this representation solves a restricted surface recognition problem. In section 3, we extend this method to more complex surfaces and present a recognition algorithm for noiseless surfaces and we introduce a correspondence prediction function to accelerate the search for the correct correspondence. The theory depends on a certain richness in the geometry of surfaces which is absent from some common surfaces, so in section 4, we extend the theory to treat some cases of those simple surfaces. In section 5, we summarize our results.

2. Parametrized curves in 2-D and surfaces in 3-D

In general, even a smooth surface must be defined as a collection of several parts[13] to permit a parametric representation that is regular on each part. The surface of a solid object will require more than one range image to completely represent it; and the different views might overlap. A tomographical surface might involve several layers where outer layers enclose inner layers. Such a surface will be represented as several parametrized parts possibly having intersecting parameter domains, but all rigidly related geometrically. The following theory easily treats such surfaces on several parameter domains since the transformation that we will introduce will always map every occurrence of a particular point to the same coordinates, where multiple occurrences can be properly exploited, say to reduce noise.

The technique developed in this paper will involve the re-parametrization of a surface, originally parametrized on sensor coordinates, perhaps. Since this new parametrization will involve domains unrelated to those of the original parametrization, we will assume for clarity that the data is originally given on one domain.

In this paper, we will make extensive use of the theory of differential geometry, and, in particular, representations for and properties of curvatures. Space does not permit a review of this theory here. The interested reader is directed to standard texts[13] for more information. Curvature is a property of surfaces in 3-space which exhibits some very attractive invariances to changes in viewpoint. In the next section, we illustrate, through a simple one-dimensional example, some of these invariant features.

2.1 Invariance of Curvature to Viewpoint: An Example

In this section we demonstrate the fact that curvature is invariant to viewpoint using a simple one-dimensional example. Suppose we are given the function in the $z=0$ plane

$$y = x^2. \tag{2.1}$$

Any point on this curve may be represented by the 3-vector

$$\mathbf{r} = [x, y, z]^T, \tag{2.2}$$

or, in parametric form (using the single parametrization $t = x$)

$$\mathbf{r} = [t, t^2, 0]^T. \tag{2.3}$$

Now, the curvature

$$\kappa = \frac{|\dot{\mathbf{r}} \times \ddot{\mathbf{r}}|}{(\dot{\mathbf{r}} \cdot \dot{\mathbf{r}})^{3/2}} = \frac{2}{[1+4t^2]^{3/2}}, \tag{2.4}$$

where a dot indicates a derivative with respect to t, so that

$$\dot{\mathbf{r}} = [1,\ 2t,\ 0]^T \qquad \ddot{\mathbf{r}} = [0,\ 2,\ 0]^T\ . \tag{2.5}$$

Now suppose this original parabola undergoes a motion in the plane consisting of a rotation of θ about the Z axis, followed by a translation of $[dx, dy]$. The new coordinate pairs (x', y') are related to the original points by

$$x' = x \cos\theta - y \sin\theta + dx \tag{2.6}$$

$$y' = x \sin\theta + y \cos\theta + dy.$$

Using the parametrization defined above, this becomes

$$x' = t \cos\theta - t^2 \sin\theta + dx, \tag{2.7}$$

$$y' = t \sin\theta + t^2 \cos\theta + dy,$$

$$\dot{\mathbf{r}} = \begin{bmatrix} \cos\theta - 2t \sin\theta \\ \sin\theta + 2t \cos\theta, \\ 0 \end{bmatrix} \tag{2.8}$$

and

$$\ddot{\mathbf{r}} = \begin{bmatrix} -2\sin\theta \\ 2\cos\theta \\ 0 \end{bmatrix}. \tag{2.9}$$

Now,

$$|\dot{\mathbf{r}} \times \ddot{\mathbf{r}}| = 2\cos^2\theta - 4t \sin\theta \cos\theta + 2\sin^2\theta + 4t \sin\theta \cos\theta \tag{2.10}$$

$$= 2$$

$$\begin{aligned}\dot{\mathbf{r}}\cdot\dot{\mathbf{r}} &= \cos^2\theta - 4t\cos\theta\sin\theta + 4t^2\sin^2\theta + \sin^2\theta + 4t\cos\theta\sin\theta + 4t^2\cos^2\theta \\ &= (\cos^2\theta + \sin^2\theta) + 4t^2(\sin^2\theta + \cos^2\theta) \\ &= 1 + 4t^2\end{aligned} \tag{2.11}$$

and we have the same result as before.

2.2 Parametrized Surfaces in 3-D

Consider a segment of a parametrized[13] surface: it is a map, X, from some set of points on the plane called the domain, $U \subseteq R^2$, to a set of points in 3-space,

$$X: U \rightarrow R^3. \tag{2.12}$$

X is specified by three parametric equations

$$\mathbf{x}(\mathbf{u}) = [\, x_1(u_1,u_2),\ x_2(u_1,u_2),\ x_3(u_1,u_2)\,]^T, \tag{2.13}$$

that are three times continuously differentiable except perhaps on a finite number of curves and points. The x's are the real Cartesian components of the point on the surface identified by the u's. This representation of a surface is both general and popular. It is, for instance, the ideal surface that is sampled in a single range image array.

We may refer to the entire surface represented by X on U as the set image, $X(U)$, of the domain set U. Since we will use several maps, this kind of notation will help keep our meaning clear. The actual surface, $X(U)$, is visible-invariant, but the details of the parametrization are not unique. The identical surface may be represented by $Y(V)$ where $Y: V \rightarrow R^3$, and Y is some different triple of functions from those of X defined on some different domain set V. Or the surface $Y(V)$ may be a rotated or translated "copy" of $X(U)$. Or it may be unrelated. If we denote the "image surface" as $X(U)$ and the "model surface" as $Y(V)$, the question of interest is whether $X(U)$ is "congruent" to $Y(V)$. That is to say, does there exist a rigid motion, M, such that $X(U) = MY(V)$?

