1993-07

Subset Warping: Rubber Sheeting with Cuts

SUBSET WARPING: RUBBER SHEETING WITH CUTS

Pierre Landau and Eric Schwartz

July, 1993

Technical Report CAS/CNS–93–050

Boston University Center for Adaptive Systems and
Department of Cognitive and Neural Systems
111 Cummington Street
Boston, MA 02215

Subset Warping: Rubber Sheeting with Cuts

Pierre Landau*and Eric Schwartz†

May 13, 1993

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

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 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.

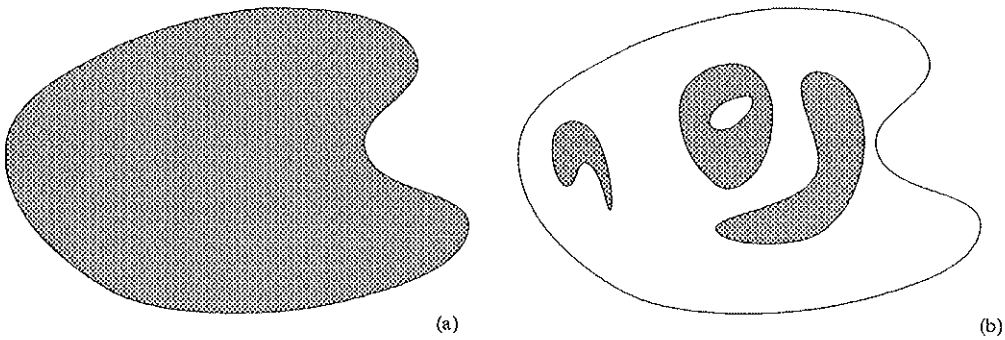

Figure 1: Subset warp of an image. The shaded area in (a) is warped to the shaded areas in (b). This is done by assuming that each polygonal area of the range draws its image from the portion of the domain image closest to itself.

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 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 [34, 37] . In these cases, analytic functions exist to describe the distortion.

One biomedical application of image warping is in modeling the mapping of the image formed on the retina onto the primary cortical representation. This is an example of "topographic mapping," which is common in the visual, auditory, and somatosensory systems of all higher vertebrates. In this application, a conformal mapping may be computed analytically, and then used to construct the image warp model [9].

For other applications, notably cinematic special effects [33, 38] where the correspondences are not analytic, two paired sets of grid points are digitized by hand, and an assortment of interpolation techniques [40] 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 model, 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 [23]. 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.

In addition to this particular biomedical system, it is now clear that much of the sensory cortex consists of various forms of "columnar" systems. All of these systems have in common the representation of multiple copies of a visual image, which has been filtered by one or another modality, with the results "warped" into a cortical columnar system. Thus, in addition to the left and right eye system outlined above, the image of a single eye is represented by "orientation columns" in cortex, which are strips of tissue which contain neurons which all respond to a specific orientation of stimulus. Many other columnar systems have been described in visual, auditory, and somatosensory cortex. It is fair to say that the cortical systems of our brains are based on a columnar architecture, and the understanding of this architecture, and the ability to simulate this form of architecture, is crucial to studies of the brain.

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 above, 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. The class of transformations described is particularly useful, as outlined above, in modeling columnar structures of the cerebral cortex. However, in generalizing the technique of image warping to arbitrary topological subsets, we feel that this algorithm may have more general area of application. In particular, a number of geometric constructs are used, and introduced here, which provide some insight into an approach which significantly generalizes the usual analytic or point-correspondence approach to image warping.

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 original image to be painted like a zebra-skin, which each connected black or white subset termed a "column". We do not restrict these columns to be either convex or simply connected. Thus, the first stage of this algorithm is to decompose the original image into a collection of regions which form a partition of the original image. In order to this, we use the geometric data structures of Generalized Voronoi Polygon and medial axis. We also introduce new structures which we term external medial axis (the dual to the medial axis), 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 and a set of columns defined on that image into a union of structures which we call annuli and proto-annuli. This result applies to images and column systems of arbitrary topological

type.

Having constructed a partition of the image/column system, we then outline a method for constructing a piecewise bilinear image warp. We introduce the use of path homotopy, which has been implicitly used in other warping methods, and a construction which we term a "ladder", which allows the performance of a local, piecewise bilinear image warp on the results of the image partition that we have constructed.

Finally, we demonstrate the application of this technique to several different "columnar" systems of the primate brain, extending earlier results described in [19].

2 Voronoi Diagrams and Medial Contours

2.1 Generalized Voronoi Diagram

Given a set $\mathbf{P}$ 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 $\mathbf{P}$:

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

where $d(x,y)$ is the Euclidean distance from x to y. The *Voronoi diagram* is the set of all such polygons[1] [27].

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 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, 7], 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 [27]. To construct the medial axis in this latter way, the Voronoi procedure is modified to compute adjacency

[1] Properly speaking, the polygons of points on the convex hull of the set of points $\mathbf{P}$ 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.

[2] a Jordan curve is a closed curve in the plane that separates a simply connected interior region from an unbounded exterior region [27].

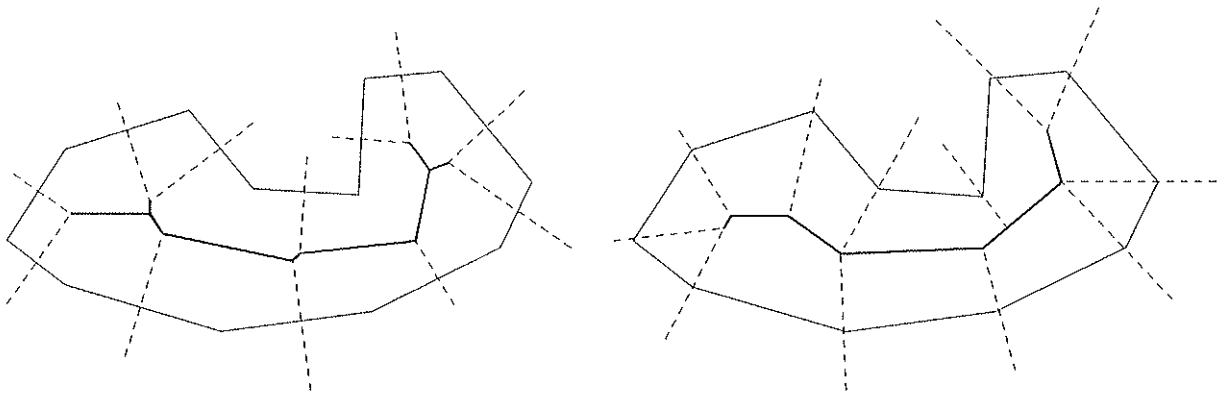

Figure 2: Two types of medial axes can be created, one using the Voronoi tessellation due to the vertices of a polygon, the other due to the edges of the polygon. As the number of vertices increases, the two medial axes become indistinguishable.

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][21, 19, 20]. 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 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 than 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. The region of the plane delimited by these two contours is called the *protoannulus*, and the

[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.

[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 [17].

portion of the protoannulus internal to the original polygon itself 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.

2.4 Construction of an annulus

An annulus is the subset of the protoannulus contained in the region 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 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, we term the mapping "inward" (see figure 9).

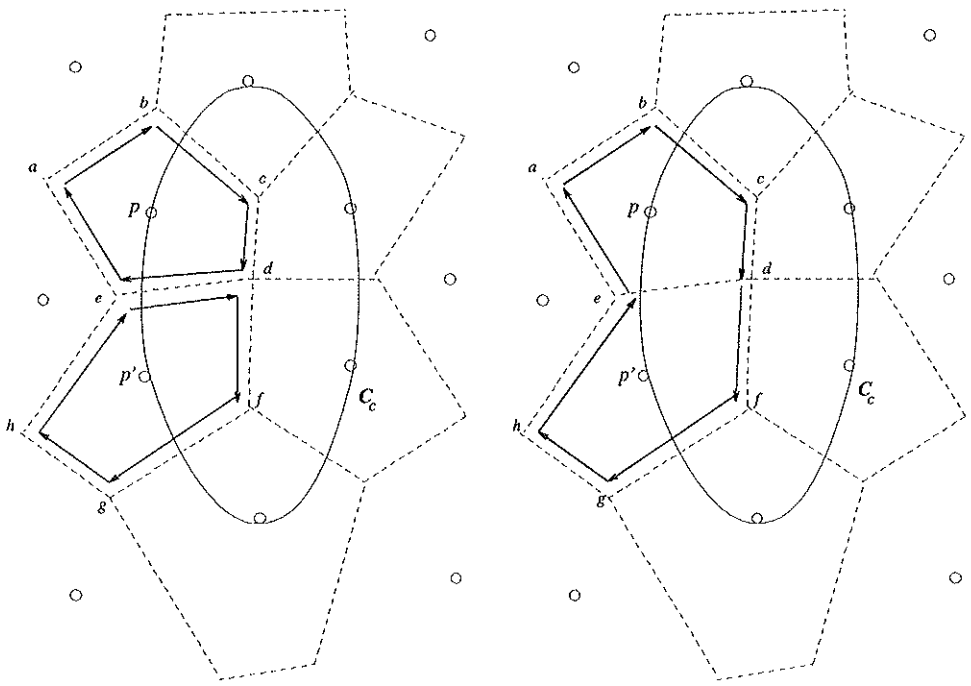

Figure 3: The two adjacent points p and p' have Voronoi polygons $abcde$ and $edfgh$. In constructing our medial and dual-medial contours, the union of these two sets is formed, edges $\vec{ed}$ and $\vec{de}$, which divide the polygons belonging to adjacent contour points on C_c are subtracted, leaving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

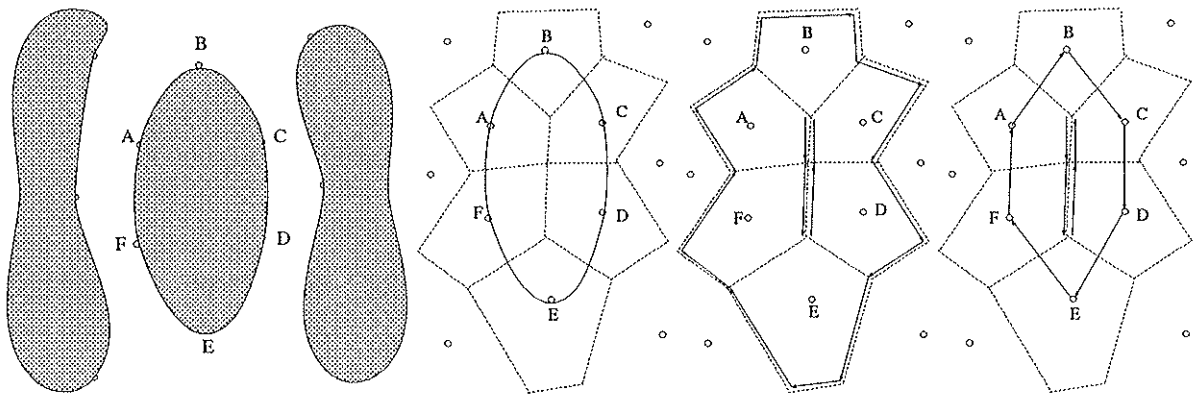

Figure 4: (a) A simple, convex region (shown with two adjacent regions). This polygonal region is delimited by a single contour. For illustrative purposes, six labeled contour points A, B, C, D, E, F were chosen, though in the algorithm, the grid is much denser. (b) the corresponding Voronoi polygons. (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.

