Subset Warping: Rubber Sheeting with Cuts

Pierre Landau*and Eric Schwartz†

February 14, 1994

Correspondence should be sent to:

Eric Schwartz

Department of Cognitive and Neural Systems

Boston University

111 Cummington Street

Boston, Mass. 02215

(617) 353-6179

fax (617) 353-6178

Running Head: "Subset Warping: Rubber Sheeting with Cuts"

*Brain Research Labs, New York University Medical Center

†Department of Cognitive and Neural Systems, Boston University. This work supported by NIMH 5R01MH45969-04.

Abstract

Image warping, often referred to as "rubber sheeting" represents the deformation of a domain image space into a range image space. In this paper, a technique is described which extends the definition of a rubber-sheet transformation to allow a polygonal region to be warped into one or more subsets of itself, where the subsets may be multiply connected. To do this, it constructs a set of "slits" in the domain image, which correspond to discontinuities and concavities in the range image, using a technique based on generalized Voronoi diagrams. The concept of medial axis is extended to describe inner and outer medial contours of a polygon. Polygonal regions are decomposed into annular subregions, and path homotopies are introduced to describe the annular subregions. These constructions motivate the definition of a *ladder*, which guides the construction of grid point pairs necessary to effect the warp itself.

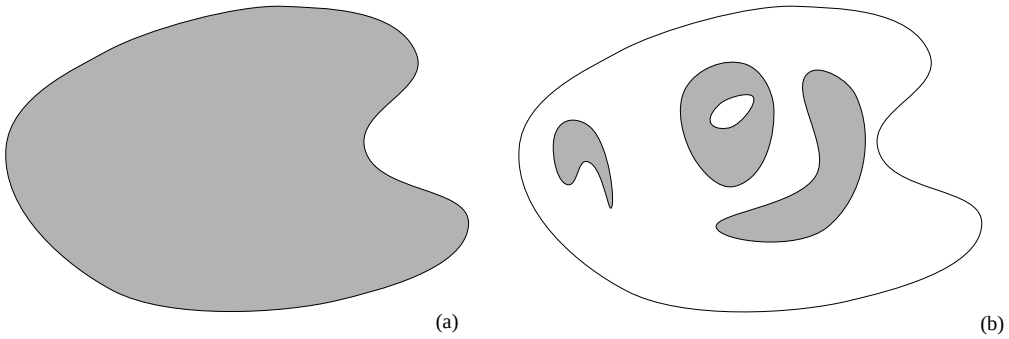

Figure 1: Subset warp of an image. The shaded area in (a) is warped to the shaded areas in (b). This is done by assuming that each polygonal area of the range draws its image from the portion of the domain image closest to itself.

1 Introduction

Image warping functions are usually known as "rubber-sheet transformations," as they represent the elastic deformation of a domain image space into a range image space. In this paper, a technique is described which extends the definition of a rubber-sheet transformation to allow a polygonal region to be warped into one or more subsets of itself, where the subsets may be multiply connected. To do this, it constructs a set of "slits" in the domain image, which correspond to discontinuities and concavities in the range image, using a technique based on generalized Voronoi diagrams.

One of the early motivations for the development of image warping was to correct images that had been distorted during acquisition. For example, in examining planetary images, it is necessary to correct the images for perspective or for known aberrations of the imaging lenses [39, 42] . In these cases, analytic functions exist to describe the distortion. In the biomedical field, one application of image warping is in modeling the mapping onto the primary cortical representation of the image formed on the retina. In this application, a conformal mapping is computed analytically, and then used to construct the image warp model [11].

For other applications, notably cinematic special effects [38, 43] where the correspondences

are not analytic, two paired sets of grid points are digitized by hand, and an assortment of interpolation techniques [45] allows texture mapping between the scenes.

Thus, existing image warping techniques are constrained to situations where the mapping between domain and range is known either as a low order polynomial or other analytic function, or can be specified by some simple heuristic or human assisted procedure of marking correspondences.

However, there are application areas of image warping which are of a more general nature, and for which existing image warping algorithms are not sufficiently powerful to provide a model. One such area is suggested by the details of the structure of the visual cortex in monkeys and humans: the left and right eyes terminate in a single cortical sheet in the form of the "ocular-dominance column pattern". This structure is most simply visualized as a kind of "zebra-skin" in which the black stripes represent the terminations of the left eye and the white stripes represent the terminations of the right eye [26]. The entire visual image recorded by the left eye is warped into its set of stripes in visual cortex, and similarly for the right eye. Thus, at a very basic level, an understanding of the architecture of visual cortex requires the ability to simulate image warps that are considerably more complex than have generally been considered. This aspect of the brain is termed "functional architecture", and provides a major current research area into the relationship of brain structure and function. Moreover, the complex patterns of spatial structure that have been observed in the brain provide a difficult challenge for the development of image warping techniques, such as the one described in this paper. This cross-disciplinary area in which problems in computational methods and in neuroscience are jointly addressed has been called "computational neuroscience" [34].

