

# Segmentation and Interpretation of MR Brain Images: An Improved Active Shape Model

Nicolae Duta and Milan Sonka,\* *Member, IEEE*

**Abstract**—This paper reports a novel method for fully automated segmentation that is based on description of shape and its variation using point distribution models (PDM's). An improvement of the active shape procedure introduced by Cootes and Taylor to find new examples of previously learned shapes using PDM's is presented. The new method for segmentation and interpretation of deep neuroanatomic structures such as thalamus, putamen, ventricular system, etc. incorporates *a priori* knowledge about shapes of the neuroanatomic structures to provide their robust segmentation and labeling in magnetic resonance (MR) brain images.

The method was trained in eight MR brain images and tested in 19 brain images by comparison to observer-defined independent standards. Neuroanatomic structures in all testing images were successfully identified. Computer-identified and observer-defined neuroanatomic structures agreed well. The average labeling error was  $7\% \pm 3\%$ . Border positioning errors were quite small, with the average border positioning error of  $0.8 \pm 0.1$  pixels in  $256 \times 256$  MR images.

The presented method was specifically developed for segmentation of neuroanatomic structures in MR brain images. However, it is generally applicable to virtually any task involving deformable shape analysis.

## I. INTRODUCTION

**A**UTOMATIC detection of brain structures is motivated by an ongoing effort to advance knowledge about relationships between anatomy and mental diseases in human brains. To better understand how human brain works and fails, an immense number of brain data sets from various sources [structural and functional magnetic resonance (MR), positron emission tomography (PET), single photon emission tomography (SPECT), etc.] must be analyzed. To date, most of the segmentation work in neuroscientific studies relies on manual tracing. As such, intraobserver and interobserver variability, and small study sample sizes remain a limiting factor. Despite the substantial interest and effort in the development of automated methods to identify brain structures [1]–[14], no reliable automated tools are available to date.

Manuscript received December 24, 1997; revised August 19, 1998. This work was supported in part by the National Science Foundation (NSF) under Grant IRI 96-16747. The Associate Editor responsible for coordinating the review of this paper and recommending its publication was L. P. Clarke. *Asterisk indicates corresponding author.*

N. Duta is with the Department of Computer Science, Michigan State University, East Lansing, MI 48823 USA.

\*M. Sonka is with the Department of Electrical and Computer Engineering, 4400 EB, The University of Iowa, Iowa City IA 52242 USA (e-mail: milan-sonka@uiowa.edu).

Publisher Item Identifier S 0278-0062(98)09425-7.

Building models describing shape and appearance of non-rigid objects, and employing them for automated identification of such objects in the analyzed images has been of substantial research interest over the past several years [15]–[18]. Among them, the point distribution model (PDM) representing the variation of a set of shapes around the average was designed by Cootes and Taylor that has many favorable shape representation properties [19]–[21].

The goal of our work reported here is twofold. The first goal was to develop a practically applicable MR brain segmentation method. The second goal was to improve the *active shape* procedure introduced by Cootes and Taylor [21] to find new examples of previously learned shapes using the PDM's. The automated brain segmentation and interpretation approach described below represents a novel approach that incorporates *a priori* knowledge about neuroanatomic structures, their context, shapes, and shape variation to provide robust segmentation and labeling. The method presented below was specifically developed for segmentation of neuroanatomic structures in MR brain images. However, it is generally applicable to virtually any task involving deformable shape analysis.


## II. PRELIMINARIES

### A. The Statistics of Shape

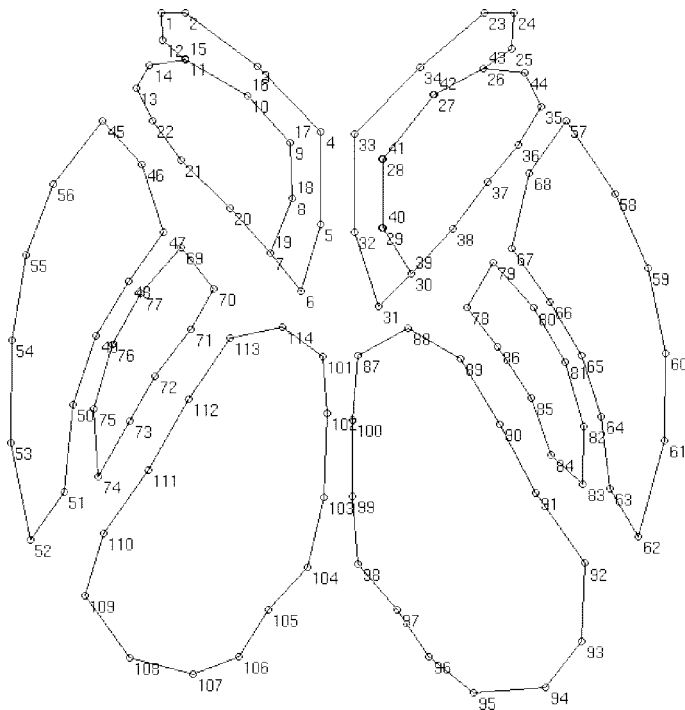
Let *shape instance* denote a set of points in the Euclidean plane describing an object contour. If there is no connectivity information, a shape instance is an amorphous set of points similar to that shown in Fig. 1(a). By adding information about the number of contours and the point sequential order, the point set may become the shape model [Fig. 1(b)]. A more detailed description of the anatomical significance of the points and the model construction methodology is given in Section IV.

*Shape* is defined to be the equivalence class of a shape instance, within the collection of all point sets of the same cardinality, under the operation of the similarity group (translation, rotation and scale). It is worth noting that an element of a shape class (a shape instance) is uniquely determined by the positions of two of its points since a planar linear transformation has exactly four parameters which can always be computed by fixing the coordinates of two points (imposing four constraints).

The *squared shape distance* between one point set (shape instance)  $A$  and another point set (shape instance)  $B$  is defined as the minimum sum of squared Euclidean distances between



(a)



(b)

Fig. 1. Landmarks of deep neuroanatomical structures. (a) MR image of the brain, coronal slice, resolution  $256 \times 256$ , slice thickness 1.5 mm, with 114 landmark points superimposed. (b) A 114-point shape model of ten brain structures.

Mathematically speaking, let  $A = \{(x_{Aj}, y_{Aj})\}_{j=1, \dots, n}$  and  $B = \{(x_{Bj}, y_{Bj})\}_{j=1, \dots, n}$ . Then a set of points  $C$  has the same shape as  $B$  ( $C \in \text{shape}(B)$ ) if and only if there exists a linear transformation (a composition of rotation, translation, and scaling) that maps the coordinates of  $B$  into the coordinates of the corresponding points in  $C$ . If we apply a rotation with an angle  $\theta$ , a scale  $s$ , and a translation  $(b, d)$  to  $B$ , its new coordinates are given by

$$x_{Cj} = s \cos(\theta)x_{Bj} - s \sin(\theta)y_{Bj} + b = ax_{Bj} - cy_{Bj} + b \quad (1)$$

$$y_{Cj} = s \sin(\theta)x_{Bj} + s \cos(\theta)y_{Bj} + d = cx_{Bj} + ay_{Bj} + d. \quad (2)$$

The *squared shape distance*  $\mu^2$  between shape instances  $A$  and  $B$  is defined as

$$\begin{aligned} \mu^2(A, B) &= \min_{C \in \text{shape}(B)} \sum_{j=1}^n [(x_{Aj} - x_{Cj})^2 + (y_{Aj} - y_{Cj})^2] \\ &= \min_{a, b, c, d \in \mathbb{R}} \sum_{j=1}^n [(x_{Aj} - ax_{Bj} + cy_{Bj} - b)^2 \\ &\quad + (y_{Aj} - ay_{Bj} - cx_{Bj} - d)^2]. \quad (3) \end{aligned}$$

One can prove that  $\mu^2$  has the following properties.

