Principal Axes and Surface Fitting Methods for Three-Dimensional Image Registration

Henry Rusinek, Wai-Hon Tsui, Alejandro V. Levy, Marilyn E. Noz and Mony J. de Leon

Departments of Radiology and Psychiatry, New York University Medical Center, New York, New York, N.S. Kline Institute for Psychiatric Research, Orangeburg, NY and Brookhaven National Laboratory, Upton, New York

We evaluated the effect of the image acquisition parameters on the accuracy of the principal axes and surface-fitting techniques for three-dimensional image registration. Using two types of phantom objects, MR brain image and a mathematically defined ellipsoid, we simulated pairs of scans with known acquisition parameters, including longitudinal coverage, magnitude of misregistration, number of sections and section thickness. Both methods are sensitive to the systematic deformation of contours. The principal axes method is also sensitive to incomplete scan coverage and to the x-axis and y-axis misangulation. Both methods are insensitive to the number of sections, section thickness and the number of points per section. Surface fitting performed well without user supervision. There is no need for routine inclusion of the scaling factors as search parameters. The results confirm the feasibility of three-dimensional multimodality registration of brain scans with accuracy 1–2 mm, with surface fitting being the method of choice.

J Nucl Med 1993; 34:2019–2024

The registration of multiple tomographic scans is becoming increasingly important for interpretation of nuclear medicine images. Registration techniques are applied to combine structural and functional modalities and monitor changes in lesion size and shape.

In repeated tomographic studies the angle of the imaging plane and the location of the image may be different. Different imaging modalities may employ different section thickness, field of view and resolution. To register two such studies, it is necessary to find a coordinate transformation that will map one image set onto another.

Algorithms for image registration include matching anatomical or externally placed landmarks, matching surfaces by minimizing distance, and matching spatial moments (centroids and principal axes) of corresponding objects. The latter two techniques may be used retrospectively and they do not require an expert observer to identify corresponding landmarks.

The surface-fitting method (SFIT) minimizes the average distance between the contours extracted from the two scans (*1*). The method has been used to construct hybrid PET-MR images of the brain to display functional abnormalities and in neurosurgical planning (*2,3*).

Two objects can also be registered by aligning their centroids and principal axes (PAX). A set of three orthogonal principal axes is inherent to each three-dimensional solid body. For simple objects (parallelopiped, ellipsoid), the principal axes coincide with the axes of symmetry. The PAX method was applied in radiation treatment planning based on combined CT/MR images (*4*) in a study of orientation of the scaphoid bone in the wrist (*5*) and in analyzing PET images of schizophrenics (*6*).

Although SFIT and PAX have been singled out as the most promising methods for automated image registration (*7*), they have not been previously compared in a systematic fashion using the same input data. The complexity of the algorithms and a large number of factors that affect their performance make it difficult to predict the magnitude of errors mathematically. Previous error analyses were based on repeated scans of physical phantoms (*1,4*). Such an approach is limited by the difficulty in controlling the orientation of the phantom. The error in measuring the angulation of the phantom with respect to scanner coordinate frame is of the order of several degrees, i.e., the same as the registration error. Furthermore, repositioning phantoms is tedious and prevents studying a large number of variables.

The goal of our study is to compare the accuracy of the two methods and to evaluate the effect of the scan acquisition parameters on registration errors. The present investigation is based on "digital phantoms". We start with a complete object, reslice it at two different orientations (simulating two scans that differ in coverage, slice thickness or number of slices), apply PAX and SFIT methods to register the two scans and finally compare the computed results with the known parameters of the transformation. Precise control over simulated images enables us to investigate the registration error empirically.

METHODS

Registration Algorithms

The PAX technique requires the computation of the centroid and the three principal axes for each of the two scans (Fig. 1). The

Received Jan. 28, 1993; revision accepted June 23, 1993.

For correspondence or reprints contact: Henry Rusinek, PhD, Dept. of Radiology, Bellevue C Bldg., Rm D120, New York University Medical Center, New York, NY 10016.

