Volume of Arbitrary Shapes

from Boundary Curvature and Medial Scale

Technical Report # TR99-001, Department of Computer Science University of North Carolina at Chapel Hill

Originally presented to the Geometry Seminar at the
Department of Computer Science, UNC, Chapel Hill, 12/2/98

submitted December 30, 1998, to
Graphical Models and Image Processing

George D. Stetten M.D., (Ph.D. expected April 1999)
Medical Image Display and Analysis Group, UNC, Chapel Hill
NSF/ERC for Emerging Cardiovascular Technologies, Duke University.

running head:

Volume from Boundary Curvature and Medial Scale

Please address correspondences to:

George D. Stetten
Assistant Research Professor
Department of Biomedical Engineering
Room 136 Engineering Bldg.
Duke University
Durham, NC 27708

phone (919) 660-5363
fax (919) 684-4488
stetten@acpub.duke.edu
http://www.stetten.com

Abstract

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. The method 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 $\Re^m$ 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.

Introduction

The relationship between an object's interior and its boundary is of fundamental concern in geometry, and provides an approach by which properties of the object's interior, such as 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*.

Brief Review of Blum Medial Manifold

For closed contours in $\Re^2$, 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 [1]. Thus for the rectangle in Fig. 1, the centers of all such *medial circles* form the branching medial manifold shown as thick dotted lines.

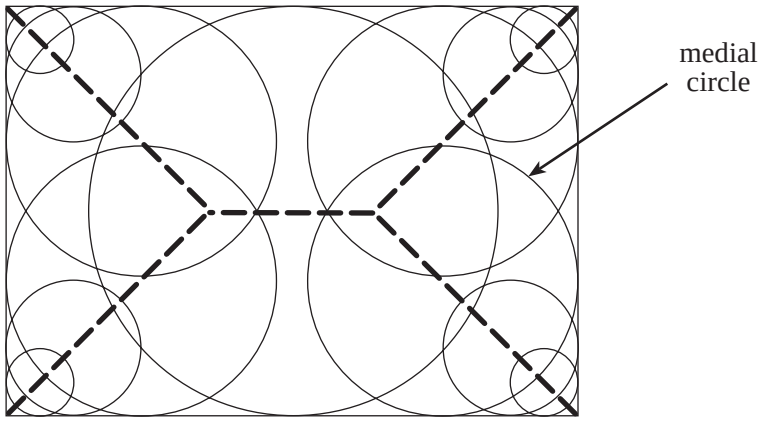

Fig. 1. The Blum medial manifold of a rectangle (thick dotted lines) is the locus of centers of all medial circles, i.e., those circles lying completely within the rectangle that touch the boundary in at least 2 places.

For objects in $\Re^3$, 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 [2, 3]. For m-dimensional objects, the generic medial manifold has $(m-1)$ dimensions, although fewer dimensions are possible in non-generic cases. In $\Re^3$, 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.

Volume of Symmetrical Figures by Integration of "Mining Rights"

