

Chapter 7 Measuring Volume

One last piece of the overall method is required to achieve the goals of this dissertation: a technique for measuring the volume of the left ventricle. As stated in the abstract, the measurement of volume is to be accomplished without exact delineation of an object's boundaries. Most segmentation routines attempt to locate complete and precise boundaries, which provide explicit parameters for measurements and which can be rendered with realism using computer graphics techniques. Perhaps explicit boundaries seems a more accurate portrayal of physical reality, because, after all, most real surfaces are defined to a microscopic scale and interact with light in a manner easily interpretable by the human visual system. A probabilistic representation of a surface may be just as useful and more robust for making geometric measurements, but it is not so easily visualized. The accuracy of measurements based on probabilistic boundaries can, however, be assessed against those from manually determined boundaries.

In previous chapters, core atoms formed the basis of a method for identifying an object. Core atoms can also be used to derive parameters of an object's shape. The ability of core atom populations to accurately measure slab thickness, cylinder diameter, and sphere diameter has already been demonstrated in Chapter 4 with parametric test objects (see Table 4.2). This approach is successfully adapted to measure the volumes of balloons scanned by RT3D ultrasound in Chapter 8, using a simple spherical model.

To analyze more complicated shapes such as the left ventricle, a general theory for measuring volume based on the medial manifold is required. Such a theory, that of *truncated wedges* is developed in Section 7A and comprises the bulk of this chapter. This takes a foray into a new concept: using local boundary curvature as an aspect of local medial dimensionality. Although very interesting and perhaps useful eventually, the idea is not actually put into practice during the remainder of the dissertation. Some of the underlying concepts are, however, important in developing an iterative Bayesian method for fuzzy segmentation of the LV in Section 7B. The method is based on an assignment of probability to voxels of being in the LV by their intensities and locations with respect to the medial node model. From these probabilities, a total volume is computed and a local estimate of error derived relative to manual traces in Section 8E.

7A. Truncated Wedges: Volumes from Boundary Curvature and Medial Scale

In this section the theoretical relationship between the boundary and the medial manifold is explored for the purpose of computing volumes. The volume contained within any closed, simple, piece-wise smooth boundary can be determined by integrating over the boundary a function whose only parameters are the principal boundary curvatures and the distance to the medial manifold. This approach is an extension of the common concept of *mining rights* by which the interior of the earth is parceled out to patches of real estate on the surface. This concept can be extended to any object in \mathfrak{R}^m for $m \geq 2$ independent of topological genus, yielding a one-to-one mapping between the boundary and the interior, which can be used to compute properties such as volume.

The relationship between an object's interior and its boundary is of fundamental concern in geometry and provides an approach by which many properties of the object's interior, including its volume, can be determined from its surface. Since an m -dimensional object has an $(m - 1)$ -dimensional boundary, the boundary may provide a more convenient domain for such calculations than the interior. One may determine a 3D object's interior volume (or a 2D object's area) by establishing a one-to-one mapping between its surface (or boundary contour) and its interior -- the *mining rights*, as it were -- and then integrating these rights over the entire surface. Vectors normal to the surface separate the mining rights between neighboring portions of the surface just as stakes driven into the ground separate the claims of neighboring miners. At the very center of the earth, the mining rights transfer to the other side of the planet. The center of the earth is just a special case of the locus of points for any shape called the *Blum medial manifold*.

For closed contours in \mathfrak{R}^2 , the Blum medial manifold is that locus of the centers of all circles completely enclosed by the boundary contour that touch the contour in more than one location (Blum and Nagel 1978). Thus for the rectangle in Fig. 7.1, the centers of all such *medial circles* form the branching medial manifold shown as thick dotted lines.

For objects in \mathfrak{R}^3 , the Blum medial manifold is the locus of all spheres completely enclosed by the boundary surface that touch the surface at more than one location (Nackman 1982; Nackman and Pizer 1985). For m -dimensional objects, the generic medial manifold has $(m - 1)$ dimensions, although fewer dimensions are possible in non-generic cases. In \mathfrak{R}^3 , for example, although the generic medial manifold has 2 dimensions, the medial manifold of a cylinder may contain portions that collapse to 1 dimension. For a sphere, the medial manifold has zero dimensions, i.e., that point at the center of the earth.

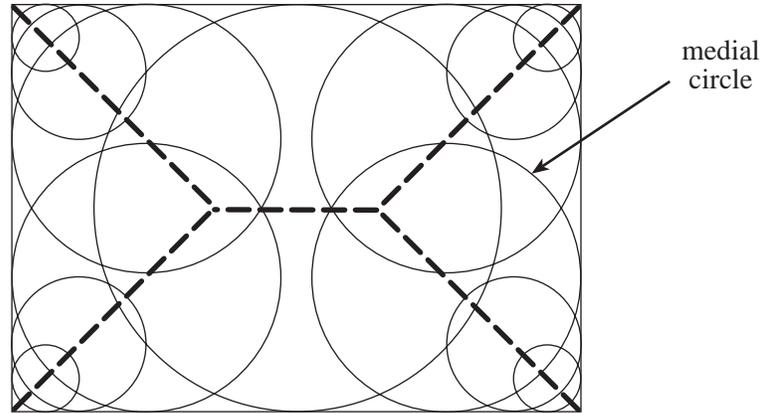


Fig. 7.1. The Blum medial manifold of a rectangle (thick dotted lines) is the locus of centers of all medial circles, i.e., those circles lying completely within the rectangle that touch the boundary in at least 2 places.