With this motivation in mind, the main thrust of the present work describes an approach to a broad class of transformations, in which an image is warped into a possibly disjoint set of smaller

subsets of its original area, each delimited by some arbitrary curve (see figure 1). We term this a "subset warp." As outlined in, e.g. [22], the columnar architecture of the brain provides a prominent example of this type of image warp, which also might result, e.g., from printing an image on a stretched rubber sheet and making arbitrary cuts in the sheet, allowing it to shrink in various ways. While initially developed for columnar modeling, this work generalizes the technique of image warping to arbitrary topological subsets, and we feel that the algorithm may have more general areas of application.

In the following discussion, we will use a number of geometric constructs, including several novel ones, in order to provide a general solution to the problem of subset image warping. There are two major sections to this algorithm. First, we outline a method for constructing a set of image correspondences between the original image and a union of some collection of disjoint subsets. We will imagine the range image to be painted on the domain image like a zebra-skin, with each connected black or white subset termed a "column". These columns need be neither convex nor simply connected. For each such column, a corresponding region of the original image (which we call the pre-column or protocolumn) will contain the pre-image of this column. Thus, the first stage of this algorithm is to decompose the original image into a collection of regions which form a partition of the image. In order to this, we use the geometric data structures of Generalized Voronoi Polygon, medial axis and external medial axis. We also introduce new structures which we term *annulus* and *proto-annulus*.

After introducing and defining these geometric structures, we prove the main theorem of the first part of this algorithm, which is a constructive partition of an image into a union of structures which we call protoannuli. This partition is entirely constrained by the structure of the range image (column system), which is partitioned by the technique into a set of annuli, each having a one-to one correspondence with a protoannulus of the domain. This result applies to images and

subsets of arbitrary topological type.

The partition of domain and range images into annular regions which must be warped to one another pairwise reduces the general warping problem to a series of smaller, more accessible ones, which can each be solved independently. As several possible methods suggest themselves to solve the annulus-to-annulus warping problem, we introduce the use of path homotopy, to provide a framework for the comparison of warping techniques. This paper suggests that a grid developed from the Voronoi tesselation leads to a minimum displacement (least deformation) piecewise bilinear warp. To build this grid, we introduce a construction which we term a "ladder".

Finally, we demonstrate the application of this technique to several different "columnar" systems of the primate brain, extending earlier results described in [22].

2 Voronoi Diagrams and Medial Contours

2.1 Generalized Voronoi Diagram

Given a set **P** of N labeled points $p \in \mathbf{P}$ in the plane, the *Voronoi polygon* (also known as *Voronoi region*) of each p is the locus of points closer to this point than to any other point in **P**:

$$v_p = \left\{ q \in \mathbf{P} | \nexists p' \in \mathbf{P}, p' \neq p, d(q,p') < d(q,p) \right\}$$

where $d(x,y)$ is the distance from x to y. The *Voronoi diagram* is the set of all such polygons[1] [30].

Similarly, let the set P consist of a set of N polygons in the plane, The *generalized Voronoi polygon* of a polygon is defined as the region containing points in the plane closer to this polygon than to any other labeled polygon, and the *generalized Voronoi diagram* as the set of all such generalized

[1]Properly speaking, the polygons of points on the convex hull of the set of points **P** extend to infinity, but they are restricted here to an image which is a subset of the plane, so that all the Voronoi polygons are finite. Voronoi regions are typically defined with Euclidean metrics, though others have been used [9].

Voronoi polygons. Generalized Voronoi polygons are not necessarily convex and may contain one or more holes.

2.2 Inner and Outer Medial Contours

Given a binary image, containing regions of one of two classes of points, the *contours* are the set of Jordan curves[2] that form the boundaries of regions. Each region has one exterior contour, and one additional contour for each internal hole. These contours are sampled by some spatial sampling process, subject to a minimal distance criteria to be specified, and are represented by their respective sets of sample points.

The medial axis, or skeleton, of a polygon [4, 8], (and see [18] for a recent review) is defined in one of two ways: the result of a morphological, "prairie fire" thinning operation applied to the digital representation of the polygonal region, leaving just the "skeleton" of the figure, or, more precisely, as the loci of proximity of the non-adjoining sides of the containing polygon [30]. To construct the medial axis in this latter way, the Voronoi procedure is modified to compute adjacency sets of edges rather than points, where the distance between a point and an edge is simply the shortest distance to the edge $d(x, E) = \inf(d(x, y)), y \in E$.