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

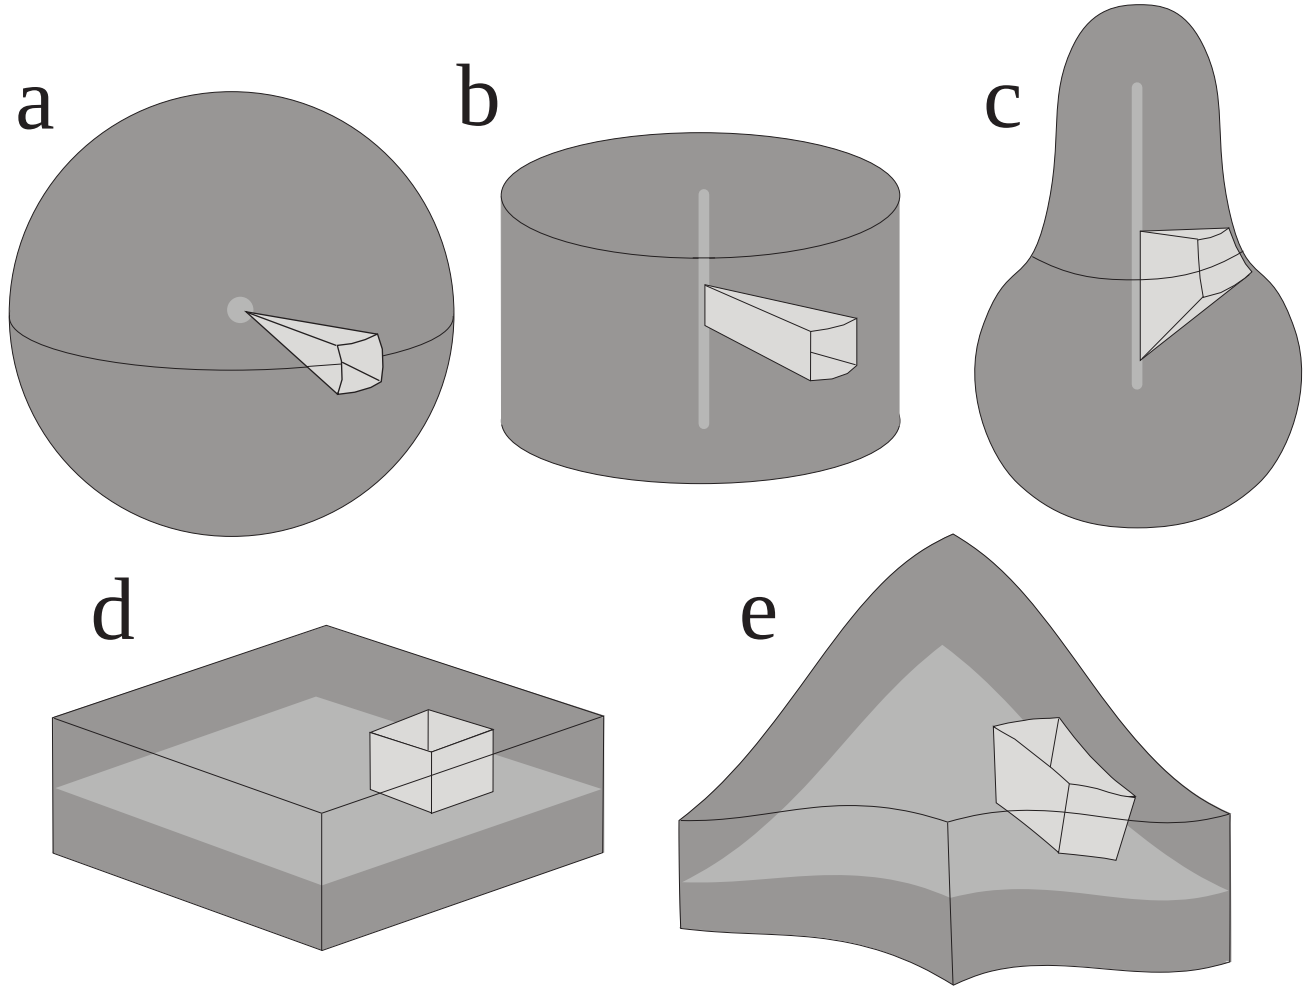

Fig. 2 Examples of 3D objects with simple shapes. The "mineral rights" for a surface patch are shown in each case extending orthogonal from the surface half-way through the object to the opposite side.

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. 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 object -- 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 [4]). Consider again the square patch on the surface of the sphere (Fig. 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. 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. 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. 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 from a wedge that widens at its cutting edge like the blade of an ax. Such wedges would fill the space within the pear and would not overlap, and thus could still be used to compute the volume from the surface, exactly as before.

Fig. 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. 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.

Consider now the closed boundary contour of the 2D object depicted in Fig. 3a. At each boundary point a coordinate system can be defined consisting of the unit [1] tangent vector $\hat{\mathbf{t}}$ and the unit normal vector $\hat{\mathbf{n}}$. As can be seen in Fig. 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 the 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.

