Gradient-Oriented Profiles for Unsupervised Boundary Classification

Robert J. Tamburo
U. Pitt. Bioengineering
rjtst21@pitt.edu

George D. Stetten, M.D., Ph.D.
U. Pitt. Bioengineering – CMU Robotics Institute
george@stetten.com

Abstract

We present a method for unsupervised boundary classification by producing and analyzing intensity profiles. Each profile is created by sampling an ellipsoidal neighborhood of voxels oriented along the image gradient. The profile is analyzed via non-linear optimization to find the best fitting cumulative Gaussian. The parameters of the cumulative Gaussian parameterize the boundary directly yielding (1) extrapolated intensity values for voxels located far inside and outside of the boundary, (2) estimates boundary location and boundary width. For these parameters, intrinsic measures of confidence are established to eliminate low-confidence parameter estimates. Neighborhoods overlap considerably, yielding sufficient high-confidence estimates for a thorough survey of the boundary. Gradient oriented profiles are demonstrated on artificially generated three-dimensional test data and proved to accurately parameterize and classify the boundary.

1. Introduction

Boundary detection is a well-established field, described in a number of general references [1-4]. Many methods of image analysis employ a boundary detection scheme that first begins with the identification of boundary candidates by simple measurement of the gradient magnitude throughout an image and selection of candidates above a predetermined threshold. Such systems are sensitive to both object and background intensity, as well as overall image contrast. These issues may be overcome by providing a more detailed parameterization of the boundaries.

The desirable property of the gradient magnitude is its inherent insensitivity to rotation. More complex parameterizations of the boundary can also be made rotationally invariant by operating along the local gradient vector. *Steerable Filters*, for example, parameterize the boundary using a basis set of Gaussian derivatives that is easily rotated into the coordinate system of the local boundary [5]. We accomplish an equivalent rotation by projecting voxels onto the gradient vector to produce a one-dimensional intensity profile. Fitting a specific function to the profile, the cumulative Gaussian, directly yields physically significant estimates that parameterize the boundary. Specifically, the (1) extrapolated intensity values for voxels located far inside and outside of the boundary, (2) estimates of boundary location and boundary width. Individual measures of confidence are calculated for these parameters to eliminate low-confidence parameters. However, initial oversampling for candidate boundary points results in a large population of overlapping neighborhoods. Thus, sufficient high-confidence parameter estimates exist for a thorough survey of the boundary.

2. Methods

Our algorithm for boundary parameterization and classification proceeds in the following steps (corresponding to subsections in the text):

1. Find boundary candidates with a Difference of Gaussian (DoG) gradient detector.
2. Define ellipsoidal neighborhoods around detected boundary points and splat voxels within each ellipsoid into bins along the major axis, yielding a profile of voxel intensity.
3. Fit a cumulative Gaussian to the intensity profile, estimating values for intensity on both sides of the boundary as well as location and width of the boundary.
4. Eliminate blatantly "bad" profiles.
5. For the remaining profiles, calculate measures of confidence for the estimated values and eliminate those according to a certain threshold.
6. Classify the boundary using the high-confidence estimated values.

0-7695-0978-9/00 $10.00 © 2000 IEEE

2.1. Finding Boundary Candidates

A population of candidate boundary points is collected on a regularized sampling grid using a Difference of Gaussian (DoG) gradient detector. There are many forms DoG gradient detectors, a common one finds the difference between two concentric Gaussian kernels of different scale [1]. In contrast, our DoG consists of three pairs of same-scale Gaussian filters displaced in location along each of the cardinal axes to measure the respective components of the gradient. We use our DoG gradient detector because it is more efficient and delivers gradient direction as well as gradient magnitude. Each boundary candidate that meets a predetermined threshold for gradient magnitude is used to generate an individual boundary profile.

