

An Optimal Three-Dimensional
Edge Operator*

Steven W. Zucker	&	Robert A. Hummel
Dept. of Elec. Eng.		Dept. of Mathematics
McGill University		Univ. of Minnesota
Montreal, Quebec		Minneapolis, MN
Canada		U.S.A.

Report No. 79-10

April 1979

An Optimal Three-Dimensional Edge Operator*

Steven W. Zucker
Computer Vision and Graphics Lab.
Department of Electrical Engineering
McGill University
Montreal, Quebec, Canada

Robert A. Hummel
Department of Mathematics
University of Minnesota
Minneapolis, Minnesota

ABSTRACT

Modern scanning techniques, such as computed tomography, have begun to produce true three-dimensional imagery of internal structures. The first stage in finding structure in these images, like that for standard two-dimensional images, is to evaluate a local edge operator over the image. If an edge segment in two-dimensions is modelled as an oriented unit line segment that separates unit squares (i.e., pixels) of different intensities, then a three-dimensional edge segment is an oriented unit plane that separates unit volumes (i.e., voxels) of different intensities. In this paper we derive an operator that finds the best oriented plane at each point in the image. This operator, which is based directly on the 3-D problem, complements other approaches that are either interactive or heuristic extensions of 2-D techniques.

* This research was supported by MRC grant no. MA-6154

1. Introduction

The development of non-invasive techniques for imaging the interior structure of three-dimensional objects is currently revolutionizing many areas of medicine and industry [1]. One of the most widely known of these techniques is computed tomography (CT), which uses either sonar or X-ray energy sources. In X-ray computed tomography, a number of X-ray projections are made from different angular positions around the object. Each of these projections yields a one-dimensional absorption profile. These profiles are then used to reconstruct a two-dimensional slice through the object [2]. After several of these reconstructed slices are made sequentially along the third axis, they can be stacked into a true three-dimensional image of the object. Sonar-based tomography yields similar three-dimensional imagery.

The development of systems for processing and displaying three-dimensional (3-D) imagery has revealed a number of new problem areas, such as the need for special hardware facilities (e.g. [3,4]) and graphical data structures [5]. Another essential problem underlying all of these systems is the need for algorithms for finding structure in 3-D images. Success in this area would improve both the presentation of the image content, and, eventually, the diagnostic usefulness of these noninvasive techniques.

Since different physical objects usually give rise to different image intensities, the first stage in finding

structure requires the location of these intensity differences. In two-dimensions this is essentially the problem of edge detection, and it can be formulated in the following way: (1) apply a local operator that responds strongly to intensity differences, such as a gradient operator, to every point in the image, and (2) interpret the response of this operator into assertions about the presence (or absence) of edge elements.

To understand our notion of a 3-D edge, recall that edge elements have an orientation associated with them, such as that shown in Fig. 1. If pixel (i,j) is part of the edge (on the dark side), then the orientation of the edge element located at (i,j) is the orientation of the line passing through the center of pixel (i,j) that best separates the intensities in the local neighborhood around (i,j) . (Note that, by symmetry, we also could have placed an oppositely oriented edge element on the light side of the edge. Or, we could have located the edge element in the interstitial space between pixels.) For an introductory discussion of many of these issues, together with techniques for grouping edge elements into longer curves, see [6].

In a geometric sense, a 3-D edge is a direct extension of the 2-D edge model just described. Instead of considering an edge as a line through a unit square of the image (i.e., through a pixel), we shall consider it as a

plane passing through the center of a unit volume (or, as Liu[7] called it, a voxel). The specific purpose of this paper is to present a local operator that defines this plane in an optimal fashion. This operator is a true generalization of the operators used in two dimensions. Also, it supplements the other approaches to 3-D edge detection, which require either interactive assistance [8] or heuristic decomposition of the 3-D surface detection problem into the maxima of the three underlying 2-D edge detection problems (namely, one along each axis) [7]. Applying our operator and interpreting its response should result in a smooth surface separating adjacent volumes in the image, where these volumes (or sub-images) correspond to different intensities. If this were an image of the abdomen, for example, one such surface could delimit the stomach.

Following the mathematical preliminaries in the next section, in which we formulate the feature detection problem as one of functional analysis, we derive our optimal operator. Surprisingly, it turns out to be a rather pleasing generalization of the (2-D) Sobel operator [9]. And finally, we present the results of applying the operator to several 3-D images.

