NEUROIMAGE **6,** 209–217 (1997)
ARTICLE NO. NI970290

Multimodal Image Coregistration and Partitioning—A Unified Framework

J. Ashburner and K. Friston

Wellcome Department of Cognitive Neurology, Institute of Neurology, London, United Kingdom

Received February 17, 1997

This paper presents a method for the coregistration and partitioning (i.e., tissue segmentation) of brain images that have been acquired in different modalities. The basic idea is that instead of matching two images directly, one performs intermediate within-modality registrations to two template images that are already in register. One can use a least-squares minimization to determine the affine transformations that map between the templates and the images. By incorporating suitable constraints, a rigid body transformation which directly maps between the images can be extracted from these more general affine transformations. A further refinement capitalizes on the implicit normalization of both images into a standard space. This facilitates segmentation or partitioning of both original images into homologous tissue classifications. Once partitioned, the partitions can be jointly matched, further increasing the accuracy of the coregistration. In short, these techniques reduce the between-modality problem to a series of simpler within-modality problems. These methods are relatively robust, address a number of problems in image transformations, and require no manual intervention.

INTRODUCTION

The coregistration of brain images of the same subject acquired in different modalities has proved useful in many areas, both in research and in clinical settings. This work concentrates on the registration of magnetic resonance (MR) images with positron emission tomography (PET) images and on coregistering MR images from different scanning sequences. The aim was to coregister images as accurately and quickly as possible, with no manual intervention.

Intermodality registration of images is less straightforward than registering images of the same modality. Two PET images from the same subject generally look similar, so it suffices to find the rigid-body transformation parameters that minimize the sum of squares difference between them. However, for coregistration between modalities there is nothing quite so obvious to minimize. AIR (Woods *et al.*, 1992) is a widely used algorithm for coregistration of PET to MR images, but it has the disadvantage that it depends on preprocessing of the MR images. This normally involves laborious manual editing in order to remove any tissue that is not part of the brain (i.e., scalp editing).

An alternative method is the manual identification of homologous landmarks in both images. These landmarks are aligned together, thus bringing the images into registration. This is also time-consuming, requires a degree of experience, and can be rather subjective.

The method described in this paper requires no preprocessing of the data or landmark identification and is still reasonably robust.

THEORY

The proposed method of image coregistration relies on images other than the images which are to be registered (**f** and **g**). These are template images of the same modalities as **f** and **g** ($\mathbf{t_f}$ and $\mathbf{t_g}$) and probability images of gray matter, white matter, and cerebrospinal fluid. These probabalistic images will be denoted by the matrix **B** (where each column is a separate image). Images $\mathbf{t_f}$, $\mathbf{t_g}$, and **B** conform to the same anatomical space.

The between-modality coregistration described here is a three-step approach:

1. Determine the affine transformations that map between the images and the templates by minimizing the sum of squares differences between **f** and $\mathbf{t_f}$ and between **g** and $\mathbf{t_g}$. These transformations are constrained so that only the parameters which describe the rigid body component are allowed to differ.

2. Segment or partition the images using the probability images and a modified mixture model algorithm. The mapping between the probability images to images **f** and **g** was determined in step 1.

3. Coregister the image partitions using the rigid body transformations computed from step 1 as a starting estimate.

AFFINE TRANSFORMATIONS

Rigid body transformations, necessary to coregister images of the same subject, are a subset of the more general affine transformations. An affine transformation is carried out by a simple matrix multiplication of the coordinates of the voxels from one image to give the corresponding coordinates of another. Here we illustrate the mapping from a voxel at coordinates x_1, x_2, x_3,

1053-8119/97 $25.00

1, to coordinates $y_1, y_2, y_3, 1$:

$$\begin{aligned} y_1 &= m_{11}x_1 + m_{12}x_2 + m_{13}x_3 + m_{14}1 \\ y_2 &= m_{21}x_1 + m_{22}x_2 + m_{23}x_3 + m_{24}1 \\ y_3 &= m_{31}x_1 + m_{32}x_2 + m_{33}x_3 + m_{34}1 \\ 1 &= m_{41}x_1 + m_{42}x_2 + m_{43}x_3 + m_{44}1. \end{aligned}$$

Note that in order to express the operation in matrix form, an extra 1 is needed to fully specify the coordinates. The fourth row of the transformation matrix always contains the elements (0 0 0 1).

The image is then resampled according to this set of transformations using one of a number of different interpolation methods. Each time an image is resampled, it is degraded slightly. The elegance of formulating these transformations in terms of matrices is that several transformations can be combined by simply multiplying the matrices together.

Often, the images **f** and **g** will have voxels which are anisotropic. The dimensions of the voxels are also likely to differ between images of different modalities. For simplicity, we work in a Euclidian space, where measures of distances are expressed in millimeters. Rather than interpolating all the image data such that the voxels are cubic and have the same voxel sizes in all images, we can simply define affine transformation matrices which map from voxel coordinates into this Euclidian space. For example, if image **f** is of size $128 \times 128 \times 43$ and has voxels which are $2.1 \times 2.1 \times 2.45$ mm, we can define the following matrix:

$$\mathbf{M_f} = \begin{pmatrix} 2.1 & 0 & 0 & -134.4 \\ 0 & 2.1 & 0 & -134.4 \\ 0 & 0 & 2.45 & -52.675 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

This transformation matrix maps voxel coordinates to a Euclidian space whose axes are parallel to those of the image and distances are measured in millimeters, with the origin at the center of the image. Similar matrices can be defined for **g** ($\mathbf{M_g}$) and the template and probability images ($\mathbf{M_t}$).

