A Feature Space for Edgels in Images with Landmarks

FRED L. BOOKSTEIN
Center for Human Growth and Development and Institute of Gerontology, University of Michigan, Ann Arbor, Michigan 48109

WILLIAM D.K. GREEN
Center for Human Growth and Development and Orthopaedic Research Laboratories, University of Michigan, Ann Arbor, Michigan 48109

Abstract. In many current medical applications of image analysis, objects are detected and delimited by boundary curves or surfaces. Yet the most effective multivariate statistics available pertain to labeled points (landmarks) only. In the finite-dimensional feature space that landmarks support, each case of a data set is equivalent to a deformation map deriving it from the average form. This paper introduces a new extension of the finite-dimensional spline-based approach for incorporating edge information. In this implementation edgels are restricted to landmark loci: they are interpreted as pairs of landmarks at infinitesimal separation in a specific direction. The effect of changing edge direction is a singular perturbation of the thin-plate spline for the landmarks alone. An appropriate normalization yields a basis for image deformations corresponding to changes of edge direction without landmark movement; this basis complements the basis of landmark deformations ignoring edge information. We derive explicit formulas for these edge warps, evaluate the quadratic form expressing bending energies of their formal combinations, and show the resulting spectrum of edge features in typical scenes. These expressions will aid all investigations into medical images that entail comparisons of anatomical scene analyses to a normative or typical form.

Key words. edges, image averaging, morphometrics, thin-plate spline, deformable templates

1 Introduction

Thin-plate splines originated in a nineteenth-century problem of continuum mechanics, the bending of a thin metal plate subject to stresses normal to its rest position. The mathematics of this problem, which involves solution of the biharmonic equation under diverse boundary conditions, was adapted for interpolation theory by Duchon [1] and Meinguet [2], and shortly thereafter it was introduced into computer graphics by Terzopoulos and his colleagues [3], [4]. Most recently, these splines have been introduced into morphometrics, the biometrics of shape, where they have revolutionized statistical analysis and visualization for one type of data peculiar to that field.

A *landmark* is a point having not only a Cartesian location in a picture plane but also a homologous (biologically corresponding) point in every other picture of a data set. One familiar example of such a point is the bridge of the nose in photographs of the human face in profile. The thin-plate-spline interpolant supplies a smooth surface through any combination of *heights* over any combination of *knots* in a base plane. Applications to morphometrics regard the coordinates of any one set of landmark locations as if they were heights above the picture of the landmark locations in any other form. The relation between any two configurations of correspondingly labeled landmark locations combines two thin-plate splines on the same set of knots, which are taken as the locations of the landmarks in one of the forms. One of these splines uses as its heights the x coordinates of the landmark locations in a second configuration, and the other spline uses for its heights the y coordinates of that same second set of landmarks. There results a model of the landmark

rearrangement as a *deformation* that takes the whole of the picture plane of one of the sets of landmarks onto the picture plane of the other, a mapping in which every pair of correspondingly labeled landmarks corresponds and that is pleasantly (indeed, optimally) smooth in between.

The splines not only supply a visualization of the interpolation, a solution of D'Arcy Thompson's old "problem of Cartesian transformation" [5], but also structure the multivariate *shape space* of the landmark locations by rotating it to a very useful new basis that is a function of the typical landmark configuration. Two dimensions of the space of landmark shapes are reserved to represent the affine transformations, or uniform shears, maps whose derivatives are the same everywhere; the remaining dimensions span the transformations having derivatives that vary from point to point of the interpolating map. Any thin-plate spline has a *bending energy*, the (idealized) physical energy required to bend an infinite, infinitely thin metal plate into the specified form from an initially flat configuration. By extension, any deformation of a landmark configuration modeled in this way has a bending energy, the sum of the energies of its two plates. A useful set of features into which deformations may be decomposed is the set of principal warps, eigenfunctions of the bending energy with respect to summed squared landmark displacement in the picture plane. Because relative vertical displacement between adjacent landmarks requires less energy as the landmarks are separated, these eigenvectors emerge ordered by an effective geometric scale for each feature. A recent monograph [6] explains this approach to biometrics in considerable detail.

In supplying a natural basis of features for the statistical space of landmark shapes near a mean form, the splines allow a tremendous expansion of morphometric technology for this special sort of data. The affine component and the principal warps of an observed sample of forms not only permit statistical tests of group differences, correlations, and the like but also supply a suggestive range of descriptions of any effects of group or other cause of deformation that might be suspected.

But the elegant algebra and multivariate statistics of the feature space yielded by the splines come at high cost: one's data must be restricted to discrete-point information, landmark locations only. In either morphometrics or image processing, one is accustomed to enriching data so discrete by information from edges or outlines. Such information typically is used to process and summarize medical images, to *ordinate* sets of biological forms and the like. Outlines have a multivariate statistics all their own, based on various primitive measures of distance between forms that do not much accord with notions of biological correspondence. Furthermore, it is typically much easier for automatic algorithms to acquire outlines than to acquire landmark locations; also, the outline methods can apply in contexts, such as pattern recognition of military or industrial scenes, to which the biological notion of homology (true anatomical correspondence) is irrelevant.

Thus although the multivariate methods for analysis of landmark locations are remarkably powerful, the data to which they apply are not typical of most applications of image analysis. It would seem desirable to bridge this gap by adapting the multivariate statistics of landmark locations to the richer data sets that are achievable by automatic or semiautomatic image processing.

The methods of this paper extend the image-warping technology of landmark locations, in particular its powerful subspaces of features and its tools for rigorous averaging, to incorporate some outline information typical of the broader range of applications of image analysis. We will show, in a variety of preliminary demonstrations, that the landmark-driven methods and the outline-driven methods may be firmly welded together by extending the thin-plate-spline formalism beyond landmark data to incorporate edge information *at landmarks*. This combination of landmark information and edge information uses the latter to augment the feature space describing the landmarks; approaches to the synthesis by the opposite route, augmenting edge descriptors by landmark information, have thus far proved unavailing. The main attraction of the landmark analyses, the linearized low-

dimensional feature sets they span, will apply to the additional features of edge information as well.

2 Thin-Plate Splines as Deformations

This section briefly reviews the formulation of the relatively more familiar thin-plate-spline (TPS) interpolant that we shall shortly perturb. For more details of this approach, see [6] and [7].

Let U be the function $U(x, y) = \frac{1}{2}(x^2 + y^2) \log(x^2 + y^2) = r^2 \log r$, and let $P_i = (x_i, y_i), i = 1, \ldots, k$, be k points in the plane. U is a fundamental solution of the biharmonic equation: we have $\Delta^2 U = \delta_{(0,0)}$, where δ here is Kronecker's function, zero everywhere except at the origin but with integral equal to 1, and Δ^2 is the iterated Laplacian $(\partial^2/\partial x^2 + \partial^2/\partial y^2)^2$. It can be shown [8] that the equation of a thin, uniform metal plate originally flat and now bent by vertical displacements at various points is $\Delta^2 U = 0$ except at points where forces are applied. The equation presumes that displacements normal to the rest position of the plate are sufficiently small that strains in the plane of the plate itself, and their energy, can be ignored. This equation has been solved under a great variety of schemes of strain and support. The application to image analysis, however, rests on a scenario with no equivalent in the world of real plates: the displacement of an infinitely extended metal plate at a finite series of discrete points in a world wholly lacking in gravity. It was a great discovery in interpolation theory [1], [2] that the equation of so unreal a plate may be integrated as a finite sum of copies of the kernel function U, along with an affine term for a tilt at infinity, by a matrix manipulation of extraordinary simplicity.

Let the knots of the spline (later to be the landmarks at which we are calibrating the deformation of one biological form into another) be at points P_i, $i = 1, \ldots, k$, in one single image. By writing $U_{ij} = U(P_i - P_j)$, build up matrices

$$K = \begin{pmatrix} 0 & U_{12} & \cdots & U_{1k} \\ U_{21} & 0 & \cdots & U_{2k} \\ \vdots & \vdots & \ddots & \vdots \\ U_{k1} & U_{k2} & \cdots & 0 \end{pmatrix}, \quad (1)$$

$$Q = \begin{pmatrix} 1 & x_1 & y_1 \\ 1 & x_2 & y_2 \\ \vdots & \vdots & \vdots \\ 1 & x_k & y_k \end{pmatrix}, \quad (2)$$

and

$$L = \begin{pmatrix} K & Q \\ Q^T & O \end{pmatrix}, \quad (k+3) \times (k+3), \quad (3)$$

where O is a 3×3 matrix of zeros. The thin-plate spline $f(P)$ having heights (values) h_i at points $P_i = (x_i, y_i)$, $i = 1, \ldots, k$, is the function

$$f(P) = \sum_{i=1}^{k} w_i U(P - P_i) + a_0 + a_x x + a_y y, \quad (4)$$

where

$$W = (w_1 \quad \cdots \quad w_k \quad a_0 \quad a_x \quad a_y)^T = L^{-1} Y \quad (5)$$

with

$$Y = (h_1 \quad h_2 \quad \cdots \quad h_k \quad 0 \quad 0 \quad 0)^T. \quad (6)$$

Note that the w's multiply copies of the kernel function $U = r^2 \log r$ evaluated with respect to each landmark in turn, whereas the coefficients a_0, a_x, a_y calibrate the function at infinity.

Then the function $f(P)$ has three crucial properties:

1. $f(P_i) = h_i$ for all i. (The function f interpolates the heights h_i at the landmarks P_i.) This is guaranteed by the first k rows of L.

2. The function f has the minimum bending energy of all functions that interpolate the heights h_i in that way: the minimum of

$$\int \int_{\mathbf{R}^2} \left(\left(\frac{\partial^2 f}{\partial x^2} \right)^2 + 2 \left(\frac{\partial^2 f}{\partial x \partial y} \right)^2 + \left(\frac{\partial^2 f}{\partial y^2} \right)^2 \right), \quad (7)$$

where the integral is taken over the entire picture plane.

3. The value of this bending energy is

$$\frac{1}{8\pi} W^T K W = \frac{1}{8\pi} W^T \cdot Y$$
$$= \frac{1}{8\pi} Y^T L_k^{-1} Y, \quad (8)$$

where L_k^{-1} is the $k \times k$ upper-left submatrix of L^{-1}. This value is zero precisely when the given heights are consistent with a linear surface $f(x, y) = a_0 + a_x x + a_y y$ over the knots; in that case the w's are all zero.

The spectrum of L_k^{-1} is of considerable practical interest in connection with biometrical analysis. It is of rank $k-3$ only; the three annihilated dimensions span the space of surfaces that are linear over the knots. The remaining $k-3$ dimensions are spanned by the *principal warps* of the landmark configuration. These come in a hierarchy of effective geometric scales. Typically, the largest represents the shift of one whole diagonal of the landmark configuration upward and the other downward, or, if the configuration is relatively elongated in one direction, this largest warp represents a bending of the splined surface along that long axis. Subsequent principal warps refer to aspects of the spline that are smaller and smaller in scale, until the very smallest encodes the discrepancy of heights in the smallest triangle of landmarks with respect to all the others. These eigenfunctions are very strongly dependent on the mean landmark configuration.

In the application to images we compute two of these splined surfaces, one for the x coordinates of the landmarks in a second form and one for the y coordinates. The resulting map $(f_x(P), f_y(P))$ is now a deformation of one picture onto the other that maps landmarks onto their homologues and has the minimum bending energy of any such interpolant. (In this context one may think of bending energy as a sort of *information*, the extent to which the affine derivative of the interpolant varies from location to location). The affine part of the map is now an ordinary shear, which can be represented by its principal strains in the usual way; the nonlinear part now comprises one multiple of each principal warp in the x direction and another in the y direction. The result is to project each principal warp into the picture plane as a *partial warp*, a 2-vector expressing its relative extent and direction, as it displaces all landmarks in parallel by multiples of the height of the corresponding principal warp.

Samples of homologous landmark configurations can be reduced to deformations of a stan-

dard, or of their average, in this way, and biostatistical analysis of effects on these shapes can proceed by any familiar multivariate technique as applied to the scores on these partial warps. Figure 1 shows a typical TPS interpolant for a configuration of six landmarks in two positions of no particular symmetry.

3 For a Single Edge-Element

3.1 An Intuitive Approach

We wish to investigate the behavior of the preceding interpolant when one landmark approaches infinitesimally close to another along a straight line. Figure 2 shows a configuration of six landmarks: four in the form of a square, invariant in locations between the panels, and two others that are displaced between left and right images. The segment between them lies horizontally in the left-hand configuration, but it is sheared to a dip of 45° at the right. The transformation grid shows how the shearing of this segment is propagated outward and is damped by the vertices of the square. In fact, it is slightly overdamped: beyond the boundary of the original square, the spline interpolant tends not to the identity map but to a gentle shear at infinity having the opposite sense (that is, upward, not downward, toward the right). But note that that boundary is not explicitly involved in the computation of the spline.

Figure 3, introducing a series of steadily shorter segment lengths δ, reduces the separation between the inner landmarks by a factor of 4 and also centers the one on the left. The separation is now $\delta = 0.1$ (in arbitrary units). The slope of the segment linking the two inner landmarks on the right is, of course, 45°.

If we now halve the separation of the inner landmarks (figure 4(a)), we see that whereas the slope of the segment on the right is still 45°, the relaxation of the induced shear has already set in at the separation of 0.1 representing the (fixed) grid spacing. The first grid intersection to the lower right of the center lies distinctly above the slope of the interlandmark segment at scale $\delta = 0.05$. To generate the correct correction at separation 0.1, we must shear the
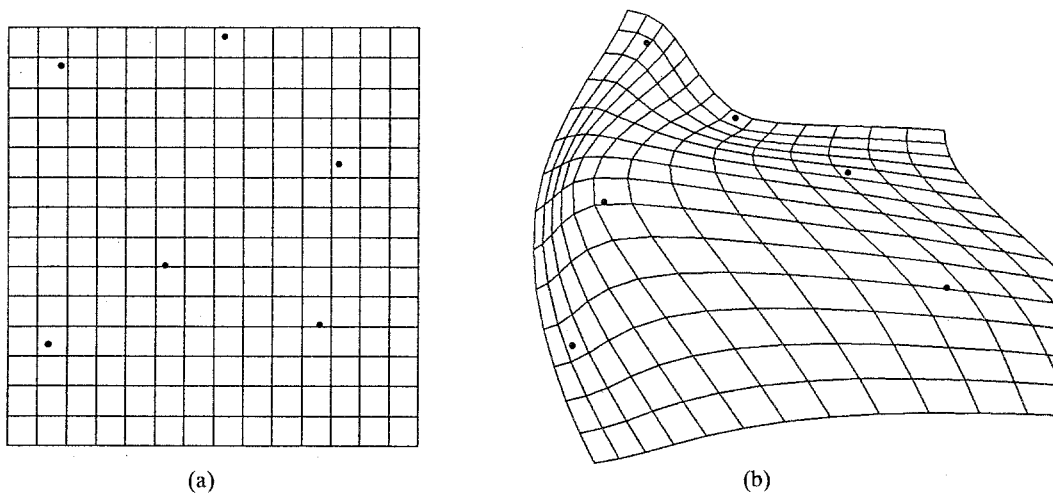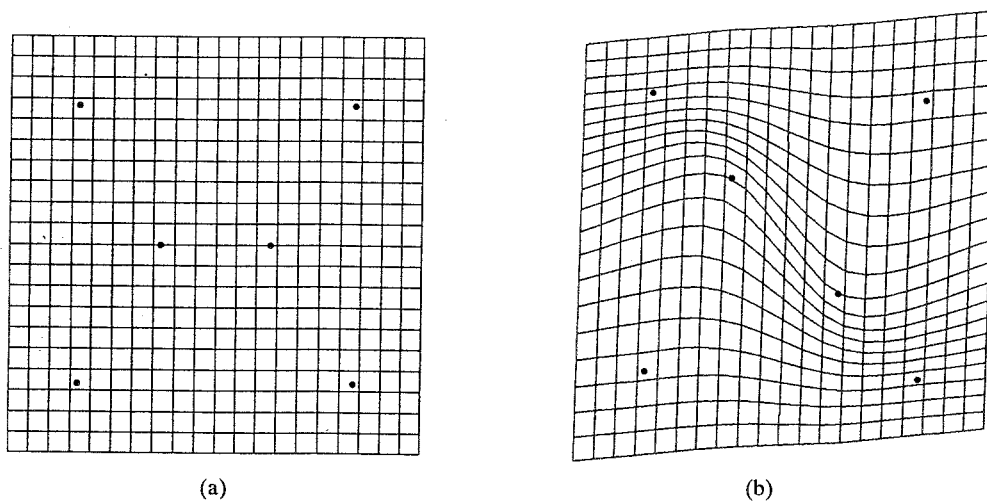
(a)

(b)

Fig. 1.