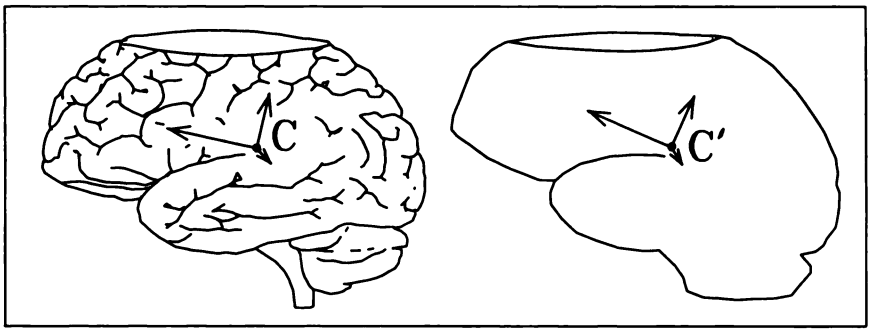

FIGURE 1. In the PAX method, the centroid and the three principal axes of the object are computed from each scan. The registration is given by the translation that maps centroid C onto C' followed by the rotations that align the principal axes.

principal axes are the eigenvectors of the symmetric 3 × 3 matrix of second moments. The registration consists of translating the centroid of the first scan onto the second, rotating it about the centroid to align the principal axes and scaling (*4,8*). Two versions of the PAX algorithm were implemented: the first treated the object as its surface contours, the other as the filled interior (*9*).

The SFIT method involves finding the intersection of a directed ray (Fig. 2) with a surface (*10*). The root mean square (RMS) distance D between the surfaces is expressed as a function of three translations, dx, dy, dz; three angles of rotation about the coordinate axes, ax, ay, az and possibly three scaling factors, sx, sy, sz. A search for the minimum value of the residual D yields the desired parameters of the matching transformation. Iterations start with the value of all parameters set to zero; they terminate when the decrease in the residual D becomes smaller than 0.001 mm, which occurred in 100–500 iterations for our test objects.

There are several ways to deal with the scaling factors when fitting surfaces. In the first variant, scaling factors can be taken as the ratio of the nominal pixel sizes of the two scans. In the second variant, the scaling factors can be computed from surface data in the same way as the translation and rotation parameters. In the presence of errors, however, computed factors may lead to non-isotropic stretching of the image. To avoid image deformation one can use a third variant in which the scaling factors are treated as parameters in the search, but are then replaced by the nominal scaling factors. Since the last two variants involve searching the nine-dimensional rather than six-dimensional parameter space, we expected them to be more sensitive to errors in input data. In order to test this hypothesis, the three variants were compared.

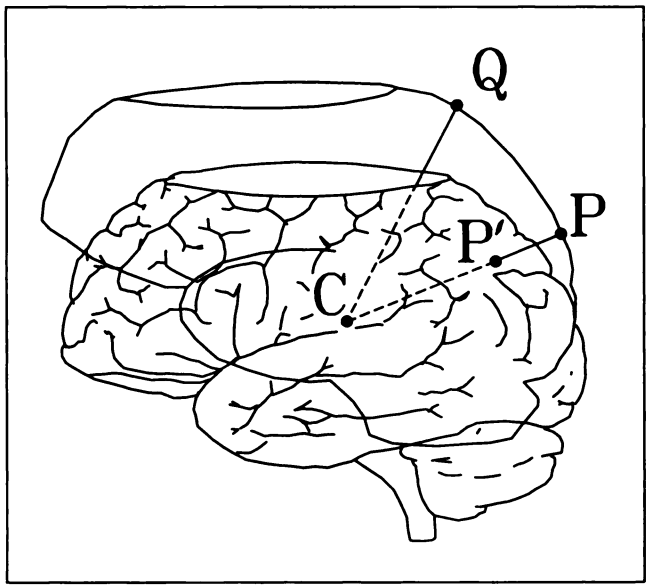

FIGURE 2. In the SFIT method, the distance PP' from each point P of the source scan to the target scan is computed. Point P' is the intersection of the target scan and the ray CP. Point C is the centroid of the target scan. Point Q, for which the intersection cannot be found, is omitted from the distance computation.