---

[1] The '^' symbol denotes normalization of the vector.

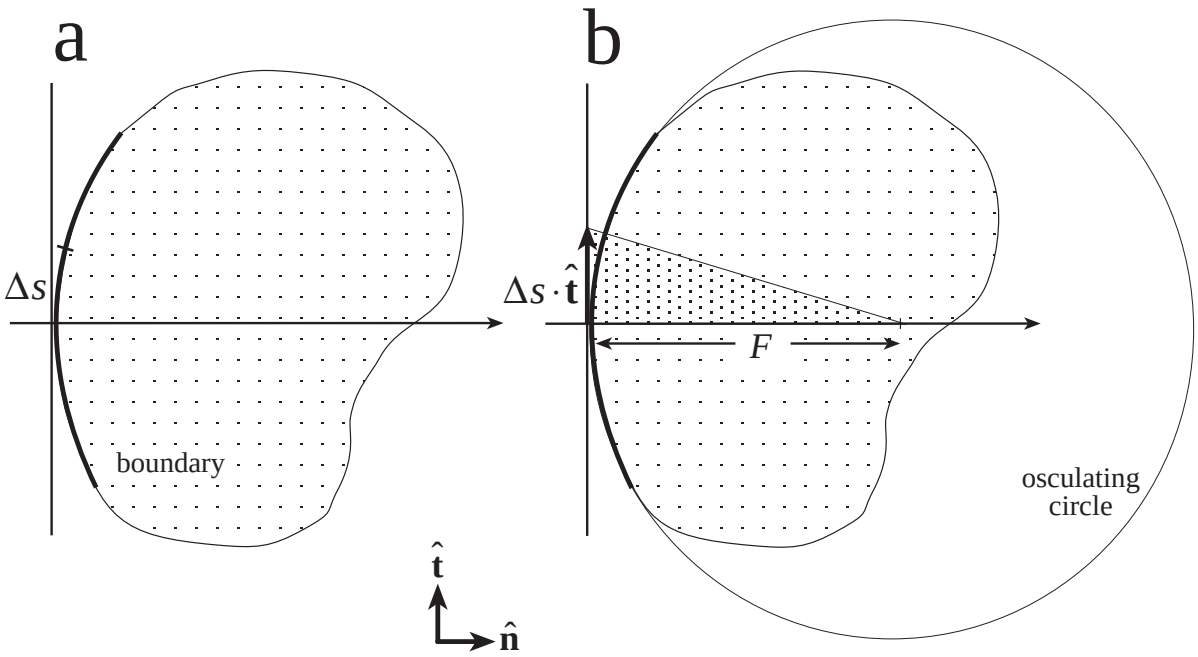

Fig. 3 **a.** Local coordinate system with $\hat{\mathbf{t}}$ tangential to, and $\hat{\mathbf{n}}$ normal to, the object's boundary. **b.** Osculating circle with radius F, and a sector of that circle corresponding to boundary interval Δs.

Area of an arbitrary 2D shape

We now develop a mathematical expression for the area of any closed smooth object in $\Re^2$ 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. 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. 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.

The relationship between the medial circle and the osculating circle is central to this paper. 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.