(a)

(b)

Fig. 2.

segment of length 0.05 downward by an angle whose tangent, as determined by numerical experimentation, should be about 1.3 (figure 4(b)).

As we continue to drive the exact spline map by shorter and shorter edges, the exact shear near the center relaxes more and more visibly. Figure 5(a) shows the spline beginning with a separation of $\delta = 0.01$, and figure 5(b) shows the shearing of the same segment by $\tan^{-1}(\log 0.01 / \log 0.1)$. It is this latter spline that reproduces the correct shearing at separation 0.1. Similarly, figure 6(a) shows the exact spline for a central segment of length $\delta = 0.0001$

(of course, we can no longer see that there are two landmarks there), and figure 6(b) shows the effect of a shearing by the angle of tangent $4 \, (= \log 0.0001 / \log 0.1)$.

This factor of the log edge length owes to the scaling of the kernel $U(r) = r^2 \log r$ to a lower order than r^2 near 0. (The following argument will be made quite rigorous in the next section). If the interpolant is to leave the corners of the square unmoved; it cannot have large coefficients for the terms of the spline centered at those landmarks. Hence the difference of δ in the y coordinates of the images of the two
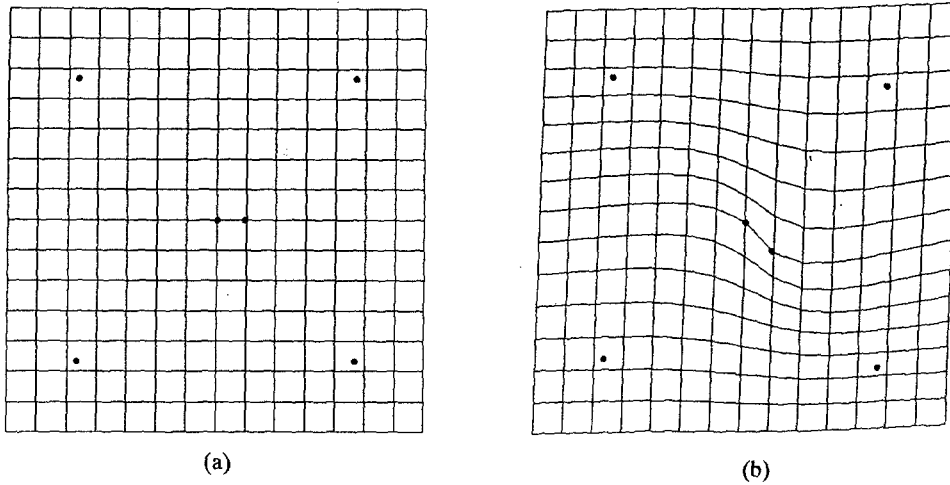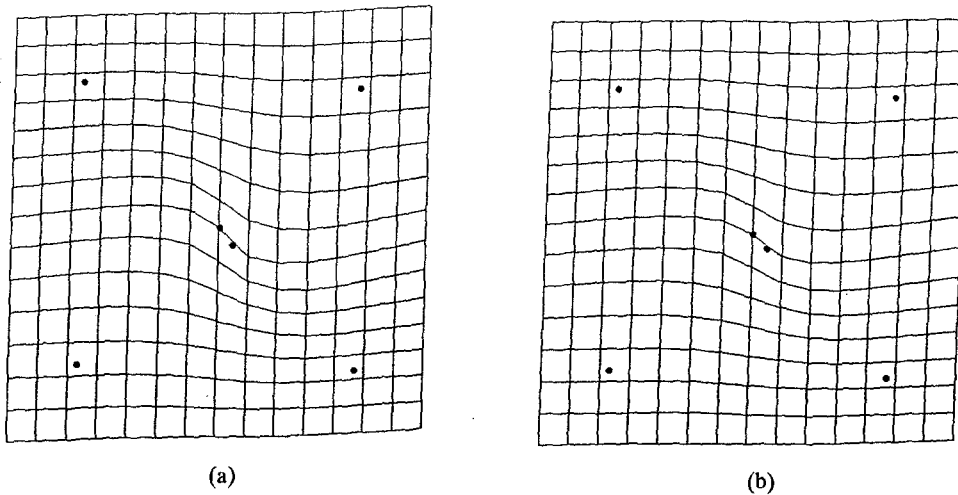
(a)

(b)

Fig. 3.

(a)

(b)

Fig. 4.

central points must be managed mainly by equal and opposite coefficients applying to the terms $U(\vec{r})$ and $U(\vec{r} - (\delta, 0))$. That is, our interpolant is dominated by $c[U(\vec{r}) - U(\vec{r} - (\delta, 0))]$ for some constant c. At $\vec{r} = (0,0)$ and $\vec{r} = (\delta, 0)$ these are $\mp c\delta^2 \log \delta$. Hence the value of c is nearly $(2\delta \log \delta)^{-1}$. Far from (0, 0) (for instance, at the end of that grid segment of fixed length 0.1), the evaluate $\delta^{-1}(U(\vec{r}) - U(\vec{r} - (\delta, 0)))$ approaches $\nabla U(r)$ in the horizontal direction. The remaining factor of $(\log \delta)^{-1}$ is the one we noticed we had to correct in figures 4(a), 5(a), and 6(a). With this factor of $\log \delta$ included, as $\delta \to 0$

the slope of the interpolant grows without limit along $(\delta, 0)$, just like that of the actual function $\nabla U(r) = r(2 \log r + 1)$, but the apparent slope along finite segments (like those composing this grid) converges quickly.

3.2 An Algebraic Development

Now let us do all this quite formally, still for the special case of a shear (now upward to simplify signs) of an originally horizontal edgel. Suppose we are using the formulas of section 2
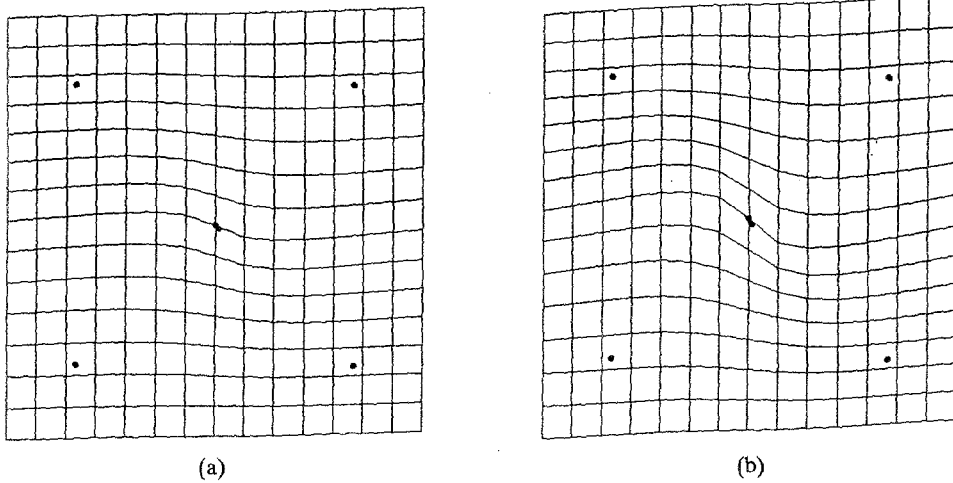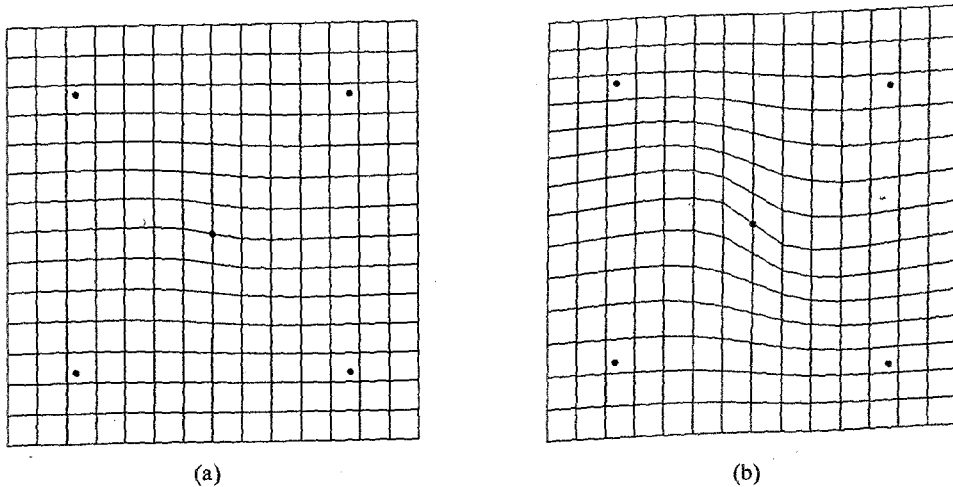
Fig. 5.

Fig. 6.

to interpolate the y coordinate of maps like that in figures 4, 5 and 6. The k landmarks $P_1 \cdots P_k$ are fixed; the $(k+1)$st is separated from P_k by a small vector $\vec{\delta} = (\delta, 0)$. The matrices K and L and the vectors Y and W are as in section 2, with $Y = (y_1 \quad y_2 \quad \cdots \quad y_k \quad 0 \quad 0 \quad 0)^T$.

To extend the spline to incorporate an additional landmark $P_{k+1} = (x_{k+1}, y_{k+1}) = (x_k + \delta, y_k)$, construct a new matrix inversion problem by bordering the matrix L rather than inserting a column internally: write

$$L' = \begin{pmatrix} L & M \\ M^T & 0 \end{pmatrix} \qquad (9)$$

for

$$M = (U_{\overline{1k+1}} \quad U_{\overline{2k+1}} \quad \cdots$$
$$U_{\overline{kk+1}} \quad 1 \quad x_{k+1} \quad y_{k+1})^T. \qquad (10)$$

We seek the $(k+1)$-landmark spline represented by $L'^{-1}Y'$ for $Y' = (Y^T \quad h_{k+1})^T$. L' is $(k+4) \times (k+4)$, Y' is a column vector of length $k+4$, and M is of length $k+3$. Here $h_{k+1} = y_{k+1} + \delta = y_k + \delta$, the height of the right end of the edgel segment in the right-hand image.

L'^{-1} may be computed from L^{-1} by using the

standard formula [9]

$$\begin{pmatrix} L & M \\ M^T & D \end{pmatrix}^{-1} = \begin{pmatrix} L^{-1} & 0 \\ 0 & 0 \end{pmatrix}$$
$$+ \begin{pmatrix} F \\ -I \end{pmatrix}(D - M^T L^{-1} M)^{-1}$$
$$\times (F^T \quad -I), \qquad (11)$$

where $F = L^{-1}M$. We apply the formula with $D = 0$, a scalar, and $-I = -1$ to get

$$\begin{pmatrix} L & M \\ M^T & 0 \end{pmatrix}^{-1} = \begin{pmatrix} L^{-1} & 0 \\ 0 & 0 \end{pmatrix} - \frac{1}{M^T L^{-1} M}$$
$$\times \begin{pmatrix} F \\ -1 \end{pmatrix}(F^T \quad -1); \qquad (12)$$

now F is a column vector of length $k + 3$.

For δ small we can approximate M by its differential:

$$M = (U_{1\overline{k+1}} \quad U_{2\overline{k+1}} \quad \cdots$$
$$U_{k\overline{k+1}} \quad 1 \quad x_{k+1} \quad y_{k+1})^T$$
$$= (U_{1k} \quad U_{2k} \quad \cdots$$
$$U_{kk} \quad 1 \quad x_k \quad y_k)^T + M^*, \qquad (13)$$

where

$$M^* \sim (\vec{\delta} \cdot \nabla_k U_1 \quad \vec{\delta} \cdot \nabla_k U_2 \quad \cdots$$
$$\vec{\delta} \cdot \nabla_k U_{k-1} \quad \delta^2 \log \delta \quad 0 \quad \delta \quad 0)^T.$$
$$\qquad (14)$$

Here, $\nabla_k U_i$ stands for the gradient of $U(P - P_i)$ evaluated at $P = P_k$, whereas $\delta^2 \log \delta$ is the value $U_{k\overline{k+1}}$ itself. (Recall that $U_{kk} = 0$.) Also, write

$$M^\circ = \lim_{\delta \to 0} M^*/\delta$$
$$= (\vec{t} \cdot \nabla_k U_1 \quad \vec{t} \cdot \nabla_k U_2 \quad \cdots$$
$$\vec{t} \cdot \nabla_k U_{k-1} \quad 0 \quad 0 \quad 1 \quad 0)^T, \qquad (15)$$

where $\vec{t} = (1, 0)$, the unit vector in the direction of the edgel.

Because $(U_{1k} \quad U_{2k} \quad \cdots \quad U_{kk} \quad 1 \quad x_k \quad y_k)^T$ is the kth column of L, we have

$$F = L^{-1}M$$
$$= (0 \quad 0 \quad \cdots \quad 0 \quad 1 \quad 0 \quad 0 \quad 0)^T + L^{-1}M^*,$$
$$\qquad (16)$$

for which we can substitute the limiting value

$$F^\circ \sim (0 \quad 0 \quad \cdots \quad 0 \quad 1 \quad 0 \quad 0 \quad 0)^T$$
$$+ \delta L^{-1}M^\circ. \qquad (17)$$

We further need to approximate that scalar reciprocal:

$$M^T L^{-1} M = M^T((0 \quad 0 \quad \cdots$$
$$0 \quad 1 \quad 0 \quad 0 \quad 0)^T + L^{-1}M^*)$$
$$= ((U_{1k} \quad U_{2k} \quad \cdots$$
$$U_{kk} \quad 1 \quad x_k \quad y_k) + M^{*T})$$
$$\times ((0 \quad 0 \quad \cdots$$
$$0 \quad 1 \quad 0 \quad 0 \quad 0)^T + L^{-1}M^*)$$
$$= U_{kk} + 2(0 \quad 0 \quad \cdots$$
$$0 \quad 1 \quad 0 \quad 0 \quad 0)M^*$$
$$+ M^{*T} L^{-1} M^*$$
$$\sim 2\delta^2 \log \delta + \delta^2 (M^{\circ T} L^{-1} M^\circ)$$
$$= \delta^2 (2 \log \delta + C), \qquad (18)$$

say. (This is how the term in $\log \delta$ arises. That the term $C = M^{\circ T} L^{-1} M^\circ$ interferes with exact proportionality to $\log \delta$ is the reason that the scaling of slopes in figures 4(b), 5(b), and 6(b) is not quite perfect.)