The objective is to determine the affine transformation which maps the space of **f** to that of **g**. To accomplish this, we wish to find a rigid body transformation matrix $\mathbf{M_r}$, such that $\mathbf{M_g}^{-1}\mathbf{M_r}\mathbf{M_f}$ will coregister the images.

DETERMINING THE MAPPINGS FROM IMAGES TO TEMPLATES

It is possible to obtain a reasonable match of images of most normal brains to a template image of the same modality using just a 12- or even 9-parameter affine transformation. One can register image **g** to template $\mathbf{t_g}$ and similarly register **f** to $\mathbf{t_f}$ using this approach. We will call these transformation matrices $\mathbf{M_{gt}}$ and $\mathbf{M_{ft}}$, respectively. Thus a mapping from **f** to **g** now becomes $\mathbf{M_g}^{-1}\mathbf{M_{gt}}\mathbf{M_{ft}}^{-1}\mathbf{M_f}$. However, this affine transformation between **f** and **g** has not been constrained to be rigid body. We modify this simple approach to incorporate this constraint by decomposing matrix $\mathbf{M_{gt}}$ into a matrix which performs a rigid body transformation ($\mathbf{M_{gr}}$) and one which performs the scaling and shearing ($\mathbf{M_{ta}}$); i.e., $\mathbf{M_{gt}} = \mathbf{M_{gr}}\mathbf{M_{ta}}$, and similarly $\mathbf{M_{ft}} = \mathbf{M_{fr}}\mathbf{M_{ta}}$. Notice that $\mathbf{M_{ta}}$ is the same for both **f** and **g**. Now the mapping becomes $\mathbf{M_g}^{-1}\mathbf{M_{gr}}(\mathbf{M_{ta}}\mathbf{M_{ta}}^{-1})\mathbf{M_{fr}}^{-1}\mathbf{M_f}$, and is a rigid body transformation:

$$\mathbf{M_{gr}} = \begin{pmatrix} 1 & 0 & 0 & q_1 \\ 0 & 1 & 0 & q_2 \\ 0 & 0 & 1 & q_3 \\ 0 & 0 & 0 & 1 \end{pmatrix} \times \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(q_4) & \sin(q_4) & 0 \\ 0 & -\sin(q_4) & \cos(q_4) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

$$\times \begin{pmatrix} \cos(q_5) & 0 & \sin(q_5) & 0 \\ 0 & 0 & & 0 \\ -\sin(q_5) & 0 & \cos(q_5) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

$$\times \begin{pmatrix} \cos(q_6) & \sin(q_6) & 0 & 0 \\ -\sin(q_6) & \cos(q_6) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

$$\mathbf{M_{fr}} = \begin{pmatrix} 1 & 0 & 0 & q_7 \\ 0 & 1 & 0 & q_8 \\ 0 & 0 & 1 & q_9 \\ 0 & 0 & 0 & 1 \end{pmatrix} \times \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(q_{10}) & \sin(q_{10}) & 0 \\ 0 & -\sin(q_{10}) & \cos(q_{10}) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

$$\times \begin{pmatrix} \cos(q_{11}) & 0 & \sin(q_{11}) & 0 \\ 0 & 0 & & 0 \\ -\sin(q_{11}) & 0 & \cos(q_{11}) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

$$\times \begin{pmatrix} \cos(q_{12}) & \sin(q_{12}) & 0 & 0 \\ -\sin(q_{12}) & \cos(q_{12}) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

$$\mathbf{M_{ta}} = \begin{pmatrix} q_{13} & 0 & 0 & 0 \\ 0 & q_{14} & 0 & 0 \\ 0 & 0 & q_{15} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \times \begin{pmatrix} 1 & q_{16} & q_{17} & 0 \\ 0 & 1 & q_{18} & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

We can now optimize the parameter set $\mathbf{q} = [q_1, q_2 \ldots]$ in order to determine the transformations which minimize the sum of squares differences between the images and the templates. The optimization method which we choose is based upon that described by Friston *et al.* (1995). The method is iterative, but generally converges within a few iterations. The chance of finding a local minimum is reduced by using smoothed data (typically images which have been convolved with an 8-mm full width at half maximum Gaussian kernel). It involves generating a linear approximation to the problem using Taylor's theorem, which is solved on each iteration. This can be expressed as computing $\mathbf{q} = \mathbf{q} + (\mathbf{A^T A})^{-1}\mathbf{A^T b}$, at each iteration. Here, $\mathbf{q}$ are the parameter estimates. $\mathbf{A}$ is a matrix in which each column contains the derivatives at various positions ($\mathbf{x}$) within the template images with respect to the parameters. The vector $\mathbf{b}$ contains the differences between the templates and the images that have been spatially transformed according to the latest parameter estimates.

For the purpose of this optimization, we define two matrices, $\mathbf{M_1} = (\mathbf{M_t}^{-1}\mathbf{M_{ft}}\mathbf{M_f})^{-1}$ and $\mathbf{M_2} = (\mathbf{M_t}^{-1}\mathbf{M_{gt}}\mathbf{M_g})^{-1}$. $\mathbf{A}$ and $\mathbf{B}$ are described in Scheme 1.