- 1) The quantity  $\mu^2$  is well defined for every two shape instances  $A$  and  $B$ ;
- 2)  $\mu^2$  is not symmetric, that is,  $\mu^2(A, B) \neq \mu^2(B, A)$  [22], [23], but if  $A$  and  $B$  have roughly the same scale then  $\mu^2(A, B) \approx \mu^2(B, A)$ ;
- 3) since  $\mu^2(A, B) = \mu^2(A, C)$  for every  $C \in \text{shape}(B)$ , it can be considered the *distance* between the fixed shape instance  $A$  and the shape class of  $B$ .

Note that the point-coordinate averaging method of adding and dividing by their count cannot be used to average shapes. However, the alternate definition can be used in which the average is defined as the quantity about which the numbers have the least sum of squared distances. In this way, a notion of shape averaging using the least squared distance characterization is inherited.

### B. Shape Averaging

There is a large number of papers dealing with averaging/aligning shapes into a common coordinate frame. While the most popular approach in the statistical/medical vision communities is the Procrustes Analysis [24]–[26], the work of Horn [22], [23] is better known to nonmedical vision researchers. Both frameworks provide analytic solutions to the shape alignment problem as well as approximate solution algorithms. The statistical shape literature mostly deals with theoretical considerations while the practical side of the problem is frequently left aside. Several iterative algorithms for aligning a set of shapes exist in

the points of  $A$  and the corresponding points in sets  $C$  as  $C$  ranges over the whole set of shape instances equivalent to  $B$ .

the medical vision literature [21], [27], [28]. However, the papers do not state how well they approximate the closed-form solution. Therefore, our approach is not iterative and is based on deriving a linear system of equations for computing the *least squares fit* to a set of shape instances.

Let  $m$  denote the number of shape instances, let  $n$  be the number of points of each instance, and let  $(x_{ij}, y_{ij})$ ,  $1 \leq i \leq m$ ,  $1 \leq j \leq n$  be the coordinates of the  $j$ th point of the instance  $i$ . The least squares fit of the  $m$  instances is sought, that is, a new shape instance  $S_0 = (x_{0j}, y_{0j})_{j=1 \dots n}$  should be determined for which the sum of the squared shape distances to each of the  $m$  shape instances is minimum. Note that even if  $S_0$  is unknown, it is a fixed shape instance. Therefore, its scale and pose must be fixed. For instance, it can be considered aligned with one of the  $m$  given shape instances, let's say the last one. Then, the distance sum  $f$  between  $S_0$  and the  $m$  instances is given by

$$f = \min_{a_i, b_i, c_i, d_i \in \mathbb{R}} \sum_{i=1}^{m-1} \sum_{j=1}^n \left[ (x_{0j} - a_i x_{ij} + c_i y_{ij} - b_i)^2 + (y_{0j} - a_i y_{ij} - c_i x_{ij} - d_i)^2 \right] + \sum_{j=1}^n [(x_{0j} - x_{mj})^2 + (y_{0j} - y_{mj})^2] \quad (4)$$

where  $a_i, b_i, c_i, d_i$ ,  $1 \leq i \leq m-1$  are the unknown parameters of the linear transformation that maps each of the first  $m-1$  instances into the instance of the corresponding shape which is closest to the least squares fit instance  $S_0$ . One can prove that if all  $m$  shape instances have roughly the same scale (for example are rescaled in such a way that the sum of the squared distances from the centroid to the vertices is equal to 1), it does not matter which of the  $m$  shape instances is chosen to align  $S_0$  with, since all the shape instances that minimize equations similar to (4) are equivalent (differ only by a linear transform). Practical results showed that this remains true even when the shape instances have quite different scales. Moreover, most of the time the square distance from  $S_0$  that minimizes (4) to the set of shape instances from which it is derived is smaller than that produced by the iterative algorithms actually employed in the literature [21], [28], [29].

Now,  $S_0$  is sought such that its coordinates minimize  $f$ . By equating the partial derivatives of  $f$  to zero, a linear system of  $4(m-1) + 2n$  equations in  $4(m-1) + 2n$  unknowns is obtained:

$$MX = T \quad (5)$$

where

$$M = \begin{pmatrix} \alpha_1 & 0 & \cdots & 0 & \beta_1 \\ 0 & \alpha_2 & \cdots & 0 & \beta_2 \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & \alpha_{m-1} & \beta_{m-1} \\ \gamma_1 & \gamma_2 & \cdots & \gamma_{m-1} & \delta \end{pmatrix}$$

with

$$\begin{aligned} \alpha_i &= \begin{pmatrix} A_i & B_i & 0 & C_i \\ 0 & -C_i & A_i & B_i \\ B_i & n & -C_i & 0 \\ C_i & 0 & B_i & n \end{pmatrix} \\ \beta_i &= \begin{pmatrix} -x_{i1} & \cdots & -x_{i,n} & -y_{i1} & \cdots & -y_{i,n} \\ y_{i1} & \cdots & y_{i,n} & -x_{i1} & \cdots & -x_{i,n} \\ -1 & \cdots & -1 & 0 & \cdots & 0 \\ 0 & \cdots & 0 & -1 & \cdots & -1 \end{pmatrix} \\ \gamma_i &= \begin{pmatrix} x_{i1} & 1 & -y_{i1} & 0 \\ \cdots & \cdots & \cdots & \cdots \\ x_{i,n} & 1 & -y_{i,n} & 0 \\ y_{i1} & 0 & x_{i1} & 1 \\ \cdots & \cdots & \cdots & \cdots \\ y_{i,n} & 0 & x_{i,n} & 1 \end{pmatrix} \\ \delta &= -mI_{2n} \\ T &= \begin{pmatrix} 0 & 0 & \cdots & 0 & -x_{m,1} & \cdots & -x_{m,n} \\ -y_{m,1} & \cdots & -y_{m,n} & & & & \end{pmatrix}^T \\ X &= \begin{pmatrix} a_1 & b_1 & c_1 & d_1 & \cdots & a_{m-1} & b_{m-1} & c_{m-1} \\ d_{m-1} & x_1 & \cdots & x_n & y_1 & \cdots & y_n \end{pmatrix}^T \\ A_i &= \sum_{j=1}^n x_{ij}^2 + \sum_{j=1}^n y_{ij}^2, \quad B_i = \sum_{j=1}^n x_{ij} \quad \text{and} \quad C_i = \sum_{j=1}^n y_{ij}. \end{aligned}$$

The *shape coordinates* of a set of shapes describe the variation of the whole set of shapes around the average in terms of separate variations at the component points. They are visualized in Fig. 2 by overlaying each individual shape over the average shape. It can be seen that some of these points show little variability around the average, while others form more diffuse clusters. Importantly, model points do not move independently within the clusters—their positions are partially correlated.

