Image segmentation using the student's t-test and the divergence of direction on spherical regions

George Stetten[1, 2], Samantha Horvath[1], John Galeotti[2], Gaurav Shukla[1], Bo Wang[1], Brian Chapman[1]

[1] University of Pittsburgh, [2] Carnegie Mellon University

ABSTRACT

We have developed a new framework for analyzing images called *Shells and Spheres* (SaS) based on a set of spheres with adjustable radii, with exactly one sphere centered at each image pixel. This set of spheres is considered optimized when each sphere reaches, but does not cross, the nearest boundary of an image object. Statistical calculations at varying scale are performed on populations of pixels within spheres, as well as populations of adjacent spheres, in order to determine the proper radius of each sphere. In the present work, we explore the use of a classical statistical method, the student's t-test, within the SaS framework, to compare adjacent spherical populations of pixels. We present results from various techniques based on this approach, including a comparison with classical gradient and variance measures at the boundary. A number of optimization strategies are proposed and tested based on pairs of adjacent spheres whose size are controlled in a methodical manner. A properly positioned sphere pair lies on opposite sides of an object boundary, yielding a direction function from the center of each sphere to the boundary point between them. Finally, we develop a method for extracting medial points based on the divergence of that direction function as it changes across medial ridges, reporting not only the presence of a medial point but also the angle between the directions from that medial point to the two respective boundary points that make it medial. Although demonstrated here only in 2D, these methods are all inherently *n*-dimensional.

1. MEDIAL RIDGES IN IMAGE ANALYSIS

A primary goal of the research presented here is to extract medial ridges from images. The lineage of the medial approach may be traced to the medial axis (otherwise known as the symmetric axis or skeleton) introduced on binary images by Blum and developed by Nagel, Nackman, and others.[1-3] A classic illustration of the Blum medial axis of a rectangle is shown in Figure 1. The dotted lines represent the locus of points equidistant from two or more boundary points of the rectangle. Also shown are a few of the circles whose centers lie on the medial axis and whose circumferences touch but do not cross the rectangle's boundary. Pizer extended the medial axis to gray-scale images, producing a graded measure called *medialness*, which links the aperture of the boundary measurement to the radius of the medial axis to produce what has been called a *core*. A core is a locus in a space whose coordinates are position, radius, and associated orientations.[4-5] Methods involving these continuous loci of medial primitives have proven particularly robust against noise and variation in target shapes.[6] Determining locations with high medialness and relating them to the core has been accomplished by analyzing the geometry of loci resulting from ridge extraction.[7] Models including discrete loci of medial primitives have also provided the framework for a class of active shape models known as *deformable m-reps* (sampled medial representations),[8,9] as well as a statistical approach using pairs of detected boundary points known as *core atoms* developed previously by author Stetten.[10]

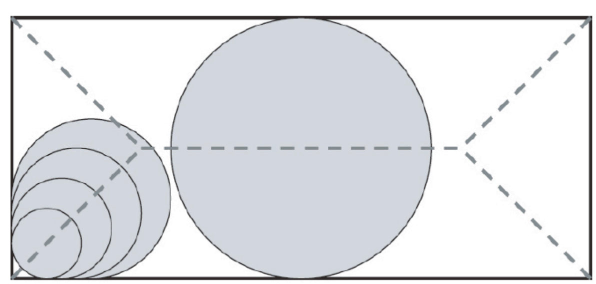

Fig. 1 Blum medial axis of a rectangle (dotted lines), being the centers of all circles that touch the boundary in more than one place, but do not cross the boundary.

2. SHELLS AND SPHERES

In our recent work with medial based image analysis, we have developed what we call the Shells and Spheres (SaS) framework.[11-12] In this framework, spheres centered at every pixel grow or shrink by adding or deleting an outer shell one pixel thick, as they use statistical measures of their contents to reach, but not cross, object boundaries in the image. Unlike conventional fixed-scale kernels, SaS operators consider as many pixels as possible to differentiate between objects and delineate boundaries. We use the word "sphere" here for convenience, since the approach is not limited to 3D and in fact is valid in *n* dimensions. We depict "spheres" in 2D for clarity in Figures 2 and 3, which show spheres with solid circles and shells with dashed circles. Note also the use of bold non-italics, **x** and **y**, to denote the *n*-dimensional (in this case 2-dimesional) locations of the sphere centers. We stay in 2 dimensions for the rest of this paper.

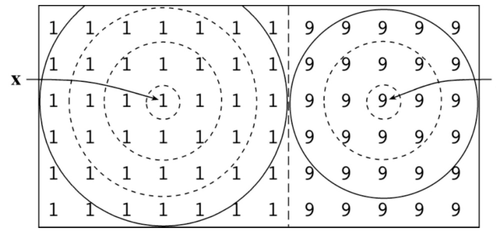

Fig. 2 Spheres centered at **x** and **y** made from concentric shells, reaching but not crossing a boundary in noiseless image.