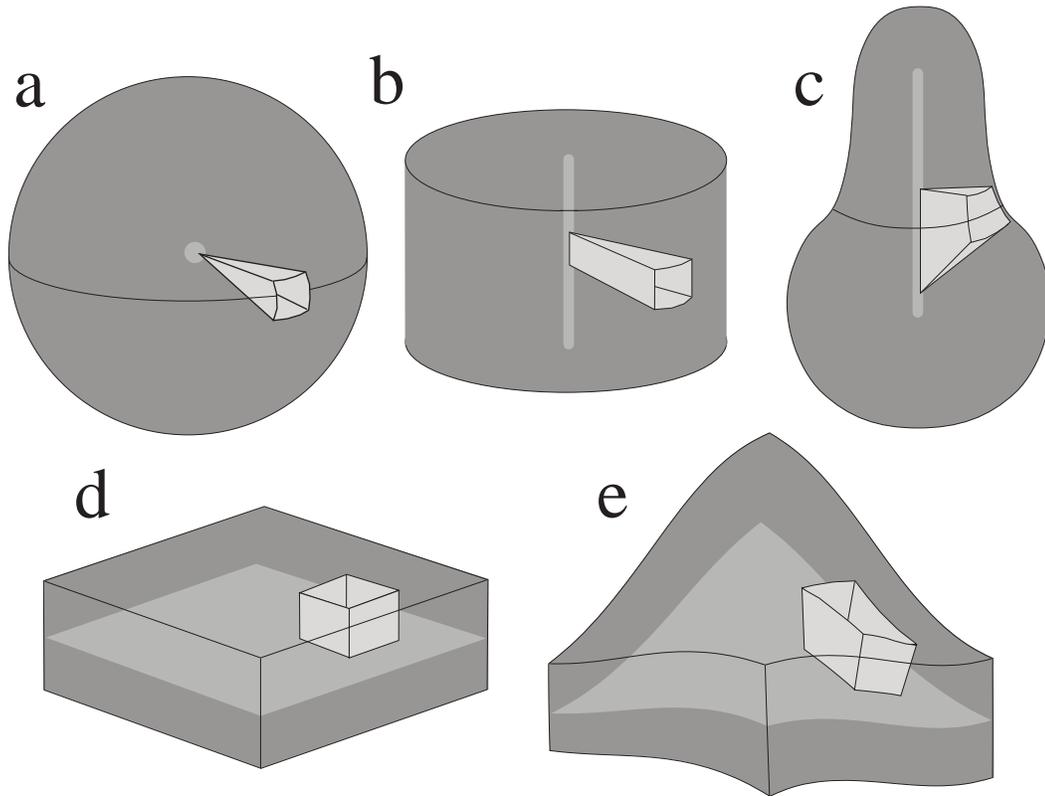


Fig. 7.2 Examples of 3D objects with simple shapes. The "mineral rights" for a surface patch are shown in each case extending orthogonal from the surface half-way through the object to the opposite side.

Corporations that extract minerals and oil divide the planet into wedge-shaped pieces extending straight down from any parcel of real estate to the center of the earth, as shown in

Fig 2a. This approach conceptually permits the measurement of the earth's volume by summing the mining rights for each patch of the earth's surface, since these rights completely fill the earth but do not overlap. This same approach could be used to establish the volume of a hypothetical cylindrical planet by extending mineral rights from the surface of the cylinder down to its central axis, as shown in Fig. 7.2b. A one-to-one correspondence exists between the surface and the interior volume that can be used to compute the volume from the surface. In these two objects -- the sphere and the cylinder -- the mineral rights converge on the medial manifold of the object, being the central point of the sphere and the central axis of the cylinder (forgetting, for now, the ends of the cylinder). Beyond the medial manifold the mining rights belong to the other side of the object.

The manner in which mining rights intersect the medial manifold is governed in part by the principal curvatures of the surface (for a review of principal curvature, see (Koenderink 1990)). Consider again the square patch on the surface of the sphere (Fig. 7.2a) and the lines orthogonal to the surface running straight down from the four corners of that patch to the center of the sphere. They meet precisely at the center because the two principal curvatures of the surface are equal and constant everywhere. The square patch on the surface of a cylinder (Fig. 7.2b) produces four straight lines orthogonal to the surface that reach the central axis in two parallel pairs, forming a wedge cut from a disk like a piece of cheese. This happens because one of the principal curvatures is constant while the other is zero.

Next consider the pear-shaped object in Fig. 7.2c. Like the cylinder, the pear's medial manifold contains a line along the central axis, although now the 2 principal curvatures on the surface patch are neither zero nor equal. For the hyperbolic surface patch shown in Fig. 7.2c, the principal curvatures have opposite signs. The mining rights are still easily defined by extending straight lines orthogonal to the surface from the 4 corners of the patch to the medial manifold, to form a wedge that widens at its cutting edge like the blade of an axe. Such wedges would fill the space within the pear and would not overlap, so they could still be used to compute the volume from the surface exactly as before.

Fig. 7.2d shows a somewhat different case in which the borders of the mining rights do not converge at the medial manifold. In this flat slab, the 2 principal curvatures of the surface are both zero and the borders of the mineral rights for the surface patch extend in parallel to the medial manifold. Half-way through the slab, those rights are simply transferred to the opposite side. Fig. 7.2e shows a curved slab with various unequal principal curvatures and a surface patch with its mining rights extending down to the medial manifold. In all these cases, it should be clear that the total volume can be found by summing the volumes of the individual mining rights.

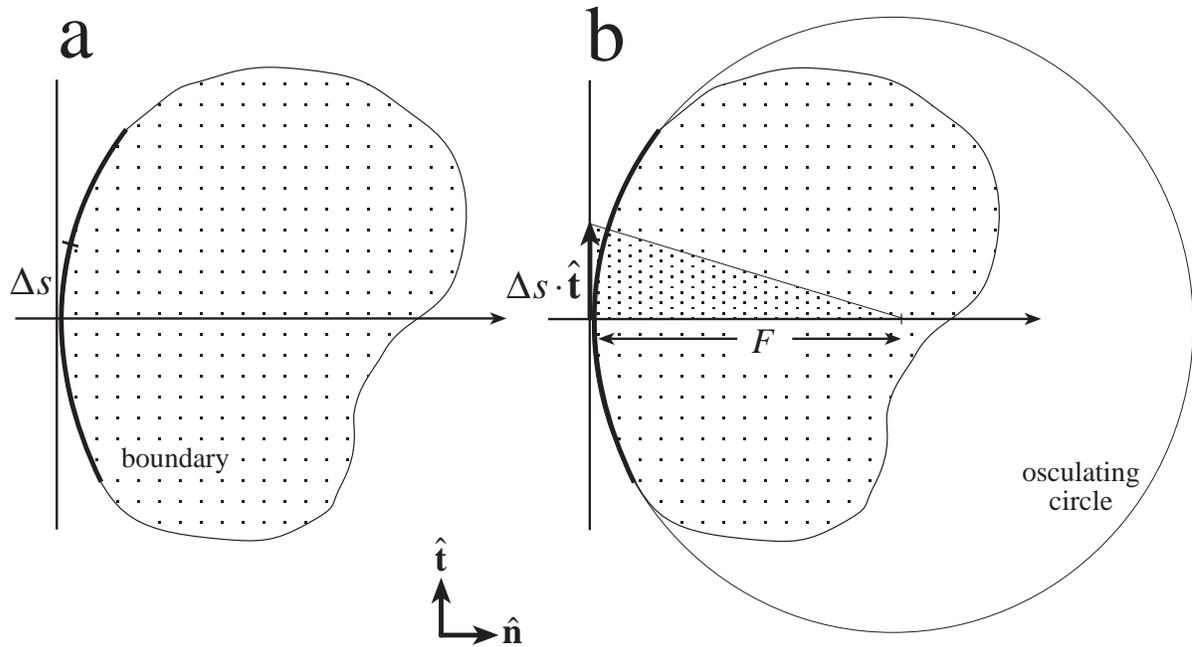


Fig. 7.3 a. Local coordinate system with $\hat{\mathbf{t}}$ tangential to, and $\hat{\mathbf{n}}$ normal to, the object's boundary. b. Osculating circle with radius F and a sector of that circle corresponding to boundary interval Δs .