The parameters describing the nonrigid transformations (q_{13} to q_{18}) could in theory be derived from either **f** or **g**. In practice, we obtain a better solution by estimating these parameters using both images and by biasing the result so that the image which fits the template better has a greater influence over the parameter estimates. This is achieved by weighting the rows of $\mathbf{A}$ and $\mathbf{b}$ which correspond to the different images. The weights are derived from the sum of squares difference between the template and the object images, obtained from the previous solution of $\mathbf{q}$. These are

$$\frac{I}{\sum_{i=1}^{I} (f(\mathbf{M_1 x}_i) - q_{19}t_f(\mathbf{x}_i))^2} \text{ and } \frac{I}{\sum_{i=1}^{I} (g(\mathbf{M_2 x}_i) - q_{20}t_g(\mathbf{x}_i))^2}.$$

Once the optimization has converged to the final solution, we can obtain the rigid body transformation which approximately maps between **f** and **g**, and we also have affine transformation matrices which map between the object images and the templates. These are used in the next step.

PARTITIONING THE IMAGES

Healthy brain tissue can generally be classified into three broad tissue types on the basis of an MR image. These are gray matter (GM), white matter (WM), and cerebrospinal fluid (CSF). This classification can be performed manually on a good-quality T1 image by simply selecting suitable image intensity ranges which

$$\mathbf{A} = \begin{pmatrix} -\frac{df(\mathbf{M_1 x}_1)}{dq_1} & \cdots & -\frac{df(\mathbf{M_1 x}_1)}{dq_6} & 0 & \cdots & 0 & -\frac{df(\mathbf{M_1 x}_1)}{dq_{13}} & \cdots & -\frac{df(\mathbf{M_1 x}_1)}{dq_{18}} & t_f(\mathbf{x}_1) & 0 \\ -\frac{df(\mathbf{M_1 x}_2)}{dq_1} & \cdots & -\frac{df(\mathbf{M_1 x}_2)}{dq_6} & 0 & \cdots & 0 & -\frac{df(\mathbf{M_1 x}_2)}{dq_{13}} & \cdots & -\frac{df(\mathbf{M_1 x}_2)}{dq_{18}} & t_f(\mathbf{x}_2) & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & \cdots & 0 & -\frac{dg(\mathbf{M_2 x}_1)}{dq_7} & \cdots & -\frac{dg(\mathbf{M_2 x}_1)}{dq_{12}} & -\frac{dg(\mathbf{M_2 x}_1)}{dq_{13}} & \cdots & -\frac{dg(\mathbf{M_2 x}_1)}{dq_{18}} & 0 & t_g(\mathbf{x}_1) \\ 0 & \cdots & 0 & -\frac{dg(\mathbf{M_2 x}_2)}{dq_7} & \cdots & -\frac{dg(\mathbf{M_2 x}_2)}{dq_{12}} & -\frac{dg(\mathbf{M_2 x}_2)}{dq_{13}} & \cdots & -\frac{dg(\mathbf{M_2 x}_2)}{dq_{18}} & 0 & t_g(\mathbf{x}_2) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \end{pmatrix}$$

$$\mathbf{b} = \begin{pmatrix} f(\mathbf{M_1 x}_1) - q_{19}t_f(\mathbf{x}_1) \\ f(\mathbf{M_1 x}_2) - q_{19}t_f(\mathbf{x}_2) \\ \vdots \\ g(\mathbf{M_2 x}_1) - q_{20}t_g(\mathbf{x}_1) \\ g(\mathbf{M_2 x}_2) - q_{20}t_g(\mathbf{x}_2) \\ \vdots \end{pmatrix}.$$

Scheme 1. Matrix $\mathbf{A}$ and vector $\mathbf{b}$, where f(x) is the intensity of image **f** at position $\mathbf{x}$, and similarly for $g(\mathbf{x})$, $t_f(\mathbf{x})$, and $t_g(\mathbf{x})$.

FIG. 1. The *a priori* probability images of GM, WM, and CS (courtesy of the Montreal Neurological Institute).

encompass most of the voxel intensities of a particular tissue type. However, this manual selection of thresholds is highly subjective.

Many groups have used clustering algorithms to partition MR images into different tissue types, either by using images acquired from a single MR sequence or by combining information from two or more registered images acquired using different scanning sequences (e.g., proton-density and T2-weighted).

The approach we have adopted here is a modified version of one of these clustering algorithms. The clustering algorithm of choice is the maximum likelihood "mixture model" algorithm (Hartigan, 1975).

We assume that the MR image (or images) consists of a number of distinct tissue types (clusters) from which every voxel has been drawn. The intensities of voxels belonging to each of these clusters conform to a multivariate normal distribution, which can be described by a mean vector, a covariance matrix, and the number of voxels belonging to the distribution.

In addition, we have approximate knowledge of the spatial distributions of these clusters, in the form of probability images (provided by the Montreal Neurological Institute (Evans *et al.*, 1992, 1993, 1994)) which have been derived from MR images of a large number of subjects (see Fig. 1). The original images were segmented into binary images of GM, WM, and CSF, and all normalized into the same space using a nine-parameter (three translations, three rotations, and three orthogonal zooms) affine transformation. The probability images are the means of these binary images, so that they contain values in the range of 0 to 1. These images represent the *a priori* probability of a voxel being GM, WM, or CSF after an image has been normalized to the same space using a nine-parameter affine transformation.

We describe here a simplified version of the algorithm as it would be applied to a single image. We use a 12-parameter affine transformation determined from the previous step to map between the space of the MR image (**f**) and that of the probability images (**B**). This allows simple "on-the-fly" sampling of the probability images into the space of the image we wish to partition.