2.2. Generating a Boundary Profile

A boundary profile is generated in the direction of the gradient by sampling the voxels in an ellipsoidal neighborhood whose center is the boundary candidate, and whose major axis is the gradient vector. We project the voxels within the ellipsoidal neighborhood onto the major axis where their values are *splatted* into bins. Splatting is a technique commonly used in computer graphics that projects the footprints of individual voxels onto a plane. Overlapping footprints are collected in bins to form a rendered image. Instead of splatting onto a plane, we splat onto a line, namely, the major axis of the ellipsoid. As shown in Fig. 1 the ellipsoid is divided into disks along the major axis, with corresponding bins collecting footprints from voxels in adjoining disks. Since the footprints are wider than the bins, each voxel contributes to a number of bins. We use a triangular footprint, with the effect of linearly interpolating the value of each voxel between neighboring bins along the axis. The weight of contribution for each voxel is separately stored so that the total value in each bin can be normalized. Thus, the profile represents the average voxel intensity for each disk within the ellipsoidal sampling region. Averaging within each disk reduces the effect of image noise on the boundary profile.

2.3. Fitting the Profile and Estimating Parameter Values

A number of different functions could potentially be used to fit to the profile. The cumulative Gaussian was chosen for reasons now described. Most anatomical boundaries are very abrupt, inherently step functions at the sub-millimeter scale where one tissue ends and another begins. Image acquisition is inevitably limited in resolution, however, with a particular device exhibiting an overall "point spread function" usually at a significantly larger scale than the actual tissue boundary. Additional blurring may be performed intentionally or inevitably, for example, during the image processing that converts the raw data to an image. The result of these sequential convolutions tends to have the effect of convolution with a Gaussian kernel because of the Central Limit Theorem. Thus we can expect the step function of the anatomical boundary to reach the image analysis stage as a cumulative Gaussian. Therefore, by fitting a cumulative Gaussian, we may hope to, in effect, reverse the blurring and parameterize the original boundary.

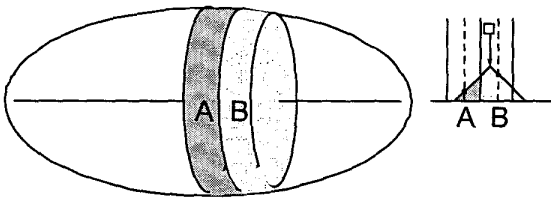

Fig. 1. Left shows an ellipsoidal neighborhood; whose center is the boundary candidate and major axis is the gradient vector, with two adjoining disks, A and B. Right shows a voxel (square) projecting onto the major axis, splatting ¾ of its value into bin B and ¼ into bin A.

Fitting a cumulative Gaussian to a boundary profile requires optimizing 4 parameters: (1) standard deviation, corresponding to the boundary width, (2) mean, corresponding to boundary location, and (3,4) two asymptotic values, corresponding to voxel intensity sufficiently far on either side of the boundary to be unaffected by blurring.

The cumulative Gaussian $C(x)$ is derived as follows: The normalized Gaussian,

$$G(x)=\frac{1}{\sigma\sqrt{2\pi}}e^{-\frac{(x-\mu)^2}{2\sigma^2}}, \qquad (1)$$

is integrated to yield the error function (*erf*),

$$\int_{-x}^{x} G(v)dv = erf\left(\frac{x-\mu}{\sigma\sqrt{2}}\right), \qquad (2)$$

which is scaled and offset to yield the cumulative Gaussian $C(x)$ as follows:

$$C(x)=I_1+\frac{I_2-I_1}{2}\left(1+erf\left(\frac{x-\mu}{\sigma\sqrt{2}}\right)\right) \qquad (3)$$

The four fixed parameters of $C(x)$ are μ (mean), σ (standard deviation), and I_1 and I_2 (asymptotic voxel values inside and out). On one side of the boundary

$$erf(-\infty) = -1,\ C(-\infty) = I_1 \tag{4}$$

and on the other side

$$erf(\infty) = 1,\ C(\infty) = I_2. \tag{5}$$

The four parameters are labeled in Fig. 2, showing a particular fit of the cumulative Gaussian along the sampled portion of a boundary profile.

