MRI and PET Coregistration—A Cross Validation of Statistical Parametric Mapping and Automated Image Registration

Stefan J. Kiebel,* John Ashburner,† Jean-Baptiste Poline,† and Karl J. Friston†

*Department of Neurology, Friedrich–Schiller University, 07740 Jena, Germany; and †The Wellcome Department of Cognitive Neurology, The Institute of Neurology, London, WC1N 3BG United Kingdom

Coregistration of functional PET and T1-weighted MR images is a necessary step for combining functional information from PET images with anatomical information in MR images. Several coregistration algorithms have been published and are used in functional brain imaging studies. In this paper, we present a comparison and cross validation of the two most widely used coregistration routines (Friston *et al.,* 1995, *Hum. Brain Map.* **2:** 165–189; Woods *et al.,* 1993, *J. Comput. Assisted Tomogr.* **17:** 536–546). Several transformations were applied to high-resolution anatomical MR images to generate simulated PET images so that the exact (rigid body) transformations between each MR image and its associated simulated PET images were known. The estimation error of a coregistration in relation to the known transformation allows a comparison of the performance of different coregistration routines. Under the assumption that the simulated PET images embody the salient features of real PET images with respect to coregistration, this study shows that the routines examined reliably solve the MRI to PET coregistration problem.   © 1997 Academic Press

INTRODUCTION

In neuroimaging, coregistration of functional PET and anatomical MR images is essential for combining information from both image modalities. The benefit of having coregistered images is that it enables the visualization of functional PET data by superimposing it on a high-resolution anatomical MR image and therefore improves localization of neural activation loci. Several algorithms have been proposed by various authors (Andersson *et al.,* 1995; Ardekani *et al.,* 1995; Friston *et al.,* 1995; Mangin *et al.,* 1994; Woods *et al.,* 1993). We chose two widely used coregistration algorithms (Woods *et al.,* 1993; Friston *et al.,* 1995) to cross validate on the same test data set. The first routine is contained in the Automated Image Registration package (*AIR,* Version

2.03) and the second is part of the Statistical Parametric Mapping package (*SPM*95). This cross validation includes an empirical validation of the *SPM* routine, which had not been undertaken before, and facilitates a comparison of the performance of both routines.

To generate a suitable PET data set, we used simulated PET images, which were generated by applying a series of transformations to some MR images. Rigid body transformations applied to the simulated PET images were used to simulate misregistrations between the PET and the MR data. After running both coregistration routines on the MR and simulated PET images, it was possible to assess and compare the errors associated with them.

METHODS

We describe the basic concepts of MRI and PET coregistration, followed by an overview and theoretical comparison of the *AIR* and *SPM* algorithms. We also describe the cross validation procedure, particularly generation of simulated PET images.

Basic Concepts

Generally, an MRI to PET coregistration algorithm transforms an MR image M in such a way that each voxel in the transformed MR image M' and in the PET image P corresponds to the same voxel in real brain space. Since images of both modalities are acquired from the same subject, a rigid body transformation T maps M to M'. We assume that the scanning processes have not introduced nonrigid spatial transformations into either the PET or the MR image or that these distortions are small. We also assume that no spatial scaling factor of the data is introduced by inaccurate voxel sizes. The required transformation can then be described by translations and rotations in three dimensions so that any coregistration algorithm computes six parameters: x, y, and z translation; pitch; roll; and yaw.

The main challenge in MRI and PET coregistration is

that, even if both images are perfectly coregistered, there is no "simple" function which maps voxel intensities from the MR uniquely to the PET image. For instance, scalp and white matter produce roughly the same voxel intensities within an MR image, but are associated with different voxel values within a PET image. Therefore, coregistration based only on the voxel intensities of both images is difficult, if not impossible. To address this issue, the images can be preprocessed to eschew this problem of nonunique mapping. An example is the often-used (manual) scalp-editing of the MR image prior to coregistration. This sort of preprocessing renders the coregistration problem solvable; however, it is desirable to automate this preprocessing task and incorporate it into a coregistration routine, since manual scalp-editing can be laborious for the user. The *SPM* algorithm is an example of this nonsupervised class of algorithm.

Since we will use the standard (4×4)-matrix format throughout the text to represent a transformation T, this matrix format is described in the following. A (4×4)-transformation matrix, which we also interchangeably denote by T, is given by

$$T = \begin{vmatrix} t_{11} & t_{12} & t_{13} & t_{14} \\ t_{21} & t_{22} & t_{23} & t_{24} \\ t_{31} & t_{32} & t_{33} & t_{34} \\ 0 & 0 & 0 & 1 \end{vmatrix}, \quad (1)$$

where the elements of the first three rows of T are chosen in a way such that

$$\begin{vmatrix} y_1 \\ y_2 \\ y_3 \\ 1 \end{vmatrix} = T \cdot \begin{vmatrix} x_1 \\ x_2 \\ x_3 \\ 1 \end{vmatrix}, \quad (2)$$