2. Mathematical Background

Feature detection can be characterized as the problem of locating instances of a set of target patterns in data. The specific target patterns that we shall consider are distributions of volume elements separated by a flat plane. This plane is oriented so that dark voxels lie on one side of the plane, and light voxels lie on the other. Mathematically, these ideal 3-D edge configurations can be described by the set of functions:

$$E_{a,b,c}(x,y,z) = \begin{cases} +1 & \text{if } ax + by + cz \geq 0 \\ -1 & \text{if } ax + by + cz < 0 \end{cases} \quad (1)$$

which are defined on the sphere:

$$S = \{(x,y,z): x^2 + y^2 + z^2 \leq 1\} \quad .$$

The vector $\bar{N} = (a,b,c)$ is the unit normal (at the origin) to the plane

$$ax + by + cz = 0 \quad . \quad (2)$$

This plane separates the dark hemisphere of the edge configuration (voxels having ideal value +1) from the light hemisphere (voxels having value -1). Thus the edge target patterns admit a parameterization through the variables defining the normal \bar{N} , and we can pose our feature-detection problem in the following way [10]:

Let $I(x,y,z)$ denote an input image defined on the unit solid sphere S that has been normalized to have zero mean and unit variance. (This requirement will be relaxed in Sec. 4.) We seek values for the parameters (a,b,c) so that

$$\|I - E_{a,b,c}\| \quad (3)$$

is minimized. A convenient norm $\|\cdot\|$ is given by the L_2 -norm:

$$\|f\|_2 = \left[\iiint_S f^2(x,y,z) \, dx dy dz \right]^{1/2}.$$

However, to obtain a practical solution to this minimization problem (3), we must consider a finite-dimensional subspace M of $L_2(S)$. If we let $\{\psi_1, \psi_2, \dots, \psi_N\}$ denote an orthogonal set of basis functions for M , then the orthogonal projection I' of I onto M is given by:

$$I'(x,y,z) = \sum_{i=1}^N c_i \psi_i(x,y,z)$$

where

$$c_i = \iiint_S I(x,y,z) \psi_i(x,y,z) \, dx dy dz.$$

Similarly, let $E'_{abc}(x,y,z)$ denote the projection of the pattern $E_{abc}(x,y,z)$ onto M . Then, if M is chosen properly, this projection will be one-to-one on the class of E-patterns, and the full minimization problem (3) can be replaced by the finite-dimensional problem of finding (a,b,c) such that

$$\| I' - E'_{a,b,c} \| \quad (4)$$

is minimized. (This is the approach adopted by Heuckel [11,12], for example, in formulating a 2-dimensional edge operator.)

More formally, the reason that a solution to (4) also provides a solution to (3) is based on the inequality:

$$\| I' - E' \| \leq \| I - E \| . \quad (5)$$

This inequality derives from the fact that the projection operation on a Hilbert space $L_2(S)$ is linear, continuous, and does not increase the norm. Its proof follows from a generalized Pythagorean theorem: Let F be an element in the Hilbert space, and let F' be its orthogonal projection onto M . Then

$$\| F \|^2 = \| F' \|^2 + \| F - F' \|^2$$

or

$$\| F' \| \leq \| F \| .$$

A special case of the inequality (5), which is of particular interest to us, is:

$$\min \| I' - E'_{a,b,c} \| \leq \min \| I - E_{a,b,c} \| .$$

Thus the parameter values for the finite-dimensional problem provide a lower bound on the full problem. Furthermore, the above inequality has three additional consequences [10]:

(i) If $\min \| I' - E' \|$ is large, then there is no pattern E' that matches the image I (i.e., $\min \| I - E_{a,b,c} \|$

is large);

(ii) if I matches one of the E -patterns exactly (i.e., if $\min ||I - E_{a,b,c}|| = 0$) then

$$\min ||I' - E'_{a,b,c}|| = 0$$

is the minimum for (4), and the selected parameters (a,b,c) are exactly those which minimize (3);

(iii) the parameters (a,b,c) that minimize (4) are close to the parameters that minimize (3) whenever I is close to an $E_{a,b,c}$. This continuity assumption is reasonable provided that the subspace M is chosen properly, and one technique for doing this is described in the next section.