Generally, we use six or seven clusters: one each for GM, WM, and CSF, two or three clusters to account for scalp, eyes, etc., and a background cluster. Since we have no probability maps for scalp and background, we estimate them by subtracting $\mathbf{b}_{\mathrm{GM}}$, $\mathbf{b}_{\mathrm{WM}}$, and $\mathbf{b}_{\mathrm{CSF}}$ from a map of all ones and divide the results equally between the remaining clusters.

We then assign initial probabilities (**P**) for each of the I voxels being drawn from each of the K clusters. These probabilities are the *a priori* probabilities (i.e., $p_{ik} = b_k(\mathbf{M_1}^{-1}\mathbf{x}_i)$. Where identical *a priori* probability maps are used for more than one cluster, the starting estimates are modified slightly by adding random noise.

The following steps (1 to 6) are repeated until convergence (or a prespecified number of iterations) is reached.

1. Compute the number of voxels belonging to each of the K clusters (**h**) as

$$h_k = \sum_{i=1}^{I} p_{ik} \text{ over } k = 1 \ldots K.$$

2. Mean voxel intensities for each cluster (**v**) are computed. This step effectively produces a weighted mean of the image voxels, where the weights are the

current belonging probability estimates:

$$v_k = \frac{\sum_{i=1}^{I} p_{ik} f(\mathbf{x}_i)}{h_k} \text{ over } k = 1 \ldots K.$$

3. Then the variance of each cluster (**c**) is computed in a way similar to the mean:

$$c_k = \frac{\sum_{i=1}^{I} p_{ik}(f(\mathbf{x}_i) - v_k)^2}{h_k} \text{ over } k = 1 \ldots K.$$

4. Now that we have all the parameters which describe the current estimate of the distributions, we have to recalculate the belonging probabilities (**P**).

Evaluate the probability density functions for the clusters at each of the voxels:

$$r_{ik} = (2\pi c_k)^{-0.5} \exp\left(\frac{-(f(\mathbf{x}_i) - v_k)^2}{2c_k}\right)$$

over $k = 1 \ldots K$ and $i = 1 \ldots I$.

5. Then utilize the *a priori* information (**B**) (this is the only deviation from the conventional mixture model algorithm which is simply $q_{ik} = r_{ik} h_k$):

$$q_{ik} = r_{ik} \frac{h_k b_k(\mathbf{M_1}^{-1}\mathbf{x}_i)}{\sum_{j=1}^{I} b_k(\mathbf{M_1}^{-1}\mathbf{x}_j)} \text{ over } k = 1 \ldots K \text{ and } i = 1 \ldots I.$$

Note that we have extended the mixture model by including an extra term $(b_k(\mathrm{M_1}^{-1}\mathrm{x}_i))/(\Sigma_{j=1}^{I} b_k(\mathrm{M_1}^{-1}\mathrm{x}_j))$. This term sums to unity over voxels and can be thought of as the probability density function of a voxel from cluster k being found at location i, irrespective of how many voxels of type k there are in the brain. Using this term, we can include the *a priori* information, without biasing the overall proportions of different tissue types.

6. And finally normalize the probabilities so that they integrate to unity at each voxel:

$$p_{ik} = \frac{q_{ik}}{\sum_{j=1}^{K} q_{ij}} \text{ over } k = 1 \ldots K \quad \text{and} \quad i = 1 \ldots I.$$

With each iteration of the algorithm, the parameters describing the distributions (**v**, **c**, and **h**) move toward a better fit and the belonging probabilities (**P**) change slightly to reflect the new distributions. The parameters describing the clusters which have corresponding *a priori* probability images tend to converge more rapidly than the other clusters—this is partly due to the better starting estimates. The final values in **P** are in the range of 0 to 1, although most values tend to stabilize very close to one of the two extremes. Examples of MR images classified in this way can be seen in Fig. 2.

Strictly speaking, the assumption that multinormal distributions should be used to model MRI intensities is not quite correct. After Fourier reconstruction, the moduli of the complex pixel values are taken, thus rendering any potentially negative values positive. Where the cluster variances are of comparable magnitude to the cluster means, the distribution deviates significantly from normal. This only really applies for the background, where the true mean voxel intensity is zero. The algorithm is modified to account for this discrepancy between the model and reality. For this background cluster, the value of v is set to zero before the variance c is computed. Also, because the background cluster is described by only a half Gaussian (and h represent the integrals of the distributions), it is necessary to double the computed values of r (step 4).

The greatest problem which the technique faces is image nonuniformity. The current algorithm assumes that the voxel values for GM (for example) have the same intensity distribution throughout the image. The nonstationary nature of MR image intensities from some scanners can lead to a significant amount of tissue misclassification.

COREGISTERING THE IMAGE PARTITIONS

The previous step produces images of GM and WM from the original images **f** and **g**. These image partitions can then be simultaneously coregistered to produce the final solution.