A variety of techniques are available for optimizing the fit of a non-linear function to a set of sample points. We chose a Quasi-Newton non-linear optimization algorithm provided as part of the *AD Model Builder* from Otter Research, Inc., [6] because it does not require explicit derivation of derivatives and is acceptably rapid and robust. To make intelligent use of the parameter estimates from the curve fitting, measures of confidence for individual parameters are established. This takes the form of two steps, as described in the following two sections.

2.4. Eliminating Bad Profiles

The first step involves the elimination of unacceptable parameters by rejecting blatantly "bad" profiles for which a reasonable fit of the cumulative Gaussian could not be found. The fit is rejected if:

1. Either extrapolated voxel intensity value I_1 or I_2 falls outside of the acceptable range of values for the imaging modality.
2. The estimated boundary location μ falls outside of the ellipsoidal sample region.

Boundary profiles that are rejected by these criteria are no longer considered. The second method determines the confidence level for the parameter estimates and eliminates low-confidence estimates.

2.5. Establishing Intrinsic Measures of Confidence

If a fit was accepted for a given profile, the next step was to determine a measure of confidence in the individual parameters of the cumulative Gaussian. We postulate that the confidence for a particular parameter is

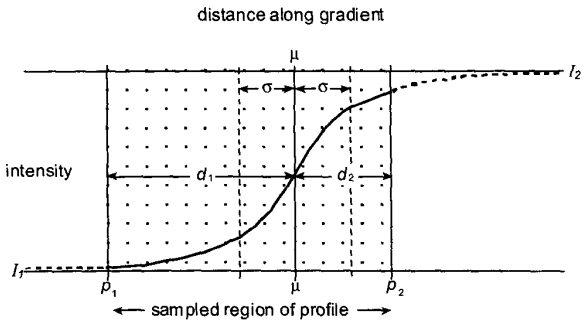

Fig. 2. A cumulative Gaussian fit to a profile from a sampled region (shaded area).

determined by the location of the boundary within the ellipsoid, as well as the boundary width.

In Figure 2, the cumulative Gaussian that best fits an intensity profile is shown. By definition the function is symmetrically distributed on either side of the mean μ (with μ representing the optimum boundary location). However, the function is not generally symmetrically distributed with respect to the sampled region of the profile, represented by the shaded area from p_1 to p_2. Given that the best fit has a mean μ and standard deviation σ, two distances d_1 and d_2 can be defined from μ to p_1 and p_2 respectively, from which two normalized distances z_1 and z_2 can be computed.

$$z_1 = \frac{d_1}{\sigma} \text{ and } z_2 = \frac{d_2}{\sigma}. \tag{6}$$

The calculations of z_1 and z_2 indicate how many standard deviations from μ in each direction the actual samples extend, and serve as measures of confidence for I_1 and I_2 respectively. For each direction, the greater the number of standard deviations the samples extend from μ, the less the effect of the boundary and the more accurately we can expect to estimate intensity. For example, in Fig 2 we have more confidence in the estimate I_1 than the estimate I_2, because $d_1 > d_2$, and therefore $z_1 > z_2$.

In addition to a measure of confidence for I_1 and I_2, we have developed one for μ. Our measure of confidence for μ determines whether sufficient samples exist on either side of μ to estimate it accurately. We define this measure of confidence as

$$z_{\min} = \min(z_1, z_2). \tag{7}$$

If on either side of μ, insufficient samples exist to anchor the cumulative Gaussian, $z_{\min}$ will be small, and we can expect difficulty in estimating μ accurately. Like the measures of confidence for profile intensity estimates, a threshold is also placed on $z_{\min}$ to determine whether the boundary location will be estimated accurately with a high confidence in μ.

2.6. Classifying the boundary

Armed with high-confidence estimates for I_1, I_2 and μ, we can now classify the local boundary in terms of these parameters. The parameters I_1 and I_2 estimate voxel intensity beyond the ellipsoid, sufficiently far from the boundary to be stable. The mean μ estimates the anatomical boundary location within the ellipsoid.