Expressed in a different way, the medial axis of a polygon is the maximal subset of the edges in the Voronoi diagram such that both endpoints of the edge are strictly contained in the polygon[3][6, 24, 23, 22]. By analogy, the *external medial axis* is the subset of the edges in the Voronoi diagram that are contained in the generalized Voronoi region[4] but have both endpoints exterior to

[2] a Jordan curve is a closed curve in the plane that separates a simply connected interior region from an unbounded exterior region [30].

[3] This is easily seen to follow from the previous definition: internal segments of the Voronoi diagram are equidistant from two sides of the polygon and therefore part of the medial axis; segments of the Voronoi diagram that are (even partly) external to the region constitute bisectors of adjoining sides of the polygon and are therefore not part of the medial axis.

[4] For this purpose the Voronoi region is assumed to be a closed interval containing its bounding polygon.

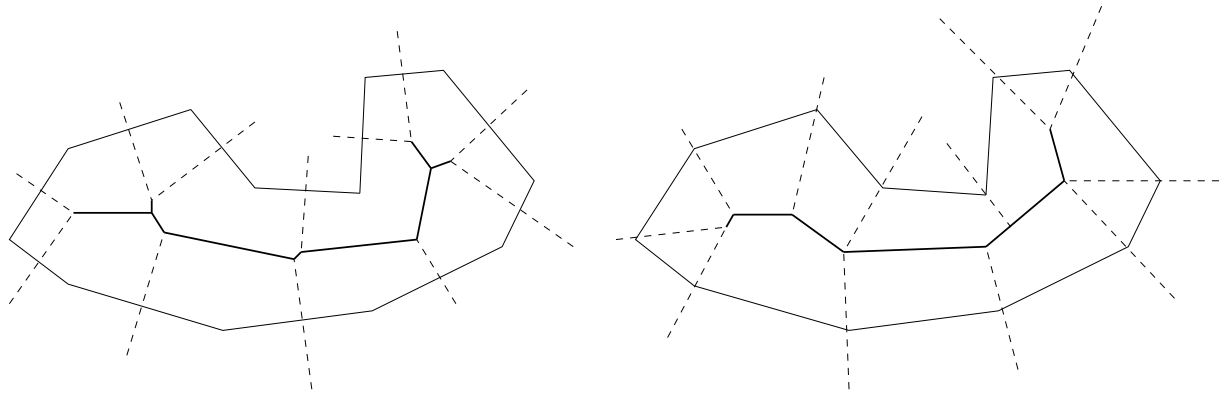

Figure 2: Two types of medial axes can be created, one using the Voronoi tessellation due to the vertices of a polygon V_P, the other due to the edges of the polygon V_E. As the number of vertices increases, the two medial axes become indistinguishable.

the polygon itself. Edges that bisect adjacent contour points are not part of the internal or external medial axes. If one disregards behavior at the image boundary, the union of the external medial axes is identical to the union of the internal medial axes of the complement of the image.

The Voronoi diagram $\mathbf{V}_E$ obtained from the edges is in a sense the dual of that obtained from the vertices of the polygon $\mathbf{V}_P$ (see figure 2). In the limit where the size of an edge becomes very small and the number of vertices large, the "medial" portion of both $\mathbf{V}_P$ and $\mathbf{V}_E$ are identical. This is because when an angle becomes more acute, its cosine approaches one, so that the distance between an interior point and its nearest contour vertex is not appreciably different from the distance between it and its nearest contour edge. In this work, $\mathbf{V}_P$ is used as it is simpler to compute.

The term *medial contour* is a generalization of the medial axis: it is the locus of *boundary points* of the medial axis[5], ordered in a counterclockwise direction. The contour lies arbitrarily close to the medial axis and is not generally convex. Analogously, the *external medial contour* is defined as the set of internal boundary points of the external medial axis, ordered in a clockwise direction.

[5]The *limit points* p of a set E are those points for which every sphere centered at p contains points of E other than p. The *boundary points* of E are all the limit points of E that do not belong to E [20].

The region of the plane delimited by these two contours is called the *protoannulus*, and the portion of the protoannulus internal to the original polygon it self is called the *annulus*.