Consider now the closed boundary contour of the 2D object depicted in Fig. 7.3a. At each boundary point a coordinate system can be defined consisting of the unit tangent vector $\hat{\mathbf{t}}$ and the unit normal vector $\hat{\mathbf{n}}$. As can be seen in Fig. 7.3b, a step along the boundary of length Δs can be approximated by the vector $\Delta s \cdot \hat{\mathbf{t}}$. The vector $F \cdot \hat{\mathbf{n}}$ establishes the center of the osculating circle, being that circle which shares the tangent and the curvature with the boundary. The focal length F is the radius of that osculating circle and the inverse of the boundary's local curvature. The mining rights of the boundary segment Δs fall within a sector formed by two such radii of the osculating circle. At some point within that sector, the mining rights may be transferred to the opposite boundary. It is worth noting here that any straight line entering an object orthogonal to its surface *must* intersect the medial manifold before exiting the other side.

Area of an arbitrary 2D shape

A mathematical expression can now be developed for the area of any closed smooth object in \mathcal{R}^2 as a function of the boundary curvature and the distance to the medial manifold. First consider a section of boundary that is convex, that is to say F is positive, as shown in Fig. 7.4a. The medial manifold (thick dotted line) is shown at a distance R from the boundary. R is the *medial scale*, the radius of the medial circle (see Fig. 7.1), while F is the radius of the corresponding osculating circle. The mining rights, approximated by the trapezoidal area ΔA , consist of a sector of the osculating circle truncated at the medial manifold.

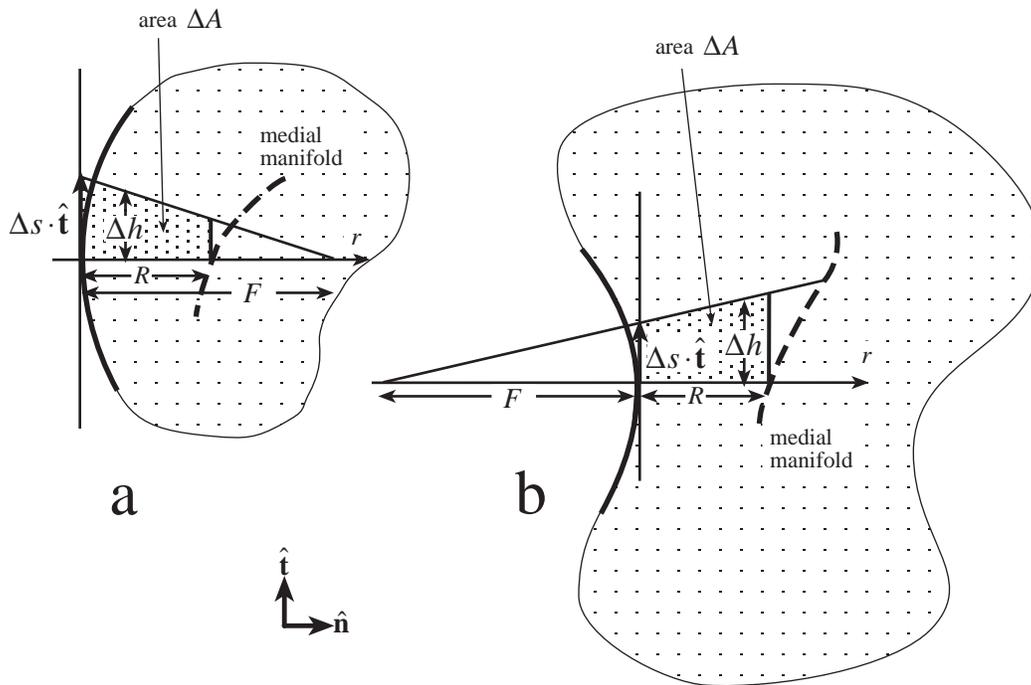


Fig. 7.4 **a.** Convex boundary segment (thick solid line) and corresponding medial manifold (thick dotted line) at a distance R from the boundary ($F > 0$). Area ΔA approximates the "mining rights" of boundary interval Δs . **b.** Same for a concave boundary segment ($F < 0$).

The relationship between the medial circle and the osculating circle is central to this dissertation. Since the medial circle must lie completely within the object, it cannot be larger than the corresponding osculating circle. Therefore, $R \leq F$, and at some point along the sector of the osculating circle (perhaps only at its vertex, if $R = F$), the sector intersects the medial

manifold and the mining rights are transferred to another location on the boundary also touched by the medial circle.

Now consider the convex boundary segment shown in Fig. 7.4b. Here the focal length F is negative and the distance to the medial manifold R can assume any positive value. This would also be true for a straight boundary segment, for which $F = \infty$. Whether concave, convex, or straight, in all cases the area of the trapezoid ΔA can be found by integrating the height of the trapezoid Δh as a function of the distance r along the $\hat{\mathbf{n}}$ axis, where

$$\Delta h = \left(1 - \frac{r}{F}\right)\Delta s. \quad (7.1)$$

As $\Delta s \rightarrow 0$, the infinitesimal area dA corresponding to the mining rights of the infinitesimal boundary interval ds may therefore be expressed as

$$dA = \left[\int_0^R \left(1 - \frac{r}{F}\right) dr \right] ds. \quad (7.2)$$

The total area A of the object can be found by integrating dA over the entire boundary contour S , assuming convergence of ΔA on the actual mining rights:

$$A = \oint_S dA \quad (7.3)$$

Several examples are illustrative here. First, consider a circle of radius R . The constant curvature of the boundary guarantees that $R = F$ everywhere, so the area dA corresponding to the segment ds can be found by substituting into Eq. (7.2) to yield

$$dA = \left(r - \frac{r^2}{2R} \right) ds \Big|_0^R = \frac{R}{2} ds \quad (7.4)$$

Integrating dA over the boundary contour S of the circle, whose length is $2\pi R$ where R is constant, yields the correct area:

$$A = \oint_S \frac{R}{2} ds = \pi R^2. \quad (7.5)$$

Another example is shown in Fig. 7.5. This rectangle serves to demonstrate the special case of the straight boundary segment whose focal length $F = \infty$. In this case, substitution into Eq. (7.2) yields

$$dA = \left[\int_0^R \left(1 - \frac{r}{\infty} \right) dr \right] ds = R \cdot ds. \quad (7.6)$$

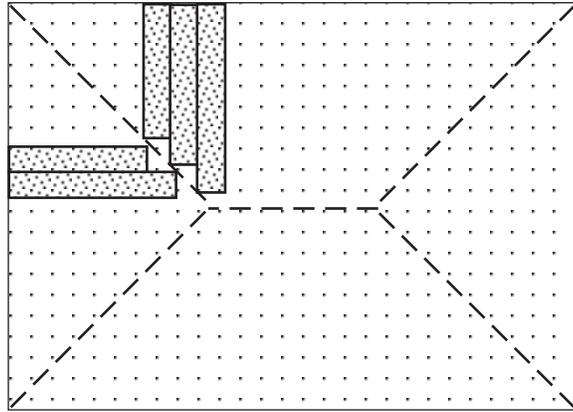


Fig. 7.5 Rectangle with mining rights extending in thin ribbons from the boundary to medial manifold.