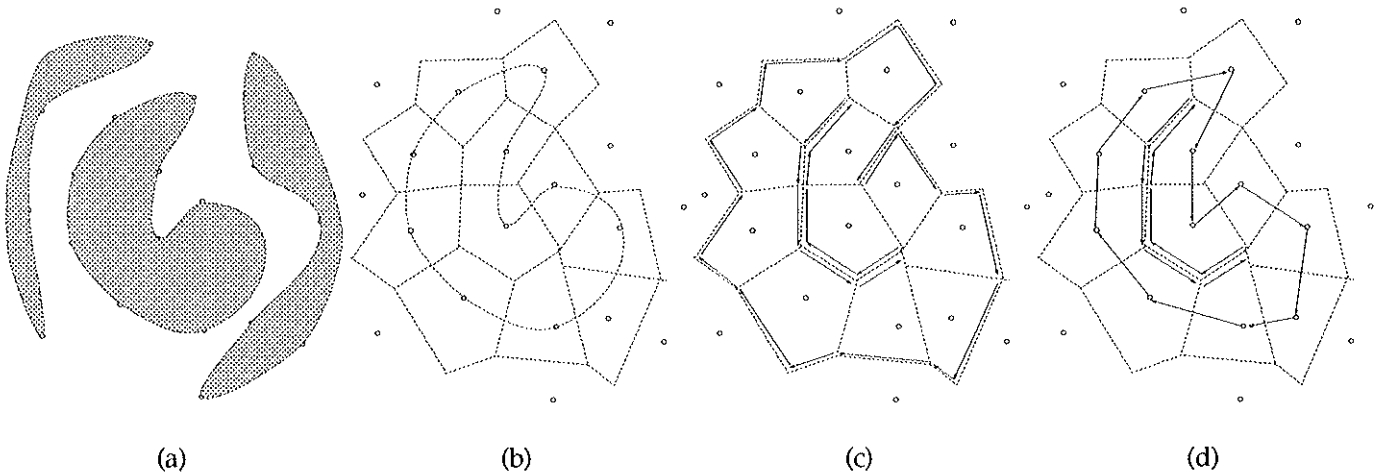

Figure 5: (a) A polygonal region which is not convex. (b) the corresponding Voronoi polygons. (c) the set $\mathbf{E}_i$ of edges that delimit the protoannulus, and (d) the set $\hat{\mathbf{E}}_i$, that delimit the annulus. Note that here part of the external medial contour is degenerate.

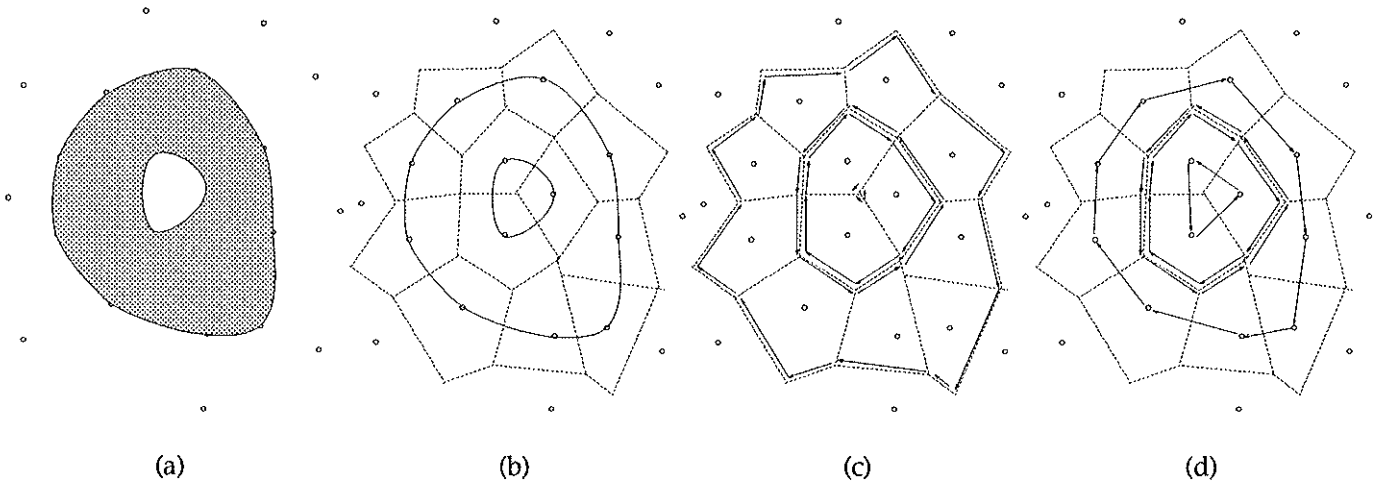

Figure 6: (a) A polygonal region with a hole in it. This region has two contours; each generates a protoannulus. (b) the corresponding Voronoi polygons. (c) the set $\mathbf{E}_i$ of edges that delimit the two protoannuli. (d) the set $\hat{\mathbf{E}}_i$, that delimit the annuli. Note that points in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

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 medial and dual-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 dual-medial contour is slit-like here as well. Figure 6 depicts a column with an embedded hole. This column has two contours, each of which gives rise to a medial and dual-medial contour. The outer contour has a medial contour that corresponds to the inner contour's dual-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

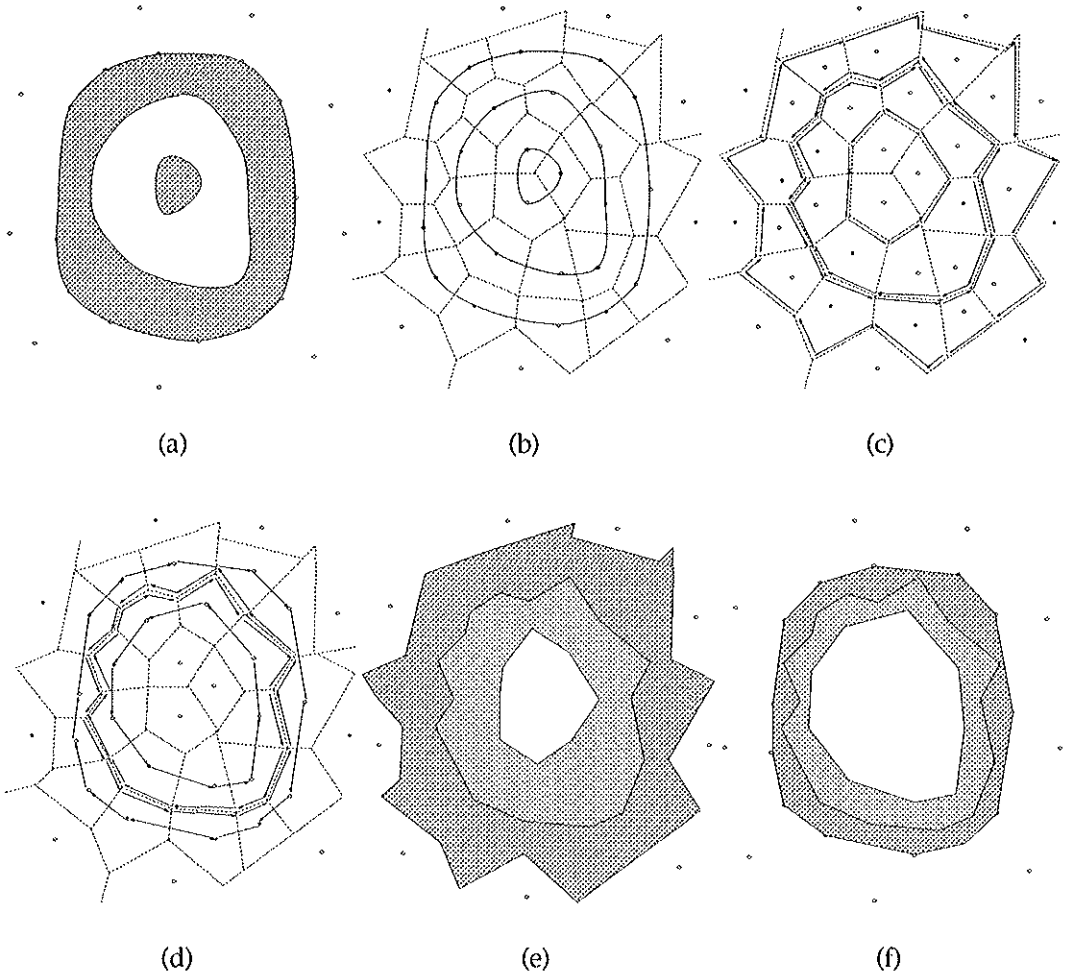

Figure 7: (a) A polygonal region with another polygonal region embedded in it. The region is decomposed by the procedure into two concentric annuli. (b) the corresponding Voronoi polygons. (c) the set $\mathbf{E}_i$ of edges form four contours that delimit the two protoannuli. (d) the set $\hat{\mathbf{E}}_i$ 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.

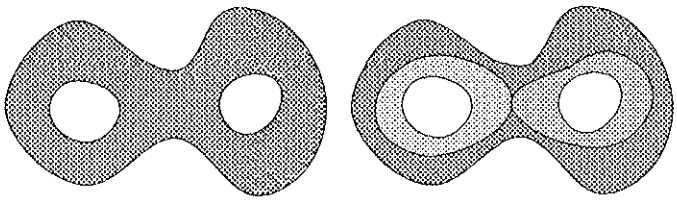

Figure 8: A planar region with holes can be decomposed into a series of annular regions, each with a single hole in it. For simplicity, regions with no holes are each assumed to have have an infinitesimal hole, allowing a region to be expressed as a sum of annular regions.