2.3 Definition and Construction of a Protoannulus

Let $\mathbf{C}_c$ represent the set of contour points of contour c. These points have been obtained by sampling the contours, both internal and external, as described below. Let $\mathbf{C} = \bigcup_i \mathbf{C}_c$ represent the union of all such sets, i.e. the union of all region and hole boundary point sets. $\mathbf{V}$ is the Voronoi diagram for $\mathbf{C}$, and $\mathbf{V}_c \subset \mathbf{V}$ the subset of $\mathbf{V}$ corresponding to points in $\mathbf{C}_c$. Each $v_{cp} \in \mathbf{V}_c$ is a Voronoi polygon associated with the contour point p of contour c. We associate with v_{cp} the set of edges $\mathbf{E}_{cp} = \{e_{cpp'}\}$ that delimit the region. These edges are oriented in a clockwise direction. Each edge is labeled by the two Voronoi polygons it separates.

Consider the set of edges associated with the c^{th} contour

$$\mathbf{E}_c = \bigcup_p \mathbf{E}_{cp} - \mathbf{E'}_c = \bigcup_{p,p'} e_{cpp'} - \mathbf{E'}_c$$

where

$$\mathbf{E'}_c = \{ e_{cpp'} \in \mathbf{E}_c | p, p' \text{ are adjacent contour points on the same contour } \}$$

(see figure 3).

The set $\mathbf{E}_c$ consists of two concentric oriented contours that delimit a region of the plane in the form of an annulus, though in many cases, the inner contour is degenerate, corresponding to the cases where points in the region are simply connected. This annular region is termed the *protoannulus* as it will define the preimage of the *annulus* to be defined below.

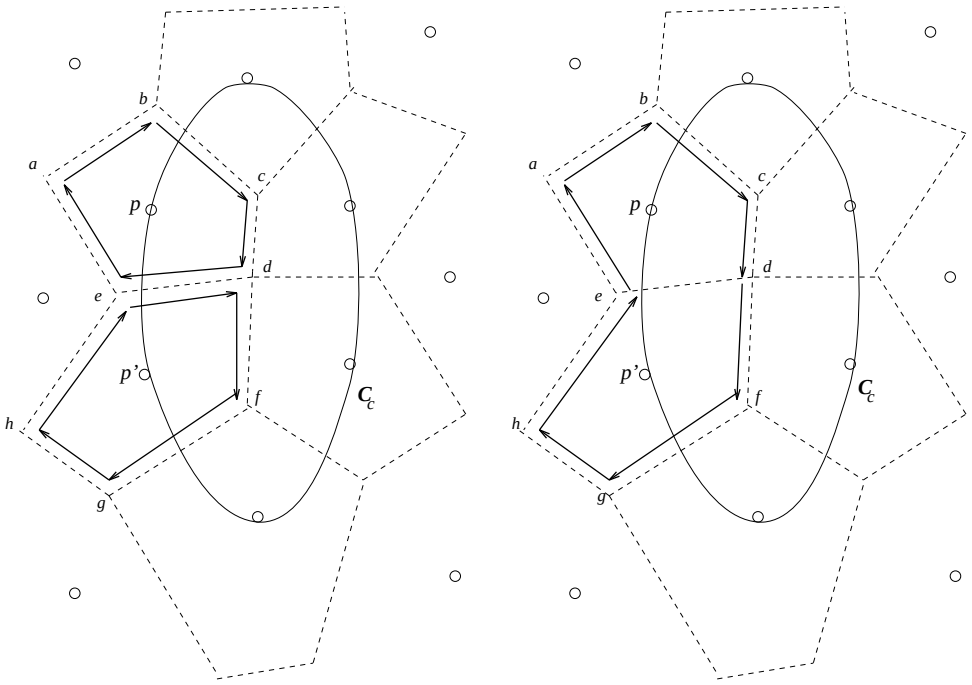

Figure 3: The two adjacent points p and p' have Voronoi polygons $abcde$ and $edfgh$. In constructing our medial and dual-medial contours, the union of these two sets, E_c, is formed, edges $\vec{ed}$ and $\vec{de}$, which divide the polygons belonging to adjacent contour points on $\mathbf{C}_c$ are subtracted, leaving polygon $abcdfgh$.