where $\mathbf{x} = (x_1, x_2, x_3)'$ is a voxel position in the original image M and $\mathbf{y} = (y_1, y_2, y_3)'$ is the according voxel position in M' after transformation with T. For instance, the matrix

$$T_t = \begin{vmatrix} 1 & 0 & 0 & t_{14} \\ 0 & 1 & 0 & t_{24} \\ 0 & 0 & 1 & t_{34} \\ 0 & 0 & 0 & 1 \end{vmatrix} \quad (3)$$

would represent a three-dimensional translation by t_{14} in x, t_{24} in y, and t_{34} in z direction (voxels). The rotations r_x, r_y, and r_z (radians) around the three axes (pitch, roll,

and yaw) are similarly coded by

$$T_{r_x} = \begin{vmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(r_x) & \sin(r_x) & 0 \\ 0 & -\sin(r_x) & \cos(r_x) & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix}, \quad (4)$$

$$T_{r_y} = \begin{vmatrix} \cos(r_y) & 0 & \sin(r_y) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(r_y) & 0 & \cos(r_y) & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix}, \quad (5)$$

$$T_{r_z} = \begin{vmatrix} \cos(r_z) & \sin(r_z) & 0 & 0 \\ -\sin(r_z) & \cos(r_z) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix}. \quad (6)$$

The advantage of this matrix representation is that several sequentially applied transformations T_i, $i = 1, \ldots, Q$, can be expressed by only one transformation matrix T:

$$T = \prod_{i=1}^{Q} T_i. \quad (7)$$

(Note that T_1 is not the first, but the last applied transformation.) In this way, we can represent a rigid body transformation with six parameters by one transformation matrix. Another important property of these matrices is that the inverse of the matrix T (denoted T^{-1}) is also the inverse of the transformation T. *SPM* and *AIR* basically use the same format except for the order of rotations, which does matter as one can see from Eqs. (4) to (6). In *SPM*, a rigid body transformation matrix is given by $T = T_t \cdot T_{r_x} \cdot T_{r_y} \cdot T_{r_z}$, and in *AIR* it is $T = T_t \cdot T_{r_y} \cdot T_{r_x} \cdot T_{r_z}$.

The *AIR* Algorithm

We used the *AIR* 2.03 routine with the cost function based upon the standard deviation of the "ratio image" for cross validation. Although Version 2.03 of the *AIR* package also offers a second routine for coregistration, we chose the first routine, since it had been described in Woods *et al.* (1993) and already cross validated with other coregistration routines (Strother *et al.,* 1994).

The approach taken by the *AIR* algorithm is to define a cost function for measuring the MRI and PET misregistration. A transformation matrix is then found by iterative, nonlinear minimization of this cost function.

The *AIR* algorithm is based on the key assumption that similar voxel intensities in the MR image correspond to the same tissue type. We refer to this assumption as the "similarity assumption" in the remainder of text. The similarity assumption is roughly valid if the T1-weighted MR image is scalp-edited, i.e., voxels not corresponding to brain tissue are set to 0. This scalp-editing has to be done prior to coregistration. Assuming that the similarity assumption holds not only for MR images but also for PET images, the idea is to find a rigid body transformation *T* so that all MR voxels with the same intensity value are mapped to a set of similar PET voxel intensity values. The similarity of a set of PET voxels can be measured by the ratio of its standard deviation and mean. Described more formally, the algorithm uses a histogram segmentation to segment the MR image into K voxel intensity partitions P_i and minimizes a weighted ratio of the standard deviations and means of the PET voxels corresponding to P_i by manipulating the transformation matrix *T*. The key formula of the algorithm is

$$f(\mathbf{PET}, \mathbf{MRI}, T) = \frac{1}{N} \sum_j^K n_j \frac{\sigma_j}{a_j}, \qquad (8)$$

where **PET** is the PET image, **MRI** is the MR image, T stands for the transformation matrix, n_j is the number of MRI voxels with an intensity of j, $N = \Sigma_j^K n_j$, a_j is the mean of all PET voxel values with a corresponding MRI voxel intensity in partition j, and σ_j is their standard deviation. Changing *T,* for example translating the MRI in relation to the PET image by 1 mm, means changing the value of σ_j/a_j for all j.

The aim of the algorithm is to find T such that the right-hand side of Eq. (8) is minimized. This computation requires a nonlinear minimization scheme, because the additive components in Eq. (8) are not independent. The minimization method chosen is Southwell's Relaxation Search (Pierre, 1986), which is an iterative, univariate Newton–Raphson search. The overall minimization is performed iteratively, each step producing a new transformation matrix *T,* which is subsequently used in the next Newton–Raphson iteration. When the coregistration is good enough, i.e., the minimization process arrives at a global or local minimum in the parameter space of *T,* or when a prespecified maximum of iteration steps is reached, the algorithm will stop and return the computed transformation matrix *T*.

The *SPM* Algorithm

The basic strategy of the *SPM* algorithm is to transform the MR image to an emulated PET image and to map this image to the real PET image.