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 C_c. Then there is no contour point $q \in \mathrm{C}_{c'} \neq \mathrm{C}_c$ on some other contour such that q is closer to p than either of p's neighbors in 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 $\mathrm{C}_c \subset \mathcal{S}$ 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 $p \in \mathrm{C}_c$. 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 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: C_c is a closed contour; the Voronoi polygon for each point along C_c is convex, and shares an edge with both of its neighbors

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

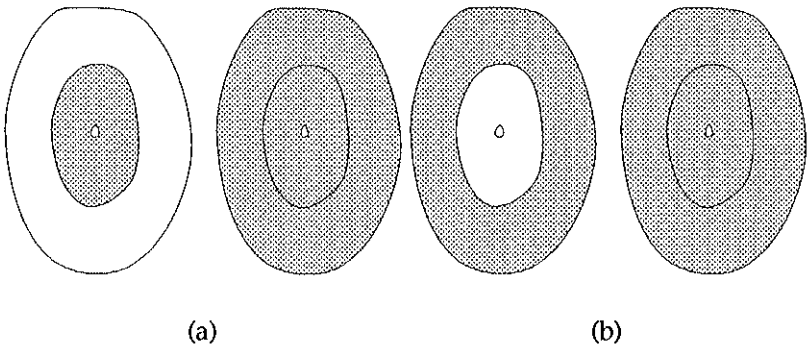

Figure 9: A circular region (*annulus*) that "grows" outward, and one that "grows" inward.

along the contour. This is always true because they are its closest neighbors. When two Voronoi polygons, each 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 where the annulus represents the external portion of the protoannulus, the mapping is termed "outward" (see figure 9).

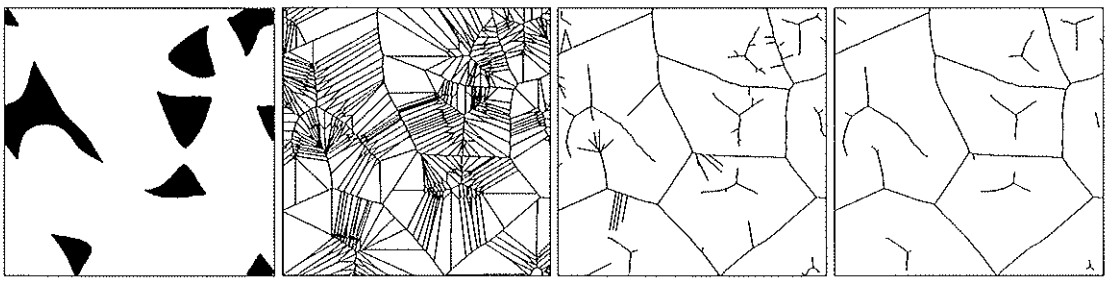

(a) (b) (c) (d)

Figure 10: Sensitivity of medial axes to boundary noise: the effect can be reduced by thresholding edges on boundary distance between the nodes the edge separates. (a) Subset of original square area. (b) Voronoi diagram of sampled contour points; (c) Medial contours, t1threshold at $\hat{d} = 0$, (d) $\hat{d} = 3$.

2.7 Sensitivity of medial contours to polygon contour shape

One 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 arm of the skeleton. This problem is often introduced by sampling a diagonal contour that occurs in an image represented as an array of square pixels. 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 and v_j, with both i and j being points on the original polygon contour. Let $d(i, j)$ represent the arc length along the contour that connects points i and j which this edge separates. A simplistic approximation to $d(i, j)$ is $\hat{d}(i, j)$, the number of points in the contour polygon separating the two points. If we assign the value $\hat{d}(i, j)$ to each edge on the medial axis, the following picture emerges. The "backbone" of the medial axis has high values of $\hat{d}$, while the small protuberances due to local fluctuations have very low values. Further, the value of $\hat{d}$ increases monotonically as one travels inward along any arm toward the core of the skeleton. If a threshold is applied to this value, and edges whose values fall below the threshold are eliminated, the remaining skeleton is less prone to local edge noise, and, due to the monotonicity of $\hat{d}$, guaranteed to still be connected. Figure 10 illustrates this effect. Naturally, all of the above applies equally to the external medial contour. One small point deserves attention, though: medial axes that initially have small numbers of points may be eliminated entirely by a strict threshold procedure, and care must therefore be taken to avoid this eventuality.

In [16, 24], the authors describe a method for computing the skeleton, or internal medial axis of a figure. They use a potential function that represents Euclidean distance from the contour, and describe an iterative hill-climbing procedure that arrives at the skeleton. While the set of points along the ridges of their potential surface are identical to those points which lie on the medial axis, the "potential" along the ridge, is different from the

one defined above. The use of a threshold on this potential function, while producing the intended effect of pruning undesirable noise in the skeleton, could easily disconnect the skeleton at a narrow point of the figure, an undesirable side-effect.

An alternative method for pruning the skeleton is used in [13], where the incident angles of bisectors of contour points are thresholded. Edges which meet other edges at too small an angle are eliminated. This approach has a serious shortcoming, as seen in the article's figures: the resulting skeletons are often not connected.

In a recent study, a significantly different approach to the computation of internal and external medial axes is proposed [20]. The authors propose a method for recursively dividing a figure into quad-trees (with appropriate extensions for higher dimensions), which efficiently computes the generalized Voronoi diagram of a figure to arbitrary resolutions. Their technique relies, however, on analytic definitions of the bounding curves of each region, which is not appropriate to our application. Appropriately adapted, this approach might provide for a more efficient computation of the medial contours than the one proposed here, though it does not provide local information needed for the computation of image warping grid points.

3 Path Homotopies provide coordinate systems for image warping

While the introduction of the topological concepts of path homotopy may appear an unnecessary diversion, it allows for a clearer discussion of the issues involved in warping between two regions, and has been used implicitly in existing applications (e.g. [39]).

The mapping $\phi : [a, b] \rightarrow P$ from the interval $[a, b] \subset \mathbb{R}$ into the plane defines a parameterized *curve*, that can be thought of as the position $\phi(t)$ of a particle as a function of time $t \in [a, b]$. A *closed curve* is a curve whose endpoints are identical: $\phi(a) = \phi(b)$.

Let ϕ_0 and ϕ_1 be two curves in a plane P defined on the same closed interval $[a, b]$. Let R be the rectangle in the $\mathbb{R}^2$ plane bounding an area defined by two variables: (t, r), with $a \leq t \leq b, 0 \leq r \leq 1$. Then a *homotopy* of ϕ_0 into ϕ_1 is a continuous mapping $: R \rightarrow P$ such that $\forall t \in [a, b]$, $(t, 0) = \phi_0(t)$ and $(t, 1) = \phi_1(t)$. If such a homotopy exists, then one can say that ϕ_0 is *path homotopic* to ϕ_1[6, 25]. In the case considered here, ϕ_0 and ϕ_1 are closed curves, and the homotopy must in addition satisfy the condition

$$\forall r \in [0, 1], (a, r) = (b, r)$$

The mapping

$$(t, r) = (1 - r)\phi_0(t) + r\phi_1(t)$$

is known as a *straight-line homotopy*. This function moves the a point $\phi_0(t)$ to a point $\phi_1(t)$ along the straight-line segment joining them (see figure 12a). The set of functions corresponding to values of r strictly between 0 and 1 define a "region" of the plane. Each point in this region is uniquely identified by two coordinates: t and r. When the two curves are closed, as with the inner and outer medial contours of a region, t can be thought of as a "circumferential" index, and r a "radial" index[7].

[7]Though one can construct homotopies for which the radial index is poorly defined, the present work

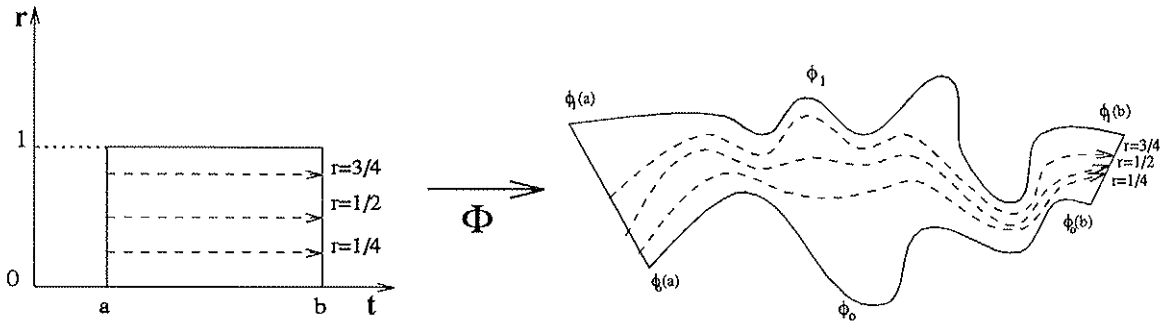

Figure 11: Homotopy of two paths $\phi_\circ$ and ϕ_1 defined on the interval $[a, b]$. There exits a continuous family of curves $(t, r), r \in [0, 1]$ such that $(x, 0) = \phi_\circ(x)$ and $(x, 1) = \phi_1(x)$. The ensemble of these functions uniquely defines all points in the region bounded by the two paths.

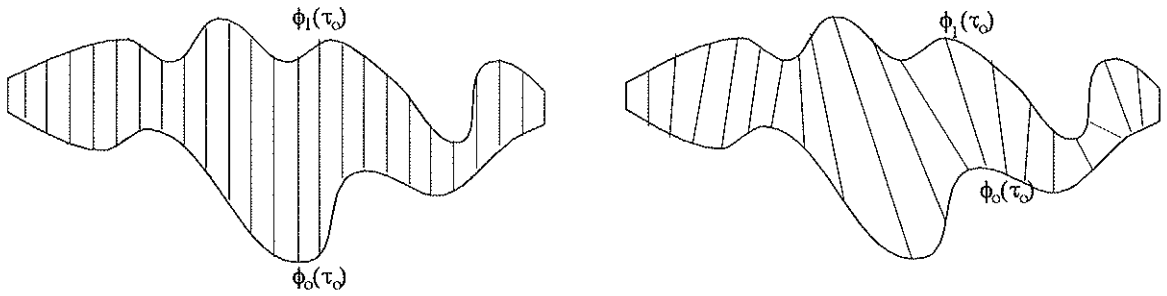

Figure 12: Two straight line homotopies. For each point $\phi_\circ(t)$ there is a corresponding point $\phi_1(t)$, and representations of t on the intermediate curves all occur on the straight line segment between these points. The straight-line homotopy assumes a parameterization of the path. The figures represent two different parameterizations of each path.