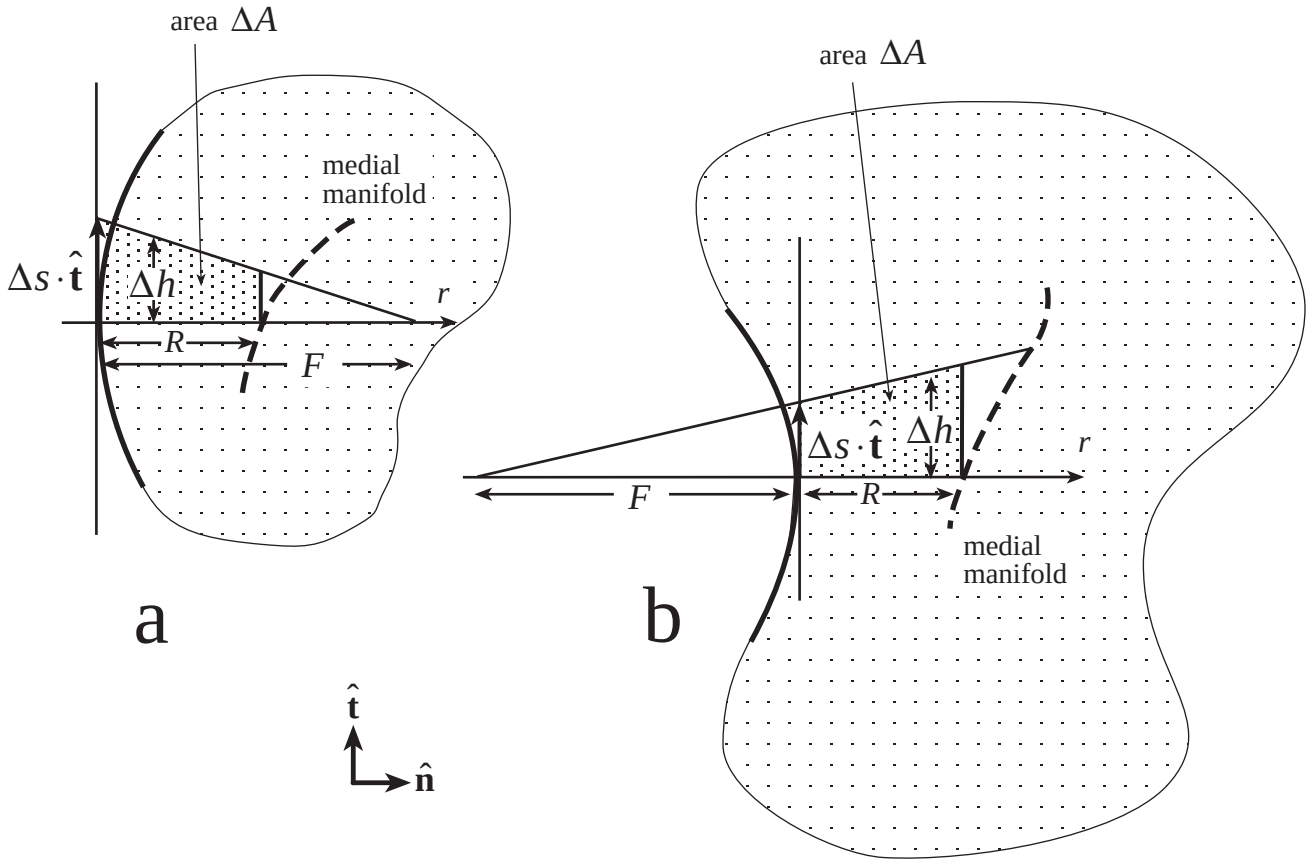

Fig. 4 **a.** Convex boundary segment (thick solid line) and corresponding medial manifold (thick dotted line) at a distance R from the boundary ($F > 0$). Area ΔA approximates the "mining rights" of boundary interval Δs. **b.** Same for a concave boundary segment ($F < 0$).

Now consider the convex boundary segment shown in Fig. 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

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

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

$$dA = \left[\int_0^R \left(1 - \frac{r}{F} \right) dr \right] ds. \tag{2}$$

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 (see below).

$$A = \oint_S dA \tag{3}$$

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. (2) to yield

$$dA = \left(r - \frac{r^2}{2R} \right) ds \Big|_0^R = \frac{R}{2} ds \tag{4}$$

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

$$A = \oint_S \frac{R}{2} ds = \pi R^2 \tag{5}$$

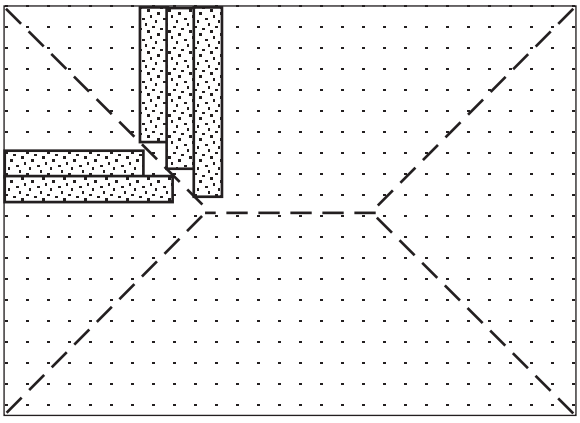

Fig. 5 Rectangle with mining rights extending in thin ribbons from the boundary to medial manifold.