The first step is to segment the MR image into four partitions: gray matter, white matter, cerebrospinal fluid (CSF), and "everything else." Because of this preprocessing segmentation it is possible to use either a scalp-edited or a non-scalp-edited MR image for coregistration: The algorithm should segment any MRI into its four partitions. In the case of a non-scalp-edited MRI, the "everything else" partition should contain scalp, skull, and meninges.

The segmentation routine is described in detail elsewhere (Friston *et al.,* 1996) and works as follows: First, the MRI is fitted to an MRI template by means of a 12-parameter affine transformation. For each voxel of the template, there exists a three-element vector specifying the *a priori* probabilities for gray matter, white matter, and CSF. This probability map is used to segment the fitted MRI into its four partitions using an iterative fuzzy clustering scheme. The next step is to generate an emulated PET image by modifying the voxel intensities so that the ratio of the mean of white matter to the mean of gray matter voxels is 1:3 and the corresponding ratio of CSF to gray matter is 1:10.

At this point, the remaining task is to find a rigid body transformation *T,* which coregisters the emulated PET to the real PET image. The method used is to linearize the optimization problem and to iteratively find a solution with linear methods. More precisely, each iteration solves an overdetermined linear equation system in a least-squares sense.

To linearize, two key assumptions are met by smoothing the emulated PET image with a Gaussian filter with 8-mm FWHM (full width at half maximum):

1. The effects of applying a sufficiently small transformation parameter (e.g., t_x) to the emulated PET image can be approximated by some linear function of the intensity values using a first-order Taylor series approximation.

2. The effects of applying sufficiently small transformation parameters to the emulated PET image are independent.

These two assumptions are the prerequisites to formulate the coregistration problem as the following linear equation system,

$$\left[\frac{\partial \mathbf{PET}}{\partial t_x} \frac{\partial \mathbf{PET}}{\partial t_y} \frac{\partial \mathbf{PET}}{\partial t_z} \frac{\partial \mathbf{PET}}{\partial r_x} \frac{\partial \mathbf{PET}}{\partial r_y} \frac{\partial \mathbf{PET}}{\partial r_z} \mathbf{1} \, \mathbf{MRI} \right] \mathbf{q}$$
$$= \mathbf{PET}, \qquad (9)$$

where t_x, t_y, and t_z stand for x, y, and z translation; r_x, r_y, and r_z for pitch, roll, and yaw; **PET** is the real PET image, **1** is a column vector of 1's, and **MRI** denotes the emulated PET image.

Note that all images (e.g., $\delta\mathbf{PET}/\delta t_x$ or **PET**) are represented as column vectors. Equation (9) assumes

that the observed **PET** can be described by a linear combination of the emulated PET image and the additive effects of the six transformation parameters applied to **PET.** The first six column vectors of the left-hand side—the numerical partial derivatives of the PET image with respect to the six transformation parameters—are known: small transformations are applied to the PET image and the change is measured. The solution, how the MR image has to be moved to approximate a good coregistration of both images, then consists of the first six elements of the vector **q**.

This step—solving for **q** and applying the determined parameters to the emulated PET image—is iterated a fixed number of times to find the best transformation matrix T (in the current implementation of SPM95 this number is 16). The last iteration step returns the final T in terms of the first six components of **q**.

Theoretical Comparison

Clearly, the main difference from a theoretical point of view between the AIR and the SPM algorithm is the applied optimization method for computing a transformation matrix T. Whereas the AIR method employs a nonlinear Newton–Raphson search to fix one parameter at a time, the SPM method solves a linear equation system to compute six transformation parameters in one step. Both algorithms have to iterate these steps to converge on a good coregistration solution. It is important to note that the SPM routine finds a solution in a nonlinear fashion, although a single iteration is only capable of finding a linear solution.

For different input images, the required computing time varies for both routines. The AIR routine is mainly dependent on the iterative optimization, whose convergence behavior is variable (Woods *et al.,* 1993). The SPM routine consists of three tasks, which are affine transformation, segmentation, and iterative solving of the linear equation system. The required computing time of the third task is only dependent on the dimensionality of the images and the fixed number of iterations and is thus rather predictable. The affine transformation and segmentation, however, use optimization routines, which are based on a quality measure of the solution found so that the required computing time varies for different input images.

The required time for coregistering an MR and a PET image is 5 to 10 min for the AIR routine (measured on a Sparc 20 workstation), whereas the SPM routine needs roughly 20 min for segmentation and 5 min for coregistration. However, note that these times are not directly comparable because the time required for the AIR coregistration does not include the necessary scalp-editing of the MR image. If this scalp-editing is done manually, the overall time of an AIR routine coregistration depends on the skill of the user. A reasonably skilled user is able to manually edit an MR image in less than an hour. There are also third-party software packages available, which can be used to edit MR images in a semiautomatic fashion so that the required time is reduced further.

Cross Validation

In this section, we describe how we cross validated the AIR and the SPM routines. Basically, the cross validation consists of three steps:

- generate simulated PET images,
- run the coregistration routines, and
- measure the error.