Thus, for a straight boundary segment, dA represents an infinitesimally thin stripe orthogonal to the surface, extending to the medial manifold along the radius of a medial circle. The rectangle is completely filled by such stripes, which are able to reach every portion of the interior. There can be no intervening branches of the medial manifold producing unreachable portions of the interior, since any such branches would represent a medial circle completely enclosed within another medial circle. This is impossible because each medial circle must touch the boundary in more than one location. Since the boundary is piece-wise smooth, the integration around it can be accomplished by sections that contain no sharp corners. Alternatively, the corners can be viewed as non-zero (but very small) minima in the focal length F , in other words, not really corners but simply maxima in the allowable curvature of a smooth boundary.

Volume of arbitrary shapes in 3 or more dimensions

The same approach can be applied to finding the volume of an object from its boundary in \mathfrak{R}^3 . All that is required for each patch of boundary surface is knowledge of the distance R to

the medial manifold and the focal lengths F_1 and F_2 of the osculating disks corresponding to the two principal curvatures (see Fig. 7.6). Whereas a single tangent $\hat{\mathbf{t}}$ and interval Δs suffice to describe the boundary contour in \mathfrak{R}^2 , two orthogonal principal directions $\hat{\mathbf{t}}_1$ and $\hat{\mathbf{t}}_2$ and a surface patch of area $(\Delta s)^2$ are required to describe the boundary surface of an object in \mathfrak{R}^3 . Given such a surface patch, the mining rights between the surface patch and the medial manifold can be found by integrating the area $(\Delta h_1 \cdot \Delta h_2)$ as a function of distance r along the $\hat{\mathbf{n}}$ axis, within an *osculating wedge* truncated by the medial manifold, where

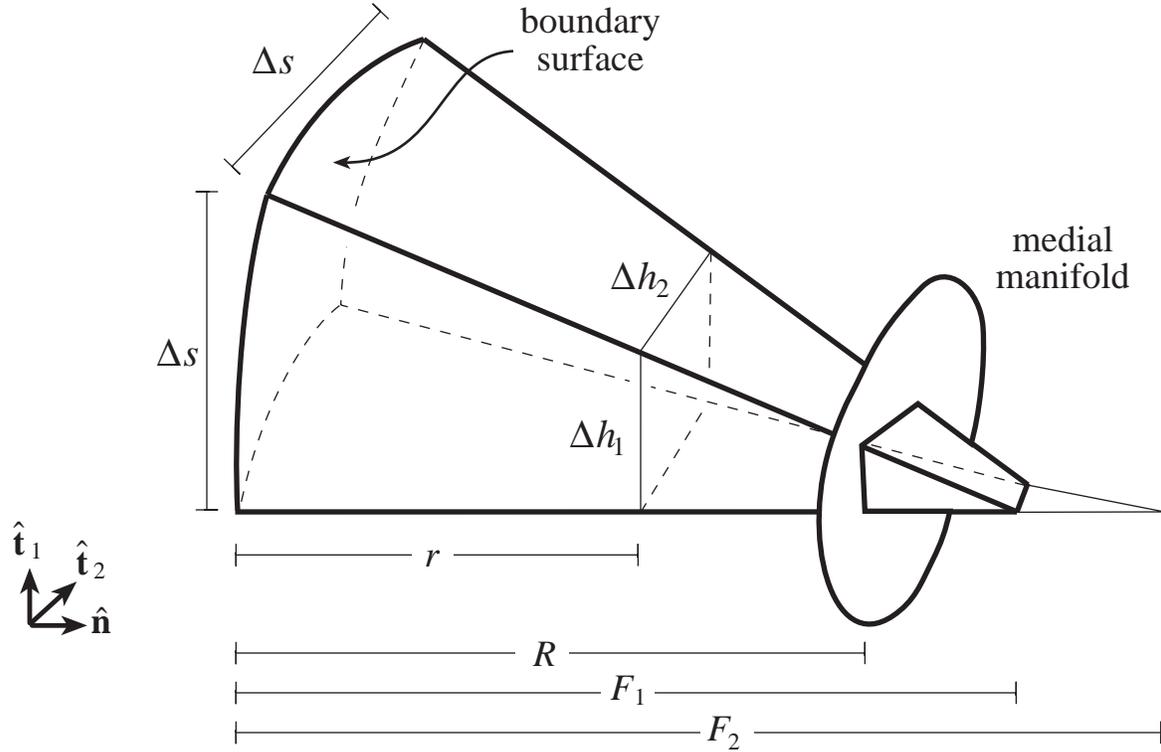


Fig 7.6 In \mathfrak{R}^3 the mining rights of a surface patch form an osculating wedge truncated at the medial manifold.

$$\Delta h_i = \left(1 - \frac{r}{F_i}\right) \Delta s \quad i = 1, 2 \quad (7.7)$$

As $\Delta s \rightarrow 0$, the infinitesimal volume dV corresponding to the mining rights for the infinitesimal boundary patch ds^2 may be expressed as

$$dV = \left[\int_0^R \left(1 - \frac{r}{F_1}\right) \left(1 - \frac{r}{F_2}\right) dr \right] ds^2. \quad (7.8)$$

The particular wedge shown in Fig. 7.6 does not converge to a point, but rather to a blade, since evidently for this surface patch, $F_1 \neq F_2$. The focal length F_1 corresponds to the principal curvature in the $\hat{\mathbf{t}}_1$ direction, and F_2 to the principal curvature in the $\hat{\mathbf{t}}_2$ direction. By the same arguments given above for osculating and medial circles in \mathfrak{R}^2 , any focal length corresponding to a convex principal curvature limits R as follows:

$$F_i \geq R, \quad \text{if } F_i \geq 0, \quad i = 1, 2 \quad (7.9)$$

The wedge must intersect the medial manifold, or at least make contact with it, within the smallest convex focal length.

The total volume V can be found by integrating dV over the boundary surface A .

$$V = \oint_A dV, \quad (7.10)$$

again assuming convergence.

Consider an example in 3D. A sphere of radius R has constant principal curvatures $R = F_1 = F_2$, so that by substituting into Eq. (7.8), the infinitesimal volume dV corresponding to the infinitesimal surface patch ds^2 is

$$dV = \left(r - \frac{r^2}{R} + \frac{r^3}{3R^2} \right) ds^2 \Big|_0^R = \frac{R}{3} ds^2. \quad (7.11)$$

Integrating everywhere on the surface $A = 4\pi R^2$, where R is constant, yields the correct volume for a sphere of

$$V = \oint_A \frac{R}{3} da = \frac{4}{3} \pi R^3. \quad (7.12)$$

In the general case of m dimensions, the hypervolume V can be found by integrating dV over the $(m-1)$ dimensional boundary, where

$$dV = \int_0^R \prod_{i=1}^{m-1} \left(1 - \frac{r}{F_i}\right) ds^{m-1} dr. \quad (7.13)$$

Thus dV is a "hyper-wedge" truncated at the medial manifold for the infinitesimal boundary patch ds^{m-1} , where F_i is the focal length corresponding to the i th principal direction on the hypersurface.