Another example is shown in Fig. 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. (2) yields

$$dA = \left[\int_0^R \left(1 - \frac{r}{\infty} \right) dr \right] ds = R \cdot ds \qquad (6)$$

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.

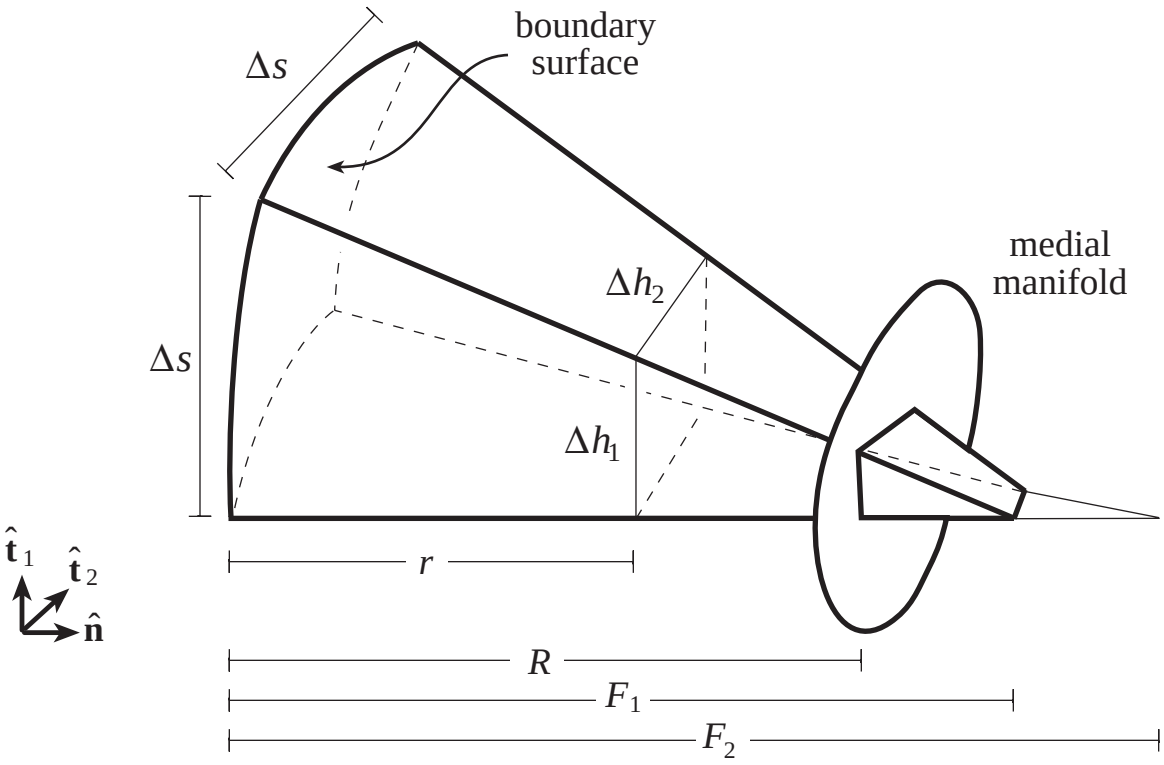

Fig 6. In $\Re^3$ the mining rights of a surface patch form an osculating wedge truncated at the medial manifold.