SFIT software developed by Pelizzari et al. provide the user with a graphical display of the surfaces and it lets the human operator to override the matching parameters (*1*). In order to study objectively two automated algorithms, our tests did not allow user interaction. In all other respects our implementation of PAX and SFIT techniques follows the published description of the algorithms. All programs were coded in C-language and executed on Sun Microsystems Sun 4 work station (Mountain View, CN).

Test Objects

Two types of test objects were employed: a MR scan of the brain and a mathematically defined ellipsoid.

Brain images of a 33-yr-old male volunteer were acquired on the 1.0 T Siemens Magnetom MR system (Siemens, Iselin, NJ) using a T1-weighted three-dimensional gradient echo sequence characterized by the repetition time 10 msec, echo time 4 msec, 15° flip angle, 23-cm field of view and 160 axial sections, each 1.15 mm thick. Since the scan covered the entire brain including the cerebellum, it allowed us to reformat the images and to simulate other acquisition geometries such as partial coverage, oblique orientation and section thickness as described below.

Independently, a computer program was written to generate arbitrary sections of an ellipsoid, a three-dimensional surface defined by the equation:

$$\frac{x^2}{r_x^2}+\frac{y^2}{r_y^2}+\frac{z^2}{r_z^2}=1. \qquad \text{Eq. 1}$$

The spatial coordinate x is defined as the left-right axis, y as the anteroposterior axis, and z as the rostrocaudal axis. With the proper choice of the three semi-axes r_x, r_y and r_z, the ellipsoid can be made to resemble the brain, heart, kidney and other body organs. Our tests used the values $2r_x$ = 144 mm, $2r_y$ = 172 mm, $2r_z$ = 120 mm and 230 mm field of view.

Data Analysis

In computing the translational and rotational parameters, errors are introduced by the discrete nature of the images and the surface points. To evaluate these errors we generated pairs of misregistered scans. The first scan, called the target scan, consisted of the phantom object in its initial position. (Brain images were translated to place the centroid at the geometric center of the three-dimensional image matrix.) The second scan, called the source scan, is derived from the target scan by rotation about the centroid, translation and reslicing at a desired section level and

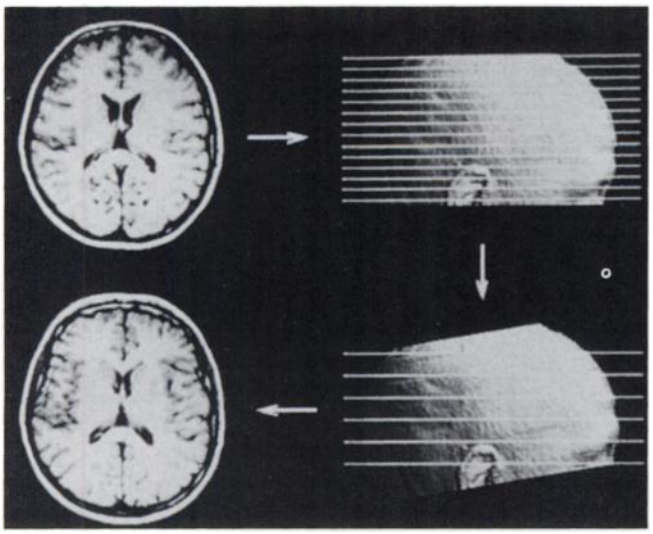

FIGURE 3. Overview of the method for generating pairs of misregistered scans.

thickness (Fig. 3). The coverage of the object is therefore different in the source and the target scans. The following parameters are analyzed: longitudinal (z-axis) coverage, magnitude of misregistration, number of sections, section thickness, sampling resolution of surface points, random and systematic deformations of the surface points and the degree of symmetry of the object.