The thin-plate spline on all $k + 1$ landmarks is the $(k + 4)$-vector $L'^{-1}Y'$ interpreted as the formula for a spline mapping. Define

$$F' = \begin{pmatrix} F^\circ \\ -1 \end{pmatrix}$$
$$= \begin{pmatrix} \delta L^{-1}M^\circ \\ 0 \end{pmatrix}$$
$$+ (0 \quad 0 \quad \cdots \quad 0 \quad 1 \quad 0 \quad 0 \quad 0 \quad -1)^T,$$
$$\qquad (19)$$

the term in Rao's formula (12) corresponding to the limit from (17). Then

$$L'^{-1}Y' \sim \begin{pmatrix} L^{-1} & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} Y \\ y_k + \delta \end{pmatrix}$$
$$- \frac{1}{\delta^2 (2 \log \delta + C)} F' F'^T \begin{pmatrix} Y \\ y_k + \delta \end{pmatrix}$$
$$= (0 \quad 0 \quad \cdots \quad 0 \quad 0 \quad 1 \quad 0)^T$$
$$- \frac{1}{\delta^2 (2 \log \delta + C)} F'(F^{\circ T} Y - (y_k + \delta)).$$
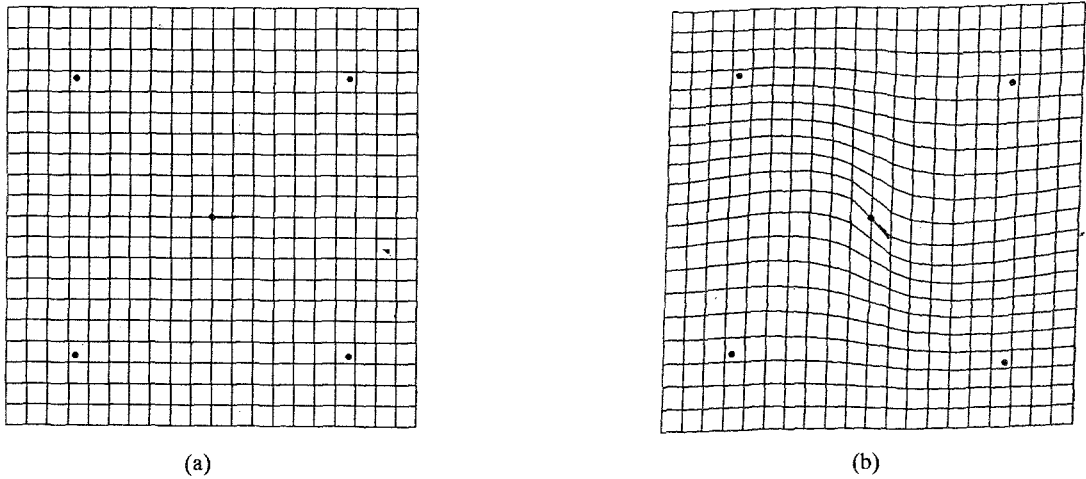$$\qquad (20)$$

(a)

(b)

Fig. 7.

The first term is the identity spline $f(x, y) = y$, and $F^{\circ T}Y = y_k + \delta M^{\circ T}L^{-1}Y = y_k + \delta M^{\circ T}(0\ 0\ \cdots\ 0\ 0\ 1)^T = y_k$. Hence the correction term by which the map $L'^{-1}Y'$ differs from the identity spline is approximately

$$\frac{1}{\delta(2\log\delta + C)}F' \sim \frac{1}{2\log\delta + C}$$

$$\times \left[(0\ 0\ \cdots\ 0\ 1\ 0\ 0\ 0\ -1)^T/\delta + \begin{pmatrix} L^{-1}M^{\circ} \\ 0 \end{pmatrix}\right].$$

$$(21)$$

The first k coefficients apply to the k real landmarks, the next three apply to the affine part $(1, x,$ and $y)$, and the last applies to the $(k + 1)$st landmark at infinitesimal distance to the right of the kth landmark.

Except in the δ-vicinity of P_k, the first term inside the square brackets in (21) is very nearly $\vec{t} \cdot \nabla U_k$, the gradient of U_k in the direction of the edge. This term perturbs each landmark except the kth to a certain extent, and it also induces an affine term (a shear at infinity). The second term inside the square brackets in (21), the spline $L^{-1}M^{\circ}$ on the first k landmarks, is the mapping that exactly cancels these unwanted landmark shifts. (Notice that the terms of M°, $\nabla_k U_i$, are opposite in sense to the $\nabla_i U_k$ of the first term.) Because C is not a function of δ, in the limit $\delta \to 0$ the multiplication by 2 log δ cancels out the first factor $(2\log\delta + C)^{-1}$ here, the factor that made the splines of figures

4(a), 5(a), and 6(a) too shallow. We may then draw the limiting form without specifying δ at all, as in figure 7. Notice three characteristic features of such deformations: the typical $r\log r$ shape of the edge through the edge element, the damping of this perturbation at different rates along and perpendicular to the edge, and the induction of changes of affine derivative at the other landmarks. The edge scale δ is not coded here at all.

Figure 7 resembles one of Timoshenko's [10]; we have seen these clamped splines in no other published source. Figure 8 shows the same spline as the vertical plate the equation (4) of which we borrowed for our in-plane work. The intuition that clamping is required to obtain the bending shown is equivalent to the observation that relation (21) is the limiting form deriving from a pair of landmarks at the middle of the plate, not a singleton.

Although the dominant component $\vec{t} \cdot \nabla U(\vec{r})$ of the edge term is a function only of the direction $\vec{t}$, the presence of the other landmarks induces a further nonlinear part, encoded in the first k rows of $L^{-1}M^{\circ}$. Figure 9 shows how these terms generally bend the edge kernel ∇U away from symmetry to leave landmarks invariant at whatever separations characterize their spacing, direction by direction.

As noted in section 2, every thin-plate-spline interpolant minimizes a global bending energy, the integrated quadratic variation (7) of the x
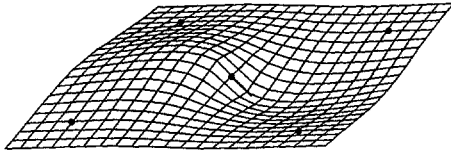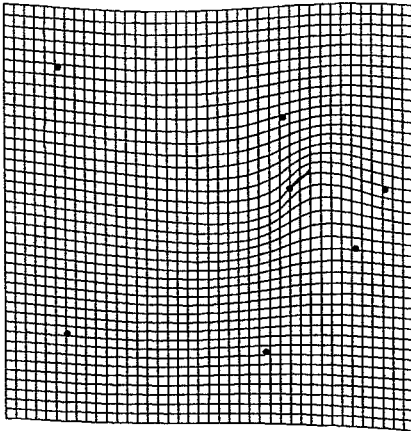
Fig. 8

Fig. 9

and y components of the mapping separately. For the case of the single edgel, we can evaluate this energy surprisingly easily. Combining (8) and (21), we see that the bending energy of F'/δ is

$$\frac{1}{8\pi}\frac{1}{2\log\delta + C}Y'$$

$$\cdot\left[(0\ 0\ \cdots\ 0\ 1\ 0\ 0\ 0\ -1)^T/\delta + \begin{pmatrix} L^{-1}M^\circ \\ 0 \end{pmatrix}\right].$$

(22)

We have

$$Y'\cdot(0\ \cdots\ 1\ 0\ 0\ 0\ -1)^T/\delta = [y_k - (y_k + \delta)]/\delta = -1$$

(23)

and

$$Y'\cdot\begin{pmatrix} L^{-1}M^\circ \\ 0 \end{pmatrix} = (Y^T\ \ 0\ \ 0\ \ 0)L^{-1}M^\circ$$

$$= (0\ \cdots\ 0\ \ 1)M^\circ = 0.$$

(24)

Hence the bending energy of this simple unit shear is proportional to $-(2\log\delta + C)^{-1}$. Rewriting that as

$$\frac{1}{-2\log\delta}\left(1 + \frac{C}{2\log\delta}\right)^{-1},\qquad (25)$$

we see the role of $C = M^{\circ T}L^{-1}M^\circ$ as an *embeddedness* — the higher the value of C, the higher the bending energy of an edge shear of fixed vertical extent, whatever the value of δ. If a shear is at angle $\tan^{-1}s$ instead of $45° = \tan^{-1}1$, the bending energy scales as the square of the tangent, approximately $-s^2/(2\log\delta + C)$.

As set forth here, formula (21) applies only to the case of unit shears in an edgel that is originally horizontal. Shearing by an angle other than $\tan^{-1}1$ is, of course, managed by an appropriate multiple of this expression. It is easily verified [7] that the thin-plate-spline interpolant relating any pair of landmark configurations is appropriately covariant under rotation of either configuration separately. The covariance extends to the limits of landmark configurations driving these edgel interpolants. Thus relation (21) actually applies to shears of edgels beginning at any orientation on the page. The expression $L^{-1}M^\circ$ transforms appropriately under rotations; one merely applies formula (21) to the Cartesian coordinate perpendicular to the edgel, leaving the coordinate along the edgel unchanged. This is shown algebraically in the context of an analysis of many edgels at once, to which we now turn.

4 For Many Edgels

The procedure of the previous section can be carried out for many edgels at the same time. Whatever their positions and orientations, the limiting form of the warp may be expressed in a simple closed form. Although the derivation is an instructive generalization of that in the preceding section, some readers may wish to skip over it. Those readers need to know mainly that the interpolant is given by the spline of (33) with coefficients from (45), (47), (50), and (52) combined according to (51) and that the edgel bending-energy matrix $\widetilde{N}$ referred to in section 5 below generalizes the embeddedness C of the single-edgel analysis according to (66).

4.1 Minimum-Energy Warp with Specified Finite Differences

Let there be k landmark locations at $P_1, \ldots, P_k$, as before, and now let there be $m > 0$ edge specifications. The jth edge specification pertains to the direction $t_j = (t_{j,x}, t_{j,y}) = (\cos\theta_j, \sin\theta_j)$ through the i_jth landmark P_{i_j}. It will be useful to arrange these edge indices i_j in a sparse matrix J having m rows of k elements, the jth row having 1 in the i_jth position and zeros elsewhere. In the analyses in this paper, no two values of i_j can be the same (but see the remark at the end of subsection 4.3). We will first solve the problem of a scalar function (vertical thin plate) with specified heights h_i at the landmarks P_i and specified slopes s_j along edges j; then, in subsection 4.4 below, we will combine the solutions for x and y coordinates separately into a single plane deformation, just as in the formalism of the landmark-driven spline [6].

In this section *only*, we are somewhat free with matters of vector and matrix notation. The letter a will stand for any of several vectors $(a_1, \ldots, a_k)^T$, b will stand for any of several m-vectors $(b_1, \ldots, b_m)^T$, and c will stand for any of several 3-vectors $(c_1, c_2, c_3)^T$. Subscripts 0 and 1 (as in splines f_0 and f_1 or energies E_0 and E_1) refer to landmark-driven terms and edge-driven modifications, respectively; thus a_0 is the k-vector a for the spline on landmarks only. Elements of vectors and matrices will often themselves be vectors or matrices. When 0 is used as a vector or matrix, its dimensions should likewise be understood from context. Subscripts or superscripts x and y will denote the usual Cartesian components of $\mathbf{R}^2$.

For any finite δ the spline we seek is an exact spline on $(k+m)$ landmarks. It therefore may be written out in full in a partitioning of equation (4): we have

$$f(z) = \sum_{i=1}^{k} a_i U(z - P_i) + \sum_{j=1}^{m} b_j U(z - P_{i_j} - \delta t_j)$$
$$+ c_1 + c_2 x + c_3 y, \qquad (26)$$

where $z = (x, y) \in \mathbf{R}^2$ and $U(z) = |z|^2 \log |z|$. Using these points and images in the standard thin-plate-spline machinery, we convert equation

(5) into one for a, b, and c:

$$\begin{bmatrix} K & K' & Q \\ K'^T & K'' & JQ + \delta V \\ Q^T & Q^T J^T + \delta V^T & 0 \end{bmatrix} \begin{pmatrix} a \\ b \\ c \end{pmatrix}$$
$$= \begin{pmatrix} h \\ h + \delta s \\ 0 \end{pmatrix}. \qquad (27)$$

In equation (27) the matrices K and Q are as in equations (1) and (2), and the new matrices V, K', K'' are defined as follows:

$$V = \begin{bmatrix} 0 & t_{1,x} & t_{1,y} \\ \vdots & \vdots & \vdots \\ 0 & t_{m,x} & t_{m,y} \end{bmatrix}, \qquad (28)$$

$$K' = \begin{bmatrix} U(P_1 - P_{i_1} - \delta t_1) & \cdots & U(P_1 - P_{i_m} - \delta t_m) \\ U(P_2 - P_{i_1} - \delta t_1) & \cdots & U(P_2 - P_{i_m} - \delta t_m) \\ \vdots & & \vdots \\ U(P_k - P_{i_1} - \delta t_1) & \cdots & U(P_k - P_{i_m} - \delta t_m) \end{bmatrix}, \qquad (29)$$

and

$$K'' = \begin{bmatrix} U(P_{i_1} + \delta t_1 - P_{i_1} - \delta t_1) & \cdots & U(P_{i_1} + \delta t_1 - P_{i_m} - \delta t_m) \\ U(P_{i_2} + \delta t_2 - P_{i_1} - \delta t_1) & \cdots & U(P_{i_2} + \delta t_2 - P_{i_m} - \delta t_m) \\ \vdots & & \vdots \\ U(P_{i_m} + \delta t_m - P_{i_1} - \delta t_1) & \cdots & U(P_{i_m} + \delta t_m - P_{i_m} - \delta t_m) \end{bmatrix}. \qquad (30)$$

(Note that the diagonal of K'' is identically zero.) It will also be helpful to rewrite the matrix K in a parallel form:

$$K = \begin{bmatrix} U(P_1 - P_1) & \cdots & U(P_1 - P_k) \\ U(P_2 - P_1) & \cdots & U(P_2 - P_k) \\ \vdots & & \vdots \\ U(P_k - P_1) & \cdots & U(P_k - P_k) \end{bmatrix}. \qquad (31)$$

The spline of equations (26) and (27) is the function with least bending energy that has the right heights and, when averaged over edges of length δ, the right slopes at the subset of landmarks specified by the i_j. From equation (8) that bending energy is

$$E = \frac{1}{8\pi}(a^T \; b^T) \begin{bmatrix} K & K' \\ K'^T & K'' \end{bmatrix} \begin{pmatrix} a \\ b \end{pmatrix}. \qquad (32)$$

We simplify the behavior of this spline as δ tends to 0 by using a first-order approximation

of $U(z - P_{i_j} - \delta t_j)$. Combine the $U(z - P_{i_j} - \delta t_j)$ terms with $U(z - P_{i_j})$ terms in the formula for $f(z)$:

$$f(z) = \sum_{i=1}^{k} a_i U(z - P_i)$$
$$+ \sum_{j=1}^{m} b_j \frac{1}{\delta} \big(U(z - P_{i_j} - \delta t_j)$$
$$- U(z - P_{i_j}) \big) + c_1 + c_2 x + c_3 y \quad (33)$$

These new coefficient vectors a and b are related to those of equations (26) and (27) by linear operations. Applying corresponding operations to the matrix in the spline and energy equations above and swapping the order of b and c gives

$$L' \begin{pmatrix} a \\ c \\ b \end{pmatrix} = \begin{pmatrix} h \\ 0 \\ s \end{pmatrix},$$

$$E = \frac{1}{8\pi} (a^T \quad 0 \quad b^T) L' \begin{pmatrix} a \\ 0 \\ b \end{pmatrix}, \quad (34)$$

with

$$L' = \begin{bmatrix} K & Q & (K' - KJ^T)/\delta \\ Q^T & 0 & V^T \\ (K'^T - JK)/\delta & V & (K'' - JK' - K'^T J^T + JKJ^T)/\delta^2 \end{bmatrix}. \quad (35)$$

4.2 Approximating L' and L'^{-1}

To expand L' as $\delta \to 0$, we need a compact notation for the derivatives of U. Its derivative in the direction t will be denoted U_t. Thus

$$U_{t_i}(z) = t_i \cdot \nabla U(z) = (2 \log |z| + 1)(z \cdot t_i). \quad (36)$$