The topological genus of the object makes no difference, since the method relies on the purely local relationship between the surface and the medial manifold. For example, one can imagine summing the mining rights for a torus, which is locally indistinguishable from a cylinder.

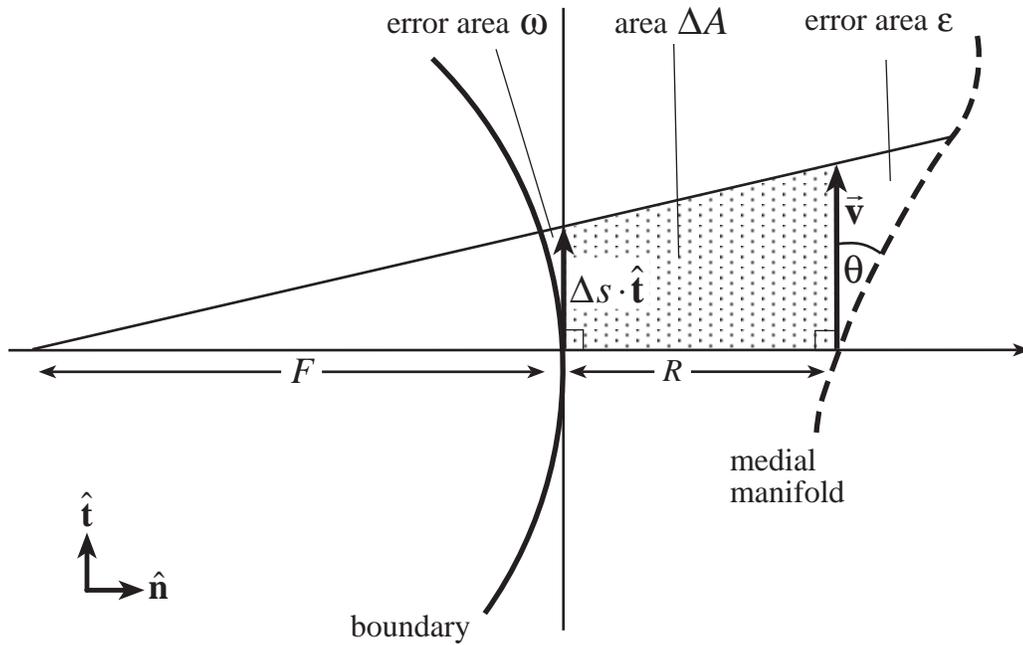


Fig. 7.7 Potential errors in the calculation of ΔA for a concave boundary segment.

Do the errors vanish?

The issue of convergence bears examination. Assume that the parametric descriptions of both the boundary and medial manifold are known. Consider an object in \mathfrak{R}^2 . Shown in Fig. 7.7 is a concave boundary segment and corresponding medial manifold. The concave boundary was chosen for this illustrations since it magnifies the area of error ε between the medial manifold and the edge \vec{v} of the trapezoid ΔA . A second area of error ω is shown between the boundary contour and the boundary tangent $\Delta s \cdot \hat{t}$. Clearly $\omega/\Delta A \rightarrow 0$ as

$ds \rightarrow 0$, so ω can be ignored in the limit. The error ε cannot so easily be ignored, however, since $\bar{\mathbf{v}}$ is generally not tangential to the medial manifold. The area ε as $\Delta s \rightarrow 0$ can be approximated for a concave boundary as

$$\lim_{\Delta s \rightarrow 0} \varepsilon = \frac{|\bar{\mathbf{v}}|^2 \tan \theta}{2} = (\Delta s)^2 \left(\frac{F+R}{F} \right)^2 \frac{\tan \theta}{2}, \quad \text{if } F < 0 \quad (7.14)$$

and for a convex boundary segment as

$$\lim_{\Delta s \rightarrow 0} \varepsilon = (\Delta s)^2 \left(\frac{F-R}{F} \right)^2 \frac{\tan \theta}{2}, \quad \text{if } F > 0, \quad (7.15)$$

which simplifies for a straight boundary segment to

$$\lim_{\Delta s \rightarrow 0} \varepsilon = (\Delta s)^2 \frac{\tan \theta}{2}, \quad \text{if } F = \infty. \quad (7.16)$$

In all cases, $\varepsilon \rightarrow 0$ as $(\Delta s)^2$ and thus the error ε vanishes in the computation of the total area A as $\Delta s \rightarrow 0$, except where $\theta \rightarrow \pi/2$ and $\tan \theta \rightarrow \infty$. This occurs when the radius vector $\bar{\mathbf{r}} = R \cdot \hat{\mathbf{n}}$ intersects the medial manifold tangentially, which can be divided into two special cases.

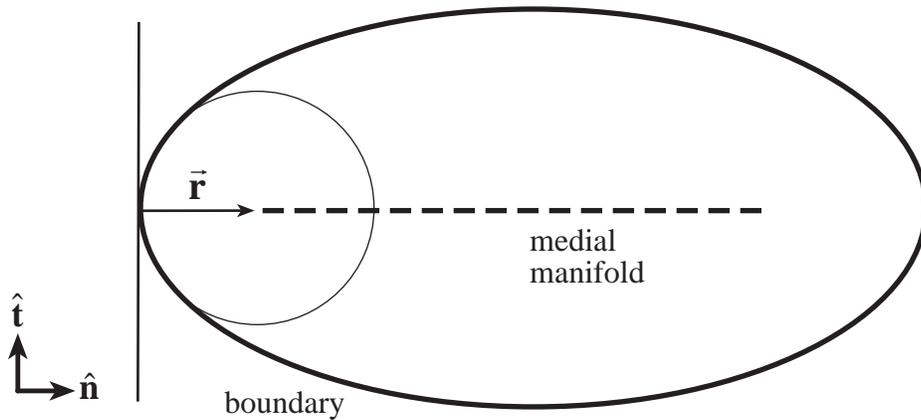


Fig. 7.8 Singular case for convex boundary.

First, consider the case for a convex boundary segment (see Fig. 7.8). In this case the medial manifold is tangential to $\bar{\mathbf{r}}$ at their intersection only for a local minimum in R along the

boundary, i.e., a smallest osculating circle, where $R = F$. This in turn implies that $\bar{\mathbf{v}}$ in Fig. 7.7 has zero length and Eq. (7.14) mandates that $\varepsilon = 0$.

Next consider the case for a concave boundary segment. The phenomenon of $\bar{\mathbf{r}}$ being tangential to the medial manifold at their intersection can only occur at a boundary point such as s_1 in the Fig. 7.9. The corresponding medial location is an end of the medial manifold and there must be at least two other contacts between its medial circle and the boundary. Thus it must represent a branch point in the medial manifold and the branches must extend in such a way as to block the expanse of error area ε , reducing the situation to the case depicted in Fig. 7.7, where $\theta \neq \pi/2$ and $\tan \theta$ is finite. The same argument can be applied to the case of the flat boundary.