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 $\Re^3$. 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. 6). Whereas a single tangent $\hat{\mathbf{t}}$ and interval Δs suffice to describe the boundary contour in $\Re^2$, 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 $\Re^3$. 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

$$\Delta h_i = \left(1 - \frac{r}{F_i}\right)\Delta s \qquad i = 1,\ 2 \tag{7}$$

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

$$dV = \left[\int_0^R \left(1 - \frac{r}{F_1}\right)\left(1 - \frac{r}{F_2}\right) dr\right] ds^2 \tag{8}$$

The particular wedge shown in Fig. 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 $\Re^2$, any focal length corresponding to a convex principal curvature limits R as follows:

$$F_i \geq R \Big|_{f > 0} \qquad i = 1,2 \tag{9}$$

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.

$$V = \oint_A dV \tag{10}$$

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. (8), the infinitesimal volume dV corresponding to the infinitesimal surface patch ds^2 is

$$dV = \left(r - \frac{r^2}{R} + \frac{r^3}{3R^2} \right) ds^2 \Bigg|_0^R = \frac{R}{3} ds^2 \tag{11}$$

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

$$V = \oint_A \frac{R}{3} da = \frac{4}{3} \pi R^3 \tag{12}$$

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

$$dV = \int_0^R \prod_{i=1}^{m-1} \left(1 - \frac{r}{F_i} \right) ds^{m-1} dr \tag{13}$$

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 *ith* 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.

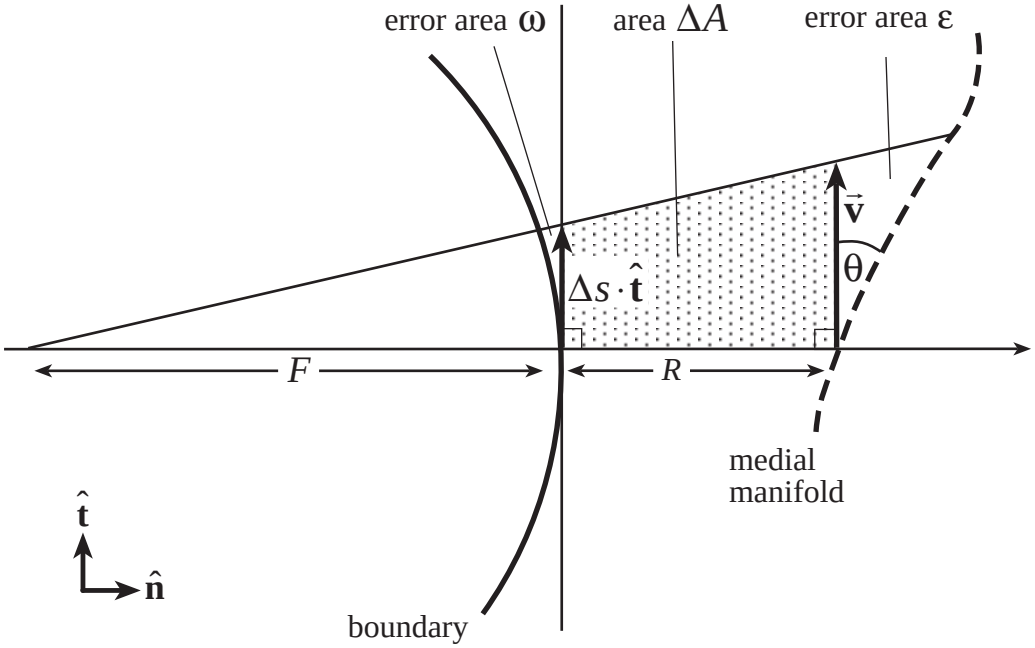

Fig. 7 Potential errors in the calculation of ΔA for a concave boundary segment.

Do the errors vanish?