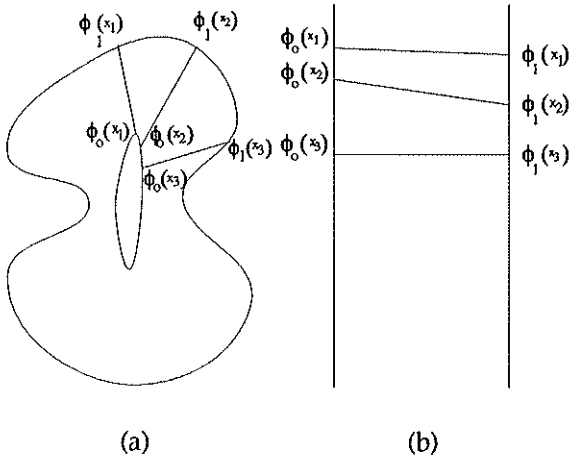

Figure 13: Ladder representation of two homotopic paths. (a) represents an annular region; (b) shows the same region represented as a "ladder", where the tops and bottoms of both ladder risers are identified, and each rung corresponds to a pair of points related by the homotopy. The ladder divides the region into a series of radial slices.

3.1 Ladder transformations

To warp the interior of an annular region to the interior of another, a sampling of each one must be defined, along with a "resampling function" that enables us to go from one to the other [40]. In the literature, *resampling* is discussed extensively [8, 11, 12, 40], both on the level of grid point correspondence, and on the level of image interpolation between corresponding grid points. These strategies can be compared by considering each as most closely approximating a particular homotopy.

The existence of a path homotopy between the inner and outer medial contours of an annular region imposes a circumferential sampling grid on the region. This is easy to show: as there exists a monotonic, nondecreasing map between points along one contour and points along the other contour, these can be associated pairwise to form a "ladder", that divides the region into "radial" slices (see figure 13). Once this circumferential sampling has been introduced, the space between each pair of rungs on the ladder is a quadrilateral, and the quadrilaterals in the two samplings can be mapped to one another in a pairwise fashion using conventional piecewise bilinear warping. Naturally, as this procedure produces only polygonal approximations to continuous contour shapes, a dense ladder is preferable to a sparse one[8].

avoids this problem: the inner and outer medial contours of a polygon can be smoothly deformed into one another using a straight-line homotopy that never crosses over itself. This is a natural result of the Voronoi construction of the medial contours, as will be shown later.

[8]In the general case of homotopy between annular regions, the quadrilaterals may have curved rungs, corresponding to concavities in one contour not matched by convexities in the other. A linear approximation to these rungs is acceptable in most cases, as the exceptions are easily detectable.

3.2 Defining a path homotopy from a Voronoi diagram

While it is conceptually simple to define a homotopy between the medial contours, it is not immediately evident how to use the Voronoi structure to define a parameterization of the medial contours and describe a homotopy. A method is described here that transforms the Voronoi structure to a ladder which implicitly defines the homotopy.

Consider the annular subset of a Voronoi diagram delimited by a single pair of medial contours. This graph contains single edges that connect the two contours; these are bisectors of adjacent points on the column contour, and are edges in the ladder. Many of the triplets of adjacent column contour points, however, give rise to a three-edged structure in the Voronoi diagram (see figure 14). This is due to the fact that their bisectors intersect before arriving at the medial contour. To construct a rung of the ladder corresponding to this bisector, one can "extend" it, replacing it by the vector sum of the the original segment and the extra edge (see figure 14).

To create the ladder, two "risers", each of which has a number of nodes on it corresponding to the number of vertices in one of the medial contours, are created. Each such node on a riser is termed a *support*. As each rung is added, it attaches to one support on each riser (this is a simple consequence of the fact that both are subgraphs of the Voronoi diagram). After adding in all the rungs, there remain two problems: unpopulated supports, corresponding to portions of the medial contour not reached by a rung, and overpopulated supports, which have more than one rung incident upon them (see figure 15). Each can be dealt with in turn using simple heuristics.

Each un-populated support can be populated by considering only its nearest populated neighbors. This preserves the monotonic nature of the homotopy function that is sought. In addition, as the two rungs that bound an un-populated point are generally the product of a single Voronoi polygon, it is natural to associate this point with the same polygon. Thus in figure 15, the edges $\overline{Cd}$ and $\overline{Ba}$ or $\overline{Bb}$ can be added. In this implementation, one of either a or b would be chosen as an attachment point, and the support would be split in the next phase. Alternatively, a node equidistant between a and b could be created at this stage.

An overpopulated support can be repeatedly "split" into several supports; the topographic ordering of the remote supports allows the new supports to be ordered correctly. The definition of a homotopy requires that there not exist regions of one function that are mapped to a single point on the other function (this would make the mapping non-invertible). The splitting operation ensures that this is always the case. To split a support, a new node is added to the riser that lies along the path between the current support and its predecessor, but very close to the current support. It suffices for the purposes of the homotopy that these supports be separated; however, the more separation is introduced, the further the ladder rungs depart from their original, Voronoi-generated positions. Thus the solution adopted here is to introduce only an infinitesimal separation between a support and the new supports that split off from it.

Procedure for ladder construction (see figure 16):

- Construct the Voronoi diagram due to the contour points, as well as the inner and outer medial contours.

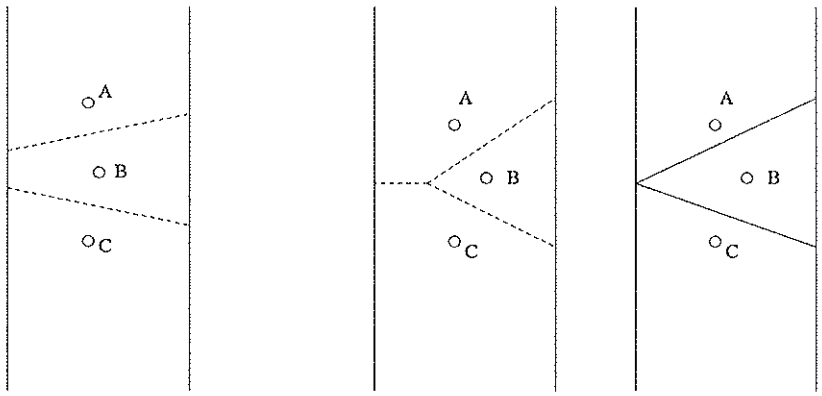

Figure 14: Voronoi edges due to bisectors of adjacent column contour points (dashed lines). In the first case, they do not intersect. In the second, new edges (dotted lines) can be created, representing the existence of a path from one riser of the ladder to the other.

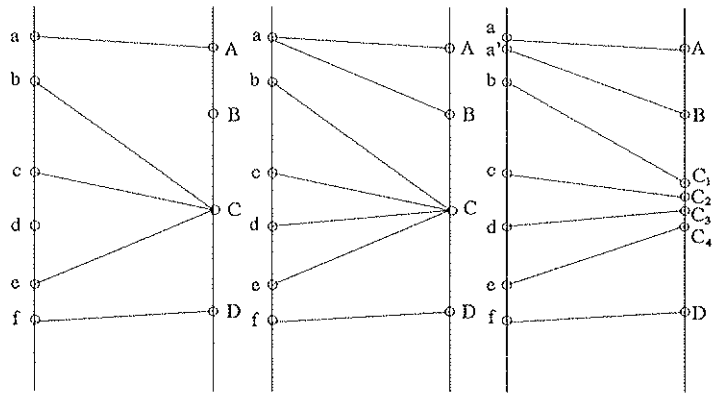

Figure 15: In the ladder portrayed above,supports B and d are un-populated, while support C is overpopulated. In a first pass (b), the un-populated supports are assigned rungs based on their neighbors. In a second pass (c), the overpopulated nodes are split.

- For each generalized Voronoi polygon,
  - Create an empty ladder, with the risers representing the inner and outer medial contours.
  - For each edge in the Voronoi diagram that connects a point along the inner medial contour to a point along the outer medial contour, add a rung to the ladder.
  - For each path of length two or three edges in the Voronoi diagram that connects a point along the inner medial contour to a point along the outer medial contour, add a rung to the ladder.
  - For each point along each riser that has no rung, add one based on the neighboring rungs.
  - For each point along each riser that has more than one incident rung, split the point (add a point to the riser).

3.3 Using a ladder to specify a warping function

One additional advantage of the ladder construction is that it defines a series of quadrilaterals, each of which is intersected by the original region contour (see figure 17). As both the medial axes were constructed using a Voronoi procedure, there cannot exist a case where the contour either fails to intersect a rung or crosses a rung more than once. This is easy to see: if the original curve had doubled back on itself, a locus of points would have existed that was equidistant from the two branches of the "U", and this locus would have been incorporated into either the internal or external medial contour by the Voronoi procedure; As a ladder rung does not cross a medial contour, it would therefore be unable to cross the original contour more than once.

Using this intersection property, it is easy to construct a new riser, consisting of, in turn, each of the intersections of the rungs with the original contour. This riser, along with the riser corresponding to the internal medial contour, forms a new ladder with exactly the same number of rungs and inter-rung quadrilaterals as the original. Corresponding to a straight-line region homotopy, the quadrilateral regions can be mapped to each other pairwise using a bilinear warp.

4 Warping subsetted images

If a polygonal image area is segmented in a binary fashion into two classes, each class occupying subsets of the image that are not necessarily connected, and the original polygonal image must be warped into the subset formed by one class (see figure 1), this is termed a *subset warp*. To achieve a subset warp, the following approach is followed:

- The boundaries of the subsets are found, and used to create a generalized Voronoi diagram. In this diagram, each generalized Voronoi polygon represents the set of