3. The Optimal Operator

In order to apply the theory of feature detection outlined in Sec. 2, we must select the finite-dimensional space M such that the patterns E' are a good approximation to E . In other words, we must find the best orthogonal basis functions $\{\psi_1, \dots, \psi_N\}$ so that the difference between E and E' is minimal. If we assume that all patterns E have zero mean and unit norm (this can be assumed without loss of generality, since the patterns can be normalized), then a formal selection criterion for M can be obtained by requiring that the expected value of $||E - E'||$ be minimal. This expectation is taken over the full set of target patterns E , and is weighted by the probability density of occurrence of the patterns. Since the set of target patterns is parameterized by (a,b,c) , we can regard each of them as the events comprising a random field with probability density $p(a,b,c)$. Thus the problem of minimizing the approximation error between E and E' can now be posed as one of finding the basis functions $\psi_1, \psi_2, \dots, \psi_N$ such that:

$$\mathcal{E} \left\| E - E' \right\| = \mathcal{E} \left\| E - \sum_{i=1}^N a_i \psi_i \right\| \quad (5)$$

is minimal, where \mathcal{E} is the expectation operator and where:

$$a_i = \iiint_S E(x,y,z) \psi_i(x,y,z) dx dy dz$$

The solution to this minimization problem is given by the Karhunen-Loeve basis functions [13,14]; i.e., the ψ_i are solutions to the integral eigenvalue problem:

$$\iiint R(x,y,z,x',y',z') \psi_i(x',y',z') dx' dy' dz' = \lambda_i \psi_i(x,y,z) \quad (7)$$

where the autocorrelation function is

$$R(x,y,z,x',y',z') = \mathcal{E} \{ E(x,y,z) \cdot E(x',y',z') \} \quad (8)$$

$$= \iiint_a \iiint_b \iiint_c E_{a,b,c}(x,y,z) \cdot E_{a,b,c}(x',y',z') \cdot p(a,b,c) \cdot da db dc$$

The autocorrelation in (7) forms the kernel of a symmetric, positive definite, compact operator, which implies that there will be a countable number of positive, real eigenvalues λ_i [15]. These can be ordered from the largest to the smallest, $\lambda_1 \geq \lambda_2 \geq \dots$, and the finite dimensional space M is formed from the eigenfunctions corresponding to the largest eigenvalues. These eigenfunctions minimize the truncation error introduced by

a finite value for N in (6). Clearly as more eigenfunctions are used, the approximation becomes better.

The mathematical computation of the autocorrelation function (8) is technically complicated, however a few geometric considerations allow us to describe R more clearly. To begin with, note that $R(x,y,z,x',y',z') = 1$ when (x,y,z) and (x',y',z') are vectors in the same direction, and $R(x,y,z,x',y',z') = -1$ when the two vectors point in opposite directions. Between these extreme points, $R(x,y,z,x',y',z')$ drops off linearly as the angle between the vectors (x,y,z) and (x',y',z') increases from 0 to π . (Recall that the dot product of two unit vectors is just the cosine of the angle between them.) Thus:

$$R(x,y,z,x',y',z') = 1 - \frac{2}{\pi} \arccos(xx'+yy'+zz')$$

Using the notation $\bar{u} = (x,y,z)$ and $\bar{v} = (x',y',z')$, observe that $R(\bar{u},\bar{v})$ depends only on $\bar{u} \cdot \bar{v}$. We can also translate the eigenvalue problem (7) into this notation:

$$\iint R(\bar{u},\bar{v}) \psi_i(\bar{v}) dS_{\bar{v}} = \lambda_i \psi_i(\bar{u}) \quad (9)$$

Now the integration is taken with respect to \bar{v} over the surface of the unit sphere, and $dS_{\bar{v}}$ denotes surface measure. We can now formulate the

Theorem 1: Let

$$\psi_1(x,y,z) = x / \sqrt{x^2 + y^2 + z^2}$$

$$\phi_2(x, y, z) = y / \sqrt{x^2 + y^2 + z^2}$$

$$\phi_3(x, y, z) = z / \sqrt{x^2 + y^2 + z^2}$$

The functions ϕ_1 , ϕ_2 , ϕ_3 are eigenfunctions of (9), with eigenvalues all equal to π . Furthermore, these are the only eigenfunctions corresponding to this eigenvalue. \square

The proof of Theorem 1 is given as an Appendix, along with a description of the full class of eigenfunctions for eq. (9). Furthermore, a full class of eigenvalues has been computed numerically (see the details in the Appendix, Table 1). This computation reveals not only that π is the largest eigenvalue, but that all of the others are smaller than $\lambda=0.25$. Thus the approximation introduced by using only the first three basis functions to define the operator should be a good one.