We will employ the surface curvature to answer this question. Although the curvature of smooth industrial surfaces is always well defined, we will limit ourselves in this section to one region of the surface on which the curvature satisfies an additional property. Consider a domain, U, on which the curvature map of the imaged surface, $K_I : U \rightarrow R^2$, has a nonsingular Jacobian matrix everywhere

$$\det J = \det \left[\frac{\partial(k_1,k_2)}{\partial(u_1,u_2)} \right] \neq 0. \tag{2.14}$$

We shall call any region on which J is never singular a "regular segment". On a regular segment, K_I therefore[13] has a local inverse everywhere. We show in Appendix A that this is sufficient for the curvature map to have also a global inverse, on the additional condition that the regular segment is part of a quadric surface. We believe that the additional condition is unnecessary but we have not been able to prove our conjecture. The conjecture simplifies the theory and we will assume it to be true. Even if it is not true, the theory can be modified by using the notion of "multiple valued functions" borrowed[8] from the study of complex variables. This complication will be mentioned in the conclusion to the paper.

If K is invertible then the composite function $X \circ K_I^{-1} : K_I(U) \rightarrow X(U)$ is a single valued, visible-invariant parametrization of the segment: For each pair of principal curvature values, a unique point on the surface is determined. A similar map can be constructed on any model segment, $Y \circ K_M^{-1} : K_M(V) \rightarrow Y(V)$. This composite map takes a pair of principal curvatures to a Cartesian triple on the surface. The entire surface can be accessed by varying the curvature pair. A small area on the surface is related to a small area on the curvature plane by

$$dA = \rho_I(\mathbf{k}) dk^2, \tag{2.15}$$

where ρ is the square root of the determinant of the metric for the image curvature parametrization,

$$\rho_I(\mathbf{k}) = \sqrt{\det G_I}. \tag{2.16}$$

The metric has a positive definite determinant[13] in this parametrization just as the metric in the original parametrization, G_{I0}, from which it is derived in the usual way:[13]

$$G_I = (J^{-1})^T G_{I0} J^{-1}. \tag{2.17}$$

And the original metric (or the "first fundamental matrix", as it is sometimes called) is defined in terms of first partials of the original parametrization of the imaged surface, $X: U \to R^3$,

$$(G_{I0})_{\alpha\beta} = \frac{\partial \mathbf{x}}{\partial u_\alpha} \cdot \frac{\partial \mathbf{x}}{\partial u_\beta} \quad 1 \leq \alpha, \beta \leq 2, \tag{2.18}$$

In a completely analogous way, a metric for the model surface, G_{M0}, can be defined from the original parametrization of the model surface; this can be transformed to the curvature parametrization for the model surface, G_M, by the appropriate Jacobian; and from this, an area density, ρ_M, can be derived for the model surface in the visible-invariant curvature parametrization.

For an imaged segment to be "congruent with" a model segment, the two must satisfy (1) a "domain test" which requires the curvature image of the former to be contained in the curvature image of the latter, $K_I(U) \subseteq K_M(V)$, with equality holding in the absence of occlusion. And (2), a "range test", $\rho_I = \rho_M$ for all $k \in K_I(U)$, that is, on the curvature image of the imaged segment. Rather than exact equality in the presence of noise, it is more sensible to require that

$$\int_{K_I(U)} dk^2 \left[\rho_I(\mathbf{k}) - \rho_M(\mathbf{k})\right]^2$$

not exceed some small threshold.

We will define the "pose" of the imaged segment as that rigid motion which brings the model segment into coincidence with the the imaged segment. Specifically, the pose is that rigid motion M for which

$$\int_{K_I(U)} dk^2 \ \rho_I(\mathbf{k}) \ |\mathbf{x}(\mathbf{k}) - M\mathbf{y}(\mathbf{k})|^2 = 0,$$

where $\mathbf{x}$ and $\mathbf{y}$ are now regarded as functions of $\mathbf{k}$, as given by the curvature parametrization. For realistic surfaces with noise, the minimum defines the proper M, but the minimum is no longer zero.

By this definition, the pose of the imaged surface depends not only on the imaging coordinate system, but also on the model coordinate system. These dependences are in opposite senses and tend to cancel, but are nevertheless essential. In this scheme, there is no AI-style conceptualization of the surface and no notion of a natural coordinate system or a preferred pose for a surface. But any definition of a parameterized surface also defines a pose by the choice of the coordinate system in which the Cartesian values of the parametrization are given. We precisely preserve this arbitrary choice.

Thus, if the imaged surface consists of exactly one regular segment, the curvature parametrization provides a transformation which will match exactly (in the absence of noise) the corresponding model segment, provided both model segment and image segment have the same pose. More specifically, a segment may be represented by an array indexed by the two principal curvatures. If two such arrays are identical, the surfaces are identical and posed identically. In Appendix B, we develop a simple technique for determining the existence (and value) of a rigid motion, M, which will make the two arrays identical or at least similar in a least squares sense. If such a motion exists, then the two surfaces are congruent and the motion, M, represents their relative pose.

If the image being viewed consists of more than one regular segment, more sophisticated techniques must be emphasized, as discussed in the next section.

2.3 Invertibility of the curvature parametrization: an example

We can eliminate the arbitrary parameter, t, from the example given in section 2.1. In equation 2.4, we stated the relationship between curvature, κ, and t which can be solved for t:

$$t = \pm\frac{1}{2}\left[\left(\frac{2}{\kappa}\right)^{2/3} - 1\right]^{1/2}. \tag{2.20}$$

Thus

$$x(\kappa) = \pm\frac{1}{2}\left[\left(\frac{2}{\kappa}\right)^{2/3} - 1\right]^{1/2} \quad y(\kappa) = \frac{1}{4}\left[\left(\frac{2}{\kappa}\right)^{2/3} - 1\right]. \tag{2.21}$$

So for a particular value of κ, $y(\kappa)$ is determined uniquely and $x(\kappa)$ is determined to within a sign. To see why x is not determined uniquely, we need to examine the derivative of $\kappa(t)$ at $t=0$,

$$\dot{\kappa} = -24t\left[1+4t^2\right]^{-5/2}. \tag{2.22}$$

Since κ is a function of only one variable, this derivative is the Jacobian of κ; we will use the two dimensional generalization of this case in the treatment of surfaces. On any interval where $\dot{\kappa}$ does not vanish, $\kappa: R \to R$ is 1-1 and therefore invertible. At $t=0$, we find $\dot{\kappa} = 0$, so that the point $t=x=0$ is a singular point of the inverted function, $t(\kappa)$, but it divides two regular regions. If we know we are in the first quadrant, $t>0$, then the parametrization on κ provides a unique solution for x and y. Furthermore, as we demonstrated in equation 2.11, this parametrization is invariant under transition and rotation. So on either of the regular intervals, any rigid motion on the curve only affects the coordinate values of the point. And the curvature of the point identifies it in a visible-invariant way.