Rather than struggling to extract the last bit of classification from each profile, our measures of confidence permit us to exclude individual parameters on a case-by-case basis. In the initial gathering of boundary candidates, we use a sampling interval well below the conventional Nyquist guidelines, to yield a comprehensive set of profiles from which a thorough examination of a given boundary can be made. By setting thresholds for the intrinsic measures of confidence, we may reject or accept individual estimates to make intelligent use of this over-sampled population. The exact method of setting these thresholds optimally is beyond the scope of this paper.

3. Results

We demonstrate our method on a 3D data set, 100 voxels on a side, consisting of 8-bit voxels. An artificial sphere was generated with a radius of 30 voxels, an interior value of 64, and an exterior value of 128. A binomial kernel was applied multiple times to the data, approximating a Gaussian and simulating the point-spread function of a real image. Using the DoG gradient detector kernel, as described in section 2.1, a total of 1,400 boundary candidates were found. According to the method described in section 2.2, voxels in an ellipsoidal neighborhood were sampled for each boundary candidate. The center of the ellipsoid was the location of the boundary kernel. The major axis of the ellipsoid had a length of 10 voxels and the minor axes had lengths of 6 voxels each. Along the major axis, 10 bins were established, each being 1 voxel wide. The intensity values of the voxels within the ellipsoid were splatted into the bins (using a triangle footprint), and normalized by the splatting weights. A cumulative Gaussian was fit to the resulting intensity profile as described in section 2.3.

Of the resulting 1,400 profiles, 240 profiles were eliminated because they did not meet the constraints defined in section 2.4. For the remaining profiles, the measures of confidence for the classification parameters were calculated as derived in section 2.5.

In summary, gradient-oriented profiles were applied to a 3D test model. The results were then compared to the know parameters of the test data, as described below.

3.1. Measuring Accuracy in Test Models

After eliminating "bad" profiles, the accuracy of the remaining profile parameters was measured using a priori knowledge of the test data. Being a sphere, the test data is parameterized by its radius (the distance R_{true} from its center to the location of the boundary), as well as the intensities on both sides of the boundary. The interior intensity of the sphere is denoted I_1, and the exterior, I_2. Assuming the estimate of intensity derived from the profile for either I_1 or I_2 is denoted $I_{estimate}$, we define I_{error} as

$$I_{error} = \left| I_{true} - I_{estimate} \right|, \tag{8}$$

where I_{true} is the equivalent known voxel intensity I_1 (64 inside the sphere) or I_2 (128 outside the sphere).

Fig. 3 shows that profiles with μ near the center of the ellipsoid are able to estimate the intensity on either side of the boundary better than when μ is located away from the center of the ellipsoid. When μ is located away from the center of the ellipsoid, intensity values are estimated less accurately in that direction. However, such a profile will estimate the intensity in the opposite direction with high accuracy. For example, in Fig. 3, I_2 can be estimated to within one unit of intensity, when μ is less than 0.

The accuracy in estimating the true boundary location of the test data was determined as follows: The distance from the center of the sphere to the location on the major axis of the ellipsoid corresponding to μ was defined as the estimated radius $R_{estimate}$. This was compared to the known radius R_{true} of the sphere, with R_{error} being defined as

$$R_{error} = \left| R_{true} - R_{estimate} \right|. \tag{9}$$

Fig. 4 shows that boundary profiles are more accurate in determining the true location of the boundary than the DoG boundary kernels. Out of the 1,160 total profiles, 1,152 (99.31%) had location estimates were within 1 voxel.

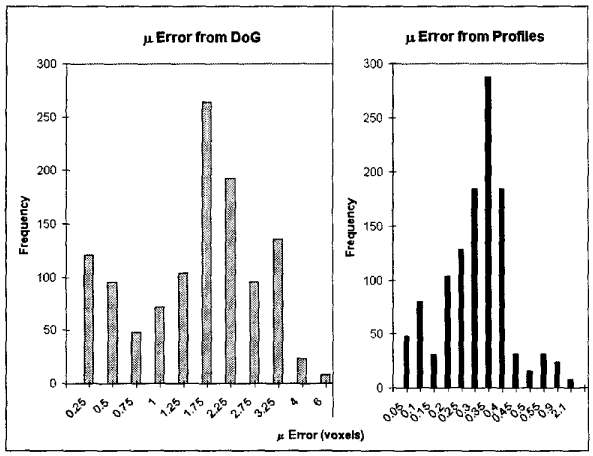

Fig. 4. Accuracy in estimating the location of the boundary by the DoG gradient detector as compared to the gradient oriented profiles.