Each image of the source and the target scans was sampled radially at fixed increments of the polar angle and the surface of the brain was detected by the edge-finding method. For the purpose of calculating parameters of the matching transformation, the images are now replaced by the surface points (Fig. 4). As seen in the figure, the surface coordinates of the ellipsoid are rounded off to fall on the finite resolution pixel grid. To simulate random noise, the x-coordinates and y-coordinates were perturbed by adding along the radial direction a normally distributed random variable.

Since the true amount of misregistration is known, the registration error can be decomposed into its translational and rotational components. (Note that such information is not available from the study of physical phantoms.) Translational error affects the entire object uniformly, while the error in rotation affects pixels in direct proportion to their distance from the center of rotation (in our case, the centroid). In studies involving centrally located structures, the effect of the rotational error is likely to be small. For the sake of conciseness we define the mean error E, measured in millimeters, to represent the average misregistration. To compute E, we first convert each of the three angular errors ax, ay and az to a mean linear displacement of the points within the object due to the rotation. The expected value of the radius of rotation about the z axis, for example, was computed using the expression:

$$r_{z,av} = \frac{3\sqrt{2\pi}}{8} \frac{r_x^2 r_y^2}{(r_x^2 + r_y^2)^{3/2}}, \qquad \text{Eq. 2}$$

and the linear displacement due to the rotation by angle az is

$$d_{az} = 2 r_{z,av} \sin\left(\frac{az}{2}\right), \qquad \text{Eq. 3}$$

with similar equations for the rotation about other two axes. The square of E is computed by adding the squares of these displacements and of the three translational errors.

RESULTS

Since our main interest is the registration of PET and MR brain images, in the first experiment we have chosen the acquisition parameters that simulate PET/MR coregistration. The source scan, approximating a routine MR brain scan, had an image resolution of 256 × 256, and a field of view of 23 cm with 18 contiguous 6.7-mm thick sections. The parameters of the target scan were typical of PET: 15 contiguous sections spaced every 7 mm. The longitudinal (z-axis) coverage of the brain was 100% for MR and 87% for PET. Brain contours were perturbed by computer-simulated Gaussian noise of mean 0 mm and standard deviation 0.8 mm. The misalignment of the two scans is given by rotations ax = 12°, ay = 3° and az = 6° and translations dx = 2 mm, dy = 4 mm and dz = 3 mm. These values represent the average misregistration between PET and MR scans at our institution.

The two versions of the PAX method were compared first. The PAX method based on boundary sampling processed 6,104 pixels and yielded translational error of 1.8 mm, a rotational error of 4.8 mm and a mean error equal to

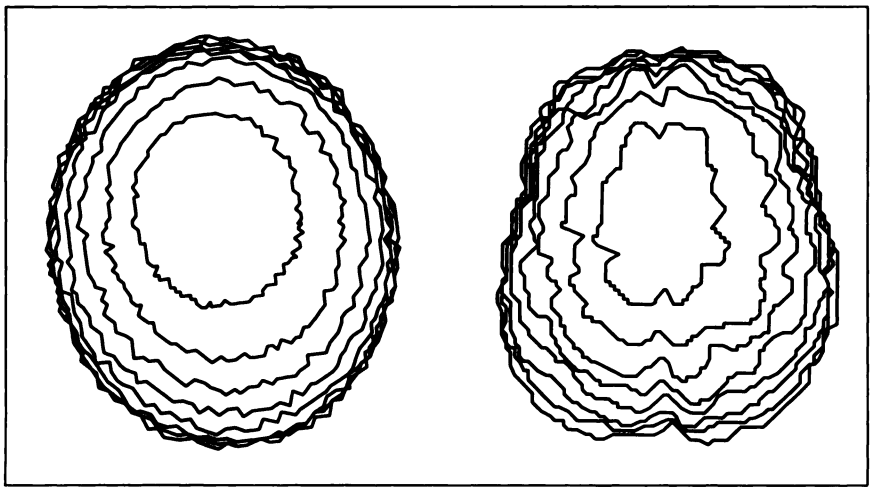

FIGURE 4. Surface points of the ellipsoid (left) and of the brain (right) show irregularities due to the finite resolution of the image. For clarity, only the top eight contours are shown.