This optimization stage only needs to search for the six parameters which describe a rigid body transformation. Again, we call these parameters **q** and define a matrix $\mathbf{M_{fg}}$ based upon these parameters (cf. $\mathbf{M_{gr}}$ as defined earlier). We define a matrix **M** as $(\mathbf{M_g}^{-1}\mathbf{M_{fg}}\mathbf{M_{gt}}^{-1}\mathbf{M_{ft}}\mathbf{M_f})^{-1}$. The way that this matrix has been formulated means that the starting estimates for **q** can all be zero, because it incorporates the results from the first step of the coregistration. Very few iterations are required at this stage to achieve convergence. No scaling parameters are needed, since the probability images derived from **f** have similar intensities to those derived from **g**. The system of equations which we iteratively solve ($\mathbf{Ax} = \mathbf{b}$) to optimize the parameters **q** is as follows (using a notation where $p_{\mathbf{g}1}(\mathbf{x}_2)$ means probability of voxel at $\mathbf{x}_2$ from image **g**

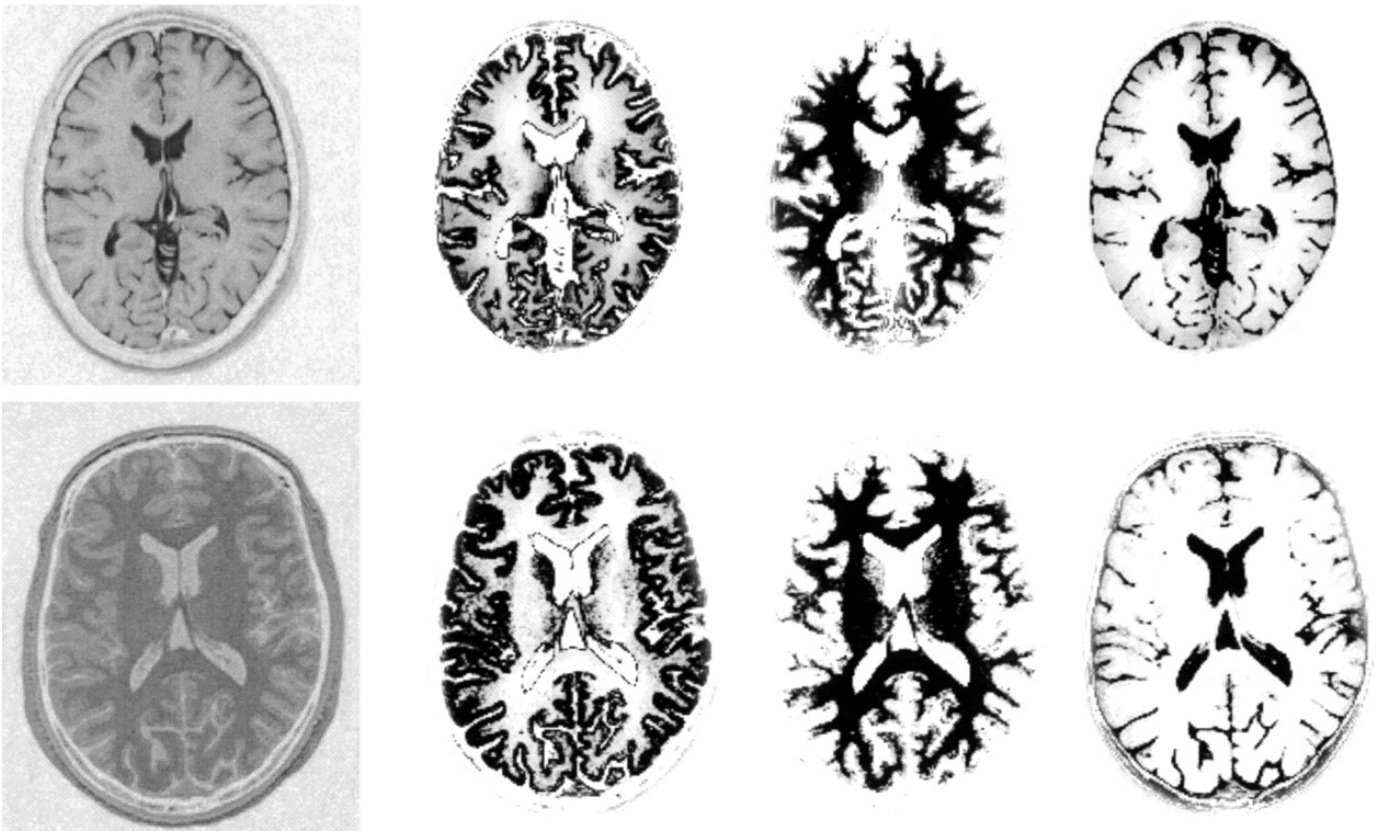

FIG. 2. Examples of MR images partitioned into GM, WM, and CSF. (Top) T2-weighted image. (Bottom) T1-weighted image.

belonging to cluster 1):

$$\mathbf{A} = \begin{pmatrix} -\frac{dp_{\mathbf{f}1}(\mathbf{M}\mathbf{x}_1)}{dq_1} & -\frac{dp_{\mathbf{f}1}(\mathbf{M}\mathbf{x}_1)}{dq_2} & \cdots & p_{\mathbf{g}1}(\mathbf{x}_1) & 0 & 0 \\ -\frac{dp_{\mathbf{f}1}(\mathbf{M}\mathbf{x}_2)}{dq_1} & -\frac{dp_{\mathbf{f}1}(\mathbf{M}\mathbf{x}_2)}{dq_2} & \cdots & p_{\mathbf{g}1}(\mathbf{x}_2) & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ -\frac{dp_{\mathbf{f}2}(\mathbf{M}\mathbf{x}_1)}{dq_1} & -\frac{dp_{\mathbf{f}2}(\mathbf{M}\mathbf{x}_1)}{dq_2} & \cdots & 0 & p_{\mathbf{g}2}(\mathbf{x}_1) & 0 \\ -\frac{dp_{\mathbf{f}2}(\mathbf{M}\mathbf{x}_2)}{dq_1} & -\frac{dp_{\mathbf{f}2}(\mathbf{M}\mathbf{x}_2)}{dq_2} & \cdots & 0 & p_{\mathbf{g}2}(\mathbf{x}_2) & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \end{pmatrix}$$