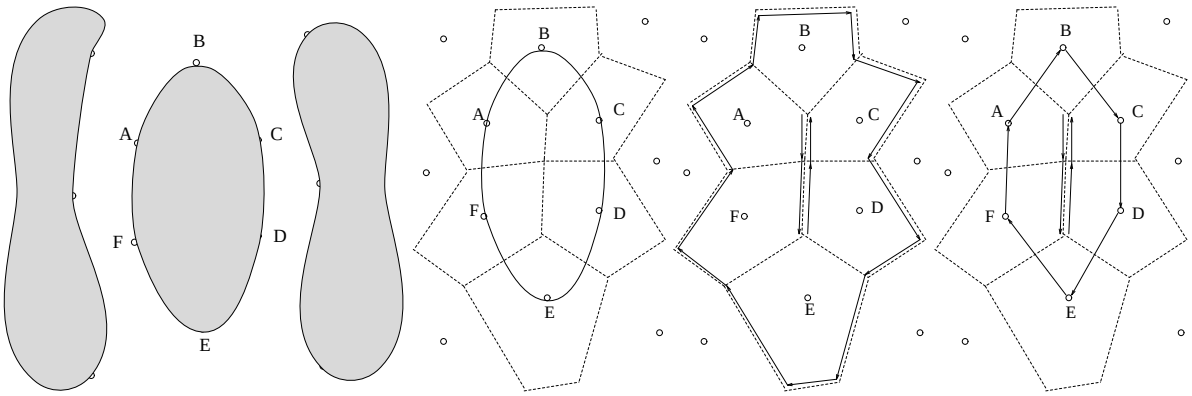

Figure 4: (a) A simple, convex region (shown with two adjacent regions). This polygonal region is delimited by a single contour. For illustrative purposes, six labeled contour points A, B, C, D, E, F were chosen, though in the algorithm, the grid is much denser. (b) the corresponding Voronoi polygons. (c) the set $\mathbf{E}_i$ of edges that delimit the protoannulus, composing two disjoint sets. termed "outer medial contour" and "inner medial contour". (d) the set $\hat{\mathbf{E}}_i$, that delimit the original region. The internal medial contour remains unchanged; the region contour is shown. The internal medial contour is effectively degenerate, but is treated as a true contour.

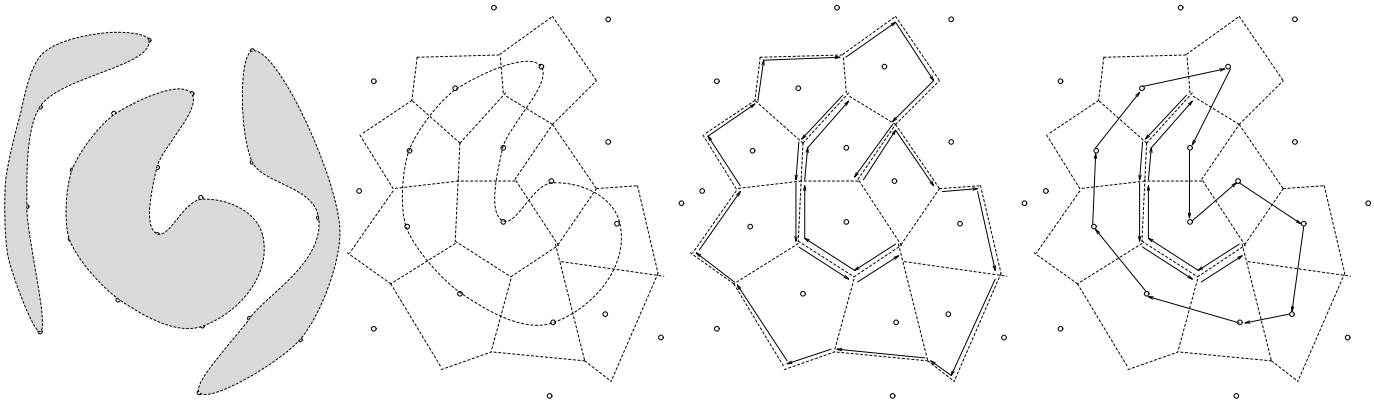

(a) (b) (c) (d)

Figure 5: (a) A polygonal region which is not convex. (b) the corresponding Voronoi polygons. (c) the set $\mathbf{E}_i$ of edges that delimit the protoannulus, and (d) the set $\hat{\mathbf{E}}_i$, that delimit the annulus. Note that here part of the external medial contour is degenerate.