### C. Point Distribution Model

Each  $n$  vertex shape instance from the training set of examples represents a single point in the  $2n$ -dimensional ( $2n$ -D) space. Thus, a set of  $m$  shape examples produces a cluster of  $m$  points in this space. The *point distribution model* (PDM) describes the coordinates variation within the clusters assuming linear dependence.

The theoretical assumptions initially stated by Cootes and Taylor are [21] as follows.

- 1) The cluster points lie within an *allowable shape domain* (ASD) which is approximately ellipsoidal.
- 2) The training examples give an indication of size of the ASD.
- 3) The cluster center is specified by the points of the least squares fit instance shape ( $\bar{\mathbf{x}}$ ).
- 4) Every  $2n$ -D point within ASD represents a vertex sequence, the shape of which is broadly similar to the shapes of examples from the training set. Thus, by moving about the ASD, we can generate new shapes in a systematic way.

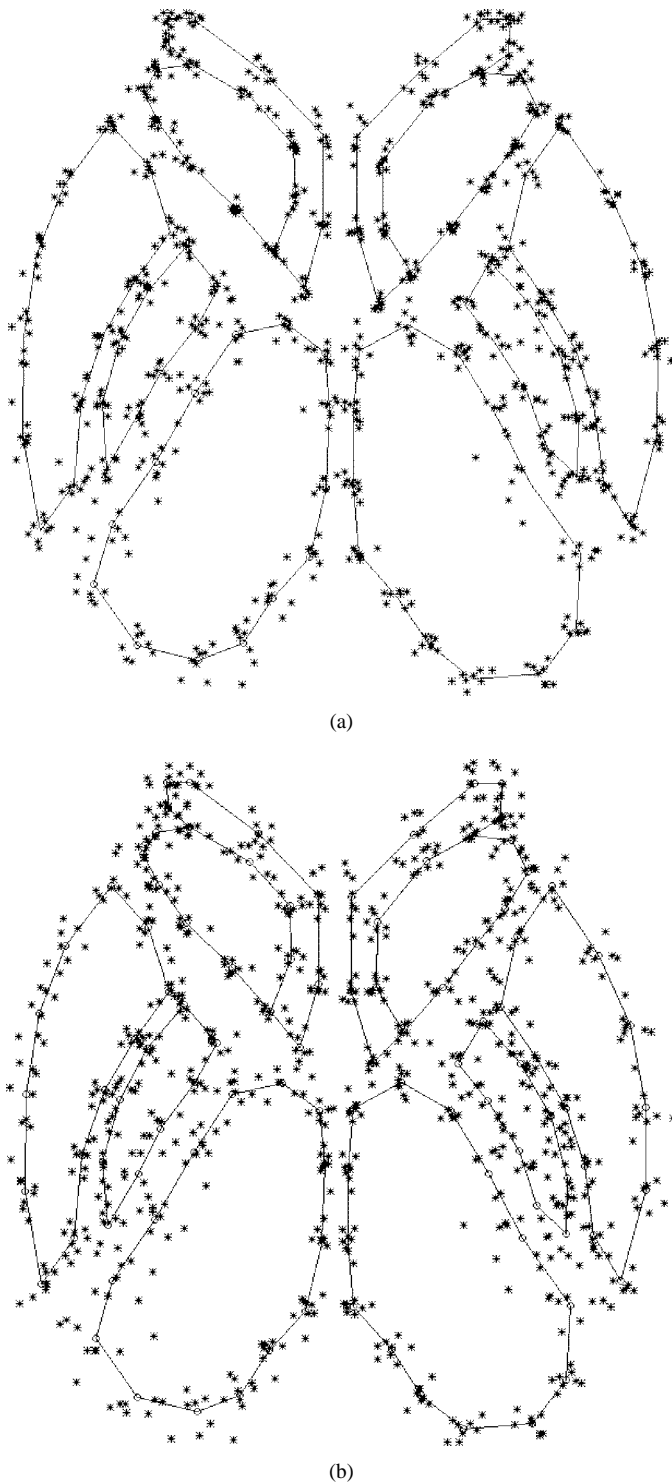


Fig. 2. Scatter of the points from aligned brain shapes, with the mean shape overlaid. (a) Each brain structure averaged separately. (b) All structures averaged simultaneously.

The principal axes of the  $2n$ -D ellipsoid are described by  $p_k$  ( $k = 1, \dots, 2n$ ), the unit eigenvectors of the covariance matrix  $S$  corresponding to  $\lambda_k$ , the  $k$ th eigenvalue of  $S$ . The principal axes can be calculated by employing the principal component analysis. Each axis gives a *mode of variation*, a way in which the landmark points tend to move together as the shape varies. The main properties of the ASD, also stated by Cootes and




Fig. 3. Effect of simultaneously varying the model's parameters corresponding to the first two largest eigenvalues (on a bidimensional grid).

Taylor, are as follows.

- 1) The eigenvectors corresponding to the largest eigenvalues describe the longest axes of the ellipsoidal ASD. Thus, they describe the most significant modes of variation in the variables used to derive the covariance matrix.
- 2) The variance represented by each eigenvector is equal to the corresponding eigenvalue.
- 3) Most of the variation can usually be represented by a small number of modes,  $t$ . This means that the  $2n$ -D ellipsoid can be approximated by a  $t$ -dimensional ellipsoid ( $t \leq 2n$ ), where  $t$  is chosen so that the original ellipsoid has a relatively small width along axes  $t + 1$  and above ( $t$  is the smallest number of modes such that the sum of their variances represents a sufficiently large proportion of the total variance  $\lambda_T$ )

$$\lambda_T = \sum_{k=1}^{2n} \lambda_k. \quad (6)$$

- 4) Any point in the ASD (any allowable shape) can be reached by taking the mean (center) cluster point and adding a linear combination of the eigenvectors. Thus, any shape in the training set can be approximated using the mean shape and a weighted sum of deviations obtained from the first  $t$  modes

$$\mathbf{x} = \bar{\mathbf{x}} + P\mathbf{b} \quad (7)$$

where

$$P = (p_1 \ p_2 \ \dots \ p_t) \quad (8)$$

is the matrix of the first  $t$  eigenvectors, and

$$\mathbf{b} = (b_1 \ b_2 \ \dots \ b_t)^T \quad (9)$$

is a vector of weights.

- 5) New examples of shapes similar to those in the training set can be generated by varying the parameters ( $b_k$ ) within suitable limits (Fig. 3).

Cootes and Taylor propose to derive limits for  $b_k$  by examining the distributions of the parameter values required




Fig. 4. Example of a globus-pallidus-like shape containing two outliers, points 1 and 9 and its approximations using: (a) eight examples and 95% of variance explained, (b) 16 examples and 95% of variance explained, and (c) 16 examples and 100% of variance explained. (d) With the absolute values of the  $b$  parameters and the corresponding  $3(\lambda)^{1/2}$  ranges for shapes given in (b)–(d) computed according to (10) and (16). The lengths of the  $b$  vectors correspond to the number  $t$  of modes used to explain the shape variance.




Fig. 5. Example of thalamus from (a) the training set, (b) average thalamus, and (c) athalamus-like shape containing two outliers (points 4 and 5), with the absolute values of the  $b$  parameters and the corresponding  $3(\lambda)^{1/2}$  ranges and maximum variance percentage values for shapes given in (b) and (c) computed according to (10), (16), and (20).

to generate the training set [21]. They recommend that  $b_k$ 's be chosen such that

$$-3\sqrt{\lambda_k} \leq b_k \leq 3\sqrt{\lambda_k} \quad (10)$$

since most of the population lies within three standard deviations of the mean.