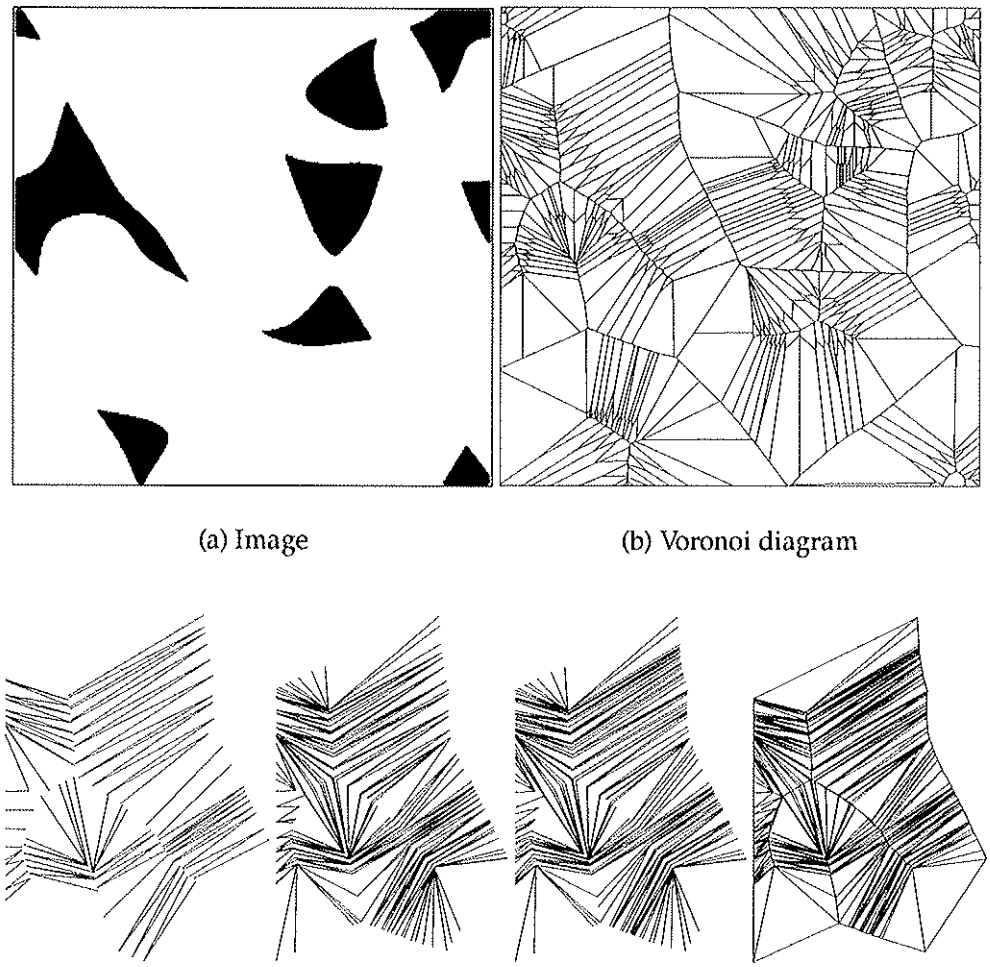

(a) Image (b) Voronoi diagram

(c) Ladder (d) Filled ladder (e) Filled,split ladder (f) Quadrilaterals

Figure 16: Construction of a Ladder from the Voronoi diagram. (b) depicts the Voronoi diagram of (a). In (c), the framework of a ladder for one region is constructed; (d) represents filling in missing rungs; (e) the effect of splitting overpopulated ladder supports. In (f), quadrilaterals obtained from (e) by removing extraneous edges are shown. Compare (f) with (b).

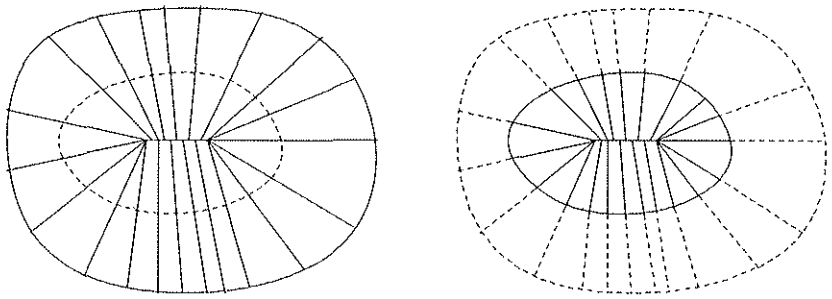

Figure 17: The intersection of a column contour and the ladder formed between its medial contours can be used to form a new ladder.

points in the original image that is warped into a particular subset. As a result of the Voronoi procedure, each polygonal area in the range is associated with the subset of points in the domain that are closest to it.

- Each generalized Voronoi polygon in the domain, as well as each polygon in the range, can be divided into protoannuli and annuli, respectively. By construction, there exists a pairing of these two sets of annuli.
- A region homotopy can be defined using ladders that maps each point in a protoannulus to a corresponding point in an annulus.

The warping defined is one to one and invertible. This is a direct consequence of the fact that the generalized Voronoi polygons tessellate the domain and of the invertibility of the homotopies described by the ladders.

5 General Warping between two objects

Warping between two arbitrary closed polygonal regions of the plane, neither containing holes, is a problem of region homotopy, as considered above. It is instructive to consider two very different approaches to this homotopy problem, in the restricted case where one polygonal region can be "fit" into the other, perhaps by scaling. The first uses distance along the contour to define its homotopy, the second is based on the subset warp described above.

5.1 Circumferential sampling

In this approach, each point in each region is considered to lie somewhere in a polar coordinate space, and described by its annular and radial indices. The radial coordinate corresponds to the shortest distance between the point and the skeleton of the region, while the annular coordinate is arrived at by measuring the perimeter of the isodistance surface and quantizing it[9].

[9]In cases where the objects are mostly convex polygons, a simplistic version of this approach might be to divide the figure into equi-angle triangles, each ending at the center of mass.

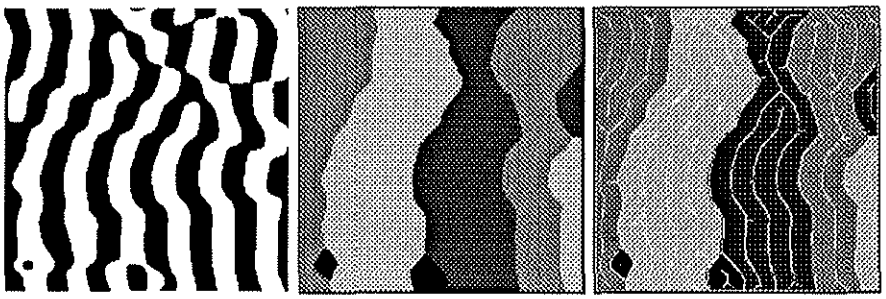

(a) (b) (c)

Figure 18: (a) A subset, from an ocular dominance column simulation. (b) The corresponding generalized Voronoi regions, each shaded differently. (c) portrays the inner and outer medial contours $\mathbf{E}_i$. In this case each column consists of a single annulus, as there are no embedded holes.

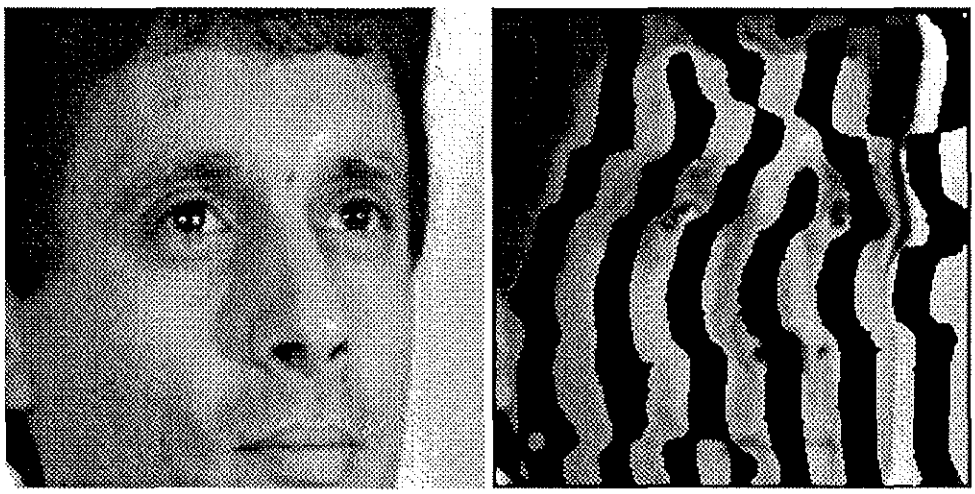

Figure 19: Texture mapping using ladders defined on the annuli in the previous figure.

In [39] a method is described in which a region, delimited by a hand-drawn curve, is decomposed into a set of "onion peels," each of which is obtained from a single iteration of a morphological erosion process. This method is clearly a quantized approximation to the one described above, although the nature of the quantization process gives rise to irregularities in the path homotopy that describes the region. To map one region onto another, the (x, y) coordinates imposed on the two regions by the path homotopies are mapped in a one to one fashion onto each other. In the discrete case, this reduces the problem to one of re-sampling one grid onto another, which can be solved by various interpolation strategies [40].

Also worthy of mention is the approach of [10], that defines a triangulation of an annular region delimited by two polygons of different order. The construction minimizes the total length of internal edges of the triangulation. Their implementation is quite fast, but defines a mapping between contours that is not strictly a homotopy unless triangle ends are split. In most cases, the resulting warping function is similar to that of circumferential distance sampling.

Circumferential distance sampling may not always yield desirable results. Consider the case of a circular domain, and a corresponding circular range into which a slit has been cut that does not extend all the way through (see figure 20a,d). The addition of the slit significantly increases the length of the second region's bounding contour, and therefore significantly distorts not only regions close to the slit, but the rest of the image as well. In addition, there is a spurious dependence on the choice of an "origin". It would appear that any other quantized approximation to the circumferential distance region homotopy would fare just as badly, reflecting an excessive sensitivity to boundary length.

5.2 Embedding one object within another using scaling

One could imagine, however, an alternative approach to this warping problem, that would attempt to preserve distance relationships. If the domain image is scaled up sufficiently to enclose the range image entirely, then the problem can be reduced to a subset warp, the solution to which was described previously. In figure 20c, a circular region has been made to enclose the circle-with-slit region. The interior polygon generates an external medial contour that consists not only of the outer circle itself, but also of an infinitesimal slit running up the center of the range image slit. This allows points that were adjacent in the domain image before the slit was made, to remain as close to one another as possible in the range image.

To see why this would be the case, consider a point p in the domain image on the slit. This point p is on the boundary between the Voronoi polygon of a contour point q and q' on opposite sides of the slit, as it is equidistant from both. In this construction, it is mapped to q and q', which are the closest possible boundary points by the definition of the Voronoi diagram. The same cannot be said of the circumferential sampling warps, where distance of a point from its image is not taken into account.