The three basis functions are radially constant, and can be written in the simpler form:

$$\phi_1 = x/r$$

$$\phi_2 = y/r$$

$$\phi_3 = z/r$$

where

$$r = \sqrt{x^2 + y^2 + z^2} .$$

It is these basis functions, or, more exactly, a discrete approximation to them, which define our local operator. Two approximations are shown, the first in which the unit sphere is partitioned into a 3x3x3 unit cube (Fig. 2), and the second into a 5x5x5 cube (Fig. 3). Since the three operators are simple rotations of one another along the different axes, only the approximation to ϕ_1 (oriented along the X-axis)

is shown. ϕ_2 looks the same but is oriented along the y-axis, and ϕ_3 is oriented along the z-axis.

4. Applying the Operator

There are two stages in the application of the operator shown in Fig. 2 or 3: (i) the unit surface normal (a,b,c) defining the best edge through each voxel (α,β,γ) must be determined; and (ii) the quality of the match in (i) must be evaluated.

The surface normal for the best edge at (α,β,γ) is obtained by convolving the ϕ_i with the (not necessarily normalized) input image; i.e.,

$$a = \langle \phi_1, I \rangle = \iiint_S \phi_1(x,y,z) I(x-\alpha, y-\beta, z-\gamma) dx dy dz,$$

$$b = \langle \phi_2, I \rangle ,$$

$$c = \langle \phi_3, I \rangle .$$

The result of this convolution is the surface normal (a,b,c) at (α,β,γ) . If the input image had first been normalized, then the surface normal would be a unit vector pointing in the same direction. Thus the effect of the image normalization is to normalize the length of the computed surface normal, i.e. to multiply the vector by a scalar. It is this unit normal that provides a precise minimum for inequality (4).

Once (a,b,c) have been computed, the quality of the edge that they define must be evaluated. The fastest measure is the norm of (a,b,c) : a large value usually indicates a strong match (i.e., a high-contrast edge), while a small value indicates a

poor match. More precisely,

$$\| (a,b,c) \|_{\text{Euclidean}} = \| I'(x,y,z) \|_{L^2(S)}$$

where I' is the projection of the (unnormalized) image I onto M . Thus, the length of (a,b,c) is a coarse measure of the image contrast after projection. Fig. 4 contains a printout of the magnitude of (a,b,c) for a slice through a 3-D image of a cube using the $3 \times 3 \times 3$ operator. Note the typical response gradient across the edge, with the maximal response directly on (the dark side of) the edge. Thus a threshold can be used to select the maximal edge responses for this example, and the unit planes (eg. 2) through each voxel can then be displayed (Fig. 5). Note also that, since each unit surface passes through the center of a voxel, the borders of these unit planes may not coincide even though orientation is varying smoothly. (This artifact of the display process could be reduced by simply smoothing the edges of these unit surfaces.)

The resolution of the 3-D object boundaries that are eventually built out of these local surface elements is directly dependent on the sampling density. Fig. 6 contains a higher resolution image of the surface of a torus that was also made by thresholding the detector's responses. This display gives a much closer description of the object in the image than the coarsely-sampled cube in Fig. 5.

Our model for edges has two possibly restrictive features. First, it was derived using continuous mathematics, but is

applied in a discrete approximation to sampled imagery. The previous examples demonstrate that this approximation does not introduce any serious problems. The second restriction is that the model is based on a 3-dimensional step edge. However, practical experience with 2-dimensional imagery has shown that other edge configurations, such as roofs and peaks, also occur [16]. Thus it is necessary to evaluate the operator over a more realistic image, and Fig. 7 contains an example from computerized tomography. Note that both the overall magnitude and the directional components of the operator's response (Fig. 8) are in apparent agreement with the original intensity distribution.

Because of the weighted combinations of intensities within our operator, it is theoretically possible for them to combine in a way such that the norm of (a,b,c) is large, but the edge quality is poor. Such situations can be detected by computing an edge quality measure after the putative edge normal (a,b,c) has been determined. A very simple quality measure, the length of (a,b,c) has already been described. A slightly more elaborate quality measure Q is given by:

$$Q \{E_{a,b,c}(\alpha, \beta, \gamma)\} = \int_S E_{a,b,c}(x,y,z) I(x-\alpha, y-\beta, z-\gamma) dx dy dz$$

where

$$E_{a,b,c}(x,y,z) = \begin{cases} +1 & \text{if } ax+by+cz \geq 0 \\ -1 & \text{if } ax+by+cz < 0 \end{cases}$$

Another quality measure could be the mean square difference between the normalized input image centered at (α, β, γ) and the ideal edge $E_{a,b,c}$. However, in the experiments that we have conducted up to this point, the simplest quality measure, $|| (a,b,c) ||$, has proved satisfactory.