To generate simulated PET images, we applied a series of transformations to six MR images. This process is not exactly the same as in the SPM routine to generate an emulated PET image. We will describe the rationale for not choosing the SPM routine for generation of simulated PET images under Discussion.

The T1-weighted MR images were generated on a 1-T scanner by a RF spoiled volume acquisition with $T_R = 21$ ms, $T_E = 6$ ms, and flip angle = 35°. The data were acquired as a 256×192 matrix. One hundred forty contiguous axial slices with a thickness of 1.3 mm were obtained, which were interpolated to generate 182 slices. The field of view was 25 cm, resulting in voxel dimensions of $1 \times 1 \times 1$ mm.

The simulated PET images which were generated by the sequence described below, had a dimension of [128 128 40] in voxels and a voxel size of [2.05 2.05 3.43] so that the whole brain including the cerebellum was covered by the simulated PET scans.

- The MR image was scalp-edited and segmented by interactive thresholding into gray matter, white matter, and CSF components.

- The gray matter, white matter, and CSF intensities were modified in such a way that the mean of the white matter voxels related to the mean of the gray matter voxels by a ratio of 1:3. The corresponding ratio of CSF to gray matter was 1:10.

- A rigid body transformation T_{orig}, which simulated misregistration between MR and PET data, was applied and the resulting image was resampled to the PET voxel size. T_{orig} was saved for future reference. A trilinear interpolation was used to implement the sampling.

- The image was smoothed in each direction to a resolution of 7 mm FWHM by applying a three-dimensional Gaussian filter.

- Gaussian white noise was added at 30% of the mean intensity value of brain voxels to simulate the worse signal to noise ratio of the PET image.

• The image was smoothed to PET resolution (~8 mm FWHM) by another three-dimensional Gaussian filter with a FWHM of 4 mm in each direction. This special order of filtering and noising was applied to visually imitate the typical PET intracranial texture.

Figure 1 shows a slice of such a transformed simulated PET image. Comparing it to approximately the same slice in the corresponding real PET of the same subject shows that the simulated PET visually resembles the real PET.

This process was repeated 32 times for each MR image to generate a series of simulated PET images. In this way, a data set consisting of 6 MR images and 192 simulated PET images was generated. Since the original transformation matrices T_{orig} are known for all simulated PET images, the error of the estimated transformation matrix T can be computed by $T_{\text{err}} = T_{\text{orig}} \cdot T^{-1}$, and the performance of the coregistration routines in terms of transformation parameters can be compared. To extract the error of the six rigid body transformation parameters t_x, t_y, t_z, r_x, r_y, and r_z from T_{err}, the inverse calculations of those described under Basic Concepts have to be applied. The parameter errors are given by

$$\text{error}(t_x) = T_{\text{err}}(1, 4)$$
$$\text{error}(t_y) = T_{\text{err}}(2, 4)$$
$$\text{error}(t_z) = T_{\text{err}}(3, 4)$$
$$\text{error}(r_y) = \sin^{-1}(T_{\text{err}}(1, 3)) \tag{10}$$
$$\text{error}(r_z) = \sin^{-1}(T_{\text{err}}(1, 2)/\cos(r_y))$$
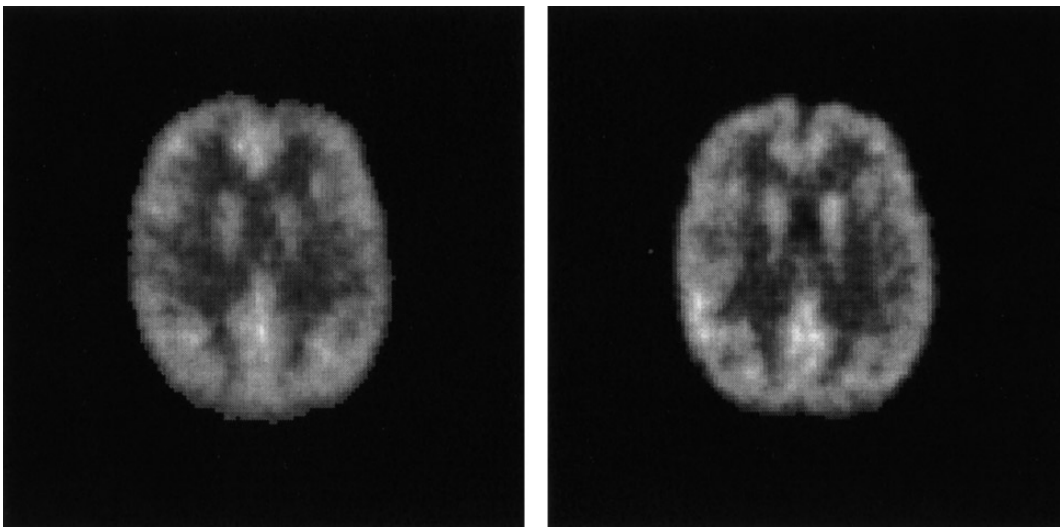$$\text{error}(r_x) = \sin^{-1}(T_{\text{err}}(2, 3)/\cos(r_y)),$$

where the order of rotation is defined as yaw, roll, and pitch (r_z, r_y, and r_x).