Similarly, the second derivative $(U_t)_u$ will be denoted $U_{t,u}$. Explicitly,

$$U_{t_i, t_j}(z) = t_j \cdot \nabla U_{t_i}(z)$$
$$= \frac{2}{|z|^2}(z \cdot t_i)(z \cdot t_j)$$
$$+ (2 \log |z| + 1)(t_j \cdot t_i). \quad (37)$$

Using this notation, we may simplify L' by expanding it to first order in δ component by

component. The entries of $(1/\delta)(K' - KJ^T)$ are

$$\frac{1}{\delta} \big(U(P_i - P_{i_j} - \delta t_j) - U(P_i - P_{i_j}) \big)$$
$$= \begin{cases} \delta \log \delta & \text{if } i = i_j, \\ -U_{t_j}(P_i - P_{i_j}) + O(\delta) & \text{otherwise.} \end{cases}$$

Hence

$$\frac{1}{\delta}(K' - KJ^T) = (\delta \log \delta)J^T - K_1 + O(\delta), \quad (38)$$

where

$$K_1 = \begin{bmatrix} U_{t_1}(P_1 - P_{i_1}) & \cdots & U_{t_m}(P_1 - P_{i_m}) \\ U_{t_1}(P_2 - P_{i_1}) & \cdots & U_{t_m}(P_2 - P_{i_m}) \\ \vdots & & \vdots \\ U_{t_1}(P_k - P_{i_1}) & \cdots & U_{t_m}(P_k - P_{i_m}) \end{bmatrix}. \quad (39)$$

The components of $(1/\delta^2)(K'' - JK' - K'^T J^T + JKJ^T)$ are

$$\frac{1}{\delta^2} \big(U(P_{i_i} + \delta t_i - P_{i_j} - \delta t_j) - U(P_{i_i} - P_{i_j} - \delta t_j)$$
$$- U(P_{i_i} - P_{i_j} + \delta t_i) + U(P_{i_i} - P_{i_j}) \big)$$
$$= \begin{cases} -2 \log \delta & \text{if } i = j; \\ -U_{t_j, t_i}(P_{i_i} - P_{i_j}) + O(\delta) & \text{otherwise.} \end{cases}$$
$$(40)$$

Hence

$$\frac{1}{\delta^2}(K'' - JK' - K'^T J^T + JKJ^T)$$
$$= (-2 \log \delta)I_m - K_2 + O(\delta), \quad (41)$$

where I_m is the $m \times m$ identity matrix and

$$K_2 = \begin{bmatrix} U_{t_1, t_1}(P_{i_1} - P_{i_1}) & \cdots & U_{t_m, t_1}(P_{i_1} - P_{i_m}) \\ U_{t_1, t_2}(P_{i_2} - P_{i_1}) & \cdots & U_{t_m, t_2}(P_{i_2} - P_{i_m}) \\ \vdots & & \vdots \\ U_{t_1, t_m}(P_{i_m} - P_{i_1}) & \cdots & U_{t_m, t_m}(P_{i_m} - P_{i_m}) \end{bmatrix}, \quad (42)$$

with the convention that $U_{t_i, t_j}(0) = 0$. Together these give

$$L' = \begin{bmatrix} K & Q & -K_1 + (\delta \log \delta)J^T \\ Q^T & 0 & V^T \\ -K_1^T + (\delta \log \delta)J & V & (-2 \log \delta)I_m - K_2 \end{bmatrix} + O(\delta). \quad (43)$$

To approximate L'^{-1}, use the trick of partitioned-matrix inversion we introduced in equation (11). The appropriate partition of L' uses the submatrix

$$L = \begin{bmatrix} K & Q \\ Q^T & 0 \end{bmatrix},$$

independent of δ as before, and also

$$M = \begin{bmatrix} -K_1 \\ V^T \end{bmatrix} + O(\delta \log \delta),$$
$$D = (-2 \log \delta)I_m - K_2 + O(\delta). \quad (44)$$

Putting

$$M \sim \begin{bmatrix} -K_1 \\ V^T \end{bmatrix},$$
$$N = K_2 + M^T L^{-1} M, \quad (45)$$

we get

$$L'^{-1} = \begin{bmatrix} L^{-1} & 0 \\ 0 & 0 \end{bmatrix} - \frac{1}{2\log\delta} \begin{bmatrix} -L^{-1}M \\ I_m \end{bmatrix}$$
$$\times \left(I_m + \frac{N}{2\log\delta} \right)^{-1} \begin{bmatrix} -L^{-1}M \\ I_m \end{bmatrix}^T + O(\delta). \quad (46)$$

4.3 Approximating the Spline Coefficients

Substituting into the spline equation, we reconstruct the zero-order terms of the spline (33) as

$$\begin{pmatrix} a_0 \\ 0 \\ c_0 \end{pmatrix}, \quad \text{where} \quad \begin{pmatrix} a_0 \\ c_0 \end{pmatrix} = L^{-1} \begin{pmatrix} h \\ 0 \end{pmatrix}, \quad (47)$$

and the zero-order term of the energy (34) as

$$E_0 = \frac{1}{8\pi}(a_0^T \quad 0)L\begin{pmatrix} a_0 \\ 0 \end{pmatrix}$$
$$= \frac{1}{8\pi}(h^T \quad 0)L^{-1}\begin{pmatrix} h \\ 0 \end{pmatrix}. \quad (48)$$

These are the coefficients and energy of the thin-plate spline f_0 resulting from the landmark locations alone. The rest of the expression is a function of

$$\Delta = \begin{bmatrix} -L^{-1}M \\ I_m \end{bmatrix}^T \begin{pmatrix} h \\ 0 \\ s \end{pmatrix} = s - M^T L^{-1} h, \quad (49)$$

a vector of differences between the desired slopes s and the (infinitesimal) slopes $v = M^T L^{-1} h$ resulting from the landmark-driven map f_0. The Cartesian components of v will be written

$$v_x = M^T L^{-1} \begin{pmatrix} h_x \\ 0 \end{pmatrix} \quad \text{and} \quad v_y = M^T L^{-1} \begin{pmatrix} h_y \\ 0 \end{pmatrix}.$$

From equations (34) and (46) the coefficients of the difference terms [the middle summation in equation (33)] are approximately

$$b_1 = \frac{-1}{2\log\delta} \left(I_m + \frac{N}{2\log\delta} \right)^{-1} \Delta. \quad (50)$$

The complete spline is then

$$\begin{pmatrix} a \\ b \\ c \end{pmatrix} = \begin{pmatrix} a_0 \\ 0 \\ c_0 \end{pmatrix} + \begin{pmatrix} a_1 \\ b_1 \\ c_1 \end{pmatrix} + O(\delta) \quad (51a)$$

with energy

$$E = E_0 + E_1 + O(\delta). \quad (51b)$$

Here

$$\begin{pmatrix} a_1 \\ c_1 \end{pmatrix} = -L^{-1}Mb_1 \quad (52)$$

and

$$E_1 = \frac{1}{8\pi}\frac{-1}{2\log\delta}\Delta^T \left(I_m + \frac{N}{2\log\delta} \right)^{-1} \Delta$$
$$= \frac{-2\log\delta}{8\pi}b_1^T \left(I_m + \frac{N}{2\log\delta} \right)b_1. \quad (53)$$

If we denote the spline with coefficients a_1, b_1, and c_1 by f_1, then it can be verified that (uniformly on compact subsets of $\mathbf{R}^2$)

$$f = f_0 + f_1 + O(\delta),$$
$$\nabla f = \nabla(f_0 + f_1) + O(-\delta \log \delta),$$
$$f_1 = O\left(\frac{-1}{\log\delta} \right),$$
$$\nabla f_1 = O(1),$$
$$E_1 = O\left(\frac{-1}{\log\delta} \right). \quad (54)$$

Within a ball of radius δ about each edgel the slopes of both f and $f_0 + f_1$ differ from what

is specified by the edgel at that landmark by an amount $O(-1/\log\delta)$.

Uniformly on compact subsets of $\mathbf{R}^2$ we have

$$\frac{1}{\delta}\big(U(z - P_{i_j} - \delta t_j) - U(z - P_{i_j})\big)$$
$$= -U_{t_j}(z - P_{i_j}) + O(-\delta\log\delta). \quad (55)$$

(The convergence is not $O(\delta)$ because $U(z - P_{i_j})$ has infinite second derivative at P_{i_j}.) Hence if we define $\tilde{f}_1$ by

$$\tilde{f}_1(z) = \sum_{i=1}^{n} a_{1,i}U(z - P_i) - \sum_{j=1}^{m} b_{1,j}U_{t_j}(z - P_{i_j})$$
$$+ c_{1,1}x + c_{1,2}y + c_{1,3}, \quad (56)$$

then $f_1 = \tilde{f}_1 + O(-\delta\log\delta)$. (The coefficients are the elements of the vectors a_1, b_1, and c_1 of equation (51a).) Note that a_1 and c_1 are the coefficients of the conventional thin-plate spline that, when added to the combination of $-U_{t_j}(z - P_{i_j})$ functions with coefficients b_1, exactly reverses the displacements of the landmarks and makes the function $O(|z|)$ at infinity. Hence the approximation by $f_0 + \tilde{f}_1$ is exact at the landmarks.

Equation (50) for b_1 can be rewritten as

$$(-2\log\delta)b_1 - K_2 b_1 - M^T L^{-1} M b_1 = \Delta. \quad (57)$$

The right-hand side is the desired change in slope at each edge. The second and third terms on the left-hand side are the slopes at each edge induced by the U_t functions at different landmarks and by the correcting thin-plate spline. Although each U_t term has infinite slope at its landmark in the direction of its edge, at scale δ (roughly speaking) it has slope $-2\log\delta$.

It can be verified that $\tilde{f}_1$ is invariant under affine transformations of the image space, but not of the domain, and that the approximations all hold uniformly in s for bounded s.

This formula is itself only a special case, in which we have assumed that no two edge specifications lie at the same landmark (that is, that $i_i \neq i_j$ for $i \neq j$). The algebra of these correction terms may be extended to include such combinations of slope specifications, allowing, in effect, a full specification of the normal to the thin-plate spline as a surface. A term J'

would then be added to the approximation to the mixed second derivative in equation (41) — we actually have

$$\frac{1}{\delta^2}(K'' - JK' - K'^T J^T + JKJ^T)$$
$$= (-2\log\delta)\mathrm{I}_m - K_2 + J' + O(\delta), \quad (41')$$

where J' is an $m \times m$ *edge-incidence matrix* that is log 2 for any pair of edges at the same landmark and is otherwise 0. The additional term continues to appear through all subsequent equations. It is a function of the specific limiting process invoked here and can be removed by a different approach to the singular perturbation.

4.4 Extension to Edgels in Two Dimensions

The development of (26)–(57) treats a single dimension of thin-plate splining. Now consider both components of the deformation. Dropping the subscript 1 from b_1, we have

$$\begin{pmatrix} b_x \\ b_y \end{pmatrix} = \frac{1}{-2\log\delta}\begin{bmatrix} Q & 0 \\ 0 & Q \end{bmatrix}^{-1}\begin{pmatrix} \Delta_x \\ \Delta_y \end{pmatrix}$$

and

$$E_1 = \frac{-2\log\delta}{8\pi}\begin{pmatrix} b_x^T & b_y^T \end{pmatrix}\begin{bmatrix} Q & 0 \\ 0 & Q \end{bmatrix}\begin{pmatrix} b_x \\ b_y \end{pmatrix}, \quad (58)$$

where

$$Q = \left(\mathrm{I}_m + \frac{N}{2\log\delta}\right).$$

Suppose that all the edges t_j are required to be mapped to directions given by unit vectors w_j but that they are allowed to grow or shrink. Let W^x denote the diagonal matrix with the components of the vector w^x along the diagonal, and similarly for W^y. It is convenient to make a change of basis at each edge to components along and perpendicular to the target directions, even though we will continue to refer to them by subscripts x and y:

$$\begin{pmatrix} b_x \\ b_y \end{pmatrix} = \begin{bmatrix} W^x & -W^y \\ W^y & W^x \end{bmatrix}\begin{pmatrix} \tilde{b}_x \\ \tilde{b}_y \end{pmatrix},$$
$$\begin{pmatrix} \Delta_x \\ \Delta_y \end{pmatrix} = \begin{bmatrix} W^x & -W^y \\ W^y & W^x \end{bmatrix}\begin{pmatrix} \tilde{\Delta}_x \\ \tilde{\Delta}_y \end{pmatrix}. \quad (59)$$

Then

$$\begin{pmatrix} \widetilde{b}_x \\ \widetilde{b}_y \end{pmatrix} = \frac{1}{-2\log\delta} \begin{bmatrix} Q^S & -Q^D \\ Q^D & Q^S \end{bmatrix}^{-1} \begin{pmatrix} \widetilde{\Delta}_x \\ \widetilde{\Delta}_y \end{pmatrix} \quad (60)$$

and

$$E_1 = \frac{-2\log\delta}{8\pi} \begin{pmatrix} \widetilde{b}_x^T & \widetilde{b}_y^T \end{pmatrix} \begin{bmatrix} Q^S & -Q^D \\ Q^D & Q^S \end{bmatrix} \begin{pmatrix} \widetilde{b}_x \\ \widetilde{b}_y \end{pmatrix}, \quad (61)$$

where $Q^S = W^x Q W^x + W^y Q W^y$ and $Q^D = W^x Q W^y - W^y Q W^x$.

If $\widetilde{s}_y$ is the assigned change of slope in this new basis, the requirement placed on the deformed image is $\widetilde{s}_y = 0$. Since $\Delta = s - v$, this means

$$\widetilde{\Delta}_y = -\widetilde{v}_y = -(-W^y v_x + W^x v_y). \quad (62)$$

The $\widetilde{\Delta}_{x,j}$ (lengths of the resulting edgels on the right) are arbitrary. To be consistent with our approach of selecting the function satisfying the constraints with least bending energy, equation (61), we should minimize over them as well. The minimum is at $\widetilde{b}_x = 0$, giving

$$\widetilde{b}_y = \frac{-1}{2\log\delta} (Q^S)^{-1} \widetilde{\Delta}_y,$$

$$E_1 = \frac{-2\log\delta}{8\pi} \widetilde{b}_y^T Q^S \widetilde{b}_y. \quad (63)$$

Note that

$$\widetilde{\Delta}_{y,j} = |v_j| \sin \Delta\theta_j, \quad (64)$$

where $\Delta\theta_j$ is the angle of rotation from v_j to w_j. If the w_j are nearly parallel to the v_j (so that the edgel-driven part of the deformation is small), to first order we may replace W^x and W^y by V^x and V^y (the components of the images under f_0 of the starting edge orientations t) in the equations for Q^S and in the change of basis for b, equations (59), and we may replace the $\widetilde{\Delta}_{y,j}$ by $\Delta\theta_j$. In this case the deformation is linear in the angles of rotation at each edgel separately, and the matrix Q^S can be simplified as $I_m + \widetilde{N}/2\log\delta$, where

$$(\widetilde{N})_{ij} = (v_i \cdot v_j) N_{ij}, \quad i, j = 1, 2, \ldots, m. \quad (65)$$

(The terms N_{ij} are characterized in (42) and (45).)

Likewise, the bending energy simplifies in the limit of small rotations. Collect the rotational deficits in an m-vector $\Delta\Theta = \{\Delta\theta_j\}$. If all landmarks are fixed ($h_x = h_y = 0$), this is just the vector of rotations, edgel by edgel, from the t's to the w's. Then the bending energy of equation (58) is very nearly

$$\frac{1}{8\pi} \frac{-1}{2\log\delta} \Delta\Theta^T \left(I_m + \frac{\widetilde{N}}{2\log\delta} \right)^{-1} \Delta\Theta. \quad (66)$$

(Compare expression (25), in which $\Delta\theta = 1$, a scalar, and the factor of 8π is suppressed.)