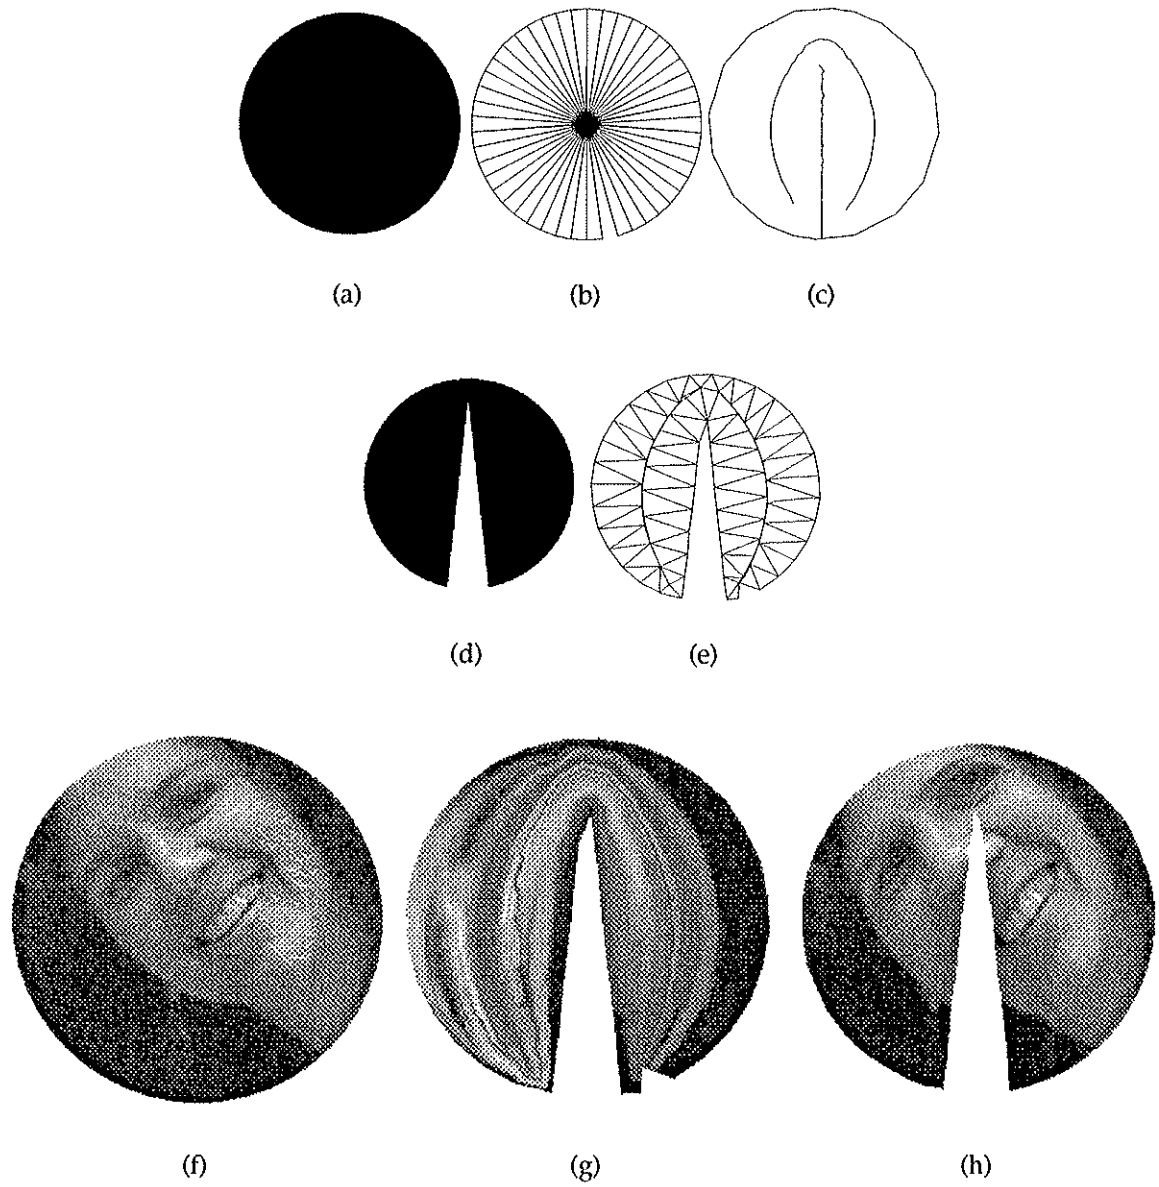

Figure 20: Warping a circular region (a) to a sliced region (d). Contour-length parameterization (e.g. [39]) is depicted in (b) and (e), and leads to the image in (g). Generalized-voronoi warping, as described in this paper, defines cut edges (c), and leads instead to (h).

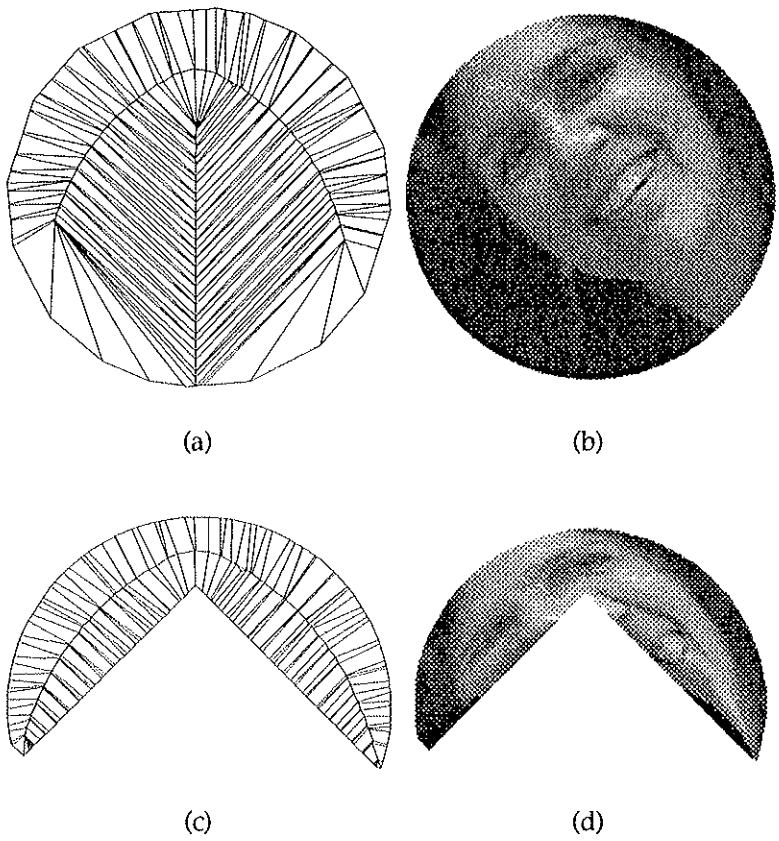

Figure 21: A more severe warp than the one in figure 20. (a) ladder representation of the domain, (c) ladder representation of the range.

6 Polymaps

6.1 Ocular Dominance Columns

This algorithm grew out of an effort to model in image terms the pattern of neural activity evoked at the level of the visual cortex by a particular set of binocular input images.

The cerebral cortex is composed of multiple layers, each organized in a highly regular fashion. A number of areas are organized as "maps" or receptotopic representations of one or another sensory modality. In each of these areas, activity in the region of a particular cortical "coordinate" is evoked by stimulation in the region of a corresponding receptor coordinate. The visual maps, occurring both in the occipital cortex (V-1, V-2) and elsewhere in the cortex, have been studied extensively in this regard[29, 31] as a cortical "image" or pattern of cellular activity that can be evoked by a visual input and measured by experimental techniques. Mathematical functions representing the distortion of retinal space into the cortical space of each of these maps have been proposed, though considerable debate continues on their exact form[28, 30, 36].

There is, however, a more fundamental structure present in each of these map areas, namely, the "column," where a population of cells with similar properties clusters together. Examples include the ocular dominance column system, consisting of a layer of cells with either a left or right ocularity preference[22] and the orientation column system, containing cells responding to contrast gradients of different orientations[14]. Columns are so pervasive that they can be said to constitute a fundamental data structure of cortical architecture. In each of these cases, the true "image" at the cortical level represents a joint occurrence of topographic mappings of multiple sub-modalities, interlaced in the form of macroscopic patches ("columns"). The term "polymap" has been introduced [19] to describe these structures, as they represent the mapping of information from several different inputs. The simulation of polymaps is of fundamental importance to the modeling and eventual understanding of cortical functional architectures.

The ocular dominance column system, located in layer IVc of primate visual area V-1, constitutes the "data input" layer, where afferent fibers from the lateral geniculate nuclei terminate. Cells in this layer appear to be members of two distinct classes, responding either to left or right eye input. The neurons of these two classes are clustered together in elongated strips one to two millimeters in width and extending over several centimeters. If one considers the cortical image of a monocular input on a global scale, it forms a topographic map of the retinal input. As one examines the image at the scale of the columns, it can be seen to alternate between areas of activity (the columns corresponding to that eye) and areas of inactivity.

To model the transformation that an image undergoes from retinal input to cortical representation, a two-step procedure is used. In a first phase, a global transformation is applied to the data to create a "protoimage", of the same size and shape as the full cortical area containing the columns; in a second phase, this protoimage is subset-warped into the regions of the cortex occupied by the columns themselves [19, 18]. The resulting two-step mapping is (theoretically) one to one and invertible. However, since the implementation discretizes both the domain and range spaces, the invertibility property no longer strictly

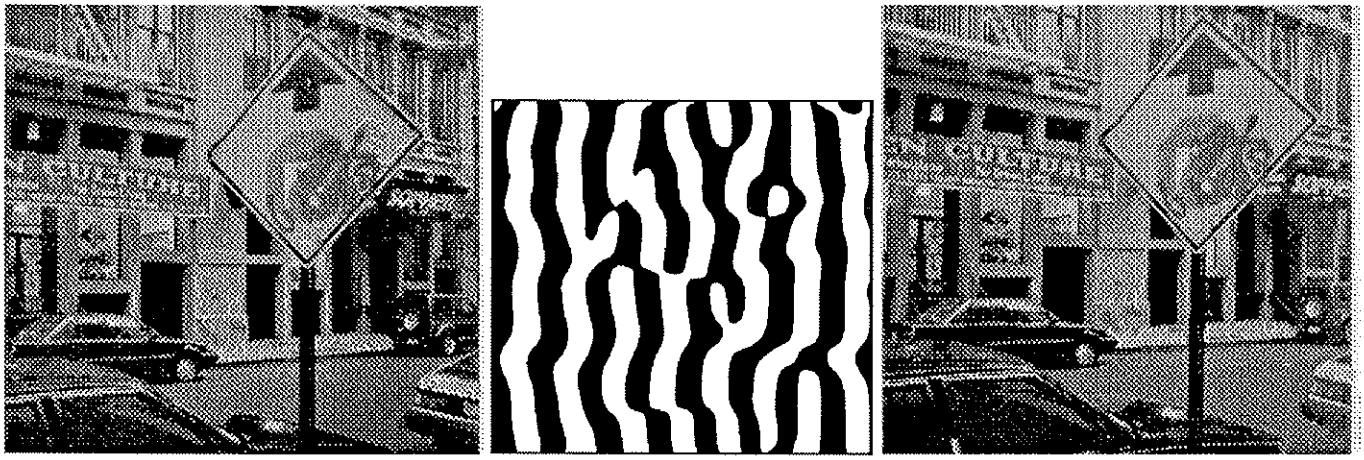

(a) (b) (c)

Figure 22: Ocular dominance columns: (a) and (c) represent corresponding patches of left and right eye inputs; in (b), the white stripes represent cortical areas which receive left-eye inputs, and conversely.

holds.