$$\mathbf{b} = \begin{pmatrix} p_{\mathbf{f}1}(\mathbf{M}\mathbf{x}_1) - p_{\mathbf{g}1}(\mathbf{x}_1) \\ p_{\mathbf{f}1}(\mathbf{M}\mathbf{x}_2) - p_{\mathbf{g}1}(\mathbf{x}_2) \\ \vdots \\ p_{\mathbf{f}2}(\mathbf{M}\mathbf{x}_1) - p_{\mathbf{g}2}(\mathbf{x}_1) \\ p_{\mathbf{f}2}(\mathbf{M}\mathbf{x}_2) - p_{\mathbf{g}2}(\mathbf{x}_2) \\ \vdots \end{pmatrix}$$

After this coregistration we have our final solution. It is now possible to map voxel **x** of image **g** to the corresponding voxel **Mx** of image **f**. An example of PET–MRI coregistration using this approach is illustrated in Fig. 3.

AN ALTERNATIVE IMPLEMENTATION FOR LOW-RESOLUTION IMAGES

Here we briefly describe an approach which may be more appropriate for the registration of SPECT or low-resolution PET images to MRI. The tissue classification model described above is not ideal for partitioning low-resolution images. It assumes that each voxel contains tissue from only one of the underlying clusters, whereas in reality, many voxels will contain a mixture of different tissue types (Bullmore *et al.*, 1995; Ashburner *et al.*, 1996).

An alternative is to partition only the MR image as described above and to generate an image from the resulting segments which resembles a PET image. This can be achieved by assigning the gray matter segment a value of 1, white matter a value of about 0.3, and CSF a value of about 0.1, followed by smoothing. It is then possible to apply the within-modality coregistration