We tested the *AIR* routine and two options of the *SPM* routine, where one option was to use scalp-edited MR images and the second was to use non-scalp-edited MR images. As a result, we produced $6 \cdot 32 \cdot 4 = 768$ transformation matrices.

To examine various aspects of the coregistration, we used three types of transformation matrices T_{orig}: The first simulated typical coregistration mismatches by using random values. In the other two types only one transformation parameter was varied, the z translation t_z or the pitch r_x. These were chosen because estimating t_z and r_x is normally the most difficult task of the coregistration routines. Estimating t_x, t_y, r_y, and r_z are in general easier because of the symmetry of the brain. To change only one parameter and hold the remaining parameters constant can reveal some information about the influence of this parameter on the overall error. Two MR images were assigned to each of these three manipulations or groups.

The random translations of the first group had a Gaussian distribution with $\sim N(0, 5)$ (mm) and the random rotations were also Gaussian with $\sim N(0, 3)$ (degrees). In the second group, we only changed the z translations t_z of the simulated PET images between 0 and 62 mm with an increase of 2 mm per image. For the simulated PET images of the third group, we changed the pitch r_x between 0 and 15.5° in 0.5-degree steps.

Since we generated a number of simulated PET images from each of the six MR images, we implicitly assumed that the variability of coregistrations between each MR image and its derived simulated PET images was the same as were individual MR–PET image pairs used.

FIG. 1. Visual comparison of a simulated PET with a real PET slice of the same subject. (Left) Real pet and (right) simulated PET.

RESULTS

The results of all coregistration runs are shown in Tables 1–6. In each of these tables the mean and the standard deviation of the errors of the coregistration routines with respect to the six transformation parameters for one group of simulated PET images are displayed so that each of the six original MR images is represented by a separate table. Since each of these six groups comprises 32 PET images, the standard deviations displayed in the tables have 31 degrees of freedom. Tables 1 and 2 show the results of the groups in which random transformations were applied to the simulated PET images. Tables 3 and 4 refer to the groups of z-translated simulated PET images and Tables 5 and 6 to the groups of x-rotated (pitch) simulated PET images.

For illustration we chose a group (corresponding to Table 6) and displayed in Fig. 2 the translational and rotational errors of both routines. The first row of the figure displays the translational errors (mm) as functions of the three translation parameters t_x, t_y, and t_z

TABLE 1

Group Random Transformations 1: Translational (mm) and Rotational (Degrees) Mean Errors for 32 Simulated PET Images, to Which Random Transformations [Translations $\sim N(0, 5)$ (mm), Rotations $\sim N(0, 3)$ (degrees)] Were Applied

	SPM with scalp-editing		SPM without scalp-editing		AIR	
	Mean	SD	Mean	SD	Mean	SD
x transl.	−0.012	0.145	0.670	0.236	0.039	0.220
y transl.	−0.422	0.175	−1.150	0.172	0.631	0.259
z transl.	1.478	0.297	2.071	0.295	−1.437	0.370
Pitch	−0.599	0.086	−1.131	0.082	0.401	0.093
Roll	0.024	0.091	−0.145	0.151	0.044	0.037
Yaw	0.120	0.060	−0.260	0.078	−0.012	0.029

TABLE 2

Group Random Transformations 2: Translational (mm) and Rotational (Degrees) Mean Errors for 32 Simulated PET Images, to Which Random Transformations [Translations $\sim N(0, 5)$ (mm), Rotations $\sim N(0, 3)$ (degrees)] Were Applied

	SPM with scalp-editing		SPM without scalp-editing		AIR	
	Mean	SD	Mean	SD	Mean	SD
x transl.	0.246	0.179	0.125	0.289	0.029	0.283
y transl.	−0.538	0.232	0.368	0.205	0.457	0.261
z transl.	1.668	0.386	1.265	0.365	−0.973	0.421
Pitch	−0.530	0.160	−0.345	0.154	0.353	0.119
Roll	0.054	0.100	0.051	0.196	0.026	0.055
Yaw	0.020	0.051	−0.012	0.069	0.009	0.039

TABLE 3

Group z Translation 1: Translational (mm) and Rotational (Degrees) Mean Errors for 32 Simulated PET Images, to Which Translations in z Direction (Ranging from 0 to 62 mm) Were Applied

	SPM with scalp-editing		SPM without scalp-editing		AIR	
	Mean	SD	Mean	SD	Mean	SD
x transl.	0.167	0.080	−1.436	0.110	−0.072	0.061
y transl.	−0.353	0.121	1.324	0.242	0.426	0.064
z transl.	1.132	0.254	−0.028	0.374	−0.872	0.106
Pitch	−0.235	0.073	0.462	0.133	0.296	0.048
Roll	−0.076	0.065	−0.469	0.080	0.026	0.052
Yaw	−0.033	0.038	0.667	0.058	0.035	0.024