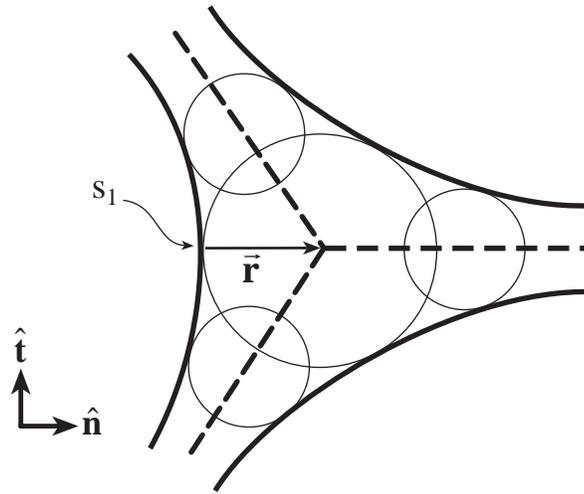


Fig. 7.9 Singular case for concave boundary.

In higher dimensions \mathfrak{R}^m , for $m > 2$, each $(\Delta s)^{m-1}$ patch of the boundary has only one radius vector $\bar{\mathbf{r}}$, but $(m - 1)$ principal curvatures. Each principal curvature F_i has its own angle θ_i with respect to the medial manifold. For a given F_i the degenerate case $\theta_i = \pi/2$ may produce a sub-manifold on the boundary of up to $(m - 2)$ dimensions in which one of the two cases in Figs. 7.8 and 7.9 apply. In all cases the errors in the calculated volume will vanish as $\Delta s \rightarrow 0$.

Practical Applications for Truncated Wedges

The method presented here uses the medial manifold to resolve ownership between boundaries across the object and it uses boundary curvature to resolve ownership between

adjoining portions of the boundary. The resulting osculating wedges, truncated by the medial manifold, map each portion of the boundary to a unique portion of the interior. The boundary of an m -dimensional object has only $(m - 1)$ dimensions and thus may present advantages in the practical problem of determining the volume of an object. While it is true that the medial manifold also has only $(m - 1)$ dimensions, it displays a more complex relationship with the interior, since a single location on the medial manifold may correspond to multiple locations on the boundary.

The truncated wedges bear some resemblance to Green's and Stokes' theorems, which relate an object's boundary to its interior by integrating a function of some underlying field (Kaplan 1973). Eberly and others have followed this approach, viewing the boundary as a level curve of some function whose value is greater inside the object than a threshold at the boundary (Eberly and Lancaster 1991, November; Eberly, Lancaster et al. 1991, November). The truncated wedges depend, instead, on a purely geometric relationship, with a deeper connection to other geometric processes. These include Delauney triangulation and Voronoi diagrams, by which the interior of a shape may be assigned to its nearest boundary and thereby broken into subunits with simpler geometric properties (O'Rourke 1998).

Any practical application of truncated wedges will depend upon prior determination of the distance to the medial manifold, which is a non-trivial problem. Several recent developments hold promise. Culver finds the medial manifold from a polygonal surface using a Voronoi approach (Culver 1998). Fritsch and Pizer have developed Deformable Shape Loci to adapt a medial model to fit objects in gray scale data (Pizer, Fritsch et al. (in press); Fritsch, Pizer et al. 1997). Furst has devised methods of tracking ridges of medialness in gray scale data (Furst and Pizer June 1996). The author has developed methods of describing the medial manifold statistically in gray scale data which are the subject of this dissertation (Stetten, Landesman et al. 1997; Stetten and Pizer 1999).

The practical application of truncated wedges will also depend upon the determination of local boundary curvature. For sampled boundary representations, boundary curvature is a function of scale. This suggests that a coarse-to-fine approach might provide estimates of volume at varying levels of precision. In situations where the derivatives of the medial locus are known, it has been shown how to compute the corresponding boundary's curvature (Nackman 1982). In situations where the medial information is derived from by core atom cluster, it may be possible to determine an appropriate scale and stabilize boundary parameters such as curvature by permitting the proper ordering of boundary points (Amenta, Bern et al. 1998).

For any of these methods, errors may arise in the subsequent volume calculation beyond those that vanish in the theoretical treatment above. Further problems may arise in

parameterization of the boundary itself, to perform the integration of volume. While parameterization of a closed contour in \mathcal{R}^2 is straightforward, it can be problematic for boundaries in higher dimensions.

Besides calculating volume, the truncated wedges approach may provide a basis for techniques, such as finite element analysis, that depend upon compartmentalizing the interior of an object and that may encounter problems when distorting the standard rectilinear coordinate system to match an object's shape. These problems might be avoided by using truncated wedges for compartmentalization.

In summary, a simple relationship exists between the volume, the boundary curvature, and the distance to the medial manifold for objects with $m \geq 2$ dimensions, piece-wise smooth boundaries, and any topological genus. The relationship generalizes the concept of mining rights, allowing the interior to be mapped from the surface, for the purpose of calculating volume and possibly determining other properties of the object.

7B. Fuzzy Segmentation of the Left Ventricle

The concept of truncated wedges is applicable in a general sense to the goal of computing the volume of the LV. Given a medial framework for a particular ventricle, an estimation of the ventricular volume and the location of the endocardial boundary can be made by organizing voxels into wedges. In the following chapter, such a medial framework (in the form of the MNM from Section 2D) will be used to locate the Apex-to-Mitral Valve (AMV) axis of the LV. This section describes how the AMV axis so produced can be used to perform a fuzzy segmentation of the LV.

Fuzzy Logic is an extension of Set Theory that permits an element to be partially included in a set. Such partial inclusion can be viewed as the probability of the element being in the set. In this manner, each voxel in an image may be assigned a probability of being in the ventricle. This probability depends on both the voxel's location and its intensity. Dark voxels near the AMV axis are likely to be in the LV. Voxels that are brighter and/or further from the axis are less likely to be in the LV. In the following sections, the probability of being in the LV due to location and that due to intensity will each be assigned an individual value. The two will then be combined into a single aggregate probability of being in the ventricle.

Probability due to location

The probability $p_L(j)$ of voxel j being within the ventricle solely because of its location is determined by using a surface model of expected ventricular shape. The surface model is located and scaled to the particular ventricle and provides a more detailed description of ventricular shape than the MNM.

Spherical rather than cylindrical coordinates are used for the surface model because spherical coordinates incorporate *medial* and *end* geometries into a single set of parameters thus avoiding the need to pre-assign end points to the ventricular cylinder. The origin chosen for the spherical coordinate system is the midpoint of the apex-to-mitral-valve (AMV) axis, as shown in Fig. 7.10, with the poles of the sphere being the end points of the AMV axis. The angle φ corresponds to latitude, with $\varphi = 0^\circ$ at the mitral valve and $\varphi = 180^\circ$ at the ventricular apex. The angle ω corresponds to longitude, or rotation around the AMV axis with $0^\circ < \omega < 360^\circ$. The third coordinate r is the radial distance to the midpoint of the AMV axis.