6.2 Orientation Columns

Orientation columns, first described by [14, 15], are groupings of cells in layer III (and also V) of the visual cortex that selectively respond to edges of a specific orientation in the retinal image. Though considerable debate continues as to the shape and organization of orientation columns [3, 2, 5, 32], there seems to be general agreement on the fact that there exists a layer of the visual cortex that is tiled by groupings of orientation tuned cells, termed *orientation columns* (see figure 25). The present work addresses the modeling of this orientation polymap, making the following assumptions which we believe to be consistent with the current level of understanding of this aspect of cortical architecture:

1. There exist n separate populations of orientation-tuned cells, each preferring a particular angle.
2. The entire visual field is visible to each of these populations of cells
3. There are no portions of the visual field for which a given orientation is multiply represented on the cortex in separate columns.

The last two assumptions are very important, as they define the orientation layer as a polymap of one cortical layer (the input from the retina) onto another cortical layer (the orientation column layer).

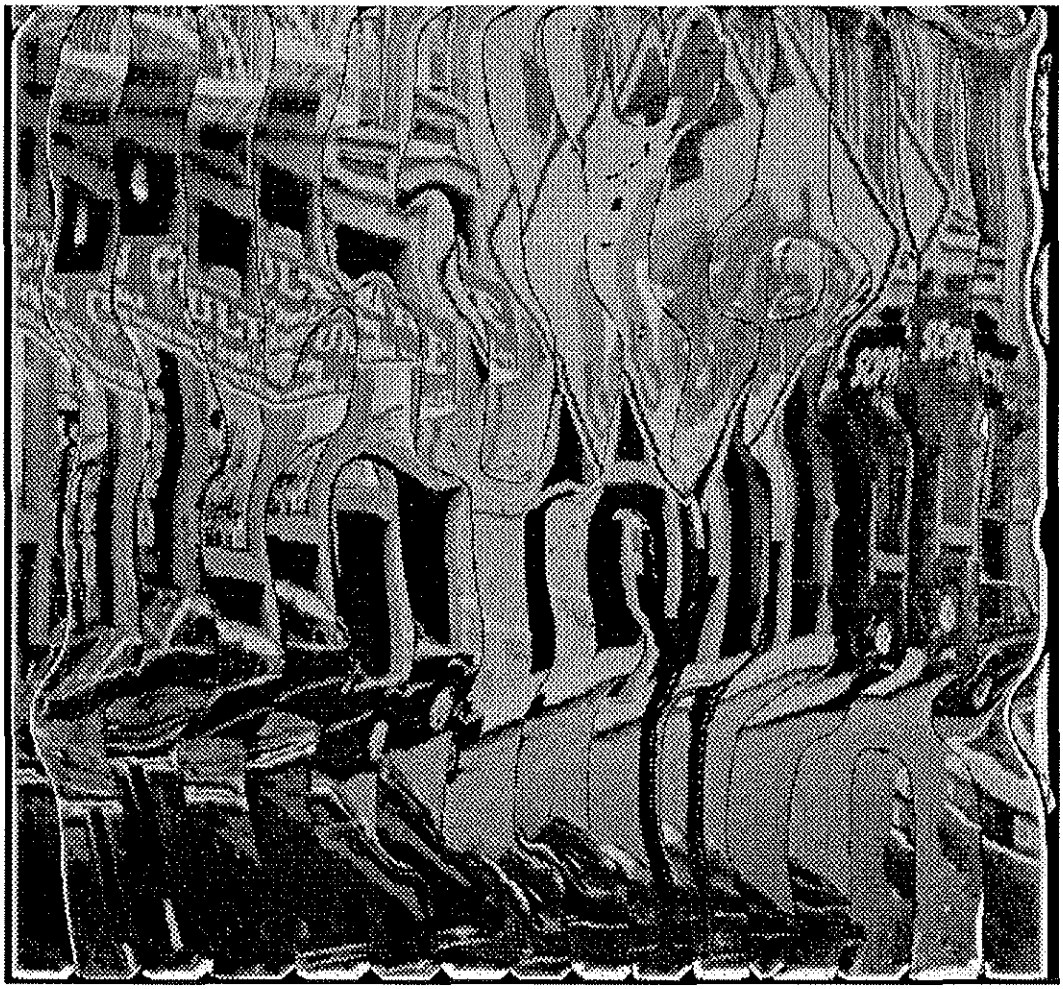

Figure 23: Ocular dominance layer image, mapped from stereo images. Note the different disparity values evident in the road sign (foreground) and the awning (background).

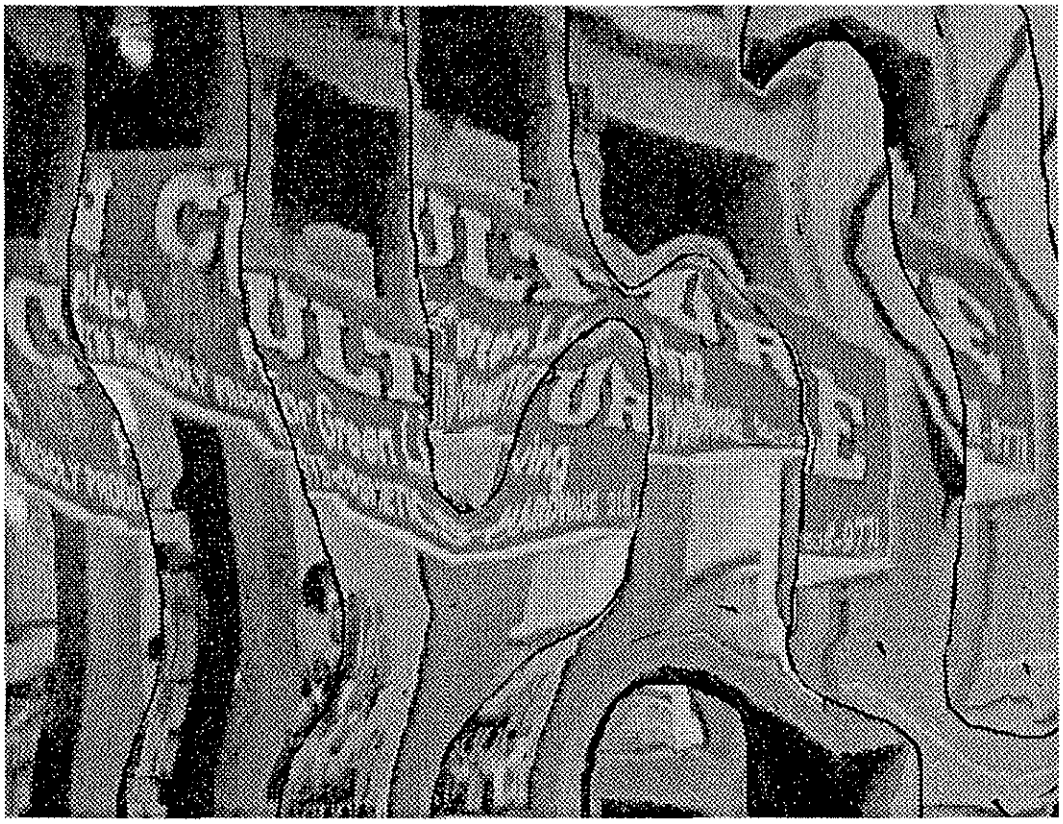

Figure 24: Detail from previous figure. Note the duplication of lettering. All detail from both complete input maps (from the left and right eyes) is presented in the final warped image.

Using the assumption that orientation columns constitute a polymap allows the use of the mechanisms described previously. Consider for example the population of cells preferring 45° degree angles. This population of cells is organized into columns comprising a subset of the cortical layer, which can be modeled by a subset warping problem. Thus one can take the input image, apply a non-isotropic spatial filter or edge detector oriented at 45° to the original image, and subset warp the result onto precisely those portions of the cortical image that contain cells preferring this orientation. An appropriately oriented edge in the input image is thus detected wherever it appears, and no edge is multiply represented. The above procedure can be repeated for each population of cells, until the "cortical image" of the ensemble of these cell emerges.

In figure 25, a planar region tessellated into regions of six classes is depicted. To obtain the initial layout of columns, an initially random grid of orientations was blurred with a Gaussian kernel, and the results binned according to preferred orientation. As these "columns" were obtained using only a lowpass filter, rather than the bandpass filter suggested by [32] and [35], their widths and shapes were relatively inhomogeneous; this example was chosen for its ability to illustrate warping techniques rather than its similarity to orientation column systems[10]. In figure 26, three of the classes of "cells" are illustrated, along with the generalized Voronoi diagrams they generate. Each of the binary images in figure 25 was used individually to generate not only generalized Voronoi diagrams, but also ladders for each column and its corresponding protocolumn. To illustrate the concept of the polymap to best effect, the same image was warped into each of the subsets, and the results of all images OR'ed together, resulting in figure 27. This figure is "globally receptotopic" while having a much more complex local structure. It is important to emphasize that this example illustrates the pattern of activity that would be evoked in a very small section of the cortex, and that the same image, viewed over the entire visual field would cover several thousand complete 180 degree sets of orientation columns.

7 Conclusion

In this paper, we have described an algorithm for warping an image into a subset of its domain, in such a way that "proximity" of source pixels to destination pixels is optimized in a sense that we make quantitative via the concept of Generalized Voronoi region. Several new geometric structures were developed in this work (outer medial contour, ladder diagram), while other well known constructions were applied for the first time in the image warping context (Generalized Voronoi region, inner medial contour).