3.2. Judging Classification Reliability Intrinsically

We have shown that gradient-oriented profiles are able to estimate the intensity values and localize the boundary location. We now determine whether these profile parameters are reliable using intrinsic measures of confidence.

The measures of confidence for the profile parameters were calculated using the methods discussed in section 2.5. Fig. 5 shows the error in estimating I_1 (interior of sphere) as a function of its measure of confidence for z_1. When z_1 is greater than 2.0, the error in estimating I_1 $I_{error} < 5$. Therefore, as seen in Fig. 5, a threshold placed on z_1 of 2.0 would guarantee this accuracy. With this threshold, 776 of the total profiles (67%) will estimate the intensity I_1 with high confidence. Similarly, Fig. 6 shows the error in estimating I_2 (exterior of sphere) as a function of its the measure of confidence z_2, with similar results using a threshold for z_2 of 1.45.

We predicted earlier that profiles with μ near the center of the ellipsoid should be able to estimate both intensities I_1 and I_2 with high accuracy. We identified 632 out of the 1,160 (54.48%) profiles that were able to reliably classify both intensities as described above. Our measure of confidence for μ is shown in Fig. 7. For z_{min} > ½ standard deviations a high confidence seems warranted for μ, yielding boundary location within 1 voxel. This threshold places high confidence in the estimates of boundary location for 1,152 out of 1,160 profiles (99.31%).

4. Conclusion

We have presented gradient-oriented profiles, a method of unsupervised boundary parameterization suitable for classification. By operating along the direction of the gradient vector, gradient-oriented profiles are insensitive to rotation.

The cumulative Gaussian provides estimates that parameterize the boundary by estimating extrapolated intensity values for voxels located far inside and outside of the boundary. The cumulative Gaussian also estimates anatomical boundary location and (although not yet validated) blurred boundary width. Intrinsic measures of confidence have been shown capable of eliminating inaccurate parameter estimates. Thresholds placed on these measures of confidence would allow classification of boundaries with high confidence.

Gradient-oriented profiles have proven capable of accurately parameterizing the boundary in our test data. To demonstrate the application of gradient-oriented profiles to classification, they must be applied to data containing multiple boundaries and noise. Also, more investigation is warranted into the effects of (1) using neighborhoods with shapes other than an ellipsoid, (2) using footprints with shapes other than a triangle, (3) altering bin size along the gradient, (4) altering relative size of the sampling neighborhood compared to the DoG gradient kernel. The application of gradient-oriented profiles to medial feature detection using core atoms [7, 8] is being investigated.