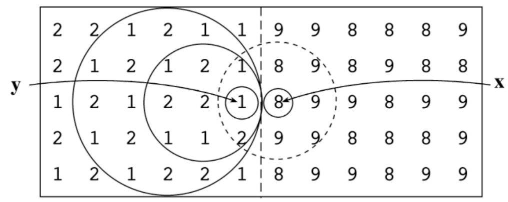

Fig. 3 In a noisy image, a set of spheres containing pixel at location **y** rejects a sphere centered at **x** from growing across boundary because its population is different.

Figure 2 shows spheres on opposite sides of a boundary between two objects with uniform intensities 1 and 9 in a noiseless image. Unlike fixed-scale kernels, the SaS framework avoids edge effects produced by conventional convolution-based operators. For example, in the left sphere (solid circle) in Figure 2, the population is truncated by the edge of the image, but it still has an unambiguous mean and standard deviation. These computations require no assumptions about pixel values outside the image.

Figure 3 depicts part of the SaS algorithm used in our previous research. The boundary between two objects is shown, now with noise in the image intensity (one object has intensities 1 and 2, the other has intensities 8 and 9). A set of spheres (solid circles) is shown to the left of the boundary, each containing pixel **y** and each reaching but not crossing the boundary. The solid circles show the correct behavior for spheres with optimal radii. The sphere at pixel **x** is attempting to add a new shell (dashed circle) and thereby join the set of spheres to its left, but it will be "repulsed" by the fact that its statistics do not match that of the other spheres. Thus the radius of the sphere at **x** will not be increased, and it will not grow across the boundary. The overall algorithm is quite complex, requiring a number of parameters that must be optimized,[12,13] but it has produced useful results.[13-16] For example, Figure 4 shows the algorithm segmenting the surface of the aorta in computerized tomography (CT) data. The image on the left is a sagittal slice through a noisy CT image of the thorax with contrast in the heart and great vessels. The image in the center shows the application of SaS removing noise while leaving sharp boundaries and uniform grayscale values for various regions. In this processed image, segmentation of the aorta and the innominate artery (pink) was achieved by a flood-fill operation through the centers of medial spheres, namely, spheres that reach more than one boundary, starting with a single manually-placed seed point in the aorta. The image on the right is a surface rendering of the union of all of the 3D spheres within the same vessel. We will not discuss this method further here, because it is described in detail elsewhere,[12] and because the present work represents a new and different approach to using the SaS framework.

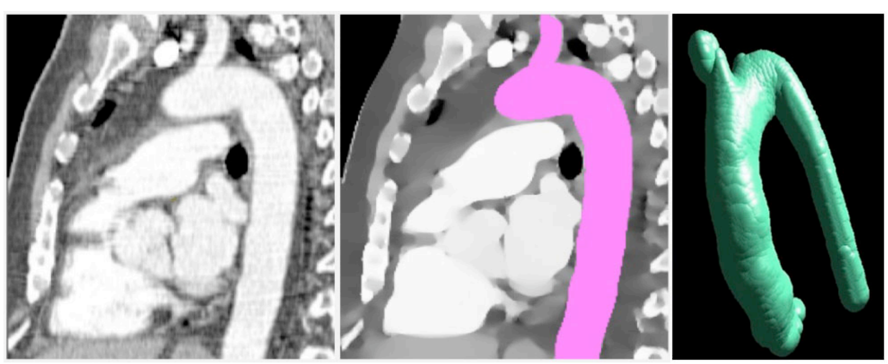

Fig. 4 *left*: noisy CT slice of thorax; *center*: 2D segmentation of aorta with Shells and Spheres; *right*: 3D segmentation with rendered surface.

3. THE STUDENT'S T-TEST

In the present work, we develop a set of simpler algorithms within the SaS framework that show promise in finding boundaries and medial manifolds in the presence of noise. We began with a classical statistic, the student's t-test for two non-pooled populations, to analyze adjacent populations of pixels. We used the standard definition

$$t = \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{\left(s_1^2/n_1\right) + \left(s_2^2/n_2\right)}} \tag{1}$$

where $\bar{x}$ is the mean intensity, s^2 the variance, and n the sample size of two spherical populations of pixels. The student's t-test statistic may be used along with the sample size to determine the P-value for pairs of spheres at arbitrary scale (radius). The goal is to find adjacent non-overlapping spherical populations of pixels that minimize the P-value and thus reject the null hypothesis that the two spheres are within the same object, given the presence of Gaussian noise.