3.0 Surfaces Composed of Several Regular Segments

The curvature images of two regular segments will not necessarily be disjoint, so that in general, K is not globally 1-1 on a surface comprising more than one regular segment. Practically, there are two cases. The first case is analogous to the undefined sign in the example of section 2.3: There may exist space curves along which the gradients of the principal curvatures are not linearly independent, either because one of the gradients vanishes along that curve, or because the two gradients are parallel along it. Secondly, there may be regions on the surface that are geometrically simple in the sense that one of these pathologies occurs. A sphere or cylinder is such a surface. We will call these unfortunately simple regions "singular segments". In this section, we will deal with the first case by partitioning the imaged surface into regular segments on each of which K is invertible. For now, we will assume that the surface contains no singular segments, and discuss the other case in section 4.

3.1 The J Partition of a Surface

We will partition the parameter domain into a union of open disjoint sets on which the curvatures are continuous and det$J$ does not vanish. There will remain an edge set, E, which contains the closure of the U's and would also contain the singular segments of more general surfaces. So

$$U = \bigcup_{i=1}^{i=n_I} U_i \cup E. \tag{3.1}$$

The partition is represented by a "label array", $L_I: U \to I$, where the image, $L_I(U)$, is the set of the first n_I integers and $L_I(\mathbf{u})$ is the number of the segment that the parameter pair $\mathbf{u}$ falls into. This partition is visible-invariant, that is, it can be reproduced uniquely and identically on the corresponding model domain, V, except that missing data may remove parts of an imaged segment without changing any boundary. The set of edge points, E, is the pre-image of zero under the Jacobian, which we here have assumed to have zero measure.

Now we will say that a surface, $X(U)$, containing n_I regular segments, is congruent to a surface, $Y(V)$, containing $n_M \geq n_I$ regular segments, if (1) there exists a permutation on the integers $\pi: I \to I$ such that for $i = 1,2,\ldots,n_I$, the i^{th} image segment is congruent with the $\pi(i)^{th}$ model segment, as discussed in section 2; and (2), there exists a single motion M that brings all the correspondent segments into simultaneous coincidence

$$\sum_{i=1}^{i=n} \int_{K_I(U_i)} dk^2 \ \rho_{i,I}(\mathbf{k}) \ |\mathbf{x}_i(\mathbf{k}) - M\mathbf{y}_{\pi(i)}(\mathbf{k})|^2 = 0, \tag{3.2}$$

where $\rho_{i,I}$ is the surface area density for the i^{th} segment of the imaged surface (we could drop the I subscript, since by the first test, the densities of the two surfaces agree). This is a straightforward generalization of the case of a single segment to a rigid body comprising several regular segments.

3.2 The Label Prediction Function

Faugeras and Hebert[6] have demonstrated that branch and bound techniques lead to efficient searches for the best of the $n!$ candidate π's. We are presently assuming noise-free surfaces that either can be made to coincide with zero mismatch or can be rejected. If the imaged surface contains no congruent segments and neither does the model surface, then branch and bound would reduce false subtrees to unit depth. This would lead to a $O(n^2)$ algorithm. In the case of noisy data, $O(n^2)$ becomes the best case time for branch and bound, because the minimal mismatch cost is no longer known, but we still might find that new minimum on our first path through the search tree. In the worst case, the time to find the global optimum becomes factorial (exponential). In practice, branch and bound probably runs fairly close to best case time because the rigid body constraint should tend to make incorrect correspondences immediately obvious (i.e. unit depth false subtrees) for reasonably asymmetric objects.

In this section, we will introduce a technique for predicting the model segment corresponding with an unmatched image segment. The prediction results from the constraints of rigid body geometry on the surfaces under consideration. If the pose of the imaged surface is known, a point in space belonging to an unmatched segment may be transformed to the corresponding point in the model space. The segment label of any model point can be determined. This segment will be the only possible match consistent with the pose. This prediction can be done in constant time, as will be shown below. The pose is defined by the motion, M, which was fixed by the previous correspondences. Each match determines M in time proportional to the number of pixels in the imaged segment. Neglecting noise, the root correspondence completely determines a motion, so that this prediction function leads to an algorithm that is asymptotically faster than any branch and bound: the only search is for the correct assignment of the root of the search tree. For each possible root, the remaining path through the search tree is determined. Evidently, this technique is of general utility: it can be used in any rigid body matching scheme that has progressed far enough to estimate M. In Faugeras and Hebert's scheme, a search path longer than three planar matches completely determines M, so all subsequent matches could be predicted. It would not be as useful there as it is here because the planar segmentation is not visible-invariant. This means that the planar image segments and model segments will not in general correspond exactly, as discussed in section 1. Of course, in any case, noise could limit the validity of the predictor to some neighborhood of known correspondences.

Define the label prediction function $\Lambda: I \rightarrow I$ from the integers to the integers so that for a image segment $i \in \{1,2, \cdots, n_I\}$, $\Lambda(i)$ is the corresponding model segment $m \in \{1,2, \cdots, n_M\}$. It is easy to implement this function in terms of previously defined functions.

$$\Lambda(i) = L_M(Y^{-1}(M^{-1} \mathbf{x})), \tag{3.3}$$

for any $\mathbf{x} \in X(U_i)$. Λ would be coded as an interpolation-and-lookup. It can be understood by working from the inside-out: M^{-1} is the inverse of the current candidate motion from model to image, so that $M^{-1}(\mathbf{x})$ is the Cartesian triple in the unmatched image segment in question, transformed to the model coordinate system. Y^{-1} is the inverse of the model surface parametrization, so that $Y^{-1}(M^{-1}(\mathbf{x}))$ is the parameter pair which identifies the transformed triple in the model parameter domain, if such point exists (if not, then either the model or the pose can be rejected). Finally, $L_M: V \rightarrow I$ is the labeling produced by the segmentation of the model image. It maps the model parameter pair to the number that was assigned to its segment.