#### D. Limitations of PDM's

While the PDM's have many favorable properties, they are not free of several limitations. The most important is given by the approximation introduced in (7) which is sensitive to both the percentage of variance explained by the  $t$  chosen modes and the number of examples in the training set. While the approximation of a shape using a number of examples that approaches the dimension of the ASD and 100% variation explained does not exhibit visible differences from the original shape, the approximations using less examples and variation have a quite different look (Fig. 4). Moreover, the approx-

imation of an unacceptable shape [Fig. 4(a)] may become acceptable if only a small number of examples were used to compute the approximation [Fig. 4(b)]. One can also provide an example of a shape instance that has all the  $b_k$  parameters smaller in absolute value than those of shape examples from the training set but are not acceptable since some of their points overlap (Fig. 5).

### III. PDM APPROACH TO MR BRAIN IMAGE INTERPRETATION

In many areas of image segmentation and interpretation, reliable *a priori* knowledge is available to help guide the image analysis process. For instance, in brain imaging, the image data are routinely represented in a normalized Talairach space [32]. Consequently, approximate positions of individual neuroanatomic structures can be determined. Knowledge about their sizes, shapes, gray-level appearance, etc. can be acquired from a training set of examples.

### A. Knowledge-Based Point Distribution Model


In order to improve the shape model and to take advantage of the available *a priori* knowledge, three additional features characteristic to MR brain images were included in the model: Gray-level appearance, border strength, and average position. We also used implicit knowledge about object context representing interrelationships of several objects.

Gray-level appearance  $g_j = (g_{j,1}, \dots, g_{j,n_p})$  is calculated in neighborhoods around each of the shape model points  $j$  [21]. It is determined for every shape model point  $j$  of each training image along a profile  $PP'$  of a constant length  $n_p$ , centered at the point  $j$  and perpendicular to the shape boundary [Fig. 6(a)]. Since the profiles vary with gray-level scaling, derivatives of the gray levels along each profile are determined and normalized. Thus, invariance to uniform gray-level scaling and gray-level offset results. If we denote by  $y_{j,1}, \dots, y_{j,n_p}$  the pixels along the profile  $PP'$ , then the  $k$ th element of the derivative profile is computed by  $g'_{j,k} = I(y_{j,k+1}) - I(y_{j,k-1})$ , and, after normalizing the profile,

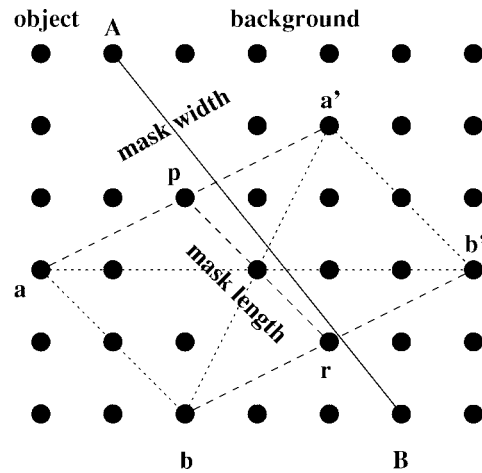
$$g_{j,k} = \frac{g'_{j,k}}{\sum_{k=1}^{n_p} |g'_{j,k}|}$$

Though the original PDM model [19]–[21] also incorporates second-order statistics of the gray-level appearance, it decreases the detection speed. At the same time, the additional variation information is only partially relevant since half of the profile spans the background which in many cases is different from case to case. Therefore, for each model point, only a mean profile  $\bar{g}$  is computed and incorporated into the model in our method.

Border strength is determined for each border segment of the model. Every two consecutive model points that lie on the object boundary define a border segment. To compute its strength, a local filtering is applied to each clique on that border segment. The designed filter was inspired by that of Geman and Jedyak [33]. If two pixels are located in close proximity and both belong to the same object, they are expected to have a smaller intensity difference than if one of them belongs to a different object or background [Fig. 6(b)]. The filter is based on a pair of close parallel profiles. Let  $aa'$  and  $bb'$  be two profiles approximately perpendicular to the border segment and centered at points  $p$  and  $r$  as shown in Fig. 6. The length of the parallel profiles (mask width) and the distance between them (mask length) are fixed when the model is defined. Due to the digitization, the actual width and length can vary slightly along the border depending on the orientation of the border segment. The overall response of the filter for a border segment is the aggregate of the individual responses of elementary cliques along that segment. An elementary clique is based on a pattern of four pixels  $(a, b, a', b')$ —see Fig. 6. The image response  $c(a, b, a', b')$  is given by (11), shown at



(a)



(b)

Fig. 6. Transversal profiles and edge cliques used to compute the gray-level appearance and edge strength: (a) contrast-enhanced to improve visibility and (b) schematic drawing of a clique.

the bottom of the page. Let  $I(t)$  denote the image intensity at pixel  $t$ , where  $\delta$  should reflect the relative border strength for a particular border segment. Therefore,  $\delta$  is determined as the largest integer such that the ratio of the overall border strength of the border segment and the maximum possible strength (twice the number of pixels in the border segment) is at least 0.66 (at least two-thirds of the cliques along an edge segment should return a nonzero response).

Average position of each shape model point that is calculated in the image coordinates is also incorporated in the model.

Interrelationships of several objects (object context) are implicitly included in the shape model. If more than one object occur in the same image, a decision must be made whether to include several of them in the same shape model

$$c(a, a', b, b') = \begin{cases} 2, & \text{if } \max(|I(a) - I(b)|, |I(a') - I(b')|) < \min(|I(a) - I(b')|, |I(a') - I(b)|) - \delta \\ 1, & \text{if only } \min(|I(a) - I(b)|, |I(a') - I(b')|) < \min(|I(a) - I(b')|, |I(a') - I(b)|) - \delta \\ 0, & \text{otherwise} \end{cases} \quad (11)$$

or whether to design an independent model for each of them. If contextual information is available, e.g., if it is known that two or more objects have a common border or are likely to appear in a prespecified relationship, combining them together may help overcome problems of their positioning or positioning of their mutual borders, especially if the borders are weak. On the other hand, it is not advisable to combine too many objects in a single object compound since the scattering of the model points in the training examples may be excessive [see, Fig. 2(b)]. Such scattering may induce a biased average and cause difficulty to recognize unallowed shapes. Therefore, a compromise solution was adopted in which some of the objects are grouped while others are treated independently. As a generalization, two objects are grouped if they have a common border (in MR brain images, examples include the ventricle and caudate nucleus) or if they are positioned in a close proximity and one of the objects exhibits a very weak boundary (e.g., putamen and globus pallidus).

To summarize, our *knowledge-based shape model* combines generally applicable parameters of the PDM and the knowledge-specific parameters appropriate for the image segmentation task in question. As such, the complete model is composed of the following.

- 1) The eigenvectors corresponding to the largest eigenvalues of the covariance matrix describing the ASD (Section II-C).
- 2) The average gray-level appearance values for each point of the model.
- 3) The average border strength for corresponding border segments and the parameters of the mask (width, length) for which the strength was computed.
- 4) The average position of the points of the average shape.
- 5) Connectivity information (the number of shapes, point ordering along contours).