TABLE 4

Group z Translation 2: Translational (mm) and Rotational (Degrees) Mean Errors for 32 Simulated PET Images, to Which Translations in z Direction (Ranging from 0 to 62 mm) Were Applied

	SPM with scalp-editing		SPM without scalp-editing		AIR	
	Mean	SD	Mean	SD	Mean	SD
x transl.	0.247	0.078	0.462	0.107	0.049	0.052
y transl.	−0.153	0.094	0.788	0.142	0.325	0.070
z transl.	0.692	0.171	−0.708	0.305	−0.824	0.120
Pitch	−0.161	0.053	0.475	0.091	0.241	0.043
Roll	−0.014	0.034	−0.295	0.056	0.060	0.022
Yaw	−0.027	0.038	−0.293	0.041	0.007	0.021

and the second row shows the rotational errors r_x, r_y, and r_z (degrees) for all axes.

The individual performance of both routines can be summarized as follows:

SPM Routine with Scalp-Edited MR Images

In Fig. 2, the results of the *SPM* routine with prior scalp-editing are displayed as a solid line. From Tables 1–6, one can see that the mean error of translation in x, y, and z direction is below 1 mm except in three instances (z translation in Tables 1, 2, and 3). The standard deviations of the translation errors are small throughout all the simulated PET data sets. Regarding the estimation of the rotations, the mean error of pitch, roll, and yaw estimates is always smaller than 0.60° and often close to 0.

SPM Routine with Non-Scalp-Edited MR Images

The results of the *SPM* routine without prior scalp-editing are displayed as a dotted line in Fig. 2. Except

for one instance (Table 2) the mean errors of translation and rotation are larger if the input of the coregistration routine is a non-scalp-edited MR image. This is most obvious in Table 6, where the mean errors of y and z translations are 2.16 and 2.91 mm, respectively, compared to 0.02 and 0.62 mm (*SPM* routine with scalp-edited MRIs). However, in most cases the increase of mean error (scalp-edited vs non-scalp-edited MRI) is at most around 1 mm in translation and 0.3° in rotation.

AIR Routine

The results of the *AIR* routine are displayed as a dashed line in Fig. 2. The mean translation errors are always smaller than 1 mm except in one instance (*z* translation in Table 1) and have a small standard deviation. The mean rotation errors are close to 0° in all experiments; the maximal error is 0.40° (Table 1 pitch estimate).

One can also see from Fig. 2 that the error of both routines in terms of the six parameters is independent on the actual *x* rotation, even if it is as high as a pitch of 15.5°. This independence of error on actual displacements also holds for the other five groups.

On a Sparc-20 computer, the time required for the *AIR* coregistration is roughly 5 to 10 min. The *SPM* routine (including noninteractive segmentation of the MR image) needs roughly 20 min for segmentation and 5 min for coregistration. These times are not directly comparable without considering the additional time required for scalp-editing the MR image, which is a necessary preprocessing step for the *AIR* routine. See Theoretical Comparison for an estimation of the required time to scalp-edit an MR image.

DISCUSSION

The basic result of our cross validation is that *AIR* and *SPM* coregistration perform well on our simulated

TABLE 5

Group Pitch 1: Translational (mm) and Rotational (Degrees) Mean Errors for 32 Simulated PET Images, to Which Rotations around the *x* Axis (Ranging from 0 to 15.5°) Were Applied

	SPM with scalp-editing		SPM without scalp-editing		AIR	
	Mean	SD	Mean	SD	Mean	SD
x transl.	0.236	0.097	0.310	0.102	0.021	0.058
y transl.	0.105	0.131	0.506	0.127	0.653	0.226
z transl.	0.165	0.199	−1.267	0.243	−0.585	0.109
Pitch	−0.068	0.067	−0.006	0.085	0.097	0.080
Roll	0.084	0.057	0.023	0.074	−0.033	0.029
Yaw	0.039	0.048	0.051	0.058	−0.044	0.023

TABLE 6

Group Pitch 2: Translational (mm) and Rotational (Degrees) Mean Errors for 32 Simulated PET Images, to Which Rotations around the *x* Axis (Ranging from 0 to 15.5°) Were Applied