The expected surface for the ventricle is determined by a function $r_0(\varphi)$, such that $r = r_0(\varphi)$ at the surface. The function $r_0(\varphi)$ is invariant to ω , exhibiting cylindrical symmetry around the AMV axis. A second surface, defining a region of interest (ROI), is

specified by the function $r_1(\varphi)$. The ROI surrounds and includes the expected ventricular shape and will be used to generate regional statistics of voxel intensity in the following section. For the purposes of this dissertation, the functions $r_0(\varphi)$ and $r_1(\varphi)$ have simply been drawn by hand to resemble a ventricle and a reasonable surrounding ROI. Actual values have been taken from the drawing in Fig. 7.10, by measuring angles and distances from the axis midpoint. These will be adjusted to fit the orientation and size information given by the MNM for a particular ventricle.

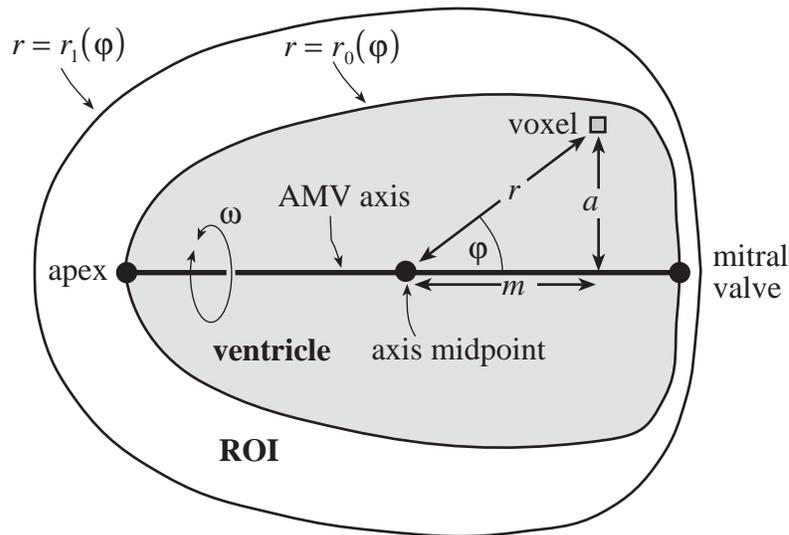


Fig 7.10 Surface model defining a region of interest (ROI) and an expected ventricular boundary.

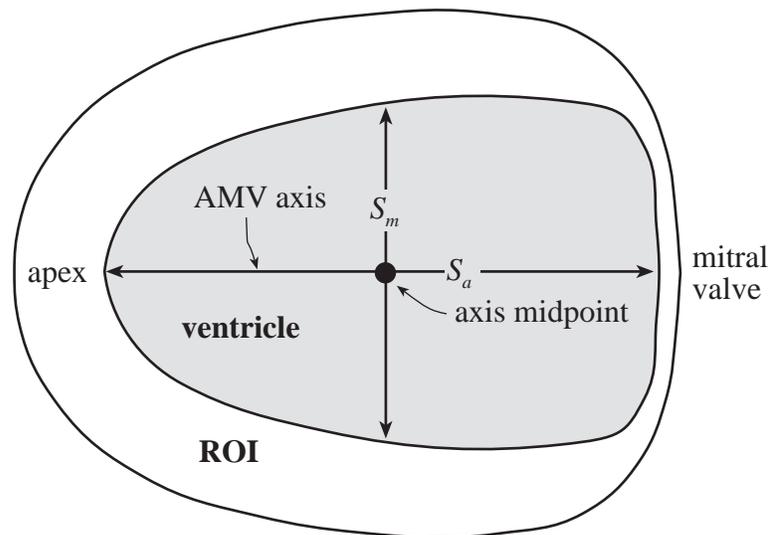


Fig 7.11 Two independent scales are determined for the ventricular boundary.

The surface model is scaled in two independent dimensions: an axial scale S_a and a medial scale S_m as shown in Fig. 7.11. The lengths of these two scales must be related to the actual ventricle in question. When the MNM finds a particular ventricle, two corresponding scales S'_a and S'_m are measured. The axial scale S'_a is the length of the particular AMV axis. The medial scale S'_m is the mean length of the core atoms in the cylindrical clusters that match the particular LV.

Voxel location is related to the surface model as follows: The AMV axis in the surface model is set to the AMV axis found by the MNM. For a given voxel, two orthogonal components of the radius to the midpoint are measured, a' and m' , corresponding to a and m in Fig. 7.10. The axial component a is signed along the AMV axis, with $a > 0$ towards the mitral valve and $a < 0$ toward the apex from the axis midpoint. The medial component m is unsigned and symmetric around the axis in 3D. The measured values a' and m' are scaled to the surface model as follows

$$a = a' \left(\frac{S_a}{S'_a} \right) \quad m = m' \left(\frac{S_m}{S'_m} \right) \quad (7.17)$$

As usual, the angle φ and radius r are determined as

$$\varphi = \tan^{-1} \left(\frac{m}{a} \right), \quad 0^\circ \leq \varphi \leq 180^\circ \quad (7.18)$$

$$r = (a^2 + m^2)^{\frac{1}{2}}.$$

Given voxel j with angle φ and radius r , the probability $p_L(j)$ of being in the ventricle is determined as a function of r using the probability function shown in Fig. 7.12 where $r_o(\varphi)$ is the expected distance to the ventricular boundary at latitude φ . Inside the boundary $r = r_o$, the probability is greater than 1/2, whereas outside it is less than 1/2. The function in Fig. 7.12, as well as that in Fig 7.13, have simply been drawn by hand for the purposes of this dissertation. Finding corresponding functions that optimize the accuracy of the resulting segmentation is beyond the scope of the present research.

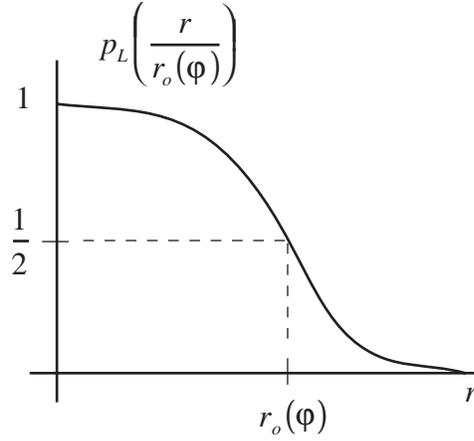


Fig 7.12 Probability p_L of a voxel being within the ventricle, given φ and r .

Probability due to intensity

Besides location, another factor governing the probability of being in the ventricle is intensity. Ventricular voxels are darker than other voxels. Let $p_I(j)$ denote the probability of voxel j being within the ventricle, given its intensity. Intensity in ultrasound images varies unpredictably from one scan to another. To compensate for this, the probability function for $p_I(j)$ is based on a statistical study of voxels in the ROI for a particular heart.