3.3 Recognition Algorithm for Surfaces Composed of Regular Segments

The following algorithm MATCH will either reject an image/model match or produce the rigid motion that brings the two into coincidence. It assumes that both image and model surfaces are segmented as discussed in section 3.1, and that $\det G$ and **u** are available in global memory perhaps as arrays indexed by **k** for both surfaces. Since we are assuming noise-free surfaces, the rigid motion M is determined by the first call to FINDPOSE, and subsequent M's are simply required to be identical; this is certainly not the optimal treatment of noisy surfaces, but we hope it is a good way to convey the gist of the technique. For an image that has no congruent segments, the algorithm will be shown to run in time that is linear in the number of pixels. The segmentation can be done [12, 5] on conventional sequential architectures in practically linear time, where N is the number of pixels in the image, and in linear time on special purpose architectures. Evidently, it may be possible to build a vision system that would operate in time that is practically linear in the number of pixels.

```
MATCH
    begin
    m:=1;
    while m < n_M
        begin
        if ISCONTAINED( 1, m ) = SUCCESS then
            begin
            M := FINDPOSE( i, m );
            if FINISH( 2 ) = SUCCESS then
                return SUCCESS;
            end
        m := m + 1;
        end
    return FAILURE
    end

FINISH( i )
    begin
    if i > n_I then
        return SUCCESS
    else
        begin
        m := Λ(i);
        if m = 0 then
            return FAILURE
        if ISCONTAINED( i, m ) = SUCCESS then
            begin
            MTEST := FINDPOSE( i, m );
            if MTEST = M then
                return FINISH( i + 1 )
            else
                return FAILURE;
            end
        else
            return FAILURE
        end
    end
```

The routine ISCONTAINED(i, m) returns SUCCESS if imaged segment U_i and model segment V_m satisfy both the domain test and the range test as defined in section 2; otherwise it returns FAILURE. The routine FINDPOSE(i, m) determines M, the rigid motion that brings U_i and V_m into coincidence, as discussed in Appendix B. Λ is defined in section 3.2.

In the worst case, MATCH will call ISCONTAINED O(n) times to find the model segment corresponding to imaged segment 1. For unsymmetric objects, there is only one possibility. If none is found, the model is rejected. If a candidate is found, then FINISH will be called. If all of the segments in the image have been dealt with, FINISH returns SUCCESS. Otherwise, it calls Λ which uses the current motion, M, to obtain the number, m, of that model segment corresponding to the imaged segment, i, if there is such a model segment. If not, FINISH fails. If FINISH fails for any reason, MATCH tries the next assignment for imaged segment 1. If no model segment corresponds with imaged segment 1, MATCH rejects the model. If FINISH finds that m and i are congruent and if the relative orientation between them is the same as for all previous matches, it calls itself to treat the next imaged segment.

In the absence of symmetry (no congruent segments in the imaged surface and none in the modeled surface), FINISH is called at most once for each imaged segment, with time complexity for the k^{th} recursion, $T_k = c_1 + T_{k+1}$, for $k = 2,...,n_I - 1$, and $T_{n_I} = c_2$, where c_1 and c_2 are constant times. The total time spent in FINISH is the time spent in the topmost call T_2, which is found from the solution of the recursion equation to be at worst O(n). Symmetric surfaces might be made to coincide by any of several rigid motions. This particular algorithm will return SUCCESS for the first completely consistent match that it finds. If the surfaces contain some congruent segments, but are only partially symmetric, MATCH may require more time. If, for instance, imaged segment 1 were separately congruent with half the model segments, MATCH would not in general find the correct one first, and each rejection might require $O(n)$ time.

MATCH will correctly treat any surface composed of regular segments, such as an ellipsoid with three distinct semi-major axes. As discussed in Appendix A, such an ellipsoid would contain eight regular segments. And because of the symmetry of the ellipsoid, some of the segments would be congruent, so there would exist four rigid motions which would make two copies of such an object coincide. This presents no problem to the algorithm. However, MATCH will fail if even two of the semi-major axes are identical. In that case, the surface is cylindrically symmetric and would would contain no regular segments, but instead only two singular segments.

4. Extension to Singular Segments

A singular segment cannot be parametrized by its inverse curvature image, since $\det J = 0$ everywhere on it. The curvature image of a singular segment is not fully two dimensional: it is either a curve or a point in the curvature plane instead of an area. We will call the first a "partially resolved" segment and say that the second type of segment is "totally unresolved". We will present a mixed parametrization that extends the theory developed in section 3 to treat these segments, but it is neither as powerful nor as simple as the preceding theory of regular segments. This mixed parametrization depends on arc length from some segment boundary. This dependence on a remote feature makes the representation of singular segments imperfect. It is not resistant to occlusion and therefore not completely visible-invariant. We will develop a treatment for unoccluded singular segments. It will turn out that the logic of the above algorithm does not change. But the domain and range tests in ISCONTAINED(i, m) become more complicated. The asymptotic time complexity will remain O(n) for noiseless unsymmetric surfaces.

Visible-invariance at the level of points on the surface can be preserved in some cases if occlusion is carefully and explicitly treated. But, in general, this approach fails when singular segments are partially occluded. MATCH assumes that M can be completely determined from any corresponding pair of segments, but occlusion can completely conceal some degree of freedom in M: For instance, when none of the natural edges of a plane segment are visible in the imaged surface and the segment is entirely bounded by occlusion edges, then that segment cannot be used to determine those degrees of freedom in M belonging to motions in the plane. Even the self occlusion of a sphere along its limb is fatal to the preceding simple algorithm. Unfortunately, all developable surfaces and all surfaces of revolution fall into this industrially interesting class, which includes planes, cylinders, and cones.

We will restrict ourselves to small singular segments that are completely surrounded by regular segments (to disallow self occlusion) and that are not otherwise occluded or incomplete. This will ensure that the entire boundary of the singular segment is available as a landmark. It is not difficult to relax these requirements somewhat, but we have not yet found a completely satisfactory treatment of singular segments.

We will refine the partition of section 3 by further analyzing the edge set to reflect partially resolved and unresolved singular segments:

$$E = U_{PR} \cup U_{UR} \cup E', \tag{4.1}$$

where U_{PR} contains all the regions of E on which $\det J$ is defined, but vanishes, and whose curvature image is a curve in the k plane. Similarly we collect all the smooth regions of E with constant curvatures into U_{UR}. This will leave E' as a final remainder.