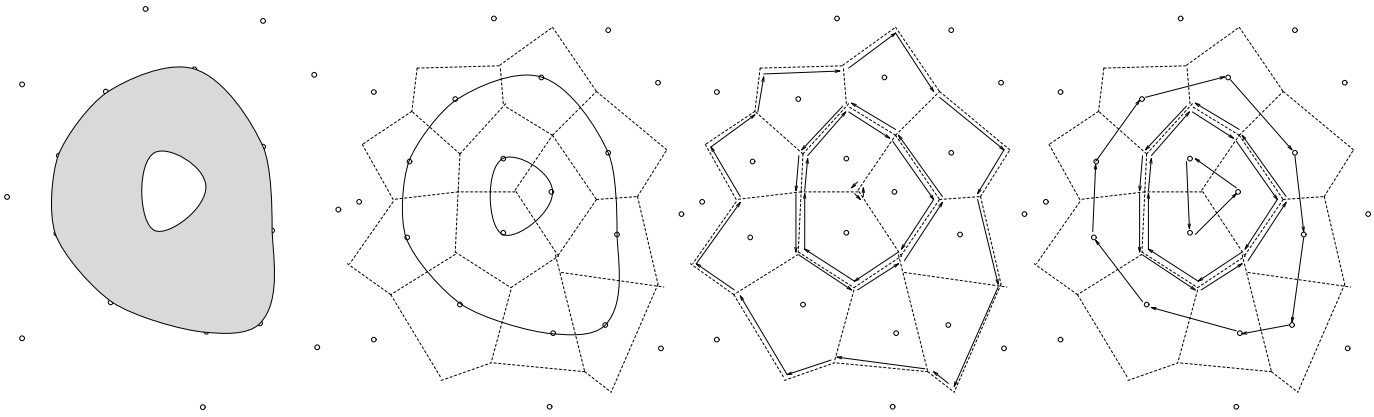

(a) (b) (c) (d)

Figure 6: (a) A polygonal region with a hole in it. This region has two contours; each generates a protoannulus. (b) the corresponding Voronoi polygons. (c) the set $\mathbf{E}_i$ of edges that delimit the two protoannuli. (d) the set $\hat{\mathbf{E}}_i$, that delimit the annuli. Note that points in the outer annulus will be mapped outward, whereas those in the inner annulus will be mapped inward.

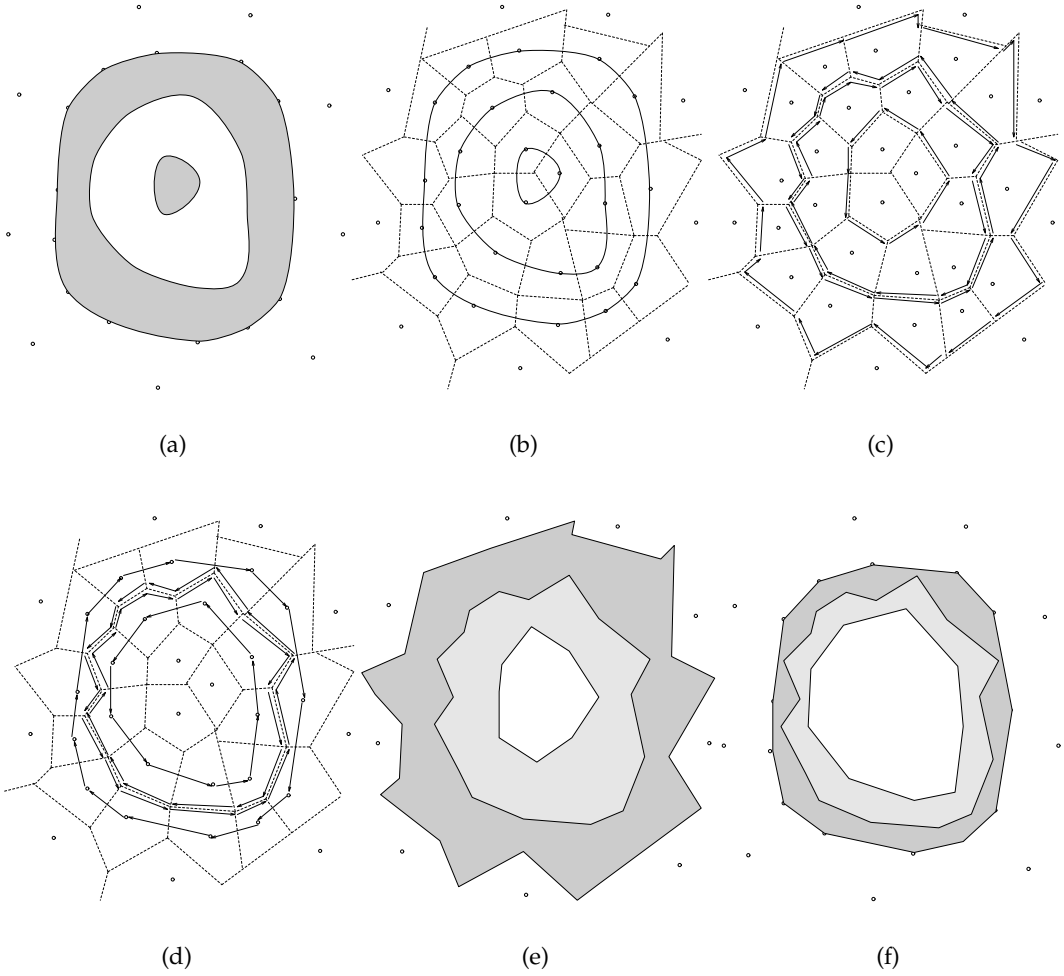

Figure 7: (a) A polygonal region with another polygonal region embedded in it. The region is decomposed by the procedure into two concentric annuli. (b) the corresponding Voronoi polygons. (c) the set $\mathbf{E}_i$ of edges form four contours that delimit the two protoannuli. (d) the set $\hat{\mathbf{E}}_i$ of edges that delimit the annuli. This resembles the previous case, except that here the innermost contour of $\mathbf{E}_i$ is no longer degenerate. (e) Protoannuli that represent the mapping from (f) annuli.