Observe that voxels in the ROI fall into 3 categories. The darkest voxels are in the ventricle and tend to have intensities near zero. The next brightest voxels are in the myocardium, located further from the AMV axis. Myocardial voxels are quite often difficult to distinguish from ventricular voxels because they are nearly as dark, often making the endocardial boundary completely indistinguishable. The epicardial/septal voxels are generally still further from the AMV axis and much brighter. Core atoms in the cylindrical clusters that match the LV node of the MNM tend to find the epicardial border rather than the endocardial border because the gradient magnitude is larger at the epicardial border.

Statistics for a particular image are computed on voxels within the ROI as follows: a mean intensity \bar{I} of all voxels in the ROI is computed weighting each voxel's intensity $I(j)$ by its $p_L(j)$,

$$\bar{I} = \frac{\sum_{j \in ROI} p_L(j) I(j)}{\sum_{j \in ROI} p_L(j)}. \quad (7.19)$$

Ventricular voxels are thus favored by their tendency to be located near the AMV axis, although other voxels are represented as well. The value of \bar{I} reflects the particular intensities of voxels in the image and is used in the function shown in Fig. 7.13 to compute the probability $p_I(j)$ for each voxel, given its intensity $I(j)$.

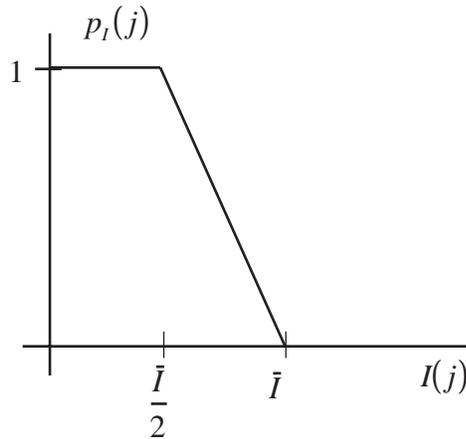


Fig 7.13 Function for computing a voxel's $p_I(j)$ from its intensity $I(j)$ given the weighted mean intensity \bar{I} over the ROI.

Aggregate probability

The probabilities due to location and intensity developed in the previous two sections are combined into a single probability for the fuzzy segmentation of the LV. Let $p_A(j)$ denote the aggregate probability that voxel j is in the ventricle, given both its location and intensity. The probability $p_A(j)$ is some combination of $p_I(j)$ and $p_L(j)$ defined by an operator denoted A

$$p_A(j) = A \left(p_I(j), p_L(j) \right). \quad (7.20)$$

The A operator is designed to exhibit certain behaviors. It is monotonic with positive slope for both $p_I(j)$ and $p_L(j)$, mapping domain $[0,1]$ for $p_I(j)$ and $p_L(j)$ into a range $[0,1]$ for $p_A(j)$. If either argument $p_I(j)$ and $p_L(j)$ is equal to $1/2$, then $p_A(j)$ equals the other argument, so that a probability of $1/2$ exerts neither positive nor negative influence. If either $p_I(j)$ or $p_L(j)$ is 0 , then $p_A(j)$ is 0 , so that either can independently exclude any voxel from the ventricle.

A linear, continuous, piece-wise smooth function that satisfies the above constraints is described by the equations

$$p_A(j) = \begin{cases} 0, & p_I(j) + p_L(j) < 1/2 \\ 2p_L(j), & p_I(j) - p_L(j) > 1/2 \\ 2p_I(j), & p_L(j) - p_I(j) > 1/2 \\ 1, & p_I(j) + p_L(j) > 3/2 \\ p_I(j) + p_L(j) - 1/2, & \text{otherwise} \end{cases} \quad (7.21)$$

and is shown graphically in Fig. 7.14.

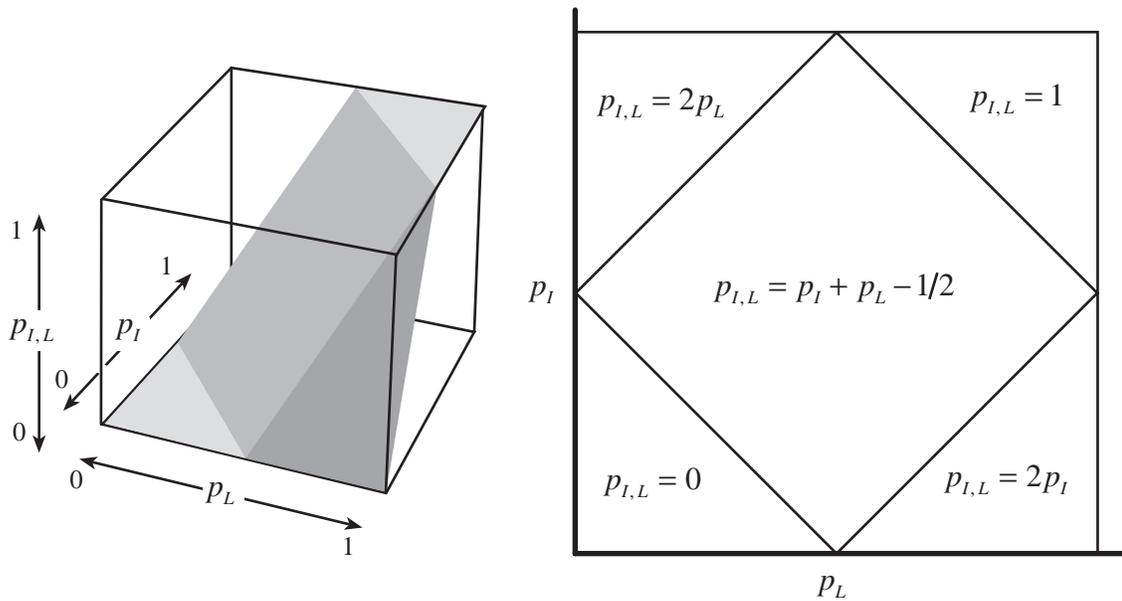


Fig 7.14 The A operator for computing aggregate probability $p_A(j)$ from $p_L(j)$ and $p_I(j)$ is linear, continuous, piece-wise smooth.

This function yields an aggregate probability $p_A(j)$ from the individual probabilities, $p_L(j)$ and $p_I(j)$, in a manner consistent with the desired behavior for the A operator.

Computing volume from aggregate probability

Once each voxel has been assigned an aggregate probability $p_A(j)$, a total volume v_A for the ventricle can be computed as the sum of the aggregate probability for each voxel in the ROI

$$v_A = \sum_{j \in ROI} p_A(j)v(j) \quad (7.22)$$

weighted by $v(j)$, the volume of that voxel.

In the following chapter, some of the methods developed thus far in this dissertation are tested on RT3D ultrasound scans of balloons and *in vivo* human cardiac ventricles. In particular, the method for fuzzy segmentation and volume computation in this section is applied to the *in vivo* left ventricle in Section 8E, with accuracy judged by comparison to manual tracings.