4.1 Partially Resolved Segments

Here the gradients of the two curvatures together only define one direction on the surface, either because one gradient is zero or because the two are parallel. The first case includes developable surfaces (except for cylinders and planes which are totally unresolved). The second includes surfaces of revolution (except for cylinders and spheres) and also surfaces whose principal curvatures obey a relation $k_1 + k_2 = constant$, such as minimal surfaces when the sum vanishes.[13]

We will let k_1 be the changing curvature or, in the case of minimal surfaces, either curvature. In either case, the k_1 partially resolves the segment: the locus of points on the segment surface with constant k_1 is a curve lying in the segment and terminated on both ends by the segment boundary. We are here tacitly assuming that each such curve is a connected set and is not broken into two or more pieces by a hole in the surface or a ripple in the segment boundary. Just as the restriction on occlusion, this assumption can also be relaxed. We will ignore this complication except to mention that it is typical of singular segments and due entirely to their imperfect visible-invariance. No such problem occurs on regular segments, where each point is uniquely identified locally by its two principal curvatures and independently of remote features.

Our singular segment can be regarded as a family of k_1 curves. Each k_1 curve can be parametrized by arc length from one of its endpoints. We have assumed that both these intersections with the segment boundary are actual features on the surface and not due to occlusion. The sense of traversal of the curve is arbitrary so long as the imaged and model curves are traversed in the same sense. Since the gradient of k_1 is perpendicular to the constant k_1 curve, we can choose the positive direction along the curve so as to make ∇k_1 point to the right. In the absence of occlusion and disjoint k_1 curves, this parametrization will agree with a similar treatment of the same k_1 curve in any congruent segment, and in particular, in the correspondent model segment.

We will partition U_{PR} into a union of open sets which are bands in the sense that they are bounded by two types of edges: those just mentioned where the curvature is evidently discontinuous to some degree (Otherwise, the segment could be extended by Taylor's theorem across the boundary). And second, by curves of constant curvature on which the k_1 gradient vanishes also. Let U be one of those segments.

On U, k_1 has a nonzero gradient and therefore can be inverted to parametrize its eigendirection. Stated differently, k_1 uniquely identifies a k_1 curve. The other curvature can now be written as a function of the first, $k_2 = k_2(k_1)$. For example, $k_2 = 0$ for every k_1 on a cone or any other developable. For minimal surfaces, $k_2 = -k_1$.

Now if $K(U)$ is the range of k_1 on U, then for each value of $k_1 \in K(U)$, we will define a second parameter along the k_1 curve. Let $S(k_1)$ be the open real interval $(0, s_{\mathbf{max}})$, where $s_{\mathbf{max}}$ is the total length of the k_1 curve. Define $s_2 \in S(k_1)$ as the arc length along the k_1 curve from one intersection with the segment boundary. This defines a parametrization of any $U \in U_{PR}$ of the form $\mathrm{x}(k_1, s_2)$. We can partition and parametrize the model set V_R in the same way. The function ISCONTAINED(i, m) tests that the two curvature domains (now simply intervals) obey $K_I(U_i) \subseteq K_M(V_m)$ and that the total arc length at each curvature is no larger in the imaged segment than in the model segment, that is $S_I \leq S_M$ for each k_1 and that

$$\int_{K_I(U)} dk_1 \int_{S(k_1)} ds_1 \left[\rho_I - \rho_M\right]^2 = 0, \tag{4.2}$$

where ρ_I is now the surface area density for the mixed parametrization of the imaged segment. As before, it is equal to the square root of the determinant of the metric for the mixed parametrization; ρ_M similarly measures the model segment. Except for the same changes in domain and metric, FINDPOSE does not change.

4.2 Totally unresolved segments

It is ironic that planes, spheres, and cylinders are conceptually the most difficult to treat in this theory. We partition the set U_{UR} into connected components on which both curvatures are constant, bounded by curves along which a surface curvature changes. Let U be one such segment of the imaged surface and V be the correspondent model segment.

The treatment of U requires two steps: Define V^* to be related to V in the following way: If V is a fragment of a plane, then V^* is that entire plane. If V is part of a sphere or cylinder, then V^* is the entire sphere or cylinder. The first step is to determine a rigid motion, M_1, that brings U into coincidence with V^*. The second step is to bring their boundaries into coincidence by a second motion, M_2. First, construct the following arbitrary coordinate system on U: in the case of planes, the system may be polar or Cartesian; in the case of spheres, it is polar; and in the case of cylinders, it may be cylindrical or helical with some prescribed pitch. In any case, construct a similar system on V^*. Now the surface can be parametrized in this coordinate system using arc length along the coordinate curves. The previous theory of section 2 can be reformulated to determine M_1 using this parametrization. The theory only requires a continuous correspondence between points on two surfaces; this second parametrization is such a correspondence.

But since these parametrizations of U and V^* are not uniquely determined, M_1 will contain a error which must be corrected by a second motion, M_2, determined by matching the actual boundaries of U and V. In the case of planes, this will be a three parameter motion in the plane. For spheres, it will involve a three parameter rotation about the center of the spheres. For cylinders, it is limited to a rotation about the axis followed by a translation along the axis. In each case, the boundaries of U and V are compared by a two dimensional analog to MATCH in which the bounding curves are parametrized by a scalar curvature.

This analog was hinted at in section 2, and now can be sketched in parallel to the theory of regular surfaces presented in the latter part of section 2 and in section 3. Let $X(U)$ be an unresolved segment, where U is the domain set of the segment. Let $X(C)$ be the boundary curve of $X(U)$. So that $X(C)$ is a space curve parametrized by the same functions as X, but whose domain is the curve C in the parameter domain.

The geodesic [13] curvature, k, of the curve is a scalar measure which reduces to ordinary plane curvature in the case of plane curves. Segment the boundary into arcs in which the derivative of the geodesic curvature with respect to arc length does not vanish. This is analogous to the J partition used in the treatment of surfaces. As before, there remains a set of constant curvature points. The domain of the boundary is now the union

$$C = \bigcup_{j=1}^{m} C_j \cup P, \tag{4.3}$$

which is the one dimensional analog of the previous partition of a surface. The model segment boundary, D, can be similarly analyzed.