2.4 Construction of an annulus

An annulus is the subset of the protoannulus contained in the region (column) proper. It is constructed, similarly to $\mathbf{E}_c$ above, from regions $\hat{v}_{cp} = \{p \in v_{cp} | p \in \mathbf{C}_c\}$. This corresponds to considering the portion of each Voronoi polygon that is internal to the original polygon.

As in the definition of the protoannulus, above, we define the set

$$\hat{\mathbf{E}}_c = \bigcup_{p,p'} \hat{e}_{cpp'} - \hat{\mathbf{E}}'_c$$

where

$$\hat{\mathbf{E}}'_c = \left\{ e \in \hat{\mathbf{E}}_c | p, p' \text{ are adjacent contour points on the same contour} \right\}$$

By construction, the annulus is entirely contained in the protoannulus. Further, it sometimes represents the inner portion of the protoannulus. In this case, we call the mapping "outward." In the case where the annulus represents the external portion of the protoannulus, the mapping is termed "inward" (see figure 9).

2.5 Examples

The simplest example of this discussion is shown in figure 4, which depicts a convex column and its two neighboring columns. The Voronoi polygons that would correspond to choosing six sample points along the contour of the column are shown, though in practice, the number of sample points used would be at least an order of magnitude greater. The internal and external medial contours are readily seen to correspond to the edges of the Voronoi polygons. The medial contour bounds a slit that is infinitesimal in width [1].

Figure 5 depicts the complications of considering a a nonconvex column. A portion of the external medial contour is slit-like here as well. Figure 6 depicts a column with an embedded

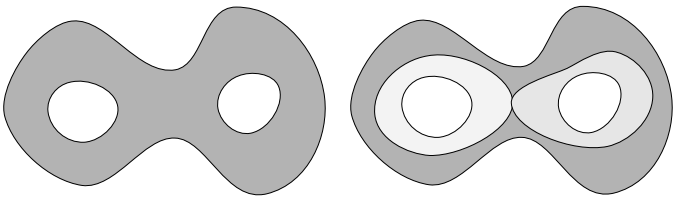

Figure 8: A planar region with holes can be decomposed into a series of annular regions, each with a single hole in it. For simplicity, regions with no holes are each assumed to have have an infinitesimal hole, allowing any region to be expressed as a sum of annular regions.

hole. This column has two contours, each of which gives rise to internal and external medial contours. The outer contour has a medial contour that corresponds to the inner contour's external medial contour. In this case, the medial contour of the inner contour is reduced to an infinitesimal circular region. In figure 7, the result of embedding a different column in the original column is considered. The results are very similar to those in the previous figure, with the exception that the medial contour of the hole boundary is no longer degenerate.

2.6 Decomposition of Planar Regions into Annuli

In order to be able to define a continuous mapping (warping) between two finite planar regions, one must in the general case guarantee that they are deformable onto each other, i.e., that they have the same Euler number, or the same topological type. To resolve this difficulty, regions are decomposed into a series of (possibly degenerate) annuli, each of which has one and only one hole in it (see figure 8).

Definitions: In the following, a discrete sampling of the plane of sufficiently high frequency not to introduce aliasing errors is assumed. Aliasing errors would use a single point to represent either two contours or two portions of the same contour; thus it is sufficient that the largest intersample distance be smaller than both the smallest inter-contour distance and the width of any concavity in a single contour. If this condition is satisfied, there is never a pixel shared between

two contours, and contours never have articulation points[6].

Let p be a point in contour $\mathbf{C}_c$. Then there is no contour point $q \in \mathbf{C}_{c'} \neq \mathbf{C}_c$ on some other contour such that q is closer to p than either of p's neighbors in $\mathbf{C}_c$. This follows directly from anti-aliasing constraints imposed on the initial sampling.