Just as the general method for a one-dimensional sheet can be extended to include specification of two edges at a single landmark, so the extension to two dimensions of interpolation allows for specification of two directions of derivatives at a landmark. That is, the method actually incorporates constraints on the general affine derivative in the δ-vicinity of landmarks, not merely these edgel shears. The rotations $\Delta\theta_j$ of the present discussion can actually be arbitrary 2-vectors applying to alter the images of the end points of the landmark-mapped edgels t_j under the linearized mapping. If the vector t_j is extended along its length, there results a localized dilatation, and so on. A computer program for this purpose is introduced in [11], and several useful special cases are demonstrated there.

The freedom afforded by these combinations of edgel shears is one of the great pleasures of our work station. Our videotape [12] demonstrates the construction of these deformations in interactive mode. In a journal such as this we can show only the final frame of one such session. Figure 10(a) is a typical mild spline on five landmarks. Figure 10(b) is the result of playfully torquing all of the edgels shown. Considerable concentration is required to construe this surface as the flat deformation grid actually being visualized.

5 Principal Edgel Warps as a Basis for the Space of Edgels

Application of these ideas to biomedical scene analysis proceeds by identifying edgels with tangents or normals to biological outlines running through landmark locations. Figure 11 shows
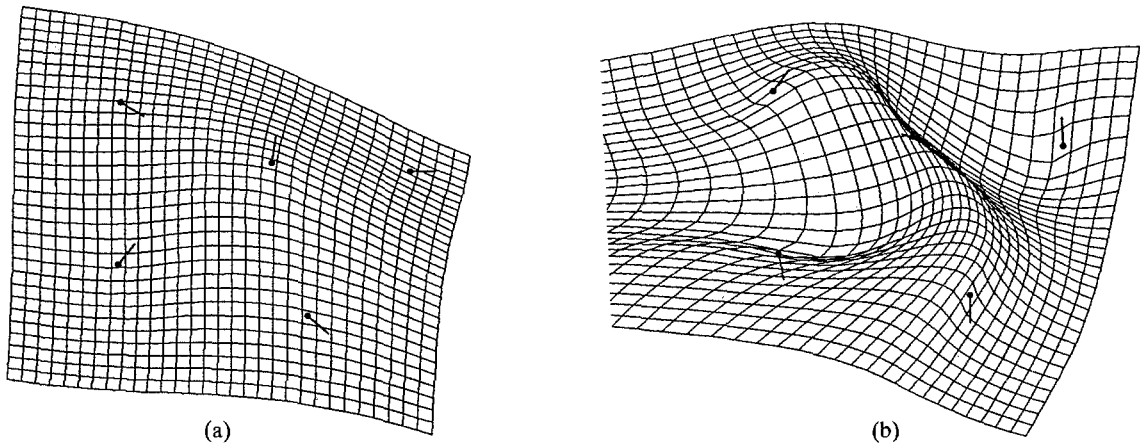
(a)   (b)

Fig. 10.

a preliminary, although still unrealistically symmetric, scene of this sort in which the edgel is tangent to a standard outline at one of four landmarks in the form of a square. The effect of the mapping $f_{\vec{t}}$ is suddenly interpretable in a different language: what is induced by the edgel shear we are abstracting is not a pseudovertical shearing, as in earlier figures, but is a *bulge* of the outline itself. Comparing figures 11(b) and 11(c), we see how the edgel-driven warps of images, when applied to outlines alone, specify changes of features of those outlines that are perceptually independent of the features of the planar maps that underlie them.

For the case of landmark data alone, the bending energy (8) is a quadratic form in the right-hand landmark coordinates, the form $Y^T L_k^{-1} Y$, having an eigenanalysis of great usefulness [6], as briefly sketched in section 1. In the vicinity of a mean form about which changes are small, the approximate bending energy for the additional deformation according with edgels, expression (66), is likewise of a very suggestive form. It is *incommensurate* with the bending energy that pertains to the landmark reconfiguration, i.e., it is weighted by an indeterminate factor $-(2 \log \delta)^{-1}$, but otherwise its spectrum is that of $\widetilde{N}$ and hence is independent of the value of δ.

For landmarks, the eigenvectors of the bending-energy quadratic form represent patterns of landmark displacement having stationary values of that energy for a unit sum of squares of Cartesian displacements. For our edgels, the eigenvectors of $\widetilde{N}$ represent, similarly, patterns of coordinated small rotations of edgels that, taken as a whole, have stationary values of bending energy for a fixed sum of squares of these rotations. They collect the interactions of these multiple rotations by a specific sort of propagation of those rotations across the interior of a form. In this matter they contrast substantially with other approaches to the interactions of edgels [13], [14] that consider only edgels aligned on a single outline and that represent localizations in arc length. The edgel warps are not restricted to outlines but can apply to any configuration of edgels inside or on the boundary of a biological image.

We shall refer to the eigenvectors of $\widetilde{N}$, and also to the deformations they specify, as the *principal edgel warps* of the schemes of landmarks P_i and P_i' and edgels t_j. These warps depend on all three of these arguments in a strongly non-linear manner. The remainder of this section exemplifies these bases for various cases characterized by diverse degrees of symmetry. In all these analyses of circles, the landmarks are presumed fixed—in the preceding formalism we take $h^x = h^y = 0$, so that the spline f_0 is the identity mapping.
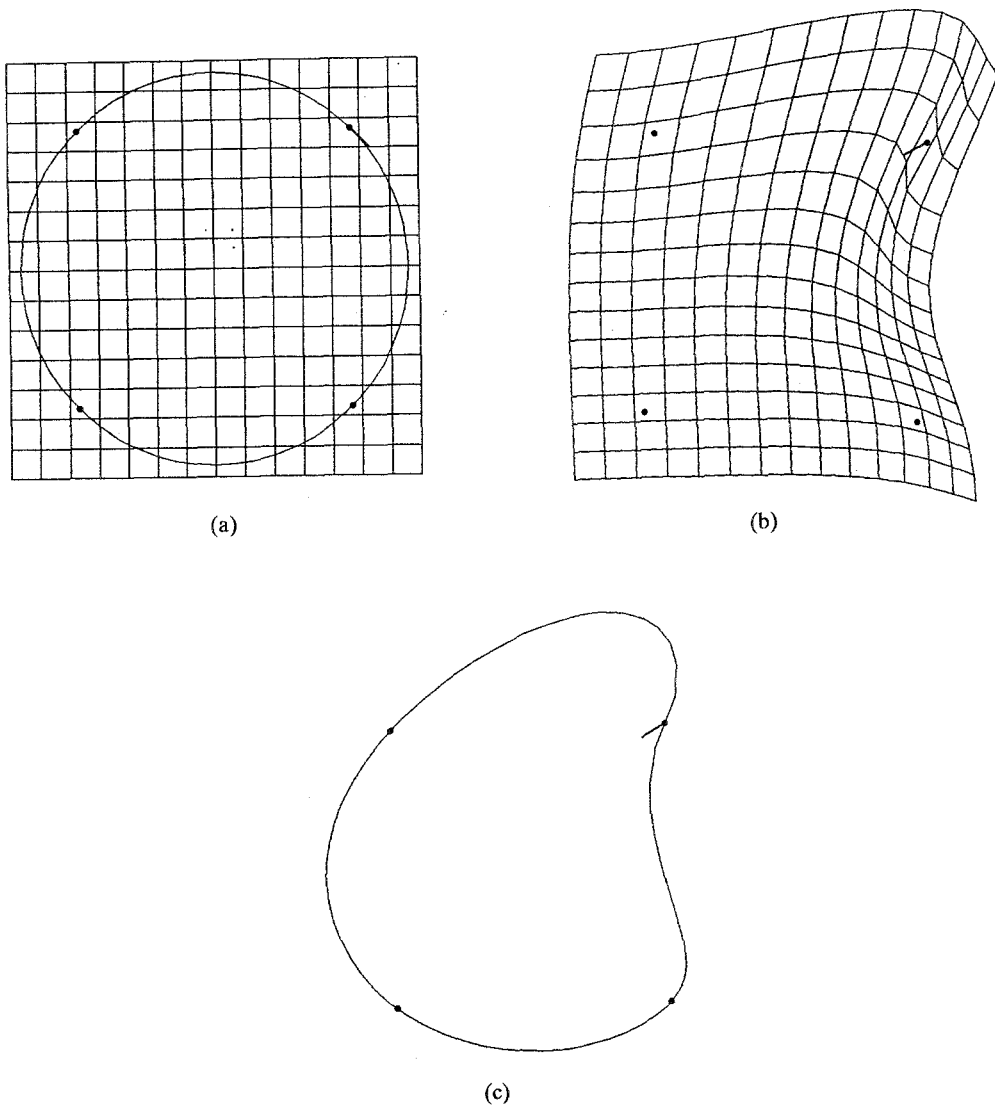
(a)

(b)

(c)

Fig. 11.

5.1 Some Symmetrical Configurations

The simplest edgel analysis after the one-edgel case of section 3 is the extension to a pair of edgels. The form of the principal edgel warps depends crucially on the angle between the edgel directions t_1 and t_2 and the angle between those directions and the segment $P_1 - P_2$ linking the landmarks at which the edgels lie. In the matrix $\widetilde{N}$ it should be noted that perpendicular edgels (those with $t_i \cdot t_j = 0$) do not interact at all. The analysis of pairs of edgels, then, can be restricted to the case of parallel edgels only. The characteristic scenarios involve edgels that are parallel to or perpendicular to the segment linking the landmarks at which they lie.

For the situation of Figure 12, to which we have added a little additional biological scenery, the matrix

$$\widetilde{N} = \begin{pmatrix} 3.63 & 2.31 \\ 2.31 & 3.63 \end{pmatrix}.$$

(The specific entries here are a mild function of the number and locations of those additional
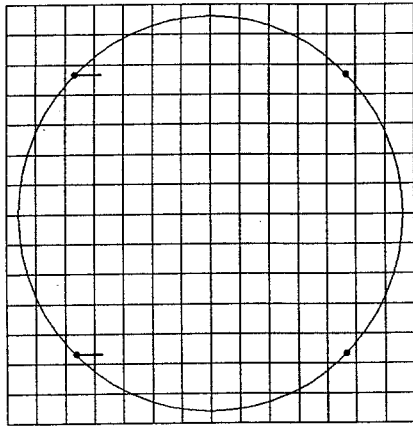
Fig. 12.

landmarks.) Its eigenvectors $(1, \pm 1)$ represent maps that shear the edgels to the same extent either in the same sense or in opposite senses. These two principal edgel warps differ considerably in specific bending energy. The deformations corresponding to a summed squared edgel shear $\|\Delta\Theta\|$ of unity are shown in figure 13, and the splined surfaces corresponding to the deformed y coordinates are shown in figure 14. In both figures part (a) looks the smoother: the impression is clearly of a virtual landmark located vaguely near the middle of the scene and dragging the interior upward or downward. (Note, however, in the surface visualization that in fact the bending has relaxed between the edgels by just a bit.) The more bent principal edgel warp appears to be an affine transformation highly localized to the interior of the configuration. The surface representation is perhaps the clearer; it shows that this warp is a pronounced twist of its coordinate sheet linking equal but opposite $r \log r$ terms at the two landmarks. The bending energy is large because the derivative $\partial f_y/\partial x$ must change sign between the landmarks, leading to a large second derivative $\partial^2 f_y/\partial x \partial y$ and thus a large quadratic variation (7).

A typical scene in which edgels point at each other might have them lying at opposite ends of an organismal axis, as in figure 15. Now the matrix

$$\tilde{N} = \begin{pmatrix} 4.70 & -1.31 \\ -1.31 & 4.70 \end{pmatrix}.$$

The principal edgel warps (figure 16) are still the same combinations of parallel or antiparallel edge shear, but the maps they drive are somewhat altered. The more bent mapping, corresponding to the eigenvector $(1, -1)$, now indeed looks the more bent as a map. The induced S-curve along the equator again requires a large variation of $\partial f_y/\partial x$, now entailing large $\partial^2 f_y/\partial x^2$; the less bent grid, as in the previous case, seems to show the effect of a virtual landmark near the middle. These impressions are consistent with the explicit representation of the deformed y coordinates as surfaces in figure 17. A pairing of two of the maps of figure 16(b) at 90° rotates the interior of the circle with respect to the exterior.

The preceding two examples were not particularly biological in that the edgels were not associated with the outline (the circle) drawn. (Indeed, in the second of the two examples the shape of that circle is hardly altered at all.) Figure 18 shows a more suggestive, although still unrealistically symmetric, scene in which the edgels lie along tangents to a standard outline at four landmarks in the form of a square. Because in this case $t_i \cdot t_j = 0$ for $i + j$ odd, the eigenanalysis of $\tilde{N}$ reduces to two copies (one for each of the pairs of parallel edgels) of the analysis of a two-edgel case, the scene of figure 15 in which the edgels are rotated by 90°. For this reduced problem

$$\tilde{N} = \begin{pmatrix} 2.61 & -0.77 \\ -0.77 & 2.61 \end{pmatrix}.$$

The principal edgel warps are not very different in energy. Each leaves two edgels clamped and twists the other two in either parallel or antiparallel fashion, as shown in figures 18(b) and 18(c); the more bent, in this case that for the eigenvector $(1, -1)$, may be referred to as a *taco bell* (after the logo of an American fast-food chain).

5.2 Principal Edgel Warps in the General Case

The degeneracy of the preceding example is broken quite easily. Figure 19 shows a small modification in which a landmark is moved and its edgel is rotated to remain tangent to the
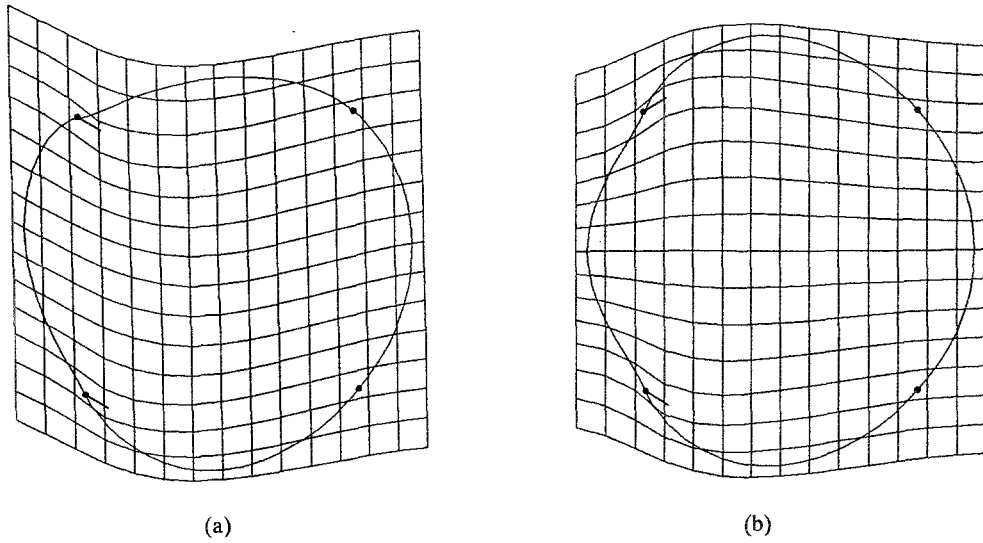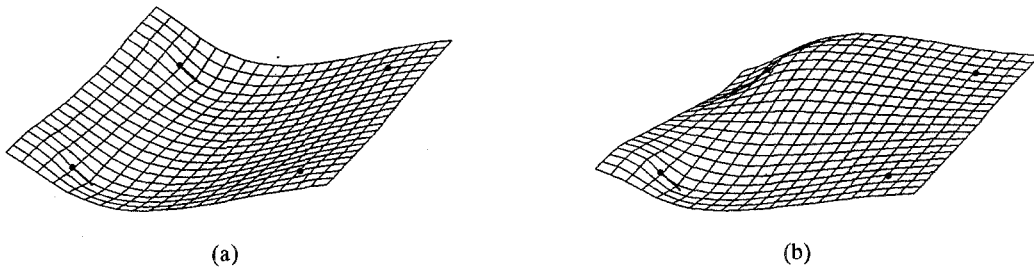
(a)

(b)

Fig. 13.