It is of course, well known that the interpretations of the responses obtained from edge operators evaluated over local portions of images are not unique, and that other kinds of processing are necessary to disambiguate them. One such possibility is to use a relaxation process analogous to that in [17], in which case these various quality measures could be used in computing the initial certainty of each possible interpretation. Or, one could evaluate a sequence of operators at different sizes, and attempt to parse the coarse through fine responses into various edge assertions [18].

5. Conclusions

The introduction of true three-dimensional images by modern scanning devices has created the need for appropriate image processing and analysis techniques. While it may be possible to extend existing 2-D techniques in certain circumstances, in others the 3-D structure may change the problem requirements. In this paper we began with a problem formulation in three dimensions, and derived a surface edge operator. Although this operator can be viewed as a generalization of a two-dimensional (Sobel) operator in retrospect, the optimal

properties of its 3-D counterpart would have been much harder to predict.

The design of an operator that responds to intensity differences is only the first stage in computing descriptions of surfaces in 3-dimensional images. Resource limitations in processing these images demand that the operator be evaluated over fairly small sub-images, and thus subsequent processing must be responsible for achieving appropriate consistencies between descriptions in neighboring sub-images. The operator described in this paper should be useful for providing input to these more global processes.

References

- [1] Stroke, G.W., Kock, W.E., Kikuchi, Y., and Tsujiuchi J., Ultrasonic Imaging and Holography-Medical, Sonar, and Optical Applications, Plenum Press, New York, 1974.
- [2] Gordon, R., Herman, G.T., and Johnson, S.A., Image reconstruction from projections, Sci. Amer., 1975, 233, 56-68.
- [3] Tanaka, K., and Tamura, S., A parallel processing system specialized in three-dimensional display based on serial tomograms, Proc. US-Japan Seminar on Research Towards Real-Time Parallel Image Analysis and Recognition, Tokyo, 1978.
- [4] Preston, K., New techniques in CT display, Proc. US-Japan Seminar on Research Towards Real-Time Parallel Image Analysis and Recognition, Tokyo, 1978.
- [5] Badler, N., and Bajcsy, R., Three dimensional representations for computer graphics and computer vision, Computer Graphics, 1978, 12, 153-160.
- [6] Rosenfeld, A., and Kak, A., Digital Image Processing, Academic Press, New York, 1976.
- [7] Liu, H.K., Two- and three-dimensional boundary detection, Comp Graph. and Image Proc., 1977, 6, 123-134.

- [8] Sunguroff, A., and Greenberg, D., Computer generated images for medical applications, Computer Graphics, 1978.
- [9] Duda, R., and Hart, P., Pattern Classification and Scene Analysis, Wiley, New York, 1973.
- [10] Hummel, R.A., Feature detection using basis functions, Comp. Graph. and Image Proc., 1979, 9, 40-55.
- [11] Heuckel, M.F., An operator which locates edges in digitized pictures, Journal of the ACM, 1971, 18, 113-125.
- [12] Heuckel, M.F., A local operator which recognizes edges and lines, Journal of the ACM, 1973, 20, 634-647.
- [13] Brown, J.L., Mean square error in series expansion of random functions, SIAM Journal, 1960, 8, 20-32.
- [14] Tou, J.T., and Gonzalez, R.C., Pattern Recognition Principles, Addison-Wesley, 1974.
- [15] Naylor, A.W., and Sell, G.R., Linear Operator Theory in Engineering and Science, Holt, Rinehart, and Winston, New York, 1971.
- [16] Herskovits, A., and Binford, T., On boundary detection, A.I. Memo 183, Artificial Intelligence Laboratory, M.I.T., July, 1970.
- [17] Zucker, S.W., Hummel, R.A., and Rosenfeld, A., An application of relaxation labeling to line and curve enhancement, IEEE Trans. Comp. 1977, C-26, 394-403 and 922-929.
- [18] Marr, D., Early processing of visual information, Phil. Trans. Roy. Soc. B., 1976, 275, 483-524.