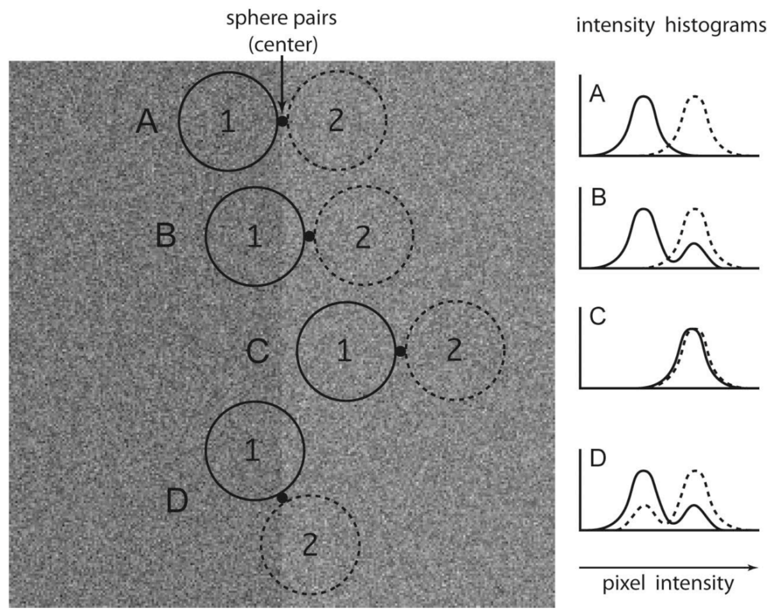

Fig. 5 Sphere pairs on an image with a boundary and noise. Each pair consists of two spheres (1 and 2) with a center point between them. Corresponding intensity histograms are shown to the right.

Figure 5 demonstrates the concept with four sphere-pairs (labeled A, B, C, and D) placed at different locations and orientations in a noisy image. The image has two regions separated by a vertical boundary, representing objects of differing mean intensity with Gaussian noise superimposed. A sphere-pair is defined in this case to be "located" at the "center" point (large black dot) between the two circles ("1" solid and "2" dashed). Corresponding intensity histograms for each sphere-pair are shown to the right. Sphere-pair A is located exactly on the boundary, so its spheres each encompass a region completely within one of the objects. Note that the corresponding histograms of pixel intensity overlap, but their means are clearly separated. Sphere-pair B is shifted to the right so that sphere B-1 now includes pixels from both objects. As seen in the corresponding histogram B, the variance of sphere B-1 (solid line) is correspondingly larger, decreasing the t-test statistic (see Eq. 1) and increasing the corresponding P-value. The two populations B-1 and B-2 are thus less likely to represent different objects than A-1 and A-2. Sphere-pair C has shifted even further to the right so that both spheres C-1 and C-2 are samples of the same population. Histogram C shows that the corresponding populations are very similar, and the difference between their means will be close to 0. Therefore the t-test statistic will be near zero as well, and the corresponding P-value will be close to 1, i.e., the null hypothesis that the samples represent the same underlying population will be very likely. Sphere-pair D shows the effect of rotating sphere-pair A to a less than optimal orientation relative to the boundary. The histogram shows that both spheres D-1 and D-2 now contain pixels from both objects, lowering the t-test

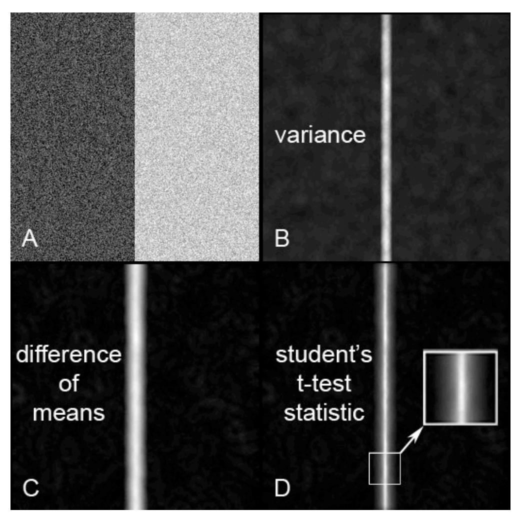

Fig. 6 Noisy image (A), variance in a single sphere (B), and statistical measures from pairs of spherical regions (C,D).

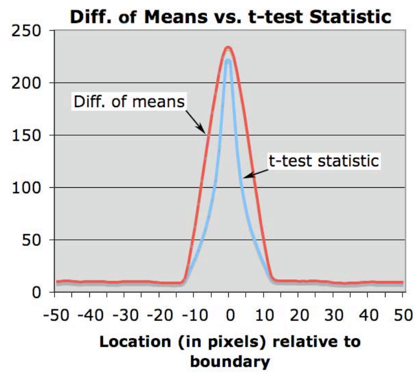

Fig. 7 Average measure for all rows in Fig. 5C and 5D, showing superior resolution for t-test.

statistic and raising the P-value. At any given location, we can find the scale (sphere size) and orientation for the sphere pair with the highest student's t-test statistic and lowest P-value, and we can use this process to identify boundaries.