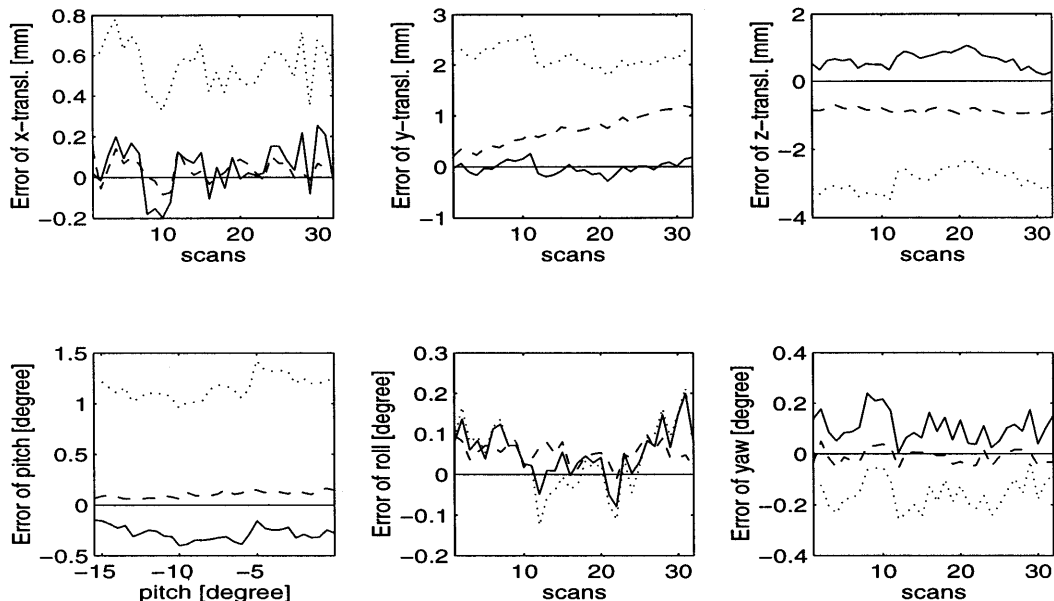|	SPM with scalp-editing		SPM without scalp-editing		AIR	
	Mean	SD	Mean	SD	Mean	SD
x transl.	0.047	0.118	0.548	0.116	0.037	0.058
y transl.	−0.017	0.124	2.161	0.199	0.719	0.289
z transl.	0.617	0.225	−2.910	0.319	−0.869	0.073
Pitch	−0.280	0.069	1.162	0.111	0.107	0.033
Roll	0.055	0.060	0.045	0.081	0.053	0.026
Yaw	0.106	0.059	−0.155	0.061	−0.010	0.030

PET data set. The maximal mean *z* translation errors made by the *AIR* routine and by the *SPM* routine are 1.44 and 1.67 mm, respectively. The maximal mean rotation error is smaller than 0.40° for the *AIR* routine and 0.60° for the *SPM* routine. The performance of the *SPM* routine decreases if the MR image is not scalp-edited. The maximal mean *z* translation error is 2.91 mm and the maximal mean rotation error is 1.16°.

Both routines provide accurate solutions to the coregistration problem on scalp-edited MR images. Moreover, the small standard deviations of the coregistration parameters suggest that the routines are robust.

For most applications, the mean error made by the *SPM* routine on non-scalp-edited MR images is still acceptable. However, if optimal accuracy is needed, manual or semiautomatic scalp-editing of the MR image is recommended when using the *SPM* routine.

The correct interpretation of these results demands some further understanding of the validation process. First of all, one has to accept that our simulated PET data is a good approximation to reality. It seems appropriate to simulate the relative voxel intensities of the PET by the ratio 1:3:10 (CSF:white matter:gray matter), scale the image to PET voxel size, smooth, add some noise, and finally smooth the resulting image to PET resolution. After these steps, the simulated PET images bear a strong resemblance to real PET images, as can be seen in Fig. 1. It could be argued that any deviation from the white:gray matter (3:10) or CSF: gray matter (1:10) ratio might decrease the performance of the *SPM* routine, since the emulated PET image is generated with these ratios. However, this is not a real limitation on the accuracy of the *SPM* routine. As one can see from Eq. (9), the polynomial fit between the emulated PET on the left side and simulated PET on the right side of the equation will compensate for most types of a deviation from the assumption about the ratios, particularly for a linear

FIG. 2. Group pitch 2: Translational (mm) and rotational (degrees) errors for 32 simulated PET images, to which rotations around the x axis (ranging from 0 to 15.5°) were applied. (Dashed line) *AIR* routine, (solid line) *SPM* routine with scalp-edited MR images, and (dotted line) *SPM* routine with non-scalp-edited images.

scaling of the ratio. Moreover, in our opinion, the most important image features for a thorough test of MRI to PET coregistration routines are the outer brain boundaries and the boundaries between gray and white matter so that the exact ratio between the three partitions only plays a minor role. Both routines rely inherently on these boundaries:

The *AIR* routine segments the MR image into K voxel intensity partitions (see above), and tries to find a transformation which maps similar PET voxel intensities into each partition. Assuming that the similarity assumption holds for the PET and MR image, it is clear that a suboptimal transformation is "punished" by the cost function, because the weighted sum of standard deviations σ_j will increase. This effect will work best if the difference in voxel intensity between roughly homogeneous regions (gray and white matter and nonbrain) is large and the transition between areas is sharp. Although the intrinsic smoothness and partial volume effects of a PET image render the boundaries less sharp, this feature enables the coregistration process to more reliably find the global minimum within the parameter space. The dependence of the *SPM* coregistration on image boundaries is also quite obvious. To find a set of coregistration parameters T the partial derivatives $\partial \mathbf{PET}/\delta T$ of the PET image with respect to the transformation parameters T [Eq. (9)] are used. As a result, the image boundaries and the relative voxel intensities of different brain tissues are important image features. The amount of added noise does not

seem to play an important role. Both routines are robust for various noise levels.