Acknowledgements

We would like to thank Dr. T. Peters of the Montreal Neurological Institute for supplying the CT images, and H. Hubschman, N. Nassif, R. Szeliski, V. Brown, and M. Henchey for their programming assistance and graphics support.

Appendix: Proof of Theorem 1.

In this appendix we prove that (i) the ϕ_i given in Sec. 3 are the eigenfunctions corresponding to the three eigenvalues equal to π in eq. (9); and (ii) that these are the largest three eigenvalues. Thus these functions define the best approximation, in the Karhounen-Loeve sense, on which to base the edge operator. We further identify all of the eigenfunctions of the integral equation as the spherical harmonic polynomials of odd degree restricted to the unit circle.

The proof of (i) begins with the

Lemma: Let $\bar{e}_1 = (1,0,0)$. Then

$$\iint_R (\bar{e}_1, \bar{v}) \phi_i(v) dS_{\bar{v}} = \pi \delta_{1,i} \quad , \quad i=1,2,3 \quad (a1)$$

$$\text{where } \delta_{1,i} = \begin{cases} 1 & \text{if } i=1 \\ 0 & \text{otherwise} \end{cases} .$$

Proof: To evaluate the integral in eq. (a1), we must first specify the surface measure in appropriate coordinates. Let (x,θ) , $-1 \leq x \leq 1$, $-\pi \leq \theta < \pi$, be a coordinatization of the surface of the unit sphere in which a point (x,θ) on the sphere corresponds to a point (x,y,z) in 3-space according to the following relations (see Fig. (9)):

$$x = x$$

$$y = (1-x^2)^{1/2} \cos\theta$$

$$z = (1-x^2)^{1/2} \sin\theta$$

Note that the point (x, y, z) is necessarily on the unit sphere; i.e., $v = 1$. In terms of the (x, θ) coordinates, surface measure is simply $d\theta dx$. Thus the integral in eq. (a1) becomes:

$$\begin{aligned} & \iint R(\bar{e}_1, \bar{v}) \phi_1(\bar{v}) dS_{\bar{v}} \\ &= \int_{-1}^1 \int_{-\pi}^{\pi} \left[1 - \frac{2}{\pi} \arccos(x)\right] \cdot x \, d\theta dx \\ &= \int_{-1}^1 (2\pi x - 4x \arccos x) dx \\ &= -4 \int_{-1}^1 x \arccos x \, dx = -4\left(-\frac{\pi}{4}\right) = \pi, \end{aligned}$$

and

$$\begin{aligned} & \iint R(\bar{e}_1, \bar{v}) \phi_2(\bar{v}) dS_{\bar{v}} \\ &= \int_{-1}^1 \int_{-\pi}^{\pi} \left(1 - \frac{2}{\pi} \arccos x\right) (1-x^2)^{1/2} \cos\theta d\theta dx \end{aligned}$$

$$= \int_{-1}^1 \left(1 - \frac{2}{\pi} \arccos x\right) (1-x^2)^{1/2} \int_{-\pi}^{\pi} \cos \theta d\theta dx$$

$$= 0.$$

Similarly using $\phi_3(\bar{v})$ yields a $\sin \theta$ instead of $\cos \theta$ in the inner integral, with the same result. \square

Using this Lemma, we can now prove that the ϕ_i , $i=1,2,3$, are eigenfunctions corresponding to eigenvalues equal to π . We first show that ϕ_1 is such an eigenfunction.

Proposition 1.

$$\iint_{\mathbb{R}} (\bar{u}, \bar{v}) \phi_1(\bar{v}) dS_{\bar{v}} = \pi \phi_1(\bar{u}).$$

Proof. Fix $\bar{u} = (x_0, y_0, z_0)^t$ and let O be any orthogonal transformation (i.e., a rotation) taking \bar{e}_1 to \bar{u} . (There are an infinite number of them.) The 3x3 matrix representation

$$O = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix}$$

has determinate 1, which implies that the Jacobean of the transformation $\bar{v} = O\bar{v}$ is 1. Further, since $O\bar{e}_1 = \bar{u}$,

$$a_{11} = x_0$$

$$a_{21} = y_0$$

$$a_{31} = z_0 \cdot$$

Now, since R depends only on $\bar{u} \cdot \bar{v}$, and

$$\bar{u} \cdot \partial\bar{v} = \partial^{-1} \bar{u} \cdot \bar{v} = \bar{e}_1 \cdot \bar{v}$$

we have, by change of coordinates, that

$$\begin{aligned} & \iint R(\bar{u}, \bar{v}) \phi_1(\bar{v}) dS_{\bar{v}} \\ &= \iint R(\bar{u}, \partial\bar{v}) \phi_1(\partial\bar{v}) dS_{\bar{v}} \\ &= \iint R(\bar{e}_1, \bar{v}) \phi_1(\partial\bar{v}) dS_{\bar{v}} \end{aligned} \tag{a2}$$