The issue of convergence bears examination. Assume parametric descriptions of both the boundary and medial manifold are known. Consider an object in $\Re^2$. Shown in Fig. 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 $\vec{\mathbf{v}}$ of the trapezoid ΔA. A second area of error ω is shown between the boundary contour and the boundary tangent $\Delta s \cdot \hat{\mathbf{t}}$. Clearly $\omega/\Delta A \to 0$ as $ds \to 0$, so ω can be ignored in the limit. The error ε cannot so easily be ignored, however, since $\vec{\mathbf{v}}$ is generally not tangential to the medial manifold. The area ε as $\Delta s \to 0$ can be approximated for a concave boundary as

$$\lim_{\Delta s \to 0} \varepsilon = \frac{|\vec{v}|^2 \tan\theta}{2} = (\Delta s)^2 \left(\frac{F+R}{F} \right)^2 \frac{\tan\theta}{2} \quad \Big|_{F<0} \tag{14}$$

and for a convex boundary segment as

$$\lim_{\Delta s \to 0} \varepsilon = (\Delta s)^2 \left(\frac{F-R}{F} \right)^2 \frac{\tan\theta}{2} \quad \Big|_{F>0} \tag{15}$$

which simplifies for a straight boundary segment to

$$\lim_{\Delta s \to 0} \varepsilon = (\Delta s)^2 \frac{\tan\theta}{2} \quad \Big|_{F=\infty} \tag{16}$$

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

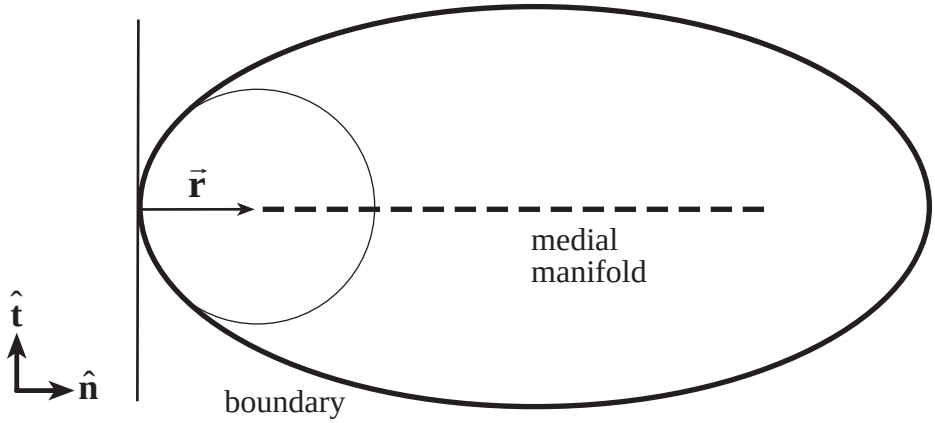

Fig. 8 Singular case for convex boundary.

First, consider the case for a convex boundary segment (see Fig. 8). In this case the medial manifold is tangential to $\vec{\mathbf{r}}$ 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 $\vec{\mathbf{v}}$ in Fig. 7 has zero length and Eq. (14) mandates that $\varepsilon = 0$.

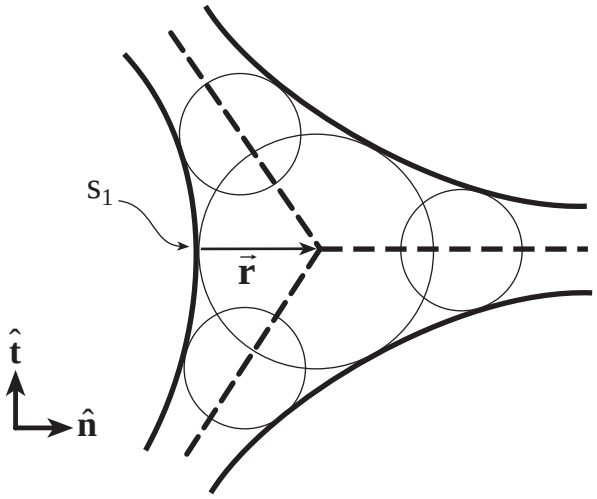

Fig. 9 Singular case for concave boundary.

Next consider the case for a concave boundary segment. The phenomenon of $\vec{\mathbf{r}}$ being tangential to the medial manifold at their intersection can only occur at a boundary point such as s_1 in the Fig. 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, where $\theta \neq \pi/2$ and $\tan\theta$ is finite. The same argument can be applied to the case of the flat boundary.