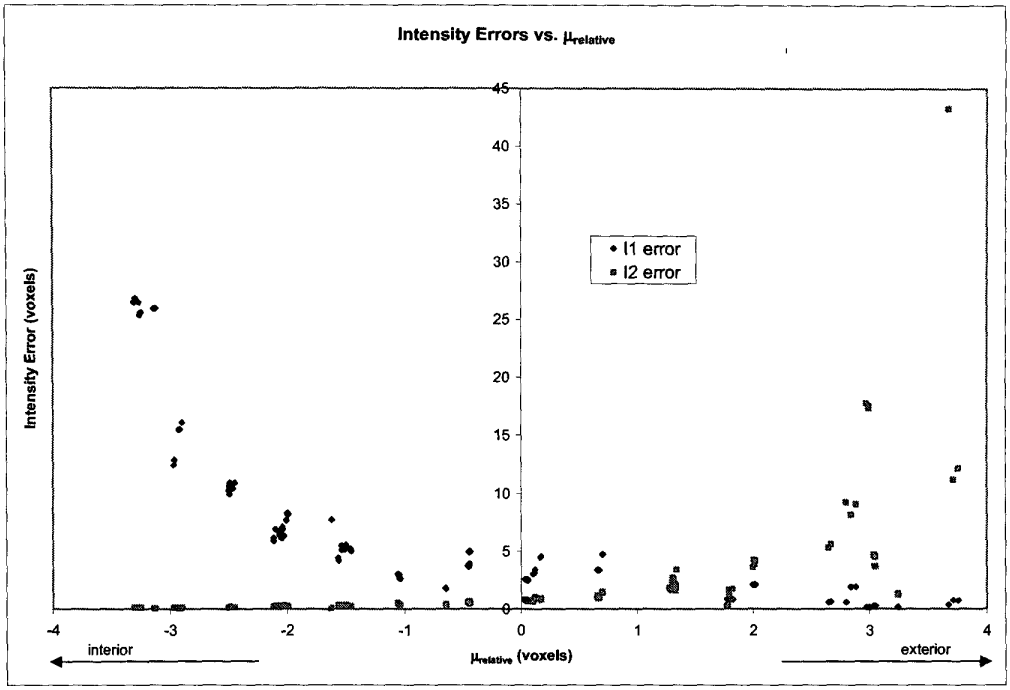

Fig. 3. The distribution of error in estimating the intensity values on either side of the boundary as a function of μ (with zero being the center of the ellipsoid).

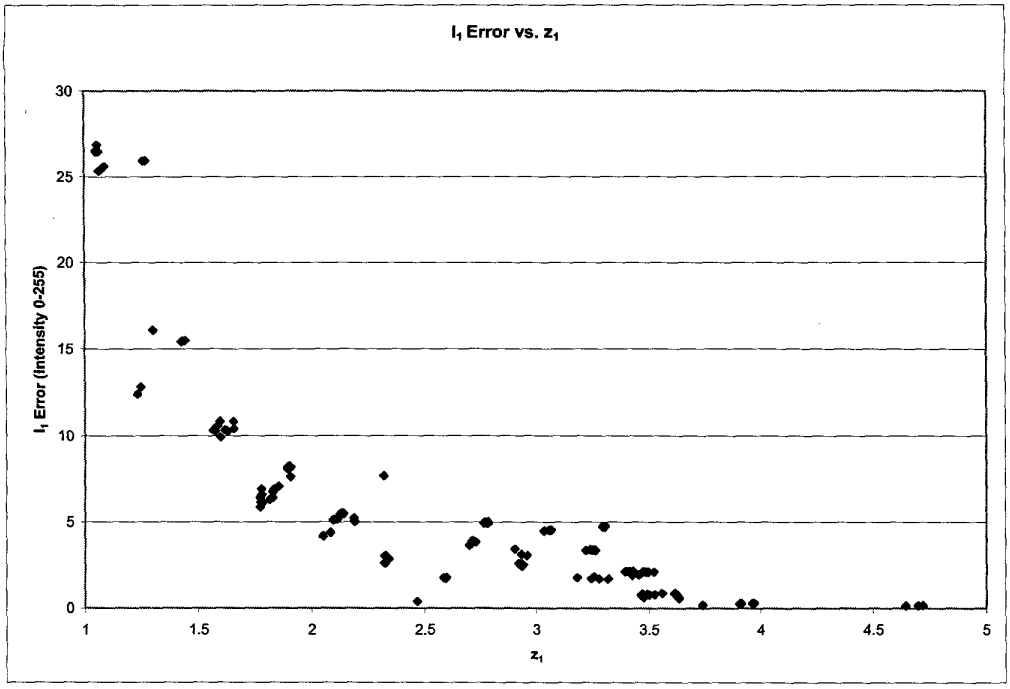

Fig. 5. Estimating the error in estimating the intensity for the interior of the sphere I_1 versus its measure of confidence z_1.