The immediate motivation for developing this algorithm was provided by a basic problem in the modeling of cortical regions of the brain: several different sub-systems of primary visual cortex provide an example of what we term here a subset warp. We have applied this algorithm to two of them, the ocular dominance column system and the orientation column system, as a demonstration; there are a number of others which we

[10]In this figure, only the spatial arrangement of orientation columns is shown. A high pass filter, at the orientation of each "column" has not been applied, since we are only attempting to illustrate the properties of the image warp which set up the orientation polymap, and not the spatial filtering properties of the polymap

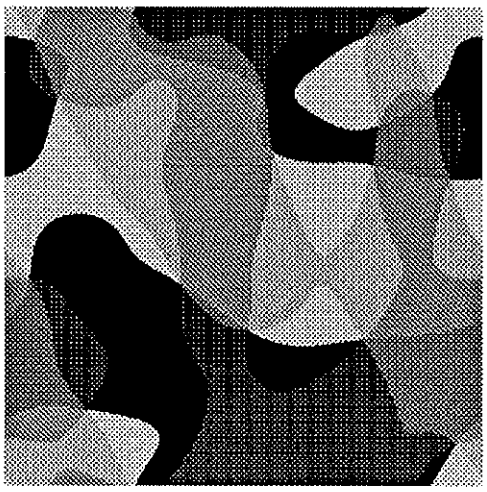

(a)

(b) (c) (d)

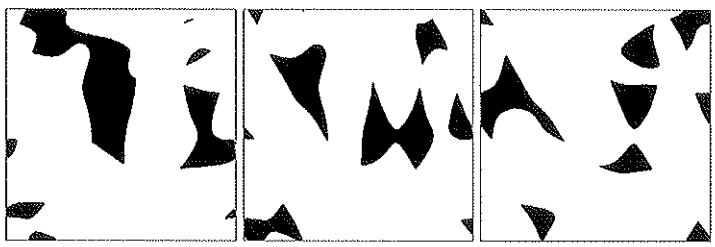

(e) (f) (g)

Figure 25: Orientation columns: (a) depicts the tessellation of the cortical plane by a series of cell populations; (b) through (g) depict each of these populations in turn

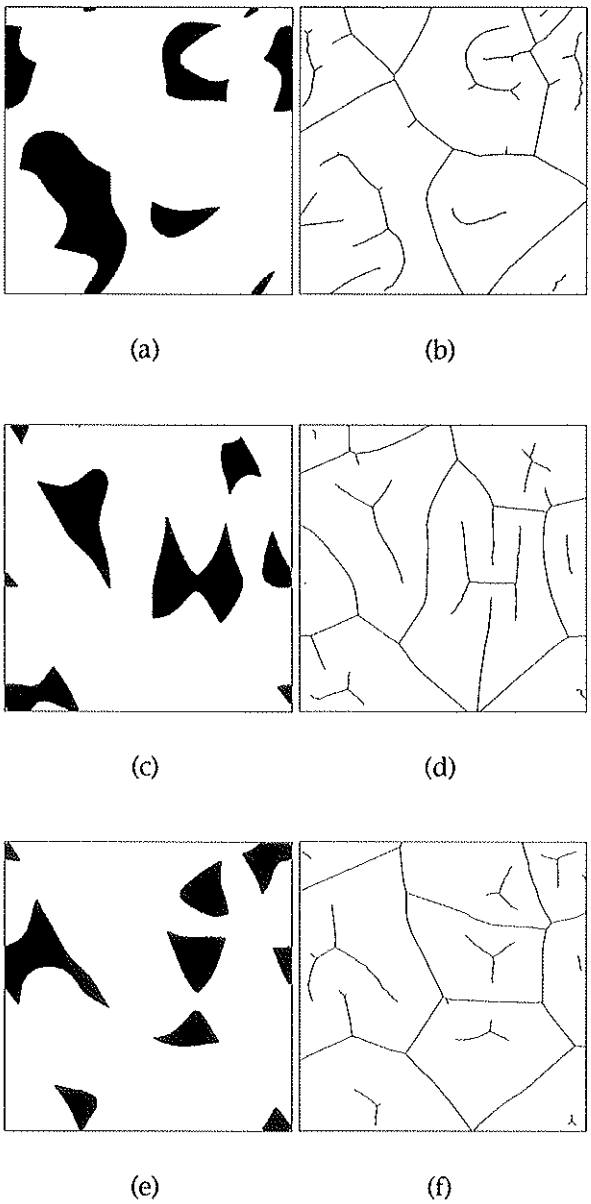

Figure 26: Medial contours: three representative orientation-tuned cell populations and their medial contours.

Figure 27: The same input (see figure 19) is warped into each subclass of cells, illustrating the mosaic-like structure of the polymap.

have not yet addressed. In this area of application, the importance of the ideas presented here is most clear. Not only for the purpose of visualization, but for a specification and description of the systems that are being modeled, the geometric structures of sub-set warping are critical to understanding the nature of systems of this type.

Usually, in applications of image warping, the warp techniques that are applied lie at one extreme, or the other of a continuum of precise knowledge of the warping function. For example, warping is applied in photometric applications, with a known polynomial warp function, to effect extremely precise corrections to image data which has been distorted in a known and well understood fashion. At the other extreme, warping is applied in visualization and entertainment applications in which the precise details of the warp function are not known, and of no great interest. Any image warp which is "visually" acceptable will do, and the relation of the specific structure of the warp function to the detailed scientific nature of the problem is neither known nor significant. The present application fills a gap for applications in which an analytic warp function is unavailable, but for which an ad hoc warp function is inappropriate. Using a small set of data structures from computational geometry, and the topological notion of "path homotopy", we have defined and implemented an image warp which is extremely general in its ability to handle range and domain shape and topology, yet is constrained in terms of proximity, and provides a warp function which is both visually acceptable, and mathematically specified. If the notion of "rubber sheeting" provides the simplest example of an affine warp function, then the method outlined in this paper might be termed "cut rubber sheeting": a local, piece-wise affine map is guided by the Voronoi Polygon construction to provide an analogue of global rubber sheeting to the problem of warping an image into multiple subsets of its support.

References

[1] L. Ahlfors. *Complex Analysis*. McGraw-Hill, New York, 1966.

[2] G. Blasdel. Orientation preference, continuity and density in monkey striate cortex. *Abstracts, Society for Neuroscience*, 17:1088, 1991.

[3] G. Blasdel and G. Salama. Voltage sensitive dyes reveal a modular organization in monkey striate cortex. *Nature*, 321:579–585, 1986.

[4] H. Blum. A transformation for extracting new descriptors of shape. In W. Wathen-Dunn, editor, *Models for the Perception of Speech and Visual Form*. MIT Press, Cambridge, Mass, 1967.

[5] Tobias Bonhoeffer and Amiram Grinvald. Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns. *Nature*, 353:429–431, 1991.

[6] W. G. Chinn and N. E. Steenrod. *First Concepts of Topology*. Math. Association America, Washington, D.C., 1966.

[7] R. O. Duda and P. E. Hart. *Pattern Classification and Scene Analysis*. Wiley, NY, 1973.

[8] R. Franke and G. Nielsen. Smooth interpolation of large sets of scattered data. *Int. J. for Numerical Methods in Engineering*, 15:1691–1704, 1980.

[9] Carl Frederick and Eric L. Schwartz. Conformal image warping. *IEEE Computer Graphics and Applications*, March:54–61, 1990.

[10] H. Fuchs, Z. M. Kedem, and S. P. Uselton. Optimal surface reconstruction from planar contours. *Communications of the ACM*, 20:693–702, 1977.

[11] Ardeshir Goshtasby. Piecewise linear mapping functions for image registration. *Pattern Recognition*, 19(6):459–466, 1986.

[12] Ardeshir Goshtasby. Piecewise cubic mapping functions for image registration. *Pattern Recognition*, 20(5):525–533, 1987.

[13] John D. Hobby. Generating automatically tuned bitmaps from outines. *JACM*, 40(1):48–94, January 1993.

[14] D. H. Hubel and T. N. Wiesel. Sequence regularity and geometry of orientation columns in the monkey striate cortex. *J. Comp. Neurol.*, 158:267–293, 1974.

[15] D. H. Hubel, T. N. Wiesel, and M. P. Stryker. Anatomical demonstration of orientation columns in macaque monkey. *J. Comp. Neurol.*, 177:361–379, 1978.

[16] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Intl. J. Computer Vision*, 1(4):321–332, 1988.

[17] Oliver Dimon Kellogg. *Foundations of Potential Theory*. Dover Publications, New York, 1953.

[18] Pierre Landau. *Computer simulation of cortical polymaps*. PhD thesis, New York University: Courant Institute, 1992.

[19] Pierre Landau and Eric L. Schwartz. Computer simulation of cortical polymaps: a proto-column algorithm. *Neural Networks*, 5:187–206, 1992.

[20] David Lavender, Adrian Bowyer, James Davenport, Andrew Wallis, and John Woodwark. Voronoi diagrams of set-theoretic solid models. *IEEE Computer Graphics and Applications*, 12(5):69–77, 1992.

[21] D.T. Lee. Medial axis transformation of a planar shape. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 4(4):363–369, 1982.

[22] S. LeVay, D. H. Hubel, and T. N. Wiesel. The pattern of ocular dominance dominance columns in macaque visual cortex revealed by a reduced silver stain. *J. Comp. Neurol.*, 159:559–576, 1975.

[23] S. LeVay, M. P. Stryker, and C. J. Shatz. Ocular dominance columns and their development in layer iv of the cat's visual cortex: A quantitative study. *J. Comp. Neurol.*, 179:223–244, 1978.

[24] Frederic Leymarie and Martin D Levine. Simulating the grassfire transform using an active contour model. *IEEE Transactions of Pattern Analysis and Machine Intelligence*, 14(1):56–75, 1992.

[25] James R. Munkres. *Topology: a first course*. Prentice-Hall, 1975.

[26] M.L. Nack. Recitification and registration of digital images and the effect of cloud detection. In *Proc. Machine Processing of Remotely Sensed Data*, pages 12–23, 1977.

[27] F. P. Preparata and M. I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, New York, 1985.

[28] E. L. Schwartz. Spatial mapping in primate sensory projection: analytic structure and relevance to perception. *Biological Cybernetics*, 25:181–194, 1977.

[29] E. L. Schwartz. Computational anatomy and functional architecture of striate cortex: a spatial mapping approach to perceptual coding. *Vision Research*, 20:645–669, 1980.

[30] E. L. Schwartz. On the mathematical structure of the retinotopic mapping of primate striate cortex. *Science*, 227:1066, 1985.

[31] E. L. Schwartz and B. Merker. Computer-aided neuroanatomy: differential geometry of cortical surfaces and an optimal flattening algorithm. *IEEE Computer Graphics and Applications*, 6(2):36–44 (March), 1986.

[32] Eric Schwartz and Alan Rojer. Cortical hypercolumns and the topology of random orientation maps. *Courant Institute of Mathematical Sciences Tech. Report 593*, 1991.

[33] Douglas Smythe. A two-pass mesh warping algorithm for object transformation and image interpolation. Technical Report 1030, Industrial Light and Magic, 1990.

[34] D. Steiner and M. E. Kirby. Geometrical referencing of LANDSAT images by affine transformation and overlaying of map data. *Potogrametria*, 33:41–75, 1977.

[35] Nicholas V. Swindale. A model for the coordinated development of columnar systems in primate striate cortex. *Biological Cybernetics*, 66(3):217–230, 1992.

[36] R. B. Tootell, M. S. Silverman, E. Switkes, and R. deValois. Deoxyglucose, retinotopic mapping and the complex log model in striate cortex. *Science*, 227:1066, 1985.

[37] Peter van Wie and Maurice Stein. A Landsat digital image rectification system. *IEEE Trans. Geoscience Electronics*, GE-15:130–147, 1977.

[38] Lance Williams. Performance-driven facial animation. *Computer Graphics*, 24(4):235–242, 1990.

[39] George Wolberg. Skeleton-based image warping. *Visual Computer*, 5:95–108, 1989.

[40] George Wolberg. *Digital Image Warping*. IEEE Computer Society Press, Los Alamitos, CA, 1990.