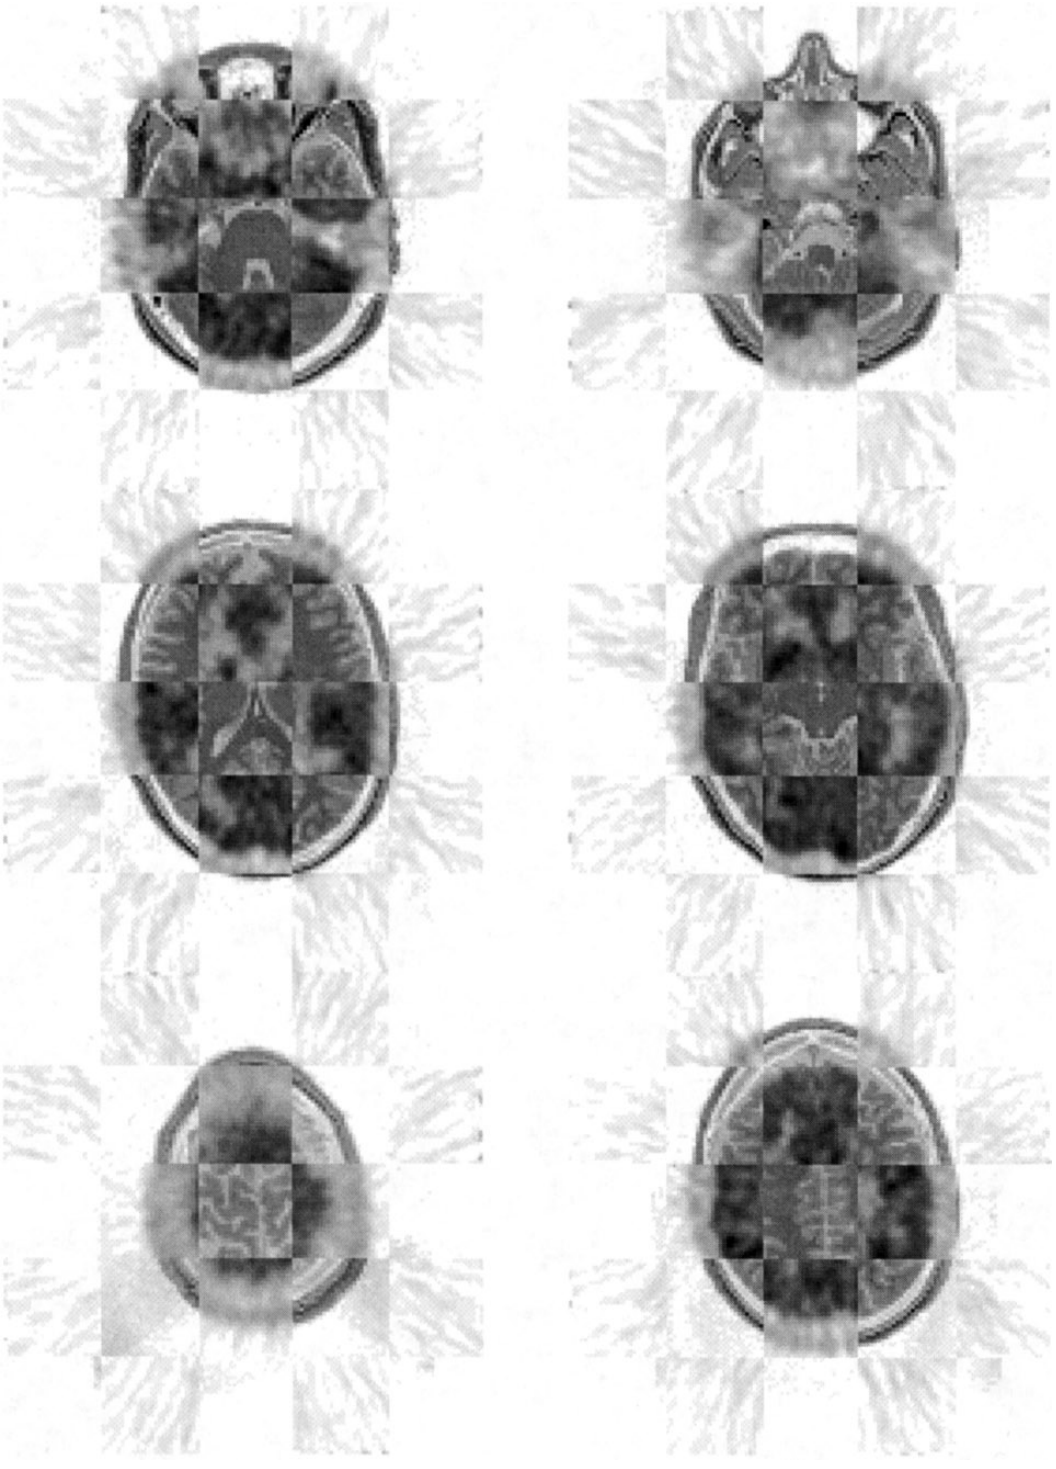

FIG. 3. An example of PET–MRI coregistration, achieved using the techniques described here.

described in the previous section to coregister the real and "fake" PET images.

EVALUATION

The coregistration methods were evaluated for PET to T1-weighted MRI, using data from the Evaluation of Retrospective Image Registration project (West *et al.*, 1996, 1997). This involved obtaining both PET and MRI data from Vanderbilt University and performing intermodality registrations on the volumes. Fiducial markers during the acquisition of these datasets enabled investigators at Vanderbilt to know the true registration parameters, but any visible traces of these mark-

ers had been removed from the images prior to their distribution to other investigators.

Registrations were performed on 11 volume images both with and without geometry distortion correction from the Vanderbilt dataset. They were done using only the first step of the registration process (constrained simultaneous affine registration) and also using all the steps.

The evaluations were performed on a Sun SPARC Ultra 2, using an implementation of the method written in C and Matlab (from Mathworks Inc., Sherborn, USA). The starting estimates for the registration parameters matched the centers of each volume together and assumed that the images were in the same orientation. No manual intervention was involved. The first step of the registration process (including the initial smoothing) took an average of 66 s, whereas the complete three-step registration required about 350 s.

The resulting parameter estimates were communicated back to Vanderbilt, where their accuracy was evaluated. The results presented in Table 1 are the mean, median, and maximum errors for the registration and can be compared directly with those in West *et al.* (1997). They show that the first step of the process quickly registered the images to within about 6 mm before the remaining steps further refined the parameters.

Although the accuracy of the registrations was found to be comparable to that of the other techniques evaluated by West *et al.*, the data used in the evaluation did violate a number of the assumptions made in the current approach. Due to the effects of field inhomogeneity and the low gray/white matter contrast, there was a considerable amount of tissue misclassification of the MR images. In addition, the assumption that brain tissue can be broadly classified as gray or white matter was complicated by the presence of tumors, which were classified as gray matter in the MR images and white matter in the PET images. This would be expected to introduce additional registration errors, since the final step is based upon matching corresponding image partitions together. The registration should be much more accurate for images where the assumptions hold. In summary, the current technique is valid in relation to existing techniques. Unlike some of the existing approaches, the present method does not require manual intervention (e.g., scalp editing).

TABLE 1

Errors for the PET-MRI Registration

	First step only		All three steps	
	Uncorrected	Corrected	Uncorrected	Corrected
Mean error (mm)	5.57	3.77	4.14	3.20
Median error (mm)	5.11	3.17	4.20	3.36
Maximum error (mm)	11.62	8.54	7.46	5.76
N	7	4	7	4

Note. Errors are presented, for both the uncorrected and the distortion-corrected MR images. The results in the left-hand column were derived after using only the first step of the registration process. The right-hand column shows the results of using all three steps of the registration process.

DISCUSSION

We have developed a strategy for the coregistration of brain images from different modalities which is entirely automatic. No manual editing of the images is required in order to remove scalp, nor does the investigator need to identify any mutual points or features or even set thresholds for morphological operations such as brain segmentation. The only occasional intervention which may be needed is to provide starting estimates to the first step of the procedure. The procedure has so far been successfully applied to the registration of T1 MRI to PET (blood flow), T1 to T2 MRI, and T2 to PET.

In addition to providing a method of coregistration, another feature of the current approach is the generation of partitioned (or segmented) images which can be used for voxel-based morphometrics (Wright *et al.*, 1995). The incorporation of the *a priori* probability images into the clustering algorithm produces a much more robust solution. However, a better result is expected when the method is applied to two (or more) exactly registered images from different scanning sequences. Although the algorithm has only been illustrated for a single image, it has actually been implemented such that the classification can be performed using any number of registered images. The mixture model clustering algorithm is described for multidimensional input data in Hartigan (1975), although the *a priori* probabilities are not included in the description.