With custom software that we constructed in Java, we have tested this concept on a 256 x 256 pixel image containing two regions with differing means separated by a vertical boundary in the presence of Gaussian noise (Figure 6A). Based on a constant radius of 6 pixels, the variance within spheres at each location is shown in Figure 6B. As would be expected, the variance is high for individual spheres that cross the boundary and lower (and nearly constant) in all other areas (since the same Gaussian noise was added to both sides). Figure 6C shows the maximum difference-of-means (absolute value) for sphere pairs centered at all locations throughout the image, over all possible orientations for a given sphere pair. This is the absolute value of the numerator in Equation 1, or $\left|\bar{x}_1 - \bar{x}_2\right|$. For sphere-pairs that cross the boundary, the highest value is with the sphere-pair oriented horizontally. Elsewhere, it is near zero. Figure 6D shows the student's t-test statistic, as defined in Equation 1, the difference-of-means divided by a denominator that increases with variance in either sphere. Thus as one moves away from the boundary even a little, as depicted by sphere-pair B in Figure 5, the variance of one of the spheres increases, reducing the t-test statistic. As detailed by the insert in Figure 6D, the t-test statistic produces a very sharp measure of boundariness, even in the presence of noise. Figure 7 shows this in graphical form, with the difference-of-means and the t-test statistic (both are absolute values and in arbitrary units) each averaged over all pixel rows. The t-test statistic clearly has a narrower peak, because it takes advantage of variance.

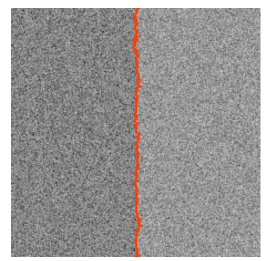

Fig. 8 Vertical boundary in noise.

Figure 8 shows an automated determination of the boundary (red line) using the same sphere-pairs that generated Figure 6D (all spheres with radius of 6 pixels). The fact that a single vertical boundary is known to exist beforehand permits a simple search for the maximum t-test statistic along each horizontal row of pixels. Granted, this is an unrealistically simple task. We will develop more general and complex methods in the following sections.

4. COMPETITIVE ELIMINATION OF SYMMETRIC SPHERE-PAIRS

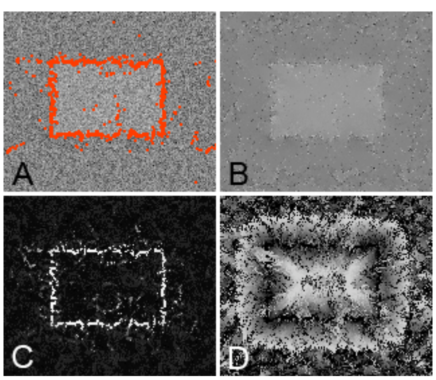

Fig. 9 Segmentation using P-values and symmetric sphere-pairs (see text).

To find boundaries when no prior information is know about their location or orientation, we have devised a method for searching through the space of sphere-pair radius and orientation to find the largest value at each location for the student's t-test of that sphere-pair. By location, we mean the center point of the sphere-pair (see Fig. 5), and we assume (for now) symmetric sphere-pairs, i.e., the radius is the same for the two spheres in a given pair. Since radius is allowed to vary in this search, we use the P-value derived from the student's t-test, instead of the t-test value itself. This is to correct for sample size, which varies with sphere radius, yielding a normalized statistic representing the probability of rejecting the null hypothesis (that the two spheres in a given pair are samples of the same population).

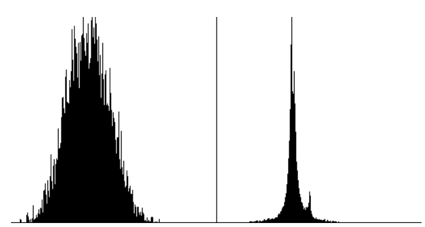

Fig. 10 Histograms (left) of raw image intensity and (right) mean intensity in spheres from sphere-pair competition.

A competitive elimination scheme is used to find local minima in P-values along putative boundaries by inactivating sphere-pairs with greater P-values, as the spheres in the sphere-pairs are allowed to grow from their smallest scale. All sphere-pairs begin active, and then at each successive scale, a sphere-pair is inactivated if the P-value at the center of either of its spheres is less than that of the sphere-pair itself. Thus stronger sphere-pairs (with lower P-values) eliminate weaker ones. The scheme has the advantage that as spheres grow larger, and thus computationally more expensive, fewer of them remain active. Figure 9A shows resulting boundary points in red (surviving sphere-pairs) for a rectangular area in noise; 9B shows the means of the constituent spheres reducing noise while leaving sharp boundaries; 9C shows the t-test statistic of every sphere pair, being highest at the boundary; 9D shows the radius of the optimum sphere pairs (black is 0 at the boundary, increasing with distance from the boundary). Figure 10 compares the histogram of the original image (Fig. 10A) to that of the means of the optimum spheres (Fig. 10B), demonstrating the extraction of the object from the noise.