### B. Searching for Objects: Model Fitting

The searching procedure developed for our PDM approach to image segmentation and interpretation is based on the model fitting strategy. Let us assume that shape models were previously constructed during a training stage in which manually traced examples of objects were utilized. Then, the goal of the image segmentation/interpretation stage is to identify objects with properties similar to those of the training objects in the analyzed images. To accomplish this goal, the shape models are fitted to the analyzed image data. Our model fitting procedure substantially differs from the *Active Shape Procedure* of Cootes and Taylor. The difference is twofold. First, our search is entirely model driven. At each step of the fitting process, several model location hypotheses are considered and evaluated. During the hypothesis generation, the actual image data play no role (no image preprocessing is done). Second, an outlier detection and replacement procedure has been developed to detect misplaced points and infer their new positions. The outlier detection improves robustness and accuracy of the shape model fitting process.

The searching procedure consists of the following steps: 1) model fitting using linear transformations, 2) model fitting

using piecewise linear transformations, 3) outlier removal, 4) final point adjustment, and 5) final outlier removal.

1) *Model Fitting Function*: As a result of the hypotheses generation processes, shape model locations are sequentially hypothesized. In order to evaluate the model location hypotheses, a *fitness function* is needed to assess the agreement between the image data and the particular model instance. We have designed a fitness function that consists of two components.

- a) *Fitness of the gray-level appearance*  $F_{GA}$  is determined as the average squared Euclidean distance between the actual gray-level appearance and the mean gray-level profile  $\bar{g}$  incorporated in the shape model (Section III-A)

$$F_{GA} = \frac{1}{n} \sum_{j=1}^n (g_j - \bar{g}_j)(g_j - \bar{g}_j)^T \quad (12)$$

where  $n$  is the number of vertices in the model,  $g_j$  is the normalized profile at point  $j$  in the current image, and  $\bar{g}_j$  is the mean normalized profile at point  $j$  computed from the training set.

- b) *Fitness of the border*  $F_B$  is calculated as the average over the contour of the individual border segment fitness. A border segment fitness is the ratio between the aggregate response of all cliques along the border segment and the maximum possible response (twice the number of cliques), where the response of a clique is given by (11) using the average  $\delta$  for that border segment as determined in the training set

$$F_B = \frac{1}{n} \sum_{j=1}^n F_j \quad (13)$$

where

$$F_j = \frac{\sum_{(a,b,a',b') \text{ along edge } j} c(a,b,a',b',\bar{\delta}_j)}{2 \times \# \text{ of cliques along edge } j}. \quad (14)$$

It is clear that the computed fitness should increase with either a smaller average profile distance or a greater border fitness. Since the border fitness is a weak qualitative property (it is satisfied at many positions in an image) and the gray-level appearance of the brain structures in the used MR images did not vary too much, one should prefer a smaller profile distance over a larger border fit. We computed two-dimensional fitness maps (similar to those in [35]-Fig. 2) over the entire image for several functional combinations of the two quantities on a large data set. The fitness map that exhibited a global, clearly delimited peak at the real position was produced by the following fitness function  $F$ :

$$F = \frac{F_B}{(F_{GA})^2}. \quad (15)$$

2) *Model Fitting Using Linear Transformations*: Shape instance hypotheses specify the locations of all model points within the analyzed image. The hypotheses are generated using linear (affine) transformations (rotation, translation, and scaling) and are applied to the model average position. The parameters of the linear transformations that contribute to the hypotheses generation are application dependent and reflect the




Fig. 7. The first two steps of a right ventricle segmentation. (a) Optimal average model position after linear transform step. (b) A five-point partial shape hypothesis is fitted to the image in a neighborhood of its actual position. (c) Optimal position of the partial ventricle hypothesis; only the position of the center vertex is kept. (d) Optimal model position after partial fitting of all model vertices. Images contrast-enhanced to improve visibility.




Fig. 8. Steps 3–5 of a right thalamus segmentation. (a) Optimal model position after piecewise linear transform. (b) Outlier detection—the two detected outliers marked by not connected white squares. (c) Final point adjustment. (d) Final outlier removal. Images contrast-enhanced to improve visibility.

*a priori* knowledge one has about the scale and pose of the objects of interest. In the brain images, the parameters were: scaling in the range  $[0.9, 1.1]$ , step 0.1; rotation  $[-8^\circ, 8^\circ]$ , step  $4^\circ$ ; and translation  $[-4, 4]$  pixels, step 1 pixel in both  $x, y$  directions. Each hypothesis represents a rigid transform of the average shape, and all generated hypotheses are sequentially evaluated using the model fitting function and the best fit is determined [Figs. 7(a) and 12(c)].


3) *Model Fitting Using Piecewise Linear Transformations:* Since nonrigid objects or objects with intersubject variability are discussed here, rigid linear transformations do not account for any potential deformations of the expected shape [Fig. 7(a)]. Therefore, in order to introduce deformations in our shape hypothesis, linear transformations (translation, rotation, and scaling) with a small range of parameters are applied to subsets of consecutive model points (subshapes of the model) [Fig. 7(b)]. Each model point is taken in turn as the center of such a subset. After the best pose of the subset is obtained (Fig. 7(c)), only the position of the center point is kept and the remaining points are discarded. The number of consecutive points that are considered for this transform is application dependent and is a function of local border strength

and length. To preserve robustness, the center point is not moved if the two border segments adjacent to it are very weak.

4) *Outlier Removal:* Under unfavorable circumstances, the previous step may introduce incorrectly positioned vertices—*outliers*. This may happen if a subshape fitted by the previous step exhibits weak edges [Fig. 8(a)] or if there exists another border of similar properties in the neighborhood. In the existing literature dealing with PDM's, the only study we are aware of to approach the outlier problem is [31]. A combination search-based solution was proposed that was based on a somewhat restricted assumption that the true position of a model point lies along the perpendicular profile through the current point position and constitutes a local peak that can be located by the gray-level model for that point. As noted by the authors, in many cases of high-curvature boundary points (and there are several such points in our brain model), the true position may not be located along the perpendicular profile and even a more time-consuming search would not help.

In most studies using PDM's, shapes that do not correspond to the allowed shape at any stage of the detection process are rejected or approximated with allowable shapes. In other






$$3 * \text{sqrt}(\lambda) = [ 6.81 \quad 4.66 \quad 3.07 \quad 2.54 \quad 1.80 \quad 1.63 ]$$

$$\text{abs}(b) = [ 1.87 \quad 1.61 \quad 0.99 \quad 1.28 \quad 1.32 \quad 1.92 ]$$

$$\text{Max. var. perc.} = [ 20 \quad 13 \quad 12 \quad 17 \quad 7 \quad 9 \quad 25 \quad 4 \quad 65 ]$$

Fig. 9. Example of a shape hypothesis rejected by the original active shape model (ASM) procedure (left). Note the one outlier responsible for rejection (marked by 9). The same shape hypothesis after our outlier removal step (right, adjusted vertex marked by \*). The average shape is shown in the middle. The values given below the figure are computed according to (10), (16), and (20) and correspond to the left-most shape.



$$3 * \text{sqrt}(\lambda) = [ 6.81 \quad 4.66 \quad 3.07 \quad 2.54 \quad 1.80 \quad 1.63 ]$$

$$\text{abs}(b) = [ 0.62 \quad 1.12 \quad 0.62 \quad 0.03 \quad 0.64 \quad 0.88 ]$$

$$\text{Max. var. perc.} = [ 58 \quad 30 \quad 23 \quad 15 \quad 6 \quad 10 \quad 13 \quad 8 \quad 25 ]$$

Fig. 10. Example of a shape hypothesis accepted by the original ASM procedure (left) despite the presence of an outlier (marked 1). The same shape hypothesis after outlier removal step (right, adjusted vertex marked by \*). The average shape is shown in the middle. The values given below the figure are computed according to (10), (16), and (20) and correspond to the left-most shape.

words, the model parameters  $b_i$  must not exceed some maximum values (10). This approach can introduce two kinds of errors.