Now we will say that the unresolved segment $X(U)$ is congruent to the unresolved model segment $Y(V)$ if $X(U)$ is congruent to $Y(V^*)$ (that is, if the surfaces are congruent without regard to boundaries) and if also the boundaries are congruent. Now the boundary $X(C)$, containing n_I regular arcs, is congruent to $Y(D)$, containing $n_M \geq n_I$ regular arcs, if (1) there exists a permutation on the integers $\pi: I \to I$ such that the intervals $K_C(C_i) \subseteq K_D(D_{\pi(i)})$ for $i=1,2,...,n$, with equality holding whenever $X(C_i)$ is not occluded, so that both endpoints actually belong to $X(C_i)$. And (2), if

$$\sum_{i=1}^{i=n_I} \int_{K_C(C_i)} dk \left[\rho_C - \rho_D\right]^2 = 0, \tag{4.4}$$

where

$$\rho = \left|\frac{ds}{dk}\right|, \tag{4.5}$$

is the absolute value of the derivative of arc length with respect to the scalar curvature. It is the one dimensional analog to the $\sqrt{\det G}$ for surfaces. And (3), if there is a single motion M_2 such that

$$\sum_{i=1}^{i=n_I} \int_{K_C(C_i)} dk \ \rho_{i,C}(k) \ |\mathbf{x}_i(k) - M_2 M_1 \mathbf{y}_{\pi(i)}(k)|^2 = 0.$$

which is the one dimensional version of the expression for surfaces. Here, M_2 corrects the earlier error in M_1, and the rigid motion that brings U into coincidence with V is the product, M_2M_1.

An added (but perhaps unsurprising) difficulty is that there may be "singular arcs" in the bounding curve, analogous to "singular segments" in surfaces. Fortunately, there is only one type, distinguished by constant curvature on planes and constant geodesic curvature on spheres and cylinders. This includes circular arcs and straight line segments on planes including their analogs on spheres and cylinders and nothing else. There is no further regress. These singular arcs can be parametrized by arc length from one end point and the remaining development is precisely parallel to the previous analysis of unresolved singular segments.

5. Conclusions

In this paper, we have introduced a representation for a surface in Cartesian 3-space that is visible-invariant except on regions that are both partially occluded and geometrically simple. This representation is an ordinary re-parametrization of the surface, with the local principal curvatures as new parameters. All points on the surface contribute uniformly to the representation. The case in which a surface is given on several domains presents no problem because a visible-invariant partition of the entire surface arises naturally from the re-parametrization. The new partition can be constructed with conventional segmentation algorithms on the originally parametrized surface.

This representation depends on a certain richness in the local geometry of the surface. We have called these rich regions "regular segments". The curvature parametrization fails on other parts of a surface which are geometrically too simple. We have called these simple regions "singular segments" and shown how they are identified by the same criterion that defines the boundaries of the partition mentioned in the preceding paragraph.

Unfortunately spheres, cones, and planes are among these singular segments. We have presented a second related representation which can be applied to these singular segments, but it depends on distance to the boundary of the region, measured in the surface. Because of its dependence on remote features, the second representation is imperfectly visible-invariant: it is sensitive to occlusion. When these difficulties can be ignored, singular segments can be treated on an equal footing with the robustly invariant regular segments.

The representation depends less severely on a conjecture concerning the single-valuedness of the curvature of a regular segment. We have shown in Appendix A that this curvature parametrization is indeed single-valued for arbitrarily complex assemblies of regular segments, if each is part of a quadric surface and if the boundaries that join the quadric segments are sharp enough to be identified. (Ultimately, the segments are treated separately). We believe that the curvature parametrization is always single-valued, but we have not yet found a proof for the conjecture. Single-valuedness is useful, but not essential. It permits each segment to be represented with an extremely simple data structure: an ordinary array. If the conjecture proves false, the representation of a segment may require more than one array. In that case, a new kind of boundary would have to be introduced, analogous to the cut lines used to treat multiple valued functions of a complex variable. This "cut line" complicates the connectivity of the domain in the same way that it does in the study of complex variables. Because we believe that this complication may not occur for surfaces in 3-D, we have assumed it as a conjecture.

We have shown that the curvature parametrization of a surface is an excellent representation for range image understanding. When an object can be identified by its regular segments, the curvature parametrization solves the problems of object recognition and pose determination: Each segment in the partition can be represented as an array indexed by the two principal curvatures. Each pair of principal curvatures identifies both a point on the surface segment and an element in an array. The array element is a Cartesian triple fixing the location of the point in (for example) the image coordinate system. Two surfaces can be compared by comparing the segments in their partitions. The comparison of segments is a correlation between arrays: it is limited to comparing elements at the same indices, because the principal curvatures are independent of rotations and translations of the coordinate system or object.

If the arrays are identical, the surfaces are identical and posed identically. If there exists a single rigid motion that makes two arrays identical, the segments are congruent and that motion is the pose of one segment relative to the other. In Appendix B, we have developed a simple least mean squares technique for determining the existence of such a rigid motion by producing the motion if one does exist. Otherwise, the segments are not "the same".

We have carefully considered the constraints presented by rigid body geometry to the problem of matching primitive surface elements with fixed relative poses. We have devised a correspondence prediction function that can be applied to geometric matching schemes. We have used this prediction function in the construction of our own recognition and pose algorithm, which embodies and organizes all the above principals. We have calculated its asymptotic time complexity to be less than standard segmentation algorithms for cases where the surfaces possess no partial symmetries and no pathological singular segments.

One reservation remains to be noted. The theory depends on reliable estimates of the second partial derivatives of the surface for the re-parametrization of a segment by its curvature. The partition that defines the segments themselves depends on the third partials of the surface. These derivatives must support enough resolution not to blur distinguishing features of the surface. This resolution requirement imposes a minimum feature size (length) of twenty, or possibly forty, pixels in the current technology of range image acquisition.

Appendix A The Curvature map is invertible on a regular segment that is also part of a quadric surface.