(a)

(b)

Fig. 14.

circle. The matrix $\widetilde{N}$ is now (with landmarks numbered counterclockwise from the lower left)

$$\widetilde{N} = \begin{pmatrix} 0.43 & 0. & -0.65 & -1.58 \\ 0. & 2.58 & 0. & -0.46 \\ -0.65 & 0. & 3.97 & 1.11 \\ -1.58 & -0.46 & 1.11 & 1.16 \end{pmatrix} . \quad (67)$$

The zero entries correspond to the perpendicularity between the edgel at the second landmark and those at the first and third.

All eigenvalues are distinct, and the different principal edgel warps now represent conceptually quite distinct aspects of the bending of this outline around its landmarks. The most-bent principal edgel warp, figure 19(b), is the induction of an S-curve between the pair of landmarks at closest separation. Taken with this sign, it grows a nose. The second-most-bent warp, figure 19(c), bulges out (when taken with this sign) rather less locally, the whole length of this shortest interlandmark arc. This is, perhaps, prognathism. The third-most-bent principal edgel warp is the same deformation on the other side, a combined protrusion and downward shear. The least-bent principal edgel warp, figure 19(e), is by far the most interesting. It represents an extrusion of the interior along the arc of the outline at the upper left, the arc least constrained by landmarks. The figure shows an uncanny resemblance to the creature called the Shmoo in the comic strip "Li'l Abner" that was popular in the United States in the 1950's and 1960's. (Its face is at the upper right, and its front and back feet are at the
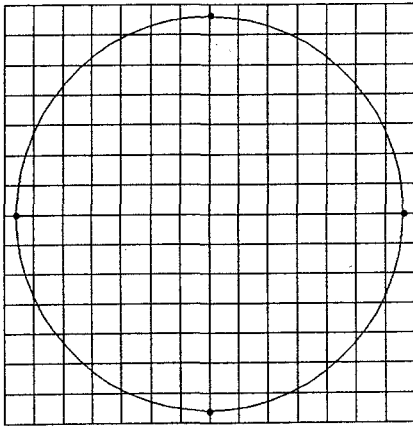
Fig. 15.

bottom.) The Shmoo (which laid eggs in cartons and produced milk in bottles) was drawn as a creature of infinite gentleness; these warps, in their curving, seem to incorporate some of that visual aspect of gentleness as well. Quite literally, one cannot imagine a twisting of this outline, leaving the landmarks alone, that is less bent than the form drawn here.

As the number of edgels increases further, the principal-edgel warps ramify among the landmarks by combining aspects of these special cases. Figure 20, for instance, shows a combination in which there is no evidence of any starting circle. The principal edgel warps of highest specific bending, such as that in figure 20(b), are those that twist single edgels that are closely spaced with respect to others lying at parallel orientation nearby. The warps of lowest and second-lowest specific bending, figures 20(c) and 20(d), pull out features like those we have been studying for the square by combination of appropriately large-scale features. In all three of these warps the landmarks have not moved. The eigenstructure of the matrix $\tilde{N}$, in fact, serves as a useful formalism for extracting scales of organization of flow in scenes as arbitrary as this.

6 Application to Medical Images

Let us turn from these abstractions to a data set of real images, as promised: 14 midsagittal MRI scans of normal UCLA medical students. The images were originally gathered and their landmarks located by Daniel Valentino of UCLA. These data have been the subject of analysis in a previous videotape [15], and the average picture the landmarks support has been published previously [16], [17]. It was generated by unwarping each image to a standard configuration, the true (statistical) average of the 14 sets of landmarks, and then averaging the unwarped pictures pixelwise. In that figure, one of us (F.L.B.) traced by hand a set of five contours, two open and three closed, corresponding to the outline drawing in figure 21. It is this outline that will bear the landmarks and edgels for the demonstration to follow. A demonstration in real time is available on a recently published videotape [12].

The machinery of the splines applies to figure 21 as to any other two-dimensional scene. We can move landmarks around, and we can define and shear edgels whether or not we choose to display the effect of the resulting mapping on a whole Cartesian grid. Figure 22 shows an artificial acrocephaly—we grabbed the vertex of the head and dragged it upward—and figure 23 shows the further evolution of one of the students away from his pithecanthropic forebears by shearing his frontal lobe forward. Corresponding to these deformations there is a warping of any underlying averaged image into that of the case and an unwarping of the case into the coordinate system of that average.

The single-edgel mapping F' of relation (21) is well defined up to a factor of proportionality. In the applications shown here, whether to grids or to pictures, that indeterminacy characterizes the geometric scale at which the finite separation vector along $\vec{t}$ displays precisely the change of slope that was specified. These scales, although individually indeterminate, may nevertheless take on empirically characterized ratios from landmark to landmark of a scene; we will investigate these in a subsequent paper. In the context of multiscale analysis of individual images, as by convolutions with derivatives of Gaussians, this implies that whether or not individual edge elements have a particular spatial scale, *rotations* (shears) of edgels have particular scales that can be estimated from image series.
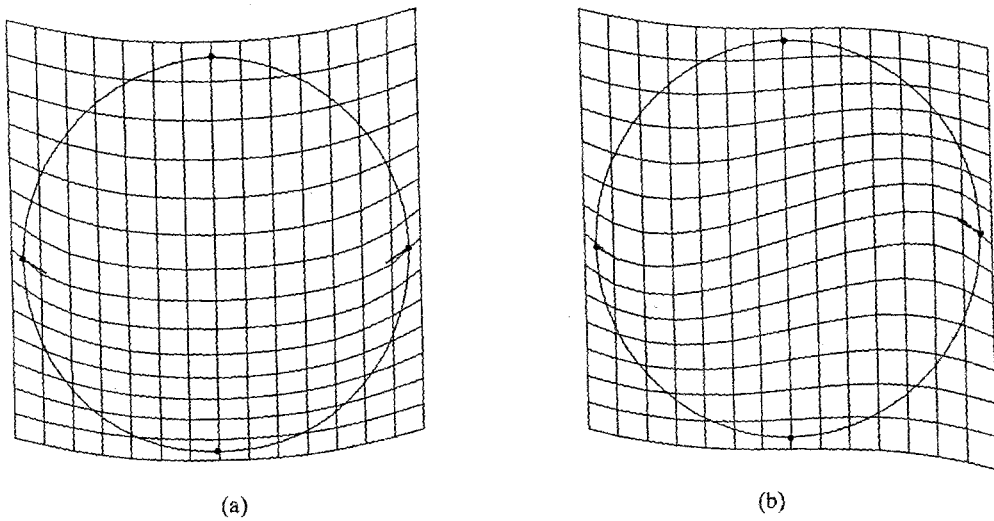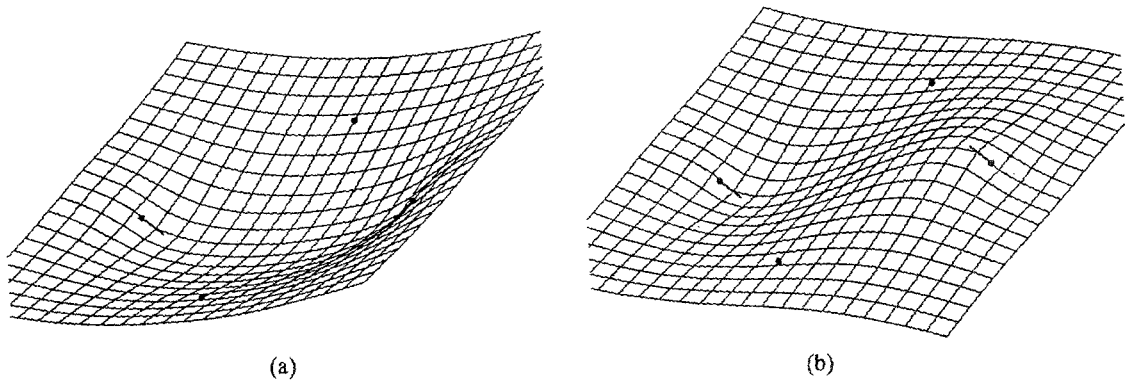
(a)　　　　　　　　　　　　　(b)

Fig. 16.

(a)　　　　　　　　　　　　　(b)

Fig. 17.

We combine these two tactics, the warping of the standard outline and the warping of the specimen images, in a new, potentially more precise approach to the problems of image averaging (the "brain atlas problem" of earlier articles [17], [18]) and of image covariances.

It is easiest to proceed by holding the patient's image fixed on the screen and then warping the standard outline interactively until it appears to match adequately the important contours of the specimen image. The landmarks and edgels of figure 21 may be considered a basis of nine landmarks and four edgels for the representation of any specimen. Figure 24 shows an arbitrary superposition over a single case; figure 25 matches the landmark locations (a step that might well be taken almost wholly automatically); figure 26 matches the four edgel slopes as well. This little interactive demonstration substitutes for a rigorous algorithm not worth developing on a sample so small: an adaptive low-dimensional version of the familiar deformable template. Notice that the edgel in the frontal lobe is a free dial corresponding to no local image information, whereas the other three all align with image gradients that themselves could be extracted by local operations. To this matching of warped standard outline to specimen contours corre-
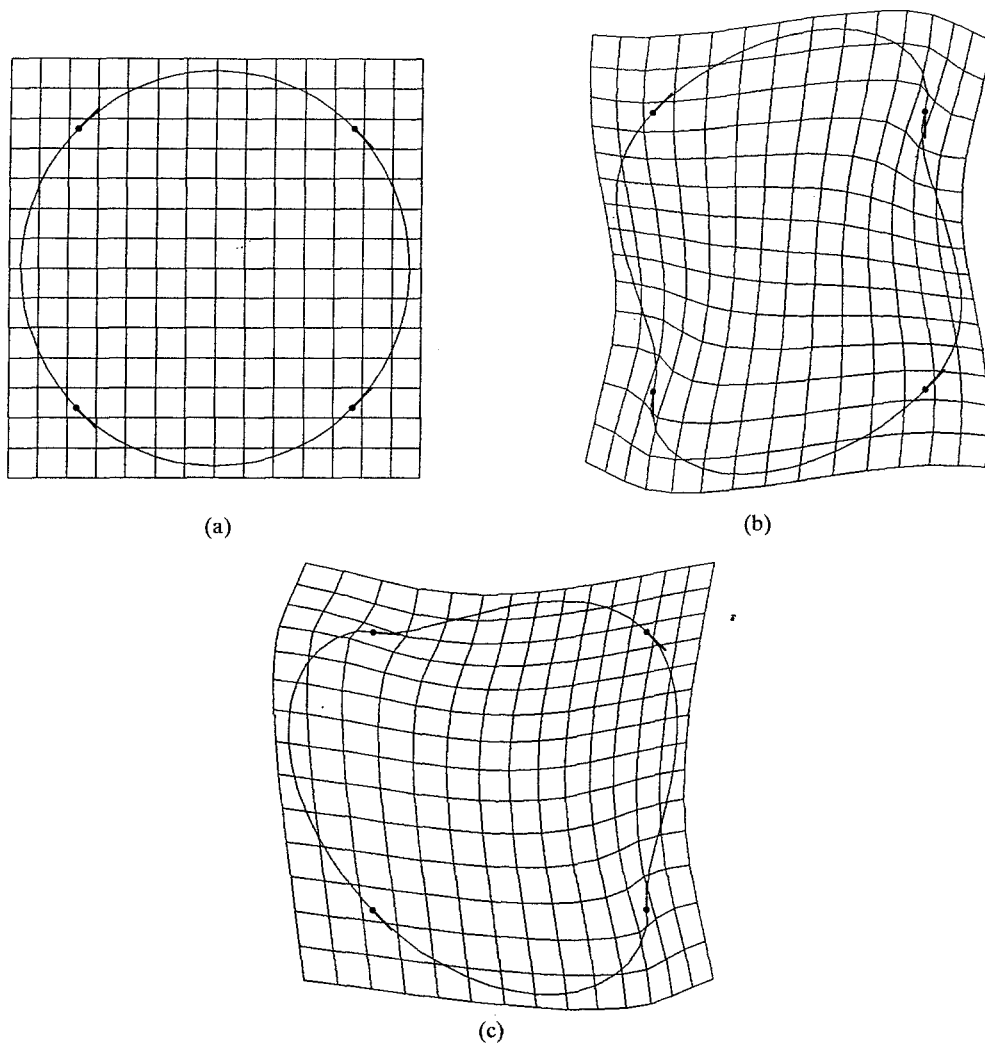
(a)

(b)

(c)

Fig. 18.

sponds a warping function (figure 26) that, inverted, converts the patient's gray levels to the standardized image of figure 27. The panels of figure 28 show other instances of this warping function.

When all specimens of the data set have been digitized as instantiations of the standard, we may proceed to average all of the unwarped images pixelwise in the newly consistent coordinate system. This summary (figure 29) shows greater sharpness in several regions of the image than does an average driven only by landmarks ([16], [17]). Whereas in one sense such improvement is surely to be expected (since the edges

were explicitly sharpened by actual geometric intervention), in another it represents a considerable enrichment of the stereotactic resources that can be turned to image averaging. The contours of this more sharply averaged picture could be traced in turn and the entire atlas computation iterated, applied to sharpen contours of specific organs, such as the pons or the cerebellum [12], or applied to new series of images, such as those with specific diseases, those undergoing specific treatments, or those that have been histologically labeled in particular ways. We might also have studied the endophrenology of this little example—the correlations between

(a)

(b)

(c)

(d)

(e)

Fig. 19.
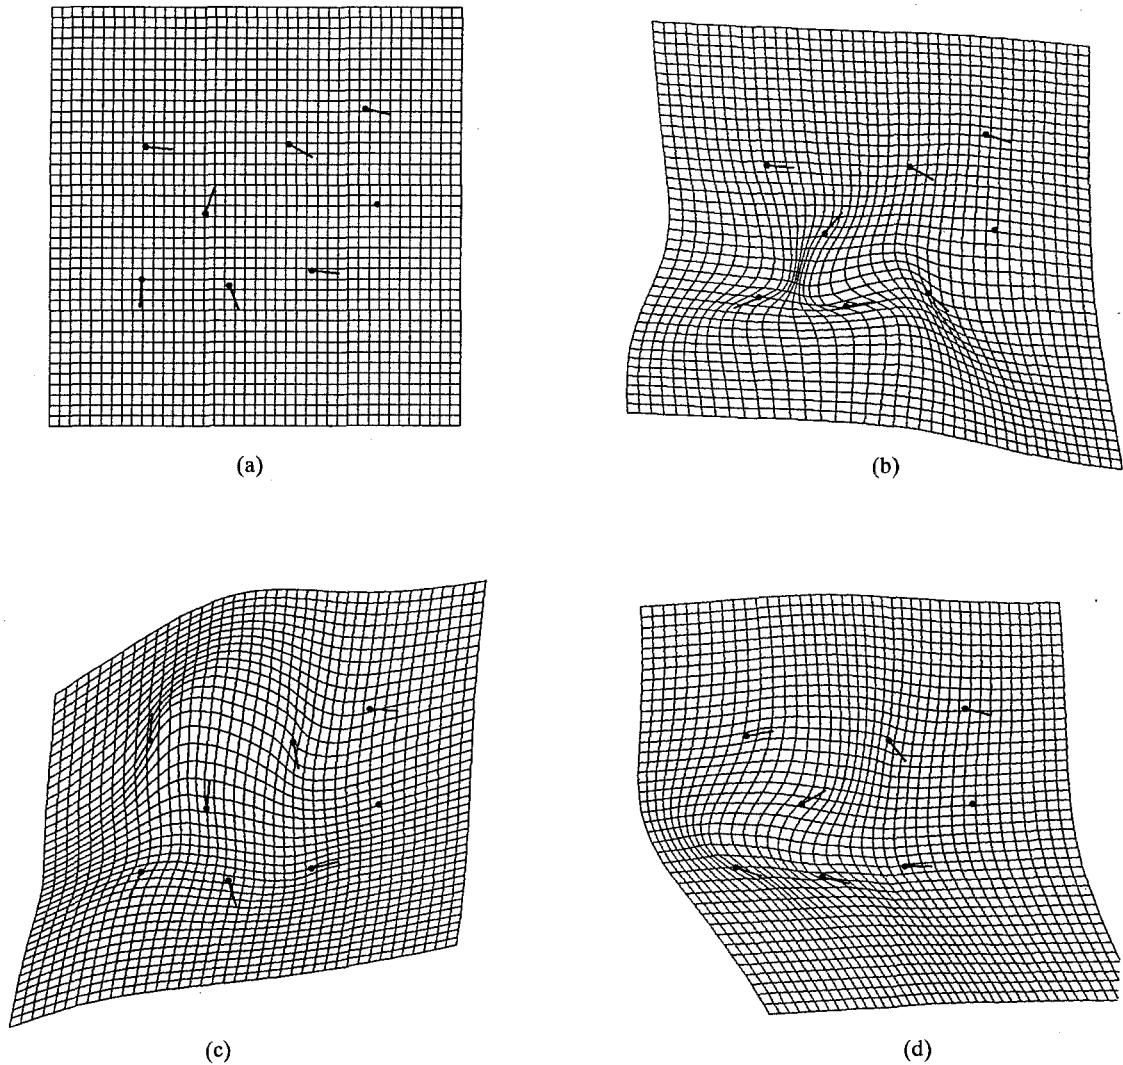
(a)