Currently, it is not possible to directly compare the speed of this method with that of others, since it has largely been implemented in Matlab (a high level programming language from Mathworks Inc.). If the algorithms were translated into efficient C code, they would compare favorably with other image registration methods. Alpert *et al.* (1996) used the Levenberg–Marquardt algorithm (LM) (Press *et al.*, 1988) for the optimization of PET–PET and PET–MRI registration, which was found to be faster than the approach of Woods *et al.* (1992). The optimization algorithms described in this paper are in fact a simplified version of LM. The nature of the linear approximations results in changes in parameter estimates that are usually slight underestimates at each iteration. Because of this, the extra precautions taken by LM are largely unnecessary and simply slow the algorithm down. This can be

demonstrated by observing the residual sum of squares at each iteration, which (almost) invariably is less than that from the previous iteration.

ACKNOWLEDGMENTS

The authors thank the group at the Montreal Neurological Institute for providing the probability images. The images and the standard transformations used in the evaluation were provided as part of the project "Evaluation of Retrospective Image Registration," National Institutes of Health, Project 1 R01 NS33926-01, Principal Investigator, J. Michael Fitzpatrick, Vanderbilt University Nashville, Tennessee. This work was supported by the Wellcome Trust.

Note. The algorithms described in this paper are available from the authors as part of the *SPM96* (Statistical Parametric Mapping) software package.

REFERENCES

Alpert, N. M., Berdichevsky, D., Levin, Z., Morris, E. D., and Fischman, A. J. 1996. Improved methods for image registration. *NeuroImage* **3:**10–18.

Ashburner, J., Haslam, J., Taylor, C., Cunningham, V. J., and Jones, T. 1996. A cluster analysis approach for the characterization of dynamic PET data. In *Quantification of Brain Function Using PET* (R. Myers, V. Cunningham, D. Bailey, and T. Jones, Eds.), Chap. 59, pp. 301–306, Academic Press, San Diego.

Bullmore, E., Brammer, M., Rouleau, G., Everitt, B., Simmons, A., Sharma, T., Frangou, S., Murray, R., and Dunn, G. 1995. Computerized brain tissue classification of magnetic resonance images: A new approach to the problem of partial volume artifact. *NeuroImage* **2:**133–147.

Evans, A. C., Collins, D. L., and Milner, B. 1992. An MRI-based stereotactic atlas from 250 young normal subjects. *Soc. Neurosci. Abstr.* **18:**408.

Evans, A. C., Collins, D. L., Mills, S. R., Brown, E. D., Kelly, R. L., and Peters, T. M. 1993. 3D statistical neuroanatomical models from 305 MRI volumes. In *Proceedings of the IEEE Nuclear Science Symposium and Medical Imaging Conference*, pp. 1813–1817.

Evans, A. C., Kamber, M., Collins, D. L., and Macdonald, D. 1994. An MRI-based probabilistic atlas of neuroanatomy. In *Magnetic Resonance Scanning and Epilepsy* (S. Shorvon, D. Fish, F. Andermann, G. M. Bydder, and H. Stefan, Eds.), NATO ASI Series A, Life Sciences, Vol. 264. Plenum, New York.

Friston, K. J., Ashburner, J., Frith, C. D., Poline, J.-B., Heather, J. D., and Frackowiak, R. S. J. 1995. Spatial registration and normalization of images. *Hum. Brain Mapping* **2:**165–189.

Hartigan, J. A. 1975. *Clustering Algorithms*, pp. 113–129. Wiley, New York.

Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T. 1988. *Numerical Recipes in C*. Cambridge Univ. Press, Cambridge.

West, J., Fitzpatrick, J. M., Wang, M. Y., Dawant, B. M., Maurer, C. R., Kessler, R. M., Maciunas, R. J., Barillot, C., Lemoine, D., Collignon, A., Maes, F., Suetens, P., Vandermeulen, D., van den Elsen, P. A., Hemler, P. F., Napel, S., Sumanaweera, T. S., Harkness, B., Hill, D. L. G., Studholme, C., Malandain, G., Pennec, X., Noz, M. E., Maguire, G. Q., Pollack, M., Pelizzari, C. A., Robb, R. A., Hanson, D., and Woods, R. P. 1996. Comparison and evaluation of retrospective intermodality image registration techniques. *Medical Imaging 1996: Image Processing, Proc. SPIE 2710*, pp. 332–347.

West, J., Fitzpatrick, J. M., Wang, M. Y., Dawant, B. M., Maurer, C. R., Kessler, R. M., Maciunas, R. J., Barillot, C., Lemoine, D., Collignon, A., Maes, F., Suetens, P., Vandermeulen, D., van den Elsen, P. A., Napel, S., Sumanaweera, T. S., Harkness, B., Hemler, P. F., Hill, D. L. G., Hawkes, D. J., Studholme, C., Maintz, J. B. A., Viergever, M. A., Malandain, G., Pennec, X., Noz, M. E., Maguire, G. Q., Pollack, M., Pelizzari, C. A., Robb, R. A., Hanson, D., and Woods, R. P. 1997. Comparison and evaluation of retrospective intermodality brain image registration techniques. *J. Comput. Assist. Tomogr.* **21**(4):554–566.

Woods, R. P., Cherry, S. R., and Mazziotta, J. C. 1992. Rapid automated algorithm for aligning and reslicing PET images. *J. Comput. Assist. Tomogr.* **16:**620–633.

Wright, I. C., McGuire, P. K., Poline, J.-B., Travere, J. M., Murray, R. M., Frith, C. D., Frackowiak, R. S. J., and Friston, K. J. 1995. A voxel-based method for the statistical analysis of gray and white matter density applied to schizophrenia. *NeuroImage* **2:**244–252.