In this section we are interested in quadric surfaces,

$$ax^2 + by^2 + cz^2 + dxy + eyz + fxz + hx + iy + jz + k = 0,$$

and to what extent the principal curvature can be used to parameterize the surface. Since the principal curvatures are invariant under rotations and translations, a given surface can undergo a rigid motion to simplify its equation. The simplified equation is of two possible types:

$$z = ax^2 + by^2 \tag{A1}$$

$$z^2 = ax^2 + by^2 + c \tag{A2}$$

Both of these equations are symmetric with respect to the xz and yz coordinate planes. The principal curvatures are identical at their corresponding symmetric points. We will see that for a large class of these surfaces the coordinate planes divide the surface into regions on which no two points have the same principal curvatures. In our calculations we will use equation (1) and leave the similar calculations for equation (2) to the reader.

The Gaussian curvature K and the mean curvature H are related to the principal curvatures k_1 and k_2 by:

$$K = k_1 \cdot k_2$$
$$H = \frac{1}{2}(k_1 + k_2) \tag{A.3}$$

Thus if two points have the same principal curvatures, then they have the same K and H. We will show that K and H never repeat in an octant and hence the same is true for the principal curvatures.

A straight forward calculation using equation (1) and the definition of the curvatures (see, for example Willmore,[13] p 98) yields

$$K = \frac{4ab}{1 + 4a^2x^2 + 4b^2y^2} \tag{A.4}$$

$$H = \frac{2b(1+4a^2x^2) + 2a(1+4b^2y^2)}{(1+4a^2x^2+4b^2y^2)^{3/2}} \tag{A.5}$$

Restricting our attention to the first quadrant, make the change of variables

$$u = 4a^2x^2 \tag{A.6}$$

$$v = 4b^2y^2 \tag{A.7}$$

Then

$$K = \frac{4ab}{1+u+v} \tag{A.8}$$

$$H = \frac{2b(1+u) + 2a(1+v)}{(1+u+v)^{3/2}} \tag{A.9}$$

The level curves $K =$ constant are straight lines in the uv space.

$$v = -u + \frac{4ab}{K} - 1 \tag{A.10}$$

On one of these line segments

$$H = \frac{2b(1+u) + 2a(1-u+\frac{4ab}{K} -1)}{\left(\frac{4ab}{K}\right)^{3/2}} \tag{A.11}$$

$$\frac{dH}{du} = (2b - 2a)\left(\frac{K}{4ab}\right)^{3/2} \tag{A.12}$$

If $a \neq b$ then this derivative is always positive or always negative. Thus H never repeats on this line segment. Therefore the pair (H,K) uniquely defines a point on the surface in the first octant. As noted earlier, this implies that the map from the surface in the first octant to the plane of principal curvatures (k_1,k_2) is one to one and hence invertible. This means that the principal curvatures may be used to parameterize the surface in the first octant. By symmetry the surfaces in each of the octants can be parameterized by the principal curvatures.

It is the symmetry of the quadric surfaces that can be used to divide the surface into pieces on which the principal curvatures may be used to parameterize the surface. These global symmetries may be hard to locate and thus a local condition may be more useful. Consider the map from xy-space to R^2 by $(x,y) \to (K,H)$. This is a smooth map and its Jacobian obeys

$$\det\left[\frac{\partial(H,K)}{\partial(x,y)}\right] = \frac{2^9\, a^3 b^3 (b-a)xy}{(1+4a^2x^2+4b^2y^2)^2} \quad .$$

This is zero when $a=0$, $b=0$, $a=b$, $x=0$, or $ay=0$. The first three cases are examples of singular segments where the principal curvatures cannot be used to parameterize pieces of the quadric surface. The last two are exactly the divisions that we need to divide the remaining quadrics into regions where the principal curvatures can be used for parametrization. Now

$$\det\left[\frac{\partial(H,K)}{\partial(x,y)}\right] = 0 \qquad \text{implies} \qquad \det\left[\frac{\partial(k_1,k_2)}{\partial(x,y)}\right] = 0 \quad ,$$

thus the subdivision produced by the second condition will be at least as fine as the first.

Appendix B. Determination of the Rigid Body Motion Relating Two Congruent Surfaces with Corresponding Parametrization

We will be interested in two surfaces in 3-space with the same parametrization in the following sense. If the two surfaces are

$$X: U \to R^3 \tag{B.1}$$

and

$$Y: V \to R^3, \tag{B.2}$$

with $U \subseteq V \subseteq R^2$, then for every $\mathbf{u} \in U$, the point located at $\mathbf{x}(\mathbf{u})$ on the first surface is the same point of the second surface, but located at $\mathbf{y}(\mathbf{u})$. In fact, we are thinking of the domain here as the principal curvatures. Our curvature parametrization provides a continuous correspondence between points on two surfaces. But this method of pose determination applies equally to discrete cases as long as the correspondence is already determined.

We will pursue a rigid motion, M, which minimizes the pointwise Euclidean squared error averaged over U, the smaller (image) domain. This mismatch is

$$E = \langle|\mathbf{d}|^2\rangle = \langle|\mathbf{x} - M\mathbf{y}|^2\rangle, \tag{B.3}$$

where the averaging operator is a weighted integral over the smaller domain, but scaled to have unit total weight, so that

$$\langle 1 \rangle = \int_K dk^2 \rho(\mathbf{k}) 1 = 1 \quad . \tag{B.4}$$

In our continuous case, ρ is the differential area of the surface. In the discrete case the integral would be replaced by a sum over pairs of corresponding points and ρ could be as simple as the reciprocal of the number of pairs. If some pairs had higher confidence or were more important, ρ would be the proper place to model that information. In our continuous case ρ would make this importance proportional to actual physical area.

The rigid motion is a linear transform of the form

$$M\mathbf{y} = R\mathbf{y} + \mathbf{t}, \tag{B.5}$$

that is, a fixed, but unknown rotation about the origin, R, followed by a fixed, but unknown translation, $\mathbf{t}$.

The translation can be expressed in terms of R by setting the gradient of E with respect to $\mathbf{t}$ equal to zero. The result is

$$\mathbf{t} = \langle \mathbf{x} - R\mathbf{y} \rangle. \tag{B.6}$$

This formal result can be substituted into the expression for the error,

$$E = \langle \, |\mathbf{x} - R\mathbf{y} - \langle \mathbf{x} - R\mathbf{y} \rangle|^2 \, \rangle, \tag{B.7}$$

or

$$E = \langle |\mathbf{x} - R\mathbf{y}|^2 \rangle - |\langle \mathbf{x} - R\mathbf{y} \rangle|^2. \tag{B.8}$$

This error can be minimized to determine the rotation R, which in turn determines the translation $\mathbf{t}$. The translation appears as a "mean" or "center of mass" term, familiar from expressions involving the expectation of quadratic functions. This unified treatment of the total motion, M, depends on a correspondence between actual points in 3-space. It was not possible for Faugeras and Hebert to obtain this unification in their seminal[6] paper because they dealt with primitives such as planes whose parameters unfortunately did not have the transformation properties of points in 3-space. Nevertheless, the following development owes much to their work.