5. ASYMMETRIC SPHERE PAIRS

We now describe a different approach, using sphere-pairs that are asymmetric. This approach has proven the most successful so far, especially in terms of its ability to extract the medial manifold, and we devote the rest of this paper to describing it. As opposed to the symmetric pairs previously described (see Fig. 5), in which the two spheres in a given pair are always the same size, each pair of adjacent circles now has an "outer" sphere (dashed circle in Fig. 11) held at a constant small radius, and an "inner" sphere (solid circle) that increases in radius. The location of the asymmetric sphere-pair is considered to be the center of its inner sphere, and thus a family of possible sizes for the inner sphere and orientations for the outer sphere are possible at each location. Two such pairs at a given location are shown in Figure 11, with the pair containing the larger inner sphere (solid circle) correctly finding the boundary with one of its outer spheres. The constant radius of the outer sphere is chosen to be small enough to provide sufficient boundary curvature while still being large enough to represent a statistically significant population. As the inner sphere is grown to each new radius, every possible orientation of outer sphere is tested, and an overall maximum student's t-test is used to identify the optimum sphere-pair. This optimum pair should have its inner sphere just touching the nearest boundary and its outer sphere just on the other side of that boundary. As will be discussed below, we actually have modified the student's t-test so as not to favor larger inner spheres. Once an optimum sphere-pair has been found for each pixel, we are ready to identify boundaries and medial ridges, as will be described next.

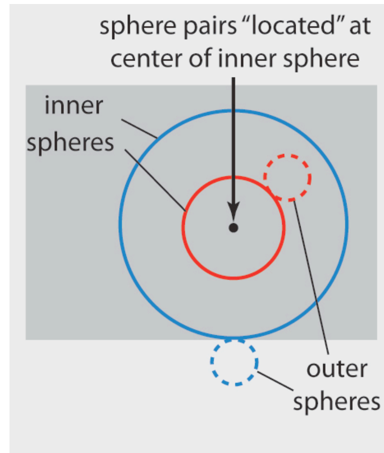

Fig. 11 Two asymmetric sphere-pairs at one location, inner = solid, outer = dashed.

6. THE DIRECTION OF THE SPHERE MAP AND ITS DIVERGENCE

Thus far we have only used a scalar radius function, which we denote $r(\mathbf{x})$, the size of the inner sphere at $\mathbf{x}$ (recall that $\mathbf{x}$ is also the location of the corresponding sphere-pair, so only one inner sphere is defined per pixel). Assuming a correct optimization of the sphere-pairs from the method described in the previous section, $r(\mathbf{x})$ represents the scalar distance to the nearest boundary. This is commonly called the Euclidian or Danielsson distance map.[17]

In addition to the scalar value $r(\mathbf{x})$, there is also a *unique* direction to the nearest boundary, which we denote $\mathbf{d}(\mathbf{x})$, defined as the unit vector in the direction from the center of the inner sphere and the center of the outer sphere for the asymmetric sphere pair at $\mathbf{x}$. Thus a vector radius function $\mathbf{r}(\mathbf{x})$ can be constructed as $\mathbf{r}(\mathbf{x}) = r(\mathbf{x})\mathbf{d}(\mathbf{x})$. The direction $\mathbf{d}(\mathbf{x})$ is not unique at the medial manifold, where at least two such directions exist. Thus $\mathbf{d}(\mathbf{x})$ changes abruptly as one crosses a medial ridge, switching from one nearest boundary to another. Detecting this switch can identify a medial ridge.

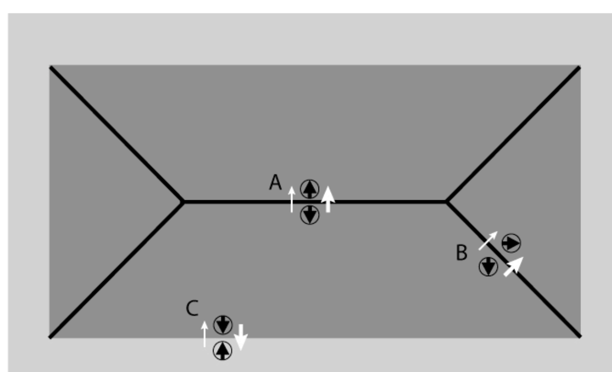

Fig. 12 Divergence of the direction: positive at the medial ridge and negative at the boundary.

When choosing how to measure a change in $\mathbf{d}(\mathbf{x})$ with respect to $\mathbf{x}$ one has a number of choices. The full first derivative of an n-dimensional vector function with respect to its n-dimensional domain is an n x n matrix of the first-order partial derivatives, the Jacobian matrix. It turns out, however, that we do not need all those partial derivatives for detecting the medial ridge (or a boundary). We note that as one moves across a medial ridge in the $\Delta\mathbf{x}$ direction, the change in the direction function $\Delta\mathbf{d}(\mathbf{x})$ will always be in the same direction as $\Delta\mathbf{x}$. This is shown in Figure 12 (the rectangle example of a medial ridge from Fig. 1). Locations A and B show the direction function on either side of a medial ridge as black arrows within circles pointing towards the nearest boundary. In each case, taking the derivative of the direction function as one moves across the medial ridge (by moving along the thin white arrow) yields a change in the direction (the thick white arrow) parallel to the direction of motion. The opposite happens at location C, where one crosses the boundary of the rectangle, with the change in the direction function being exactly in the opposite direction. Thus the only partial derivatives that are non-zero as one crosses a medial ridge or a boundary are those appearing in the divergence operator. The divergence of $\mathbf{d}(\mathbf{x})$ is defined in n-dimensions as

$$\nabla \cdot \mathbf{d}(\mathbf{x}) = \sum_{i=1}^{n} \frac{\partial d_i(\mathbf{x})}{\partial x_i} \tag{2}$$

The divergence will be positive at the medial ridge, as if the ridge acts as a source of direction flow, and negative at the boundaries, as if it acts as a trough or sink for direction flow.