- a) The shape/location hypothesis is completely rejected because one or two points are misplaced, such situation is documented in Fig. 9.
- b) The hypothesis is accepted although one or two points are misplaced as seen in Fig. 10.

To treat the problem of outliers in a systematic fashion, an approach to outlier detection and position adjustment was developed. The misplaced points are identified using the information about the relative positions of the shape model vertices that are implicitly included in the shape model. Let  $\mathbf{z} = (x', y')^T$  be the model point positions after the piecewise linear transformations were applied and the resulting shape was aligned with the shape average. According to the shape model, the hypothesized shape should satisfy (7), specifically  $\mathbf{z} = \bar{\mathbf{x}} + P\mathbf{b}$  for some parameter vector  $\mathbf{b}$ . Usually,  $P$  is not a square matrix but its columns are orthonormal vectors (they are eigenvectors of the covariance matrix). Therefore

$$\mathbf{b} = P^T(\mathbf{z} - \bar{\mathbf{x}}) = P^T \mathbf{d}\mathbf{x} \quad (16)$$

or by components

$$b_j = \sum_{i=1}^n P_{i,j} \mathbf{d}\mathbf{x}_i = \sum_{i=1}^n v_{i,j} \quad (17)$$

where

$$|v_{i,j}| = |P_{i,j} \mathbf{d}\mathbf{x}_i| \quad (18)$$

is the absolute variation induced by point  $i$  in parameter  $b_j$ . Let the percentage of variation induced by point  $i$  in parameter  $b_j$  be defined as

$$V_{i,j} = \frac{|v_{i,j}|}{\sum_{i=1}^n |v_{i,j}|} 100 \quad (19)$$

and let the maximum percentage of variation induced by point  $i$  in any of the parameters  $b_j$  be defined as

$$u_i = \max_{j=1 \dots t} V_{i,j}. \quad (20)$$

If all the points were to generate an equal amount of variation, then all the percentages  $u_i$  were approximately  $100/n$ ,  $n$  being the number of object points. However, since outliers may be present, larger variation may be associated with some points—the outliers. A point is considered to be an outlier if the percentage of variance generated by its position in any of the parameters  $b_i$  of the model is more than  $4\times$  greater than the average amount of variation. Once such misplaced points are detected, they must be moved to a new position that can be inferred from the alignment of the rest of the shape instance with the average shape. If several outliers are present, the variance is distributed among them and (perhaps) the well placed points. As a result, it may be difficult to identify outliers if more than a few occur simultaneously.

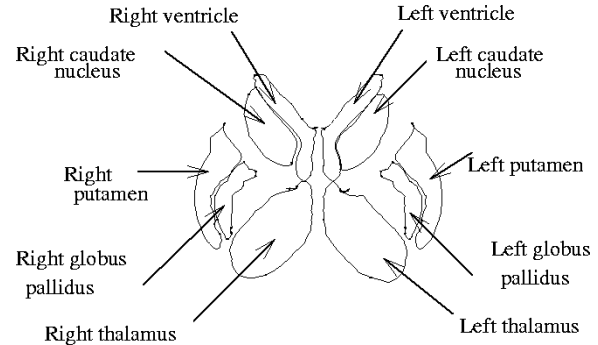


Fig. 11. Neuroanatomic structures of interest.

5) *Final Point Adjustment*: Some of the shape model points may have been declared outliers in the previous step. Consequently, their position may have been adjusted solely considering the average shape appearance and not considering the image data [Fig. 8(b)]. Therefore, they must be subjected to the position optimization step to better correspond with the image data. A hill-climbing procedure that searches a small ( $3 \times 3$ ) neighborhood of each point is employed, the model fitting function is utilized as the optimization criterion. This point adjustment process is performed twice for each point of the shape model, not just for the adjusted outliers.

6) *Final Outlier Removal*: Resulting from the previous steps or newly introduced during the final point adjustment, outliers may remain present in the shape model [Fig. 8(c)]. Following the same outlier detection procedure applied in the first outlier removal step, the outliers are identified and removed, no adjustment is performed in this final step of model fitting [Fig. 8(d)].

One might be tempted to alternate the final point adjustment with the outlier removal several times what may not be a good idea. The two procedures involved have somehow opposite effects: the outlier removal tries to reduce the point scatter of the current shape making it closer to the average shape, while the adjustment tries to better match it with the image data which usually enlarges the point variance. While there are no restrictions regarding the new position of a point that is considered an outlier, the adjustment step can move a point only in a small neighborhood of its current location. Therefore, if we alternate these two steps several times, the current shape would converge toward the average shape. On the other hand, if we allow the adjustment procedure to search a large neighborhood of each point, then the algorithm would be much slower and may have a tendency to oscillate.

#### IV. EXPERIMENTAL METHODS

The method presented above was employed to design a PDM shape model for ten brain structures (Fig. 11) and its performance was assessed in *in vivo* MR brain images. The following ten neuroanatomic structures were identified in individual MR image slices (Fig. 11): left/right ventricles, left/right caudate nucleus, left/right putamen, left/right globus pallidus, left/right thalamus. The brain image data, independent standard, and quantitative validation indexes were identical to those used in our previous study [12].

### A. Data

The image set consisted of individual T1-weighted contiguous MR images of the human brain, imaged in the coronal plane and acquired from a GE 1.5-T Signa scanner with in-slice resolution of  $256 \times 256$  pixels, slice thickness 1.5 mm, and a field-of-view of 26 cm, MR imaging parameters  $TE = 5$  ms, and  $TR = 24$  ms. The images were acquired from 27 subjects and normalized using a three-dimensional proportional Talairach grid system [32]. Normalization using the Talairach grids—while not perfect—increased the likelihood that slice  $n$  in each data set represented the same anatomic slice of the brain. Each image from the set contained the brain structures to be identified, as shown in Fig. 11.

### B. Knowledge from Observer-Traced Contours

Observer-defined contours identified by a neuroanatomist were available for all the images in the brain data set. A training subset of eight images served for the shape model construction. The remaining 19 images formed the test set that was used for validation.

Using the observer-defined contours of the individual neuroanatomic structures of interest for guidance, the landmark points were manually identified for all structures in the training set of MR images (Fig. 1). The landmarks were positioned to consistently represent the shape and shape variation of individual neuroanatomic structures. A prespecified number of landmarks was determined for each neuroanatomic structure of interest. If true landmarks were available (tips of putamen, tips of globus pallidus, etc.—numbered as 45, 47, 52, 69, 70, 74, etc. in Fig. 1), these landmarks were consistently identified. Other points of the shape models did not correspond to true anatomic landmarks (e.g., points numbered as 53–56, 71–73, etc. in Fig. 1). These points were identified to partition the border segments between true landmarks in border elements of approximately constant length.