(b)

(c)

(d)
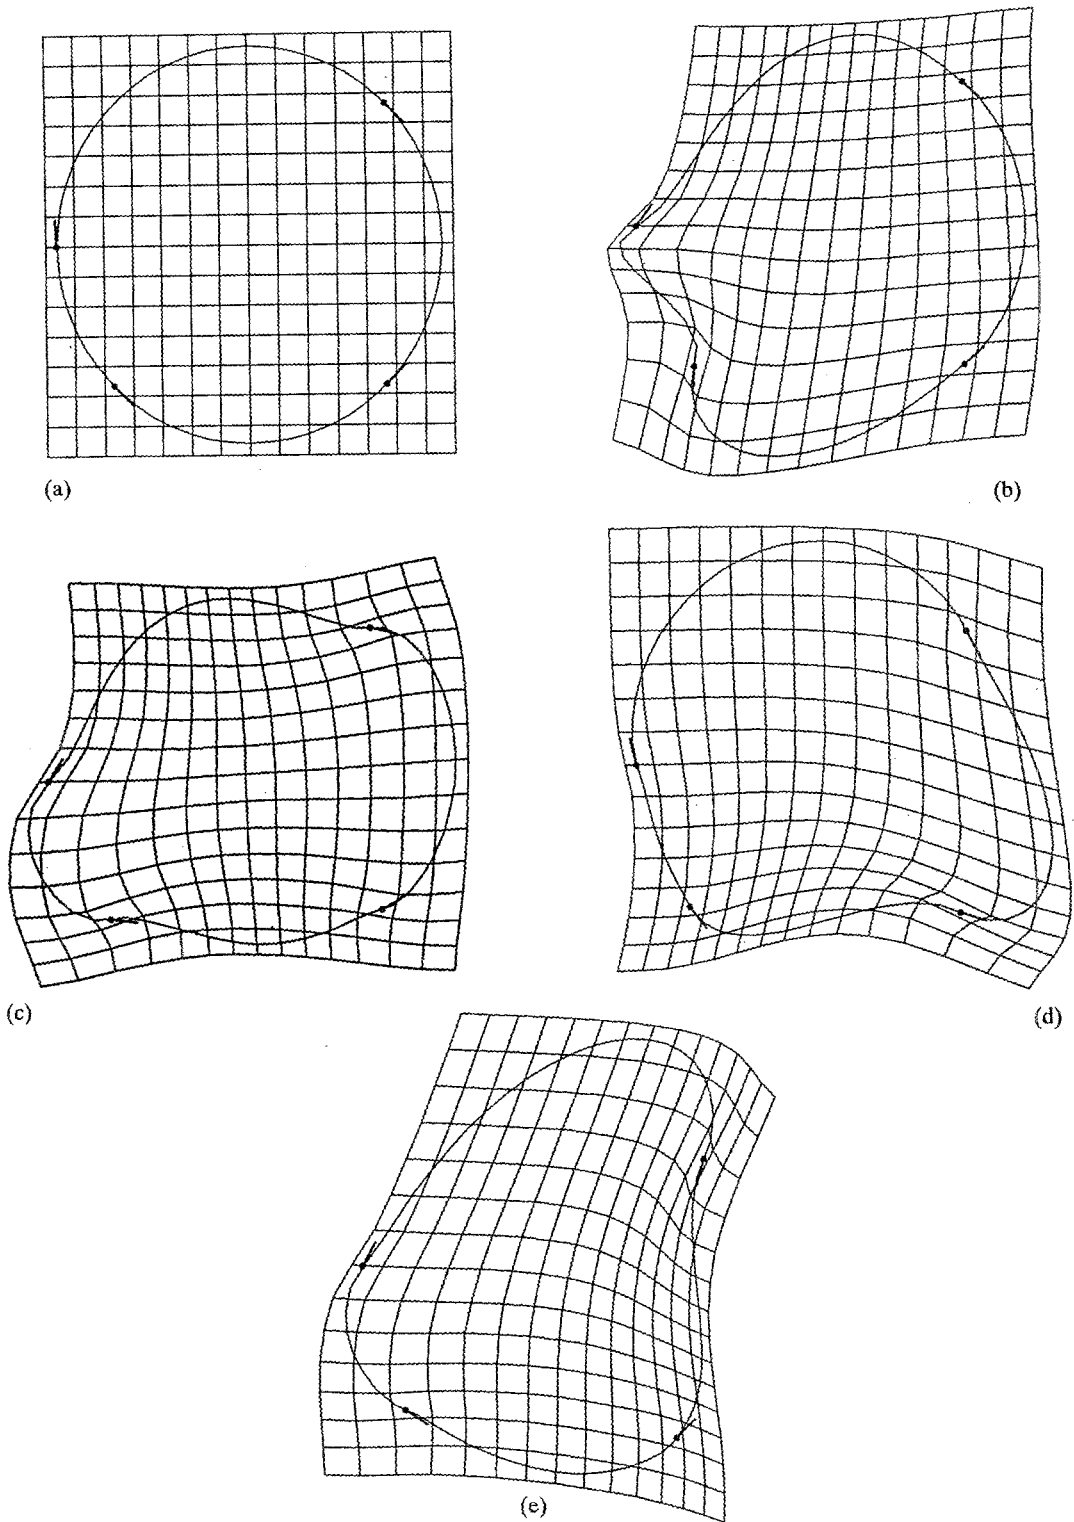
Fig. 20.

landmark position and edge orientation — or we might have explored other structural–functional associations, as between MRI shape and PET activation patterns.

7 Discussion

7.1 Warps, Templates, and Edgels

For a collection of images having a well-characterized mean landmark configuration P_1 $\cdots$ P_k and typical orientations of edge elements $t_1 \cdots t_m$ at some subset of these landmarks, the elementary edgel warps F' of relation (21), appropriately reoriented, serve as a finite basis spanning the feature space for edgels promised in our title. Easily evaluated by finite matrix operations, they apply quite flexibly for the parametric instantiation of arbitrarily detailed templates deriving from typical or validly averaged images rather than from simple geometric models. The principal edgel warps provide an alternative basis, an orthogonal rotation of the
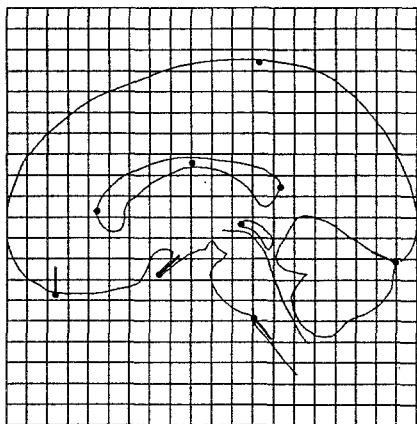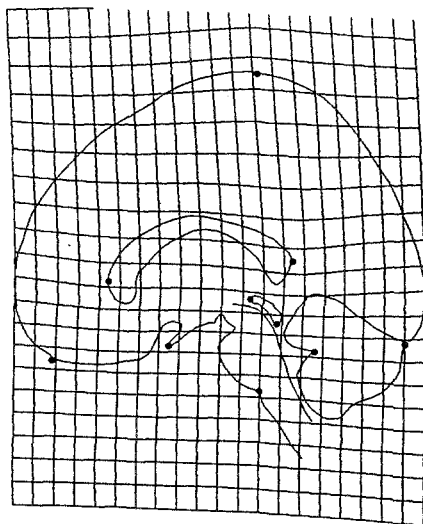
Fig. 21.
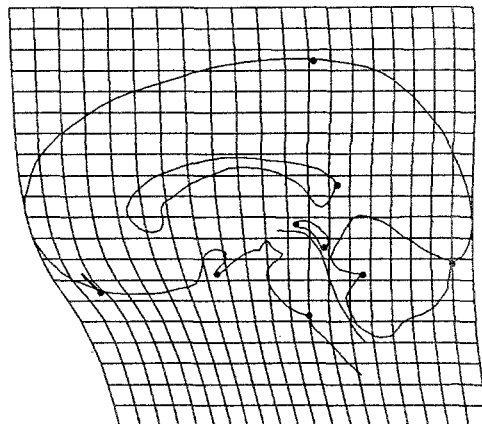
Fig. 22.

Fig. 23.

collection of separate edgels organized by effective geometrical scale.

The same technology that warps templates can be applied to unwarp specimen images into a normative coordinate system in which gray-level features can be extracted with great precision of biological localization. Furthermore, all the information used in these unusually flexible warping operations may be expressed in a finite-dimensional linearized feature space, spanned by principal warps and principal edgel warps; with the aid of this vector space any warping may be run in reverse to standardize the geometry of a specimen image without otherwise adjusting its gray levels. (Of course, if a physical theory calls for the adjustment of those gray levels, so as to conserve total counts, such adjustments can be incorporated in the course of the unwarping.) The feature spaces for geometry meanwhile archive that part of the information in a form usable for conventional multivariate analyses: these include assessments of correlations between landmark locations and outline form as well as the more usual case-oriented group comparisons, discriminant-function computations, structural–functional correlations, and the like.

Equations (33) and (51) allow any statistical signal extracted as a vector of coefficients of these warps—that is, any linear combination of principal warps and principal edgel warps—to be visualized as an incremental deformation all its own, rearranging landmarks and shearing edgels at those rearranged positions. For applications to automatic feature extraction from images processed individually, these same parameters supply a convenient basis in which to express posterior distributions that summarize the expected values and the expected covariances of these descriptors. Sample covariance matrices among the principal edgel warp scores may be summarized by a set of empirical orthogonal functions of their own (principal com-
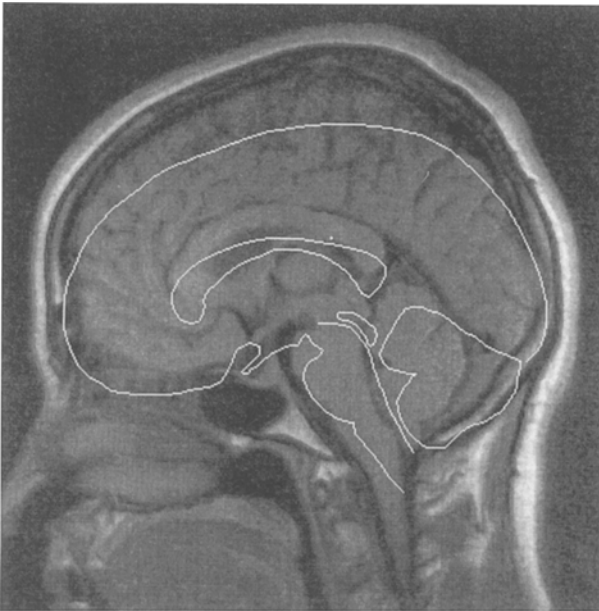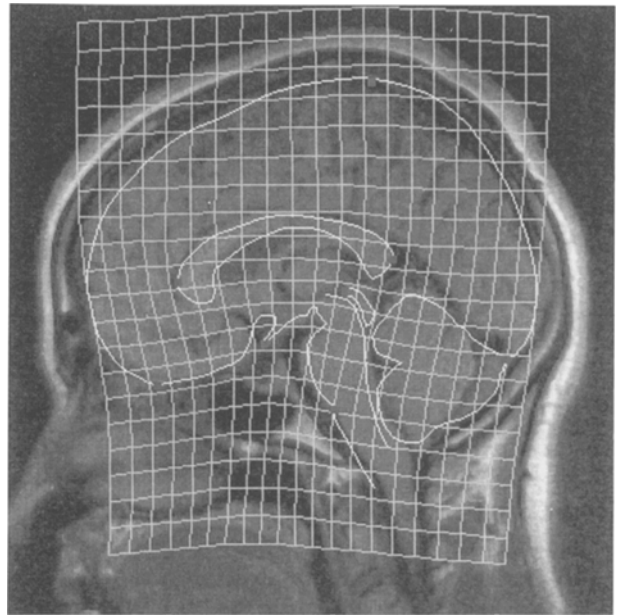
Fig. 24

Fig. 25.

ponents with respect to bending energy) in a manner parallel to that for any other system of empirical orthogonal functions.

At the conclusion of section 6 we reported the match of the original landmark-labeled outline to each of the 14 specimen images, together with a pooled feature space for the representation of the requisite deformations (a composite list of all the landmarks and all the edgels that had to be used). Although many other methods of deformable templates exist for instantiating a template to arbitrary accuracy, the method of warps demonstrated here is the only one known to us that allows for a *parametric reversion*, the computation of a rigorously defined inverse mapping that restores the geometry of the specimen image to the normative configuration up through terms linear in all the relevant parameters of those deformations. Likewise, we are not aware of any other method of freely deformable templates that, while permitting maps of quite diverse and spatially distributed complexity, nevertheless is capable of expressing all the members of a sample of these maps in a shared linearized feature space of relatively

low dimension. The method of edgel-driven splines here supports not only a flexible algebra of deformations but also a complete multivariate statistical method.

The role of the landmark-driven thin-plate spline in image normalization has been noted previously by us [16], [17] and by others [18]. When such unwarping splines are calibrated only by landmarks, however, there is variation remaining in the directions of edge elements through those landmarks. When equation (51a) is used to sharpen the unwarping mapping, the further calibration afforded by the edgels can only improve the precision of these averaged images, thus further reducing the sample sizes required to show statistical significance for clinical or experimental group differences. These edge-driven splines are much more flexible for such purposes than are families of transformations fixed by functional form in advance of encountering data, such as the polynomial warping functions of Sneath [19] or Greitz et al. [20] We are very enthusiastic about their suitability as a routine or utility unwarping method for a great variety of quantitative investigations into
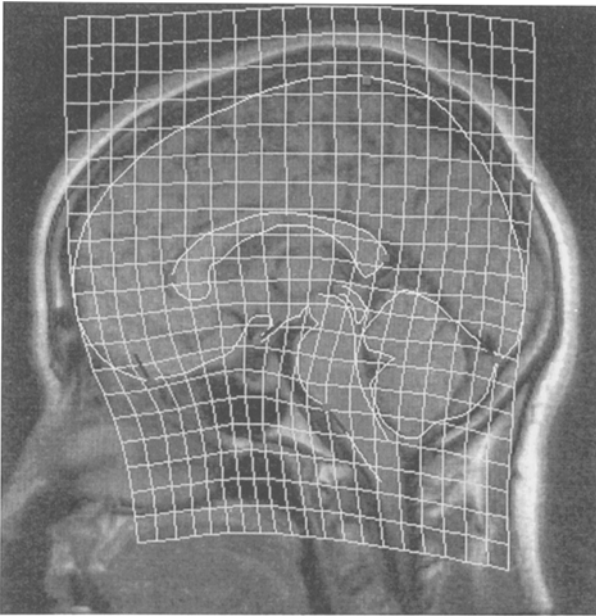
Fig. 26

Fig. 27.

medical images.