In the continuous domain of an object with perfectly defined boundaries, the direction function will change instantaneously at the medial ridges and boundaries, and its divergence will thus be infinite. Singularities are notoriously difficult to handle in computational systems, but we can neatly avoid the problem by never stepping directly on a singularity. Instead, we force a choice at each pixel as to the direction to the nearest boundary. We basically assume that the singularities, i.e. the medial or boundary points, lie between pixels, and we sense those singularities by changes in the direction between neighboring pixels. This is possible because we are processing a discrete, not continuous, image. We therefore use the discrete version of the unit direction function $\mathbf{d}[\mathbf{x}]$ with an unambiguous value at each pixel. To compute the discrete divergence of $\mathbf{d}[\mathbf{x}]$ in n-dimensions we use a difference function,

$$\nabla \cdot \mathbf{d}[\mathbf{x}] = \sum_{i=1}^{n} d_i[\mathbf{x} + \mathbf{u}_i] - d_i[\mathbf{x}] \tag{3}$$

where d_i is the ith component of $\mathbf{d}[\mathbf{x}]$ and $\mathbf{u}_i$ is the unit vector in the i^{th} cardinal direction (in other words, $\mathbf{x} + \mathbf{u}_i$ is the neighboring pixel along the i-axis direction). Operating thus on adjacent pixels with a simple difference function maximizes spatial resolution, although it also shifts the computation by half a pixel in each of the cardinal directions.

A useful feature of the divergence of the direction function is that it can serve as a measure of the angle between the direction to the nearest boundary on either side of a medial ridge. For example, in Figure 12 at point A, that angle is 180°, whereas at point B it is only 90°. This will result in a lower positive value for the divergence at point B than at point A. At point C, along the boundary, the divergence will be negative. These phenomena will be demonstrated in the next section.

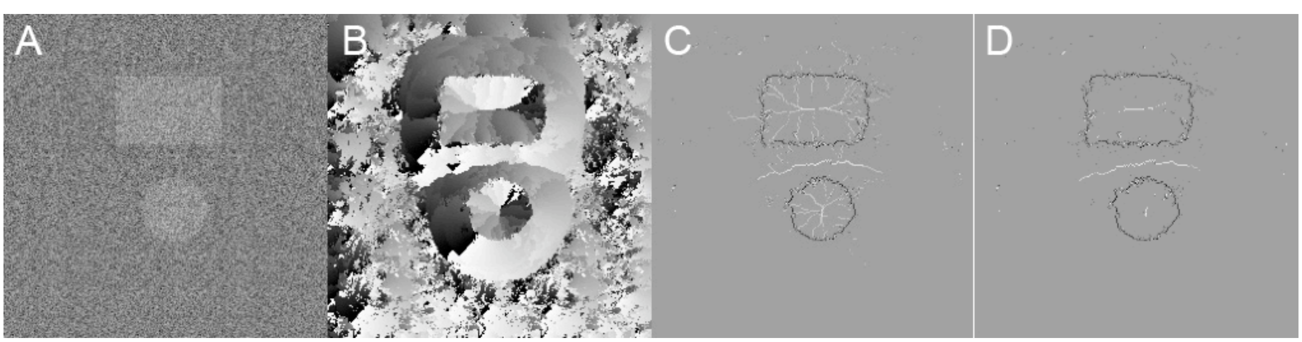

Fig. 13 **(A)** Shapes in Gaussian noise, **(B)** unit direction map $\mathbf{d}[\mathbf{x}]$ with intensity wrapping at 360°, **(C)** divergence of the distance map $\nabla \cdot \mathbf{d}[\mathbf{x}]$ with white = positive (medial ridges) and black = negative (boundary troughs), **(D)** divergence with positive values thresholded near 180° degrees.

7. RESULTS WITH ASYMMETRIC SPHERE-PAIRS AND THE DIVERGENCE OF DIRECTION

A series of experiments were performed using the methods just described with asymmetric sphere-pairs and the divergence of the direction function. Figure 13A shows a rectangle and a circle in Gaussian noise. Asymmetric sphere-pairs were optimized up to a maximum possible radius greater than the diameter of the circle or the height of the rectangle, yielding the direction function shown in Fig. 13B, which maps angle to intensity (wrapping around from black to white at +45°). The divergence of this direction function is shown in Fig. 13C, normalized over its range between white for positive divergence and black for negative. The classical medial axes of the rectangle are clearly demonstrated as white lines, as well as branches extending to the boundaries that are typical with minor fluctuations in boundary direction. The circle also shows a central medial manifold, as does the space between the circle and the rectangle ("inside" spheres form on both sides of any boundary). Boundaries are shown in black (negative divergence). Fig. 13D shows the same data with positive divergence truncated near the equivalent of 180° degrees, so that the 90° ridges to the corners of the rectangle (B in Fig. 12) and minor ridges to boundary fluctuations with even more acute angles are no longer visible. Only the major axes, or what Pizer calls *cores*, remain.[4]