### C. Quantitative Validation

Performance of the PDM approach to image interpretation was assessed in the test set by quantitatively comparing the labels, areas, and border positions of identified neuroanatomic structures with their respective true labels, areas, and border positions as determined by the observer-defined independent standard. Signed area error, labeling error and border positioning error were used as measures for comparing the observer-defined and computer-determined contours of neuroanatomic structures. The area errors resulting from automated object identification are reported in a normalized way in percents of true object areas. Labeling error was defined as a ratio of the number of incorrectly labeled pixels (both missing and extra) and the number of pixels of each object according to *a priori* knowledge and is expressed in percents. A one-pixel labeling insensitivity was used in the labeling error assessment performed here.

To evaluate the overall accuracy of the contours of the computer-identified brain structures, a distance function along object contours was derived. Considering a pair of contours, the distance function for a pixel on one contour was defined

as the minimum Euclidean distance between that pixel and all other pixels on the second contour. For each border point on the computer identified contours, the average distance errors were calculated with respect to the observer traced contours. A maximum distance error measure, between pairs of contours, was also computed. This measure was calculated as the maximum of  $\{\max_1, \max_2\}$  where  $\max_1$  was the maximum of the distance function of a computer detected contour with respect to an observer traced contour and  $\max_2$  was the maximum of the distance function of the observer traced contour with respect to the computer detected contour.

## V. RESULTS

The PDM approach to brain image segmentation and interpretation correctly interpreted the brain neuroanatomic structures in all images from the test set. Fig. 12 shows the observer-traced and computer-detected contours together with the intermediate results. In the test set, the neuroanatomic structures were identified with the signed percent area error of  $12\% \pm 5\%$  and labeling error of  $7\% \pm 3\%$ . Border positioning errors were quite small, with the average border positioning error of  $0.8 \pm 0.1$  pixels, and maximum border positioning error of  $4.3 \pm 1.2$  pixels. The detection time was about 1 min using a combination of C and MATLAB running on a Sun UltraSparc 1 workstation.

## VI. DISCUSSION

Proper design of the objective function used for model fitting is important for the success of the method. Complexity of the objective function is inversely related to the processing speed. Similarly, the number of tested model hypotheses is related with the matching success. Therefore, selection of the objective function must be a compromise between the function complexity and the processing time that is also reflected in the number of model positioning hypotheses that may be explored in the process. In the past, several sophisticated fitness functions were proposed that were based on combining the Mahalanobis distance of the current shape hypothesis to the average shape with the Mahalanobis distance of the corresponding image profiles to the average (model) profiles [34] and even including a background model in a Bayesian framework [35]. These objective functions, however, lack a way of assessing the object properties between the model points. Clearly, if two consecutive model points lie on a boundary, this information together with the boundary strength may be used in the detection procedure. This is particularly needed when the number of model points is small. When an accurate-enough initial guess of the object pose is not available, it may be necessary to test a large number of hypotheses. In the reported study, about 10 000 hypotheses were tested per image. In such case, computing Mahalanobis distances becomes a computational burden and in the tradeoff between the number of hypotheses and complexity of the objective function, evaluation of a larger number of hypotheses was considered more important. Therefore, our fitness function




Fig. 12. Example of automated brain image segmentation and interpretation. Upper row from left to right: Manual tracings. Initial average position of the shape model. Optimal shape model position after linear transform step. Bottom row from left to right: Optimal shape model position after piecewise linear transform step. Outlier detection—outliers marked by dark dots. Final point adjustment. Images contrast-enhanced to improve visibility.

combines the average *Euclidean* distance of the current profiles to the mean ones with the average boundary strength of the current object hypothesis.

Identification of illegal shapes is another important problem considered in this paper. Mathematically speaking, an acceptable shape is a shape for which its complete vector of principal component analysis parameters (computed using a large set of examples by taking the  $b$  parameters corresponding to all  $2n$  eigenvalues, Section II-C) is inside the ellipsoid defined by the parameters of the shapes in the training set. In our work, we applied a much weaker and informal condition—a shape was considered acceptable if it did not autointersect and the points along the contour were relatively equally spaced and if it was similar to the shapes present in the training set. The problems of illegal shapes come from inaccurate approximations due to a small training set and are unaffected by the model used.

The automatically determined borders of the ten neuroanatomic structures exhibit a high level of accuracy when compared to the manually identified independent standard. Putamen or caudate may serve as examples of consistently well determined structures with the average labeling errors ranging from 1% to 2%. However, the borders of the neuroanatomic structures that lack consistent border appearance like globus pallidus or thalamus represent a much bigger challenge both to the computer algorithm and human observers, average labeling errors for these structures ranged from 8% to 16%. The latter structures also exhibited larger shape variation compared to the former ones. The shape variations unmatched by the appropriate parameters of the model fitting transformations seem to be responsible for most of the gross errors in the incorrectly traced structures. Specifically, errors in detection of the lower tip of the globus pallidus and the outside border of the thalamus suggest




Fig. 13. Segmentation of globus pallidus using artificial deformations and extended scaling and rotation ranges. Columns from left to right: Original image, manual tracing, automated segmentation using computer-generated shapes, and automated segmentation using the average shape only; contrast-enhanced to improve visibility.

that the searching scaling range used is less than the real variation in scale. This observation will serve for further improvement.

Another possible explanation of the imperfect border detection in structures with weak borders is based on the implemented strategy that does not apply the piecewise linear transform to vertices with weak border segments. While this approach improves the method's robustness, it may limit its accuracy. If the true positions of such vertices are not in the neighborhood of the average position, the above strategy may forbid the point's movement toward the correct location. A possible solution to this problem may be to find optimal deformations and positions not only for the average shape but also for a number of other computer generated example shapes. The examples can be generated by discretizing the parameter space for few parameters accounting for the largest variation modes of the PDM. Preliminary results using artificially gen-

erated deformations and extended parameters for the model fitting transforms are presented in Fig. 13. Although only eight computer-generated examples were used, Fig. 13 demonstrates a substantial improvement over the current approach.

Another improvement may result from incorporating better structure positioning information. At this moment, no statistical information is used concerning the relative positions of the neuroanatomic structures. The only constraint that is enforced after each step is that no two structures overlap. In case of overlapping, the faulting structures have to be grouped together. If a structure is misplaced but does not overlap another one, the hypothesis is accepted and a wrong detection occurs. While we did not experience such behavior, the possibility is present. Also, no statistical information is used now concerning the variation of model points. Incorporation of such information may be useful for the outlier detection procedure.

## VII. CONCLUSION

A new fully automated segmentation and interpretation method has been presented. The method was utilized to identify ten neuroanatomic structures in individual MR images. The method was trained in a training set of eight MR brain images and tested in a test set of 19 MR brain images. The method correctly segmented and interpreted neuroanatomic structures in the test MR brain image set. Using quantitatively assessed comparison indexes, our method yielded brain image segmentation and interpretation with small labeling and small border positioning errors. Area measurements using computer-detected neuroanatomic structure borders correlated well with those derived from observer-defined tracings. Compared to our previously reported brain image segmentation and interpretation approach [12], our new method described here offers comparable performance while the interpretation process speed increased sixtyfold.

## ACKNOWLEDGMENT

The authors would like to thank N. Andreasen, M.D., Ph.D., M. Flaum, M.D., R. Rajarethinam, M.D., and T. Cizadlo, M.S. for providing both original and manually traced MR brain data utilized in the presented study. They would also like to thank Prof. G. Stockman for computer facility support.