7.2 Vertical and Horizontal Visualizations of Images with Landmarks

The problems we are pursuing originate in data that are already visualized. The source of information was a medical image, indicating a physical property (some sort of interaction with radiation) within each of a grid of little volumes inside a region of tissue. Our scientific concern is to investigate aspects of this sort of data many sets at a time. The visualization we seek is to concentrate certain features of particular interest out of the rather dilute, hard-to-compare information that is each original brain scan. To emphasize this task of concentration rather than the pursuit of arbitrary detail is to ask a different sort of question from that the original visualization was designed to answer. The goal now is to retrieve not what is unique of each instance but what is common to all, what is most variable among them, what typically covaries with exogenous causes or effects, etc.

The data under discussion are, in general, vectors at each point of a domain organized on Euclidean principles — multispectral pixels or

voxels. To this basis, linear operators deal most comfortably with variances and covariances pixel by pixel; it is rather difficult to express adjacencies, and it is nearly impossible to denote families of warps, even with the aid of shift operators. Let us refer to this orientation of linear statistics, that which uses pixel locations for its index set, as *vertical*. There is always additional information, then, in the *horizontal* part of this imagined figure — the information about where the labeled locations and gradients of the ground plane actually lie with respect to the pixels and how their configurations covary with height(s) of the surface(s) above them. This horizontal part, born in and borne on the non-linear world of the retina, is what we primates are used to processing. Indeed, the abstract linear visualization of the vector space indexed by pixels hardly deserves the name *visualization* at all. Whenever data are originally visual, and especially if they are originally gridded, the linear machinery *must* be supplemented, if not wholly supplanted, by a semantics of deformation.

Note that the labels attach to the points and directions of the ground plane (or ground space in the three-dimensional applications), not to the points of the imaginary data surfaces float-
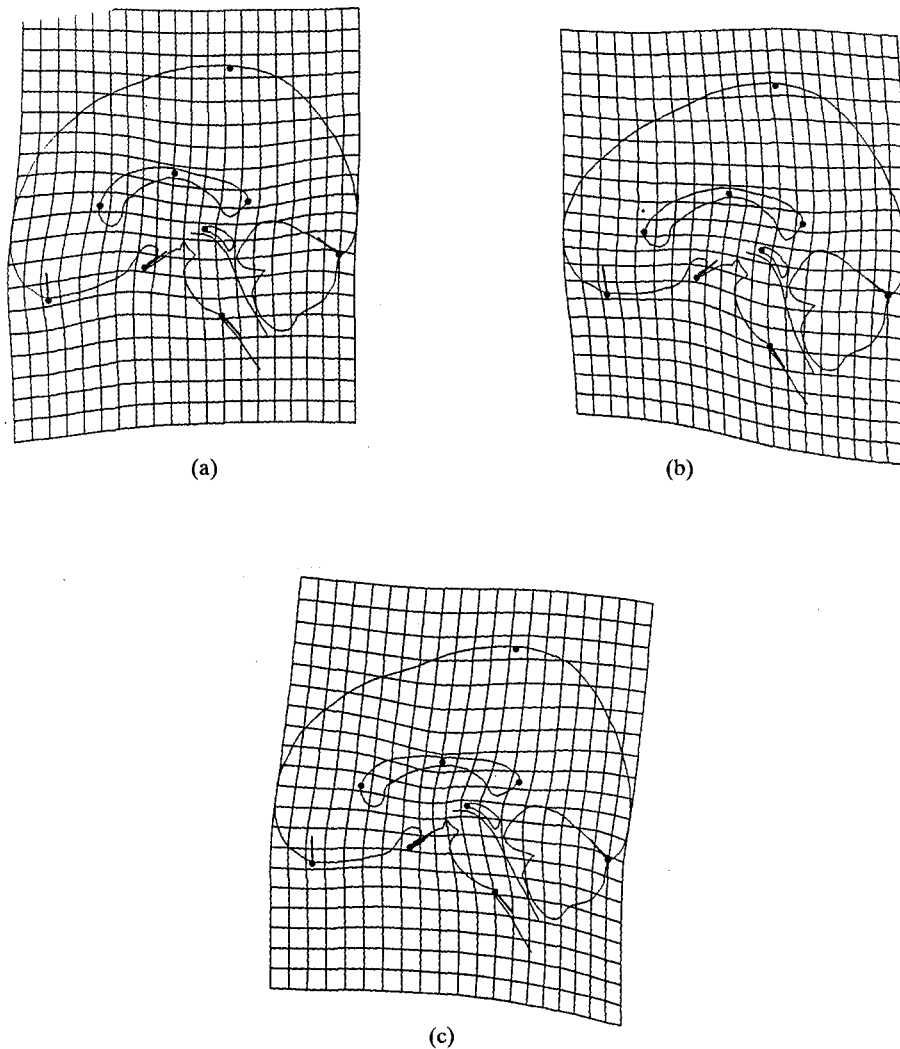
(a)   (b)

(c)

Fig. 28.

ing above them. These labels, which, of course, identify the landmarks and edgels of the earlier discussion, have their own metrics, including the versions of bending energy reviewed here, that complement the usual covariance-based metric of the multivariate observations above them. That is, the labeled points and directions can move about in their Euclidean domain at the same time that images change above them, leading to decompositions of the variance at a point that are very interesting both scientifically and statistically.

For instance, a vertical analysis may be best if one wants to use the geometry of the labeled image rather as one uses a covariate in a classic experimental design. In this case it is as if the shape of the configuration of labeled points — fluctuation in biological meaning of the basis for the vector space underlying the data — is to be treated as nuisance variation. Controlling this variation increases the precision with which other effects can be addressed. That is, one analyzes vertically — examining the gradients of the picture, for instance, or its correlations with physical or biological processes — only after unwarping horizontally to a more focused

feature space in which processes more nearly stay put to have their statistically standardized pictures taken. In multivariate language we are projecting out a complicated nonlinear feature co-space. The experience of generations of anatomists shows how this maneuver improves the power of subsequent multivariate tactics, such as discrimination or analysis of covariance. When averaging pictures of brain activity over brains of different shapes, for instance, the landmarks serve as guides to the correspondence of regions (the atlas) before averaging. It is the atlas, not the squares of the grid of a PET reconstruction, that represents the true coordinate system for valid biometric analyses.

This much is often conceded. Less often acknowledged is that in most applications this horizontal variation is not noise or nuisance but is instead a signal in its own right. In the vicinity of their mean configuration, the labeled points and edgels induce a very powerful low-dimensional feature space. With the aid of a convenient basis for shape variation this information may be concentrated into linear features of its own. The variables of this block ought to be paired with less delicately crafted descriptors of the original vertical scalar or vector content for prediction of other images, such as later images of the same system, and for the joint evocation of shape and content as a bispectral signal in a detection or classification problem, such as locating tumors or quantifying their recession under treatment.

We believe we have shown that visualization of vectors in this feature space of edge information at landmarks is at least as straightforward as visualization of changes in surfaces above the planes or volumes tagged by those points. The best visualizations are suggestive of the process explanations automatically familiar to any sentient organism that ever navigated a binocular landscape. We read the grids of figures 7, 8, 10, and the like as easily as we can read simulations of surfaces by gray-scale shading, the cues of barrels or pincushions, or any other pictorial metaphor of a spatial field. The combination of features of labeled point shape with features of the image at the average shape — the careful separation of vertical from horizontal variation in these mixed feature spaces, along with the care-
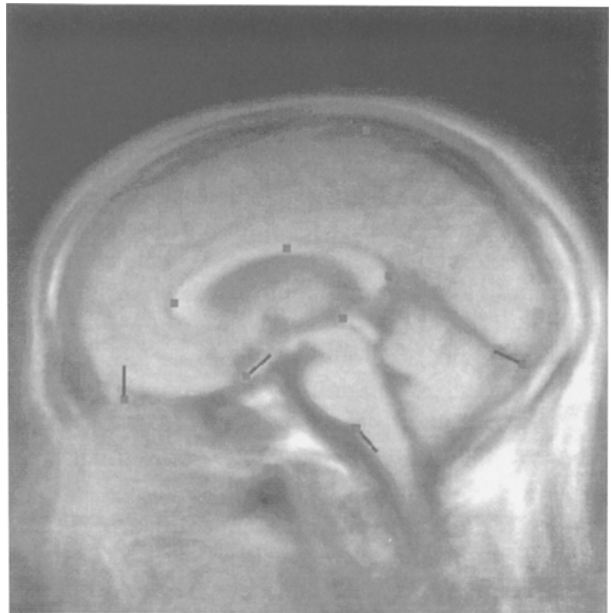

Fig. 29.

ful, specialized visualization of the horizontal — is, in our view, the most powerful discipline available for scientifically effective analyses of biometrical images. For instance, many transformations that appear hopelessly nonlinear in terms of the multivariate space oriented vertically turn out to be linear, or nearly so, in aspects of the same space viewed and measured horizontally. The averaging of biological images is made vastly more effective when these nonlinear transformations are executed first, each in its own linearized domain. In practice, the two domains of linearization, and the resulting two complementary styles of visualization, combine for composite analyses of changes over a deforming scene, as when a growing tumor changes texture while it deforms surrounding tissues or when a heart contracts as it bounces on its tether within the chest wall. The same duality applies also in analysis of image variability, as of structure–function correlations, unrelated to physical dynamics. Whether the problem is dynamic or cross sectional, the proper visualization of quantitative biometrics involves both kinds of linear operators.

The power for image analysis and scientific

insight of the new methods that exploit labeled point and edge data to enrich the conventional multivariate metaphor is thus immense, and it is mostly unexploited. In our view, the interaction of real nonlinear geometry with linear feature extraction is the key to many presently intractable problems of pattern detection and display across the medical imaging sciences. The crux of this combination of vertical and horizontal descriptions is their careful separation to begin with: separation of change of image content from deformation of the index set of pixels or voxels. The separation proceeds best with the aid of a unique intermediate linearized structure, the nonstandard multivariate technology of labeled point and edge configurations. In its peculiar finite-dimensional elegance, this space affords two bases at once: one for linear features of the arbitrarily nonlinear transformations that we see in the real world and the other for the pictorial features of the normalized images that result from these very helpful nonlinear transformations.

Acknowledgment

References

1. J. Duchon, "Interpolation des fonctions de deux variables suivant la principe de la flexion des plaques minces," *RAIRO Anal. Numér.*, vol. 10, pp. 5–12, 1976.
2. J. Meinguet, "Multivariate interpolation at arbitrary points made simple," *Z. Ange. Math. Phys.* vol. 30, pp. 292–304, 1979.
3. D. Terzopoulos, "Multilevel computational processes for visual surface reconstruction," *Comput. Vis., Graph., Image Process.*, vol. 24, pp. 52–96, 1983.
4. A. Witkin, D. Terzopoulos, and M. Kass, "Signal matching through scale space," *Intl. J. Comput. Vis.*, vol. 1, pp. 133–144, 1987.
5. F.L. Bookstein, *The Measurement of Biological Shape and Shape Change*, Springer-Verlag: Berlin, 1978, chap. 5.
6. F.L. Bookstein, *Morphometric Tools for Landmark Data*, Cambridge University Press: New York, 1991.
7. F.L. Bookstein, "Principal warps: thin-plate splines and the decomposition of deformations," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. PAMI-11, pp. 567–585, 1989.
8. S. Timoshenko and S. Woinowsky-Krieger, *Theory of Plates and Shells*, 2nd ed., McGraw-Hill: New York, 1959.
9. C.R. Rao, *Linear Statistical Inference and Its Applications*, 2nd ed., John Wiley: New York, 1973, p. 33.
10. S. Timoshenko and S. Woinowsky-Krieger, *Theory of Plates and Shells*, 2nd ed., McGraw-Hill: New York, 1959, pp. 287–289.
11. F.L. Bookstein and W.D.K. Green, "A feature space for derivatives of deformations," in *Proceedings of the Thirteenth International Conference on Information Processing in Medical Imaging*, H.H. Barrett and A.F. Gmitro, eds., Lecture Notes in Computer Science, Springer-Verlag: Berlin, 1993, to appear.
12. F.L. Bookstein and W.D.K. Green, "Edge Information at Landmarks in Medical Images," Biomedical Communications, Univ. of Michigan, Ann Arbor, MI, videotape BMC 818, 1992; also presented at the 1992 Conf. on Visualization and Biomedical Computing, Chapel Hill, NC, October 1992.
13. K.V. Mardia, J. Kent, and A. Walder, "Statistical shape models in image analysis," in *Computing Science and Statistics: Proc. 23rd Conf. on the Interface*, E.M. Keramidas, ed., Interface Foundation of North America: Fairfax Station, VA, 1991, pp. 550–557.
14. U. Grenander, Y. Chow, and D. Keenan, *Hands: A Pattern Theoretic Study of Biological Shapes*, Springer-Verlag: Berlin, 1991.
15. F.L. Bookstein and W. Jaynes, "Thin-Plate Splines and the Analysis of Biological Shape," Biomedical Communications, Univ. of Michigan, videotape BMC 650, 1990; also presented at 1st Conf. on Visualization in Biomedical Computing, Atlanta, GA, May 1990.
16. F.L. Bookstein, *Morphometric Tools for Landmark Data*, Cambridge University Press: New York, 1991, appendix 1.3.
17. F.L. Bookstein, "Thin-plate splines and the atlas problem for biomedical images," in *Information Processing in Medical Imaging*, A.C.F. Colchester and D. Hawkes, eds., Springer-Verlag: Berlin, 1991, pp. 326–342.
18. A.C. Evans, W. Dai, L. Collins, P. Neelin, and S. Marrett, "Warping of a computerized 3-D atlas to match brain image volumes for quantitative neuroanatomical and functional analysis," *Proc. Soc. Photo-Opt. Instrum. Eng.*, vol. 1445, pp. 236–246, 1991.

19. P.H.A. Sneath, "Trend-surface analysis of transformation grids," *J. Zool.*, vol. 151, pp. 65–122, 1967.
20. T. Greitz, C. Bohm, S. Holte, and L. Eriksson, "A computerized brain atlas: construction, anatomical content, and some applications," *J. Comput Assist. Tomog.*, vol. 15, pp. 26–38, 1991.
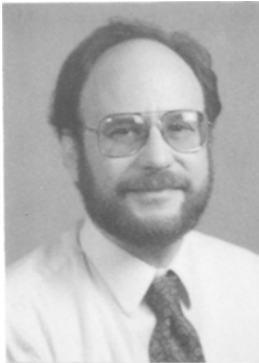
Fred L. Bookstein, Distinguished Research Scientist at the University of Michigan, received his Ph.D. in statistics and zoology from the University of Michigan in 1977. He is the author of more than 100 scientific papers, including two on videotape, and is author, co-author, or editor of five books dealing with aspects of image measurement, morphometrics, multivariate statistics, and measurement of complex systems in the biological sciences.
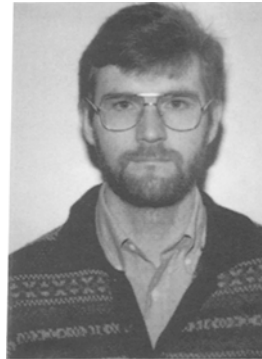
William D.K. Green received his Ph.D. in algebraic geometry in 1988 from Princeton University. He is presently a research mathematician at the Center for Human Growth and Orthopaedic Research Laboratory at the University of Michigan. Bookstein and Green have collaborated since 1991 on extensions of the thin-plate spline to incorporate differential information and on the software that produced all the pictures of this article by straightforward screen interactions.