Assuming that our simulated PET data are suitable for testing coregistration routines, we can consider the rationale for using a manual segmentation to generate the simulated PET images. As described above, *SPM* uses its own noninteractive segmentation routine to generate an emulated PET image, whereas *AIR* segments the MR image into different intensity partitions. We did not choose the same noninteractive segmentation of *SPM* for generating the simulated PET images, but chose a manual one, since we wanted to carry out a validation of the whole *SPM* routine, including its preprocessing segmentation step. Moreover, the cross validation would have probably been biased in favor of *SPM* if we used the same method to create the simulated and emulated PET images.

Concerning scalp-edited MR images, the performance of the *AIR* and the *SPM* routines seem to be comparable. This shows that both routines reliably solve the coregistration problem, if the MR image is scalp-edited. It also indicates that the segmentation process of the *SPM* routine works accurately on scalp-edited MR images.

As stated above, the performance of the *SPM* routine worsens if non-scalp-edited MR images are coregistered. Visual inspection of the MR images shows that the internal segmentation of the *SPM* routine misclassifies a small part of the meninges as gray matter. This most likely causes the observed decrease in perfor-

mance compared to coregistration based on scalp-edited MR images.

From Fig. 2, one can also see that at least some of the parameters are highly dependent on each other. This is because the primary aim of both coregistration routines is to match the outer brain boundaries. In the figure, it can be seen that the *SPM* routine based on a non-scalp-edited MR image makes relatively large errors on the y and z translation and the pitch. This combination of a positive pitch, a positive y translation, and a negative z translation error is a suitable parameter combination to preserve the matching of the outer brain boundaries, although the overall error is relatively large.

CONCLUSION

In this paper, we showed that the routines examined reliably solve the MRI to PET coregistration problem for simulated PET data. Although the study is limited in the sense that we used simulated PET data, this work gives to the functional imaging community a sense of the coregistration errors that should be expected using the two methods reviewed and provide a cross validation of these algorithms and their implementation.

ACKNOWLEDGMENTS

REFERENCES

Andersson, J. L., Sundin, A., and Valind, S. 1995. A method for coregistration of PET and MR brain images. *J. Nucl. Med.* **36:**1307–1315.

Ardekani, B., Braun, M., Hutton, B., Kanno, I., and Iida, H. 1995. A fully automatic multimodality image registration algorithm. *J. Comput. Assisted Tomogr.* **19:**615–623.

Evans, A., Collins, D., Neelin, P., MacDonald, D., Kamber, M., and Marrett, S. 1994. Three dimensional correlative imaging: Application in human brain mapping. In *Advances in Neuroimaging: Multimodal Registration,* (R. Thatcher, M. Hallett, T. Zeffiro, E. John, and M. Huerta, Eds.), pp. 145–162.

Evans, A., Marret, S., Torrescorzo, J., Ku, S., and Collins, L. 1991. MRI-PET correlation in three dimensions using a volume-of-interest (VOI) atlas. *J. Cereb. Blood Flow Metab.* **11:**A69–A78.

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

Friston, K. J., Ashburner, J., Holmes, A., Poline, J.-B., Büchel, C., Price, C. J., Turner, R., Howseman, A., Rees, G. E., Greene, J. D., and Josephs, O. 1996. *SPM Short Course.* Available at http://www.fil.ion.ucl.ac.uk.

Ge, Y., Fitzpatrick, J., Votaw, J., Gadamsetty, S., Maciunas, R., Kessler, R., and Margolin, R. 1994. Retrospective registration of PET and MR brain images: An algorithm and its stereotactic validation. *J. Comput. Assisted Tomogr.* **18:**800–810.

Mangin, J., Frouin, V., Bloch, I., Bendriem, B., and Lopez-Krahe, J. 1994. Fast nonsupervised 3d registration of PET and MR images of the brain. *J. Cereb. Blood Flow Metab.* **14:**749–762.

Pelizzari, C., Evans, A., Neelin, P., Chen, C.-T., and Marret, S. 1991. Comparison of two methods for three-dimensional registration of PET and MR images. *Proc. IEEE-EMBS* **1:**221–223.

Pierre, D. A. 1986. *Optimization Theory with Applications.* Dover, New York.

Strother, S. C., Anderson, J. R., Xu, X.-L., Liow, J.-S., Bonar, D. C., and Rottenberg, D. R. 1994. Quantitative comparisons of image registration techniques based on high-resolution MRI of the brain. *J. Comput. Assisted Tomogr.* **18:**954–962.

Turkington, T., Hoffman, J., Jaszczak, R., MacFall, J., Harris, C., Kilts, C., Pelizzari, C., and Coleman, R. 1995. Accuracy of surface fit registration for PET and MR brain images using full and incomplete brain surfaces. *J. Comput. Assisted Tomogr.* **19:**117–124.

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

Woods, R. P., Mazziotta, J. C., and Cherry, S. R. 1993. MRI-PET registration with automated algorithm. *J. Comput. Assisted Tomogr.* **17:**536–546.