THEOREM: Let $\mathcal{S}$ be a sampling of the plane of high enough frequency to avoid aliasing. Let $\mathbf{C}_c \subset \mathcal{S}$ describe a closed (interior or exterior) contour of a region. Let $\{e_{cpp'}\}$ represent the clockwise-oriented set of edges of the Voronoi polygon associated with point $p \in \mathbf{C}_c$. Each such edge separates the Voronoi polygon of point p from that of point p'. Then the set

$$\mathbf{E}_c = \bigcup_{p \in C_c, p' \in C_c} e_{cpp'} - \mathbf{E'}_c$$

where

$$\mathbf{E'}_c = \{ e_{cpp'} \in \mathbf{E}_c | p, p' \text{ are adjacent contour points on the same contour } \}$$

is composed of two closed contours, one interior to (termed *inner medial contour*), one exterior to (termed *external medial contour*), the closed contour $\mathbf{C}_c$. These two concentric contours bound a region of the plane called a *protoannulus*. Further, none of the protoannuli generated by different contours in the plane overlap, and the set of all protoannuli form a partition of the plane.

PROOF: By (modified) induction. Initial case: $\mathbf{C}_c$ is a closed contour; the Voronoi polygon for each point along $\mathbf{C}_c$ is convex, and shares an edge with both of its neighbors along the contour. This is always true because they are its closest neighbors. When two Voronoi polygons, each

[6]An articulation point of a contour is a point on the contour that, when deleted, breaks the contour into two disconnected parts.

delimited by a clockwise contour, are juxtaposed, and the shared edges are eliminated, there remains a clockwise path delimiting both regions (see figure 3). Induction: if there exists a set of adjacent contour points of length n constructed by such a procedure, then the property holds true for a set of length $n + 1$, unless this closes the contour (termination). The regular case is similar to the initial case. For the termination case, when the last point on the contour is added, closing the contour, the two "end" segments of the contours are removed, leaving two closed contours. As all shared boundary elements that are eliminated from $\mathbf{E}_i$ are perpendicular to the contour $\mathbf{C}_i$ and necessarily present in the original set, the contours formed are on either side of the initial contour and thus cannot cross it. Moreover, as the Voronoi polygons due to points along the contour are all pairwise adjacent with no gaps, the termination case forms two closed contours.

The protoannulus generated cannot overlap any other protoannulus because, by its nature, the Voronoi diagram is a partition of the plane into discrete regions of points, and any partition of the original partition is still a partition of the plane.

COROLARY: Any set of multiply connected regions in the plane can be partitioned into a set of protoannuli, as defined above.

PROOF: Simply consider each interior and exterior contour of a column separately. Each such contour generates a single protoannulus, a result that follows from above.

These regions have been termed *protoannuli* with a specific purpose. Each protoannulus can be partitioned into two annuli, one internal and one external to the original region, by a construction similar to that described above. In the image warping to be described below, a mapping is defined between each protoannulus and its corresponding annulus. The annulus sometimes represents the inner portion of the protoannulus. In this case, the mapping is termed "inward." In the case

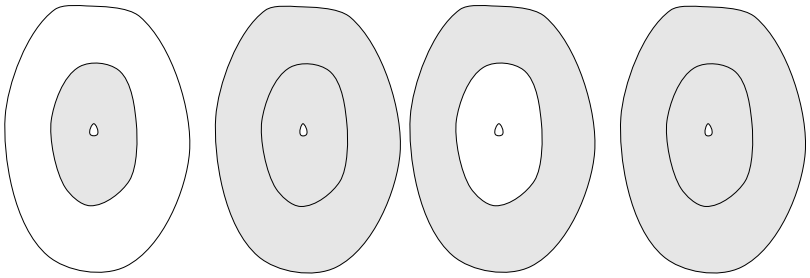

(a) (b)

Figure 9: Figure (a) shows a circular region (*annulus*) that "grows" outward (before and after), and (b) shows the case "grow" inward(before and after)

where the annulus represents the external portion of the protoannulus, the mapping is termed "outward" (see figure 9).

2.7 Sensitivity of medial contours to polygon contour shape

As the algorithm described in this paper makes extensive use of sampled contours, the question of sampling frequency becomes important. As we have shown above, a minimum sampling frequency is dictated by aliasing concerns; an increase in sampling frequency gives a better fit to the original subset shapes, at the expense of increased running time. Another potential difficulty with increasing sampling rate is the increasing branching that is observed in the medial contour. Much of this branching is due to "unimportant" features of the coutour.

A common difficulty with the use of medial axes in object recognition and other applications is instability in the presence of small amplitude noise in the original polygon boundary. The addition of a small bump in the polygon creates an entire new (and undesireable) arm of the skeleton. There is, however, a simple technique for reducing the magnitude of the problem. Recall that it has been shown previously that all edges on the medial axis are edges in the Voronoi diagram of the polygon. Each such edge forms the boundary between two Voronoi regions v_i