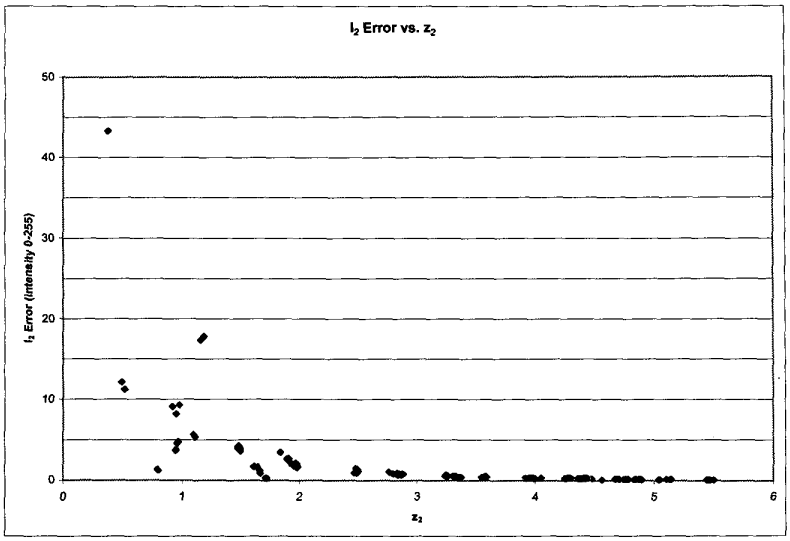

Fig. 6. The error in estimating the intensity exterior to the sphere I_2 as a function of its measure of confidence z_2.

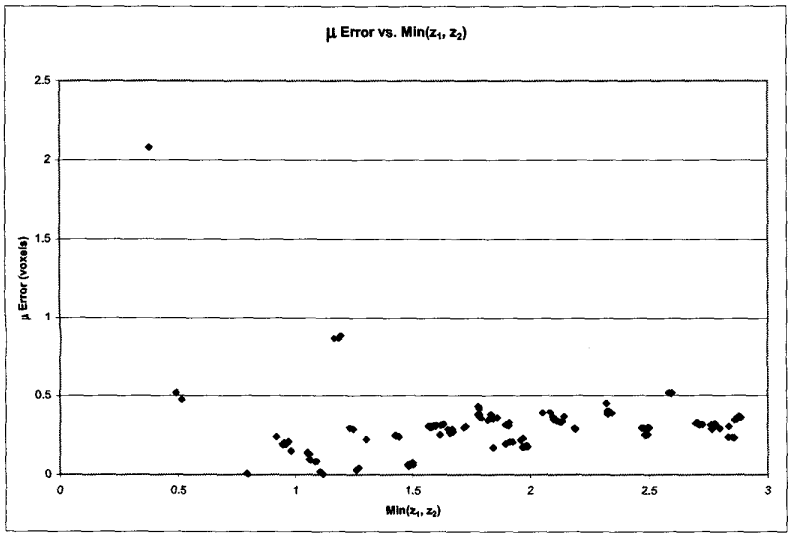

Fig. 7. Error in boundary location from gradient-oriented profiles versus measure of confidence.

5. References

[1] Mlsna, P. and J. Rodriguez, *Gradient and Laplacian-Type Edge Detection.* From Handbook of Video and Image Processing, ed. Al Bovik, Academic Press, San Diego: pp. 415-447, 2000.

[2] Petrou, M. and P. Bosdogianni, *Image Processing: The Fundamentals,* John Wiley and Sons, New York, 1999.

[3] Pratt, W., *Digital Image Processing*, John Wiley and Sons, New York, 1991.

[4] Jain, A., *Fundamentals of Digital Image Processing*, Prentice Hall, New Jersey, 1989.

[5] Freeman, W. and E. Adelson, "The Design and Use of Steerable Filters". IEEE Transactions in Pattern Analysis and Machine Intelligence, 1991, 13(9): pp. 891-906.

[6] AD Model Builder, *Otter-Research Ltd.,* http://otter-rsch.com/admodel.htm.

[7] Stetten, G. and S.M. Pizer, "Medial-Node Models to Identify and Measure Objects in Real-Time 3-D Echocardiography", IEEE Transactions on Medical Imaging, 1999. 18(10): pp. 1025-1034.

[8] Stetten, G., R. Landesman, and S. Pizer, "Core-Atoms and the spectra of scale", SPIE Medical Imaging Conference, 1997. 3034, part 2: pp. 642-652.