To evaluate the integral (a2), we need to calculate $\phi_1(\partial\bar{v})$.

Recalling that $\phi_1(x, y, z) = x/r$, and noting that $\partial\bar{v}$ is a unit vector, we obtain

$$\begin{aligned} \phi_1(\partial\bar{v}) &= \frac{a_{11}x^1 + a_{12}y^1 + a_{13}z^1}{|\partial\bar{v}|} \\ &= a_{11}\phi_1(\bar{v}) + a_{12}\phi_2(\bar{v}) + a_{13}\phi_3(\bar{v}) \end{aligned}$$

Substituting this into the integral (a2) yields three terms:

$$\begin{aligned}
& a_{11} \iint_{\bar{v}} R(\bar{e}_1, \bar{v}) \phi_1(\bar{v}) dS_{\bar{v}} + a_{12} \iint_{\bar{v}} R(\bar{e}_1, \bar{v}) \phi_2(\bar{v}) dS_{\bar{v}} \\
& + a_{13} \iint_{\bar{v}} R(\bar{e}_1, \bar{v}) \phi_3(\bar{v}) dS_{\bar{v}}
\end{aligned}$$

which become using the lemma,

$$a_{11}\pi + 0 + 0 = \pi x_0 = \pi \phi_1(\bar{u}),$$

and the Proposition is proven. \square

The proof that ϕ_2 and ϕ_3 are also eigenfunctions follows directly from the observation that if $\phi(\bar{v})$ is an eigenfunction, then $\Psi(v) = \phi(\partial\bar{v})$ is also an eigenfunction. This completes the first part of the proof of Theorem 1; namely, that ϕ_1 , ϕ_2 , and ϕ_3 are the eigenfunctions corresponding to the eigenvalue π .

By using techniques similar to those just described, it can be shown that the full class of eigenfunctions for the integral equation (9) are the spherical harmonic polynomials of odd degree, restricted to the unit circle. More precisely, we have

Proposition 2. Let

$$Y_{\ell}^k(\theta, \phi) = \begin{cases} A_{\ell}^k \cos k\phi P_{\ell}^k(\cos\theta), & k \geq 0 \\ A_{\ell}^k \sin k\phi P_{\ell}^k(\cos\theta), & k < 0 \end{cases}$$

where Y_ℓ^k is the k^{th} spherical harmonic of order ℓ ,
 $-\ell \leq k \leq \ell$, $\ell \geq 0$, and for which A_ℓ^k are normalizing constants,
 $P_\ell^k(x)$ is the associated Legendre polynomial of degree ℓ , and

$$\begin{aligned}x &= r \cos \theta \sin \theta \\y &= r \sin \theta \sin \theta \\z &= r \cos \theta\end{aligned}$$

Then there exist constants $\{\lambda_\ell : \ell \text{ odd}\}$ such that

$$\{Y_\ell^k\}_{k=-\ell}^\ell = \{\phi_\ell\}_{k=-\ell}^\ell$$

are $2\ell + 1$ distinct eigenfunctions of (9), with eigenvalues all equal to λ_ℓ . Furthermore, the collection of eigenfunctions $\{\phi_\ell^k\}_{k,\ell}$ is complete. \square

We will only remark briefly about the proof of Proposition 2, because of its similarity to the proof of Proposition 1. We have already observed that an orthogonal transformation of coordinates carries an eigenfunction to an eigenfunction belonging to the same eigenvalue. Furthermore, the spherical harmonics transform into one another, with constant degree ℓ , under orthogonal coordinate changes, since the class of harmonic homogeneous polynomials is preserved under rotations. Thus, as was the case in the proof of Proposition 1, it suffices to check the integral equation for the $2\ell + 1$ functions at a

single point. All but one of these integrals turn out to be zero.