R is conveniently constrained to be a rotation by expressing it as a quaternion product,[7]

$$R\mathbf{y} = q \cdot \mathbf{y} \cdot q^{-1}. \tag{B.9}$$

A quaternion q is any real 4-vector along with special definitions for multiplication and conjugation. Quaternion multiplication, $\cdot$, is a certain bilinear operand for 4-vectors. The components of a quaternion are thought of as a "scalar" part and a "vector" part, $q = q_0 + \mathbf{q}$. A 3-vector is taken to have a vanishing "scalar" part. The product of two quaternions is

$$q \diamond r = q_0 r_0 - \mathbf{q} \cdot \mathbf{r} + q_0 \mathbf{r} + r_0 \mathbf{q} + \mathbf{q} \times \mathbf{r}. \tag{B.10}$$

To uniquely specify the quaternion by the rotation, q is conventionally taken to have unit length (norm) and positive zeroth component,

$$|q|^2 = \sum_{i=0}^{i=3} q_i^2 = 1 \quad \text{and} \quad q_0 > 0. \tag{B.11}$$

Such a quaternion can be written

$$q = \cos\frac{\theta}{2} + \mathbf{n} \sin\frac{\theta}{2}, \tag{B.12}$$

where $\mathbf{n}$ is the axis of rotation and θ is the angle. Only one additional observation is essential about the quaternion product: $|q \diamond r| = |q|\ |r|$, for any two quaternions.

So now, we want to minimize E with the constraint $|q|^2 = 1$. With this constraint, the minimum of Eq^2 coincides with the minimum of E, but leads[6] to a linear equation. Using Lagrange multipliers, we will minimize

$$E' = E|q|^2 - \lambda|q|^2, \tag{B.13}$$

or, since the norm of the product of two quaternions is the product of the norms,

$$E' = \langle|\mathbf{x} \diamond q - q \diamond \mathbf{y}|^2\rangle - |\langle\mathbf{x} \diamond q - q \diamond y\rangle|^2 - \lambda|q|^2. \tag{B.14}$$

Any linear function of q can be written as the ordinary product[6] of a matrix that is independent of q with the components of q, so

$$\mathbf{x} \diamond q - q \diamond \mathbf{y} = Aq, \tag{B.15}$$

for some A. Solving for the elements of A, we find

$$A = \begin{bmatrix} 0 & -(\mathbf{x}-\mathbf{y})^T \\ (\mathbf{x}-\mathbf{y}) & \Delta(\mathbf{x}+\mathbf{y}) \end{bmatrix}, \tag{B.16}$$

where $\Delta(\mathbf{x})$ is the 3 by 3 matrix equivalent to cross product premultiplication: $\Delta(\mathbf{x})\mathbf{v} = \mathbf{x} \times \mathbf{v}$. It can be written with the Levi-Civita symbol on three indices, ϵ_{ijk}, which is unity if i,j,k is an even permutation of 1,2,3; is -1 if i,j,k is an odd permutation; and is zero otherwise.

$$\Delta_{ik}(\mathbf{x}) = \sum_{j=1}^{j=3} \epsilon_{ijk} x_j, \tag{B.17}$$

as is easy to verify.

Then

$$E' = <(Aq)^T(Aq)> - <(Aq)^T><Aq> - \lambda q^T q \tag{B.18}$$

attains an extreme value at

$$Bq = \lambda q, \tag{B.19}$$

where $B = <A^T A> - <A^T><A>$ is a real symmetric matrix of correlation constants, and the best solution is the eigenvector belonging to the minimum eigenvalue. Now we can evaluate the translation

$$\mathbf{t} = <A> \diamond q^{-1}, \tag{B.20}$$

which is a quaternion with a vanishing scalar part (a vector), and the inverse of a quaternion is

$$q^{-1} = \frac{q_0 - \mathbf{q}}{|q|^2}. \tag{B.21}$$

1. P. Besl and R. Jain, "Surface Characterization for three-dimensional object recognition in depth maps," *University of Michigan, RSD-TR-20-84,* (Dec 1984).

2. P. Besl and R. Jain, "Range Image Understanding," *Invited paper, CVPR,* (1985).

3. B. Bhanu, "Surface representation and shape matching of 3-d objects," *Proc. Pat. Re. and Image Processing Conference,* (June 1982).

4. M. Brady, J. Ponce, A. Yuille, and H. Asada, "Describing surfaces," *Proc 2nd Int. Symp. Rob. Res., Harajusa and Inoue (eds),* MIT Press, (1985).

5. A. W. Cowart and W. E. Snyder, "An Iterative Approach to Region Growing Using Associative Memories," *IEEE Trans on PAMI,* (May 1983).

6. O. D. Faugeras and M. Hebert, "A 3-D Recognition and Positioning Algorithm Using Geometrical Matching Between Primitive Surfaces," *Proc. IJCAI,* pp. 996-1002 (1983).

7. C. J. Joly, *A Manual of Quaternions,* Macmillan and Co., Ltd, London (1905).

8. Z. Nehari, *Introduction to Complex Analysis,* Alyn and Bacon, Inc., Boston (1968).

9. J. Ponce and M. Brady, "Toward a Surface Primal sketch," *IEEE International Conf. on Robotics,* (Mar 1985).

10. A. Rosenfeld, R. Hummel, and S. Zucker, "Scene labeling by relaxation operators," *IEEE SMC-6,* (June 1976).

11. B. Snyder, *Three Dimensional Analysis for Range Images of Polyhedra,* MS Thesis, Dept of Electrical and Computer Engineering, NCSU, Raleigh, NC (1984).

12. W. E. Snyder and C. D. Savage, "Content-Addressable Read/Write Memories for Image Analysis," *IEEE Trans on Computers* C-31(Oct 1982).

13. T. J. Willmore, *An Introduction to Differential Geometry,* Oxford University Press, London (1959).