## REFERENCES

- [1] R. Bajcsy and S. Kovacic, "Multiresolution elastic matching," *Comput. Vision, Graphics, Image Processing*, vol. 46, pp. 1–21, 1989.
- [2] S. P. Raya, "Low-level segmentation of 3-D magnetic resonance brain images: A rule-based system," *IEEE Trans. Med. Imag.*, vol. 9, pp. 327–337, 1990.
- [3] H. H. Ehrlicke, "Automated 3D MR image analysis with a rule-based segmentation system," in *Proc. CAR '91 Computer Assisted Radiology*, Berlin, Germany, Springer-Verlag, pp. 543–548, 1991.
- [4] R. C. Hallgren, "Computer-based system for quantitative analysis of physiological structures from MRI data," *J. Clin. Eng.*, vol. 16, pp. 49–55, 1991.
- [5] C. Li, D. B. Goldgof, and L. O. Hall, "Knowledge based classification and tissue labeling of MR images of human brain," *IEEE Trans. Med. Imag.*, vol. 12, pp. 740–749, 1993.
- [6] H. Li, R. Deklerck, B. D. Cuyper, A. Hermanus, E. Nyssen, and J. Cornelis, "Object recognition in brain CT-scans: Knowledge-based fusion of data from multiple feature extractors," *IEEE Trans. Med. Imag.*, vol. 14, pp. 212–229, 1995.
- [7] D. Vandermeulen, R. Verbeeck, L. Berben, and P. S. ad G. Mar.al, "Continuous voxel classification by stochastic relaxation: Theory and application to MR imaging and MR angiography," in *Proc. IPMI '93*, Berlin, Germany, Springer-Verlag, 1993, pp. 487–506.
- [8] D. N. Kennedy, P. A. Filipek, and V. S. Caviness, "Automatic segmentation and volumetric calculations in nuclear magnetic resonance imaging," *IEEE Trans. Med. Imag.*, vol. TMI-8, pp. 1–7, 1989.
- [9] H. Li, R. Deklerck, and J. Cornelis, "Integration of multiple knowledge sources in a system for brain CT-scan interpretation based on the blackboard," in *Proc. 10th Conf. Artificial Intelligence for Applications*, Los Alamitos, CA, IEEE Comput. Soc. Press, 1994, pp. 336–343.
- [10] D. L. Collins, A. C. Evans, C. Holmes, and T. M. Peters, "Automatic 3D segmentation of neuro-anatomical structures," in *Information Processing in Medical Imaging*, Y. Bizais and C. Barillot (Eds.), Dordrecht, the Netherlands: Kluwer Academic, pp. 139–152, 1995.
- [11] M. Kamber, R. Shinghal, L. Collins, G. S. Francis, and A. C. Evans, "Model-based 3-D segmentation of multiple sclerosis lesions in magnetic resonance brain images," *IEEE Trans. Med. Imag.*, vol. 14, pp. 442–453, 1995.
- [12] M. Sonka, S. K. Tadikonda, and S. M. Collins, "Knowledge-based interpretation of MR brain images," *IEEE Trans. Med. Imag.*, vol. 15, pp. 443–452, 1996.
- [13] M. Flaum, M. Sonka, S. Arndt, T. Cizadlo, S. Stoneall, and N. C. Andreasen, "A new automated method of identifying neuroanatomic regions of interest on imaging data (abstract)," in *Proc 34th Annu. Meet. American College of Neuropsychopharmacology*, 1995, p. 233.
- [14] G. Szekely, A. Kelemen, C. Brechbuhler, and G. Gerig, "Segmentation of 2-D and 3-D objects from MRI volume data using constrained elastic deformations of flexible Fourier contour and surface models," *Med. Image Anal.*, pp. 19–34, 1996.
- [15] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," in *Proc. 1st Int. Conf. Computer Vision*, London, U.K., 1987, pp. 259–268.
- [16] L. H. Staib and J. S. Duncan, "Boundary finding with parametrically deformable models," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 14, no. 11, pp. 1061–1075, 1992.
- [17] A. Jain, Y. Zhong, and S. Lakshmanan, "Object matching using deformable templates," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 18, no. 3, pp. 267–277, 1996.
- [18] M. Sonka, V. Hlavac, and R. Boyle, *Image Processing, Analysis, and Machine Vision*. London, U.K.: Chapman Hall, 1993.
- [19] T. F. Cootes, A. Hill, C. J. Taylor, and J. Haslam, "Use of active shape models for locating structures in medical images," *Image, Vision Computing*, vol. 12, no. 6, pp. 355–366, 1994.
- [20] A. Hill, T. F. Cootes, C. J. Taylor, and K. Lindley, "Medical image interpretation: A generic approach using deformable templates," *Med. Informatics*, vol. 19, no. 1, pp. 47–59, 1994.
- [21] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active shape models—Their training and application," *Comput. Vision, Image Understanding*, vol. 61, pp. 38–59, 1995.
- [22] B. K. P. Horn, "Closed form solution of absolute orientation using unit quaternions," *J. Opt. Soc. Amer. A*, vol. 4, pp. 629–642, 1987.
- [23] B. K. P. Horn, H. M. Hilden, and S. Negahdaripou, "Closed form solution of absolute orientation using orthonormal matrices," *J. Opt. Soc. Amer. A*, vol. 5, pp. 1127–1135, 1988.
- [24] C. Goodall, "Procrustes methods in the statistical analysis of shape," *J. Royal Stat. Soc. B*, vol. 53, no. 2, pp. 285–339, 1991.
- [25] F. L. Bookstein, *Morphometric Tools for Landmark Data*. New York: Cambridge Univ. Press, 1991.
- [26] J. Kent, "The complex Bingham distribution and shape analysis," *J. Roy. Stat. Soc. B*, vol. 56, pp. 285–299, 1994.
- [27] T. F. Cootes and C. J. Taylor, "A mixture model for representing shape variation," in *Proc. British Machine Vision Conf.*, BMVA, 1997, pp. 110–119.
- [28] F. L. Bookstein, "Shape and the information in medical images: A decade of the morphometric synthesis," *Comput. Vision, Image Understanding*, vol. 66, pp. 97–118, 1997.
- [29] ———, "Morphometrics," *Math Horizons*, pp. 28–31, 1996.
- [30] T. F. Cootes and C. J. Taylor, "Combining point distribution models with shape models based on finite-element analysis," in *Proc. British Machine Vision Conf.*, BMVA, 1994, pp. 419–428.
- [31] ———, "Data driven refinement of active shape model search," in *Proc. British Machine Vision Conf.*, BMVA, 1996, pp. 383–392.
- [32] J. Talairach and P. Tournoux, *Co-Planar Stereotaxic Atlas of the Human Brain*, New York: Thieme Medical Publishers, Inc., 1988.
- [33] D. Geman and B. Jedynak, "An active testing model for tracking roads in satellite images," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 18, no. 1, pp. 1–14, 1996.
- [34] T. F. Cootes and C. J. Taylor, "Active shape model search using local grey-level models: A quantitative evaluation," in *Proc. British Machine Vision Conf.*, BMVA, 1994, pp. 639–648.
- [35] J. Haslam, C. Taylor, and T. Cootes, "A probabilistic fitness measure for deformable template models," in *Proc. British Machine Vision Conf.*, BMVA, 1994, pp. 33–42.