In higher dimensions $\Re^m$, for $m > 2$, each $(\Delta s)^{m-1}$ patch of the boundary has only one radius vector $\vec{\mathbf{r}}$, 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. 8 and 9 apply. In all cases the errors in the calculated volume will vanish as $\Delta s \to 0$.

Discussion

The method presented here uses the medial manifold to resolve ownership between boundaries across the object, and 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 [5]. 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 [6, 7]. The truncated wedges depend, instead, on a purely geometric relationship, with a deeper connection to other geometric processes such as 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 [8].

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 [9]. Fritsch has developed Deformable Shape Loci to adapt a medial model to fit objects in gray scale data [10]. Furst has devised methods of tracking ridges of medialness in gray scale data [11]. Stetten has developed methods of describing the medial manifold statistically in gray scale data [12, 13].

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 course-to-fine approach might provide estimates of volume at varying levels of precision. Medial approaches can determine an appropriate scale, and stabilize boundary parameters such as curvature, by permitting the proper ordering of boundary points [14].

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 $\Re^2$ 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 which 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.

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 the object.

Acknowledgments

Supported through a Whitaker Biomedical Engineering grant, NSF grant CDR8622201 NIH grants 1K08HL03220, P01CA47982, and HL46242. Special thanks to the Geometry Group in the Department of Computer Science, UNC, Chapel Hill for valuable feedback, and to Stephen M. Pizer, Ph.D. for tutelage, advice, and editing.

References

[1] H. Blum and R. N. Nagel, Shape description using weighted symmetric axis features, *Pattern Recognition*. **10**, 1978, 167-180.

[2] L. R. Nackman, Curvature relations in 3D symmetric axes, *CVGIP*. **20**, 1982, 43-57.

[3] L. R. Nackman and S. M. Pizer, Three-Dimensional shape description using the symmetric axis transform: I. Theory, *IEEE Trans. PAMI*. **2**, 2, 1985, 187-202.

[4] J. J. Koenderink, *Solid Shape*, MIT Press, Cambridge, MA, 1990.

[5] W. Kaplan, *Advanced Calculus*, Addison-Wesley, Reading, Mass., 1973.

[6] D. Eberly and J. Lancaster, On gray scale image measurements: I. Arc length and area, *CVGIP: Graphical Models and Image Processing*. **53**, 6, 1991, 538-549.

[7] D. Eberly, J. Lancaster and A. Alyassin, On gray scale image measurements: II. Surface area and volume, *CVGIP: Graphical models and image processing*. **53**, 6, 1991, 550-562.

[8] J. O'Rourke, *Computational Geometry in C*, Cambridge University Press, Cambridge, 1998.

[9] T. Culver., Accurate computation of the medial axis of a polyhedron. *Technical Report TR98-034*, Department of Computer Science, UNC, Chapel Hill., 1998.

[10] D. Fritsch, S. Pizer, L. Yu, V. Johnson and E. Chaney, Segmentation of Medical Image Objects Using Deformable Shape Loci, *Lecture Notes in Computer Science. 1230. Information Processing in Medical Imaging, Poultney, Vermont, 1997*, pp. 127-140.

[11] J. D. Furst and S. M. Pizer, Marching Cores: A method for extracting cores from 3D medical images., *Proceedings of the Workshop on Mathematical Methods in Biomedical Image Analysis*, San Francisco, CA, June 1996.

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

[13] G. Stetten and S. Pizer, Automated Identification and Measurement of Objects via Populations of Medial Primitives, with Application to Real Time 3D Echocardiography, *Submitted to Information Processing in Medical Imaging, 1999*, (*Technical Report TR98-035*, University of North Carolina, Department of Computer Science 1998).

[14] N. Amenta, M. Bern and M. Kamvysselis, A new Voronoi-based surface reconstruction algorithm, *SIGGRAPH Computer Graphics Proceedings*, Orlando, Florida, 1998., pp. 415-421.