The particular correspondence between the three eigenfunctions given in Theorem 1 and the spherical harmonics is given by the following equations:

$$Y_1^0(\theta, \phi) = z/r$$

$$Y_1^{-1}(\theta, \phi) = x/r$$

$$Y_1^1(\theta, \phi) = y/r.$$

Now that we have determined the eigenfunctions of eq. (8), the proof of Theorem 1 requires a proof that π is the largest eigenvalue. This can be done by using an analytic expression for the spherical harmonics and actually performing the integration, or by using the known sequency properties of the spherical harmonics to conclude that eigenvalues decrease as the sequency increases. Then the above result follows from the known eigenvalue of π for the lowest harmonic. However, it is much more illuminating to numerically compute the eigenvalues for a discrete approximation to the kernel $R(\bar{u}, \bar{v})$, and the results of one such experiment are shown in Table 1. In this example $R(\bar{u}, \bar{v})$ has been represented by a 90x90 matrix, which is sufficiently fine for the largest

three eigenvalues to approximate π to within 1%. Furthermore, all other eigenvalues are less than $\lambda = .25$, which indicates that the approximation introduced by using only the first three eigenfunctions is a good one.





Fig. 1 A 2-D edge segment modeled as a unit line passing through the center of a square pixel (on the dark side).



$-\frac{\sqrt{3}}{3}$	$-\frac{\sqrt{2}}{2}$	$-\frac{\sqrt{3}}{3}$
$-\frac{\sqrt{2}}{2}$	-1	$-\frac{\sqrt{2}}{2}$
$-\frac{\sqrt{3}}{3}$	$-\frac{\sqrt{2}}{2}$	$-\frac{\sqrt{3}}{3}$

0	0	0
0	0	0
0	0	0

$\frac{\sqrt{3}}{3}$	$\frac{\sqrt{2}}{2}$	$\frac{\sqrt{3}}{3}$
$\frac{\sqrt{2}}{2}$	1	$\frac{\sqrt{2}}{2}$
$\frac{\sqrt{3}}{3}$	$\frac{\sqrt{2}}{2}$	$\frac{\sqrt{3}}{3}$

Fig. 2

3x3x3 operator



TOP SLICE



MIDDLE-SLICE BOTTOM-SLICE

Figure 7: Three consecutive CT images of a human head through the eyes.




Figure 8: The response of the 3x3x3 operator over the image in Fig. 7; the magnitude x-, y-, and z-projections of the surface normals are shown.




Figure 5: Display of the unit planes comprising the surface of a coarsely sampled cube.

FORTAN STOP
↓





Figure 6: Display of the unit planes comprising the surface of a torus.




Fig.3

5x5x5 operator

0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:		
0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:		
0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:		
0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:		
0:	5:	14:	21:	25:	25:	25:	25:	21:	14:	5:	0:	14:	35:	43:	49:	49:	49:	49:	43:	35:	14:	5:
0:	14:	35:	43:	49:	49:	49:	49:	43:	35:	14:	0:	14:	35:	43:	49:	49:	49:	49:	43:	35:	14:	5:
0:	21:	43:	25:	25:	25:	25:	25:	25:	43:	21:	0:	21:	43:	25:	25:	25:	25:	25:	43:	21:	0:	0:
0:	25:	49:	25:	0:	0:	0:	0:	0:	25:	49:	25:	25:	49:	25:	0:	0:	0:	0:	25:	49:	25:	0:
0:	25:	49:	25:	0:	0:	0:	0:	0:	25:	49:	25:	25:	49:	25:	0:	0:	0:	0:	25:	49:	25:	0:
0:	25:	49:	25:	0:	0:	0:	0:	0:	25:	49:	25:	25:	49:	25:	0:	0:	0:	0:	25:	49:	25:	0:
0:	21:	43:	25:	25:	25:	25:	25:	25:	43:	21:	0:	21:	43:	25:	25:	25:	25:	43:	21:	0:	0:	0:
0:	14:	35:	43:	49:	49:	49:	49:	43:	35:	14:	0:	14:	35:	43:	49:	49:	49:	43:	35:	14:	0:	0:
0:	5:	14:	21:	25:	25:	25:	25:	21:	14:	5:	0:	5:	14:	21:	25:	25:	25:	21:	14:	5:	0:	0:
0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:	0:

Fig. 4 Cross section of gradient vector magnitudes.




Figure 9: (x, θ) coordinates on the unit sphere.