The same procedure was applied to simulated branching vessels with noise (Fig. 14 A). Now we also show the t-test values (modified as described below) in Fig. 14B, which are high within a certain range of the boundaries (the inner sphere was not allowed to grow beyond that range). Elsewhere in the image, sphere-pairs found no significant variation. Figure 14C shows the resulting direction map, with ridges at the medial axes and troughs at the boundaries. Figure 14D shows the resulting divergence of the direction function, again thresholded near 180°, so that only the major medial axes are shown. Branch points are clearly handled effectively.

Finally, we applied the procedure to a maximum intensity projection (MIP) of cerebral vasculature imaged using time of flight Magnetic Resonance Angiography (MRA). The volume includes the Circle of Willis at the base of the brain. Figure 15A shows the original projection 2D image, with Fig. 15B showing the modified student's t-test. Given the very small diameter of the vessels, the outer sphere radius was reduced, as was the maximum range of the inner spheres, to be appropriate for the particular structures. Figures 15C and 15D show, respectively, the direction function and its divergence.

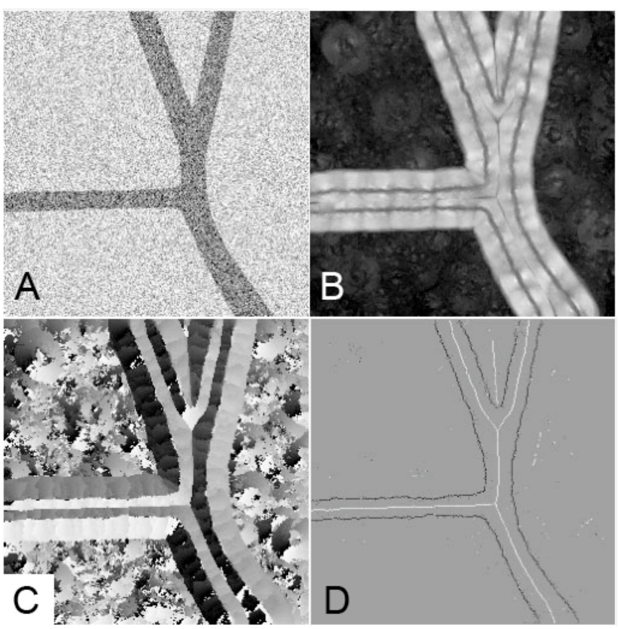

Fig. 14 (A) Simulated bronchi in noise, **(B)** altered student's t-test, **(C)** unit direction map, **(D)** divergence with positive values thresholded near 180° degrees.

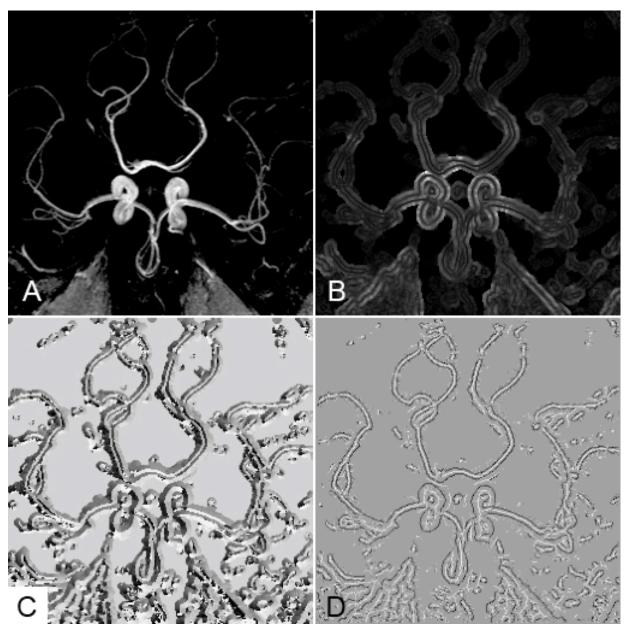

Fig. 15 (A) MRI of cerebral vasculature, **(B)** modified student's t-test, **(C)** unit direction map, **(D)** divergence with positive values thresholded near 180° degrees.

8. PROBLEMS WITH THE T-TEST

During our experiments with asymmetric sphere-pairs, we found problems with the student's t-test that were not apparent with our earlier experiments with symmetric sphere-pairs elimination. In particular, the student's t-test (and the subsequent P-value) resulted in unduly favoring larger spheres. Empirically we found that our optimization of sphere size improved dramatically if we dropped the "*n*" from Equation 1, so as to make the measure independent of sample size. The *n* basically accounts for the lower standard error of the mean as sample size increases. Thus we have adopted the following *modified* t-test,

$$t' = \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{s_1^2 + s_2^2}} \tag{4}$$

simply because it works better empirically. The theoretical basis for this is an open question requiring further study.

9. EXTENSIONS TO 3D AND CLINICAL APPLICATIONS

Nothing in the mathematics developed in this paper is specific to 2 dimensions. In particular, by using divergence instead of a direct measure of angle, we can monotonically convert what amounts to the change in orientation of the direction function in *n*-dimensions to a single scalar value.

We are presently transferring our system from its present 2D implementation in Java to the Insight Toolkit (ITK) so that it can be applied efficiently in 3D to segmenting the vasculature of the lung, as imaged with computerized tomography (CT) for the detection of pulmonary embolism. More information about our research can be found at our website, http://www.vialab.org.

ACKNOWLEDGEMENTS

This work was supported by grants from the National Institutes of Health: R21EB007721 and R01HL087119.

REFERENCES

1 Blum, H., and Nagel, R. N., "Shape description using weighted symmetric axis features," Pattern Recognit., vol. 10, pp. 167–180, (1978).

2 Nackman, L. R., "Curvature relations in 3D symmetric axes," CVGIP, vol. 20, pp. 43–57, (1982).

3 Nackman, L. R., and Pizer, S. M., "Three-dimensional shape description using the symmetric axis transform I: Theory," IEEE Trans. Pattern Anal. Machine Intelligence, vol. 2, pp. 187–202, Mar. (1985).

4 Burbeck, C. A., Pizer, S. M., "Object representation by cores: Identifying and representing primitive spatial regions," Vision Research, vol. 35, no. 13, pp. 1917–1930, (1995).

5 Pizer, S. M., Eberly, D. H., Morse, B. S., and Fritsch, D. S., "Zoom invariant vision of figural shape: The mathematics of cores," Comp. Vision Image Understanding, vol. 69, no. 1, pp. 55–71, (1998).

6 Morse, B. S., Pizer, S. M., Puff, D. T., and Gu, C., "Zoom-Invariant vision of figural shape: Effect on cores of image disturbances," Comp. Vision Image Understanding, vol. 69, pp. 72–86, (1998).

7 Furst, J. D., Pizer, S. M., "Marching cores: A method for extracting cores from 3D medical images," in Proc. Workshop Mathematical Methods Biomedical Image Analysis, San Francisco, CA, June (1996).

8 Fritsch, D., Pizer, S., Yu, L., Johnson, V., and Chaney, E., "Segmentation of medical image objects using deformable shape loci," Information Processing in Medical Imaging (Lecture Notes in Computer Science). Berlin, Germany: Springer-Verlag, (1997).

9 Pizer, S. M., Fritsch, D. S., Yushkevich, P. A., Johnson, V. E., and Chaney, E. L., "Segmentation, registration, and measurement of shape variation via image object shape," IEEE Trans. Med. Imag., vol. 18, pp. 851–865, Oct. (1999).

10 Stetten, G., Pizer, S., "Medial Node Models to Identify and Measure Objects in Real-Time 3D Echocardiography," *IEEE Transactions on Medical Imaging*, Vol. 18, No. 10, pp 1025-1034, Oct. (1999).

11 Cois, A., Rockot, K., Galeotti, J., Tamburo, R., Stetten, G., "Shells and Spheres: A Framework for Variable Scale Statistical Image Analysis," CMU Robotics Tech Report #CMU-RI-TR-04-19, April, 19, (2006).

12 Cois, A., "Variable Scale Statistics for Cardiac Segmentation and Shape Analysis," Ph.D. Dissertation, Bioengineering, U. Pitt., October, advisor G. Stetten (2007).

13 Cois, A., Rockot, K., Galeotti, J., Tamburo, R., Gottlieb, D., Mayer, J., Powell, A., Sacks, M., Stetten, G. "Automated segmentation of the right heart using an optimized shells and spheres algorithm," *International Society of Biomedical Imaging,* Arlington VA., (2007).

14 Tamburo, R., Siegle, G., Stetten, G., Cois, C. A., Reynolds, C., Aizenstein, H., "Amygdala Morphometry in Late-Life Depression," special issue Neuroimaging in Geriatric Psychiatry," in *International Journal of Geriatric Psychiatry* Vol. 24, pp. 837–846, PubMed 19085964, (2009).

15 Tamburo, R., Siegle, G., Stetten, G., Cois, C.A., Rockot, K., Galeotti, J., Reynolds, C., Aizenstein, H., "Localizing amygdala structure differences in late-life depression," ISBI (2007).

16 Fata, B., Sacks, M., Cois, A., Stetten, G., Gottlieb, D., and Mayer, J., "In Vivo Geometric Characterization of the Right Ventricular Outflow Tract," *Biomedical Engineering Society* (BMES) Annual meeting, St. Louis, Oct 2, (2008).

17 Danielsson, Per-Erik. "Euclidean Distance Mapping." *Computer Graphics and Image Processing* 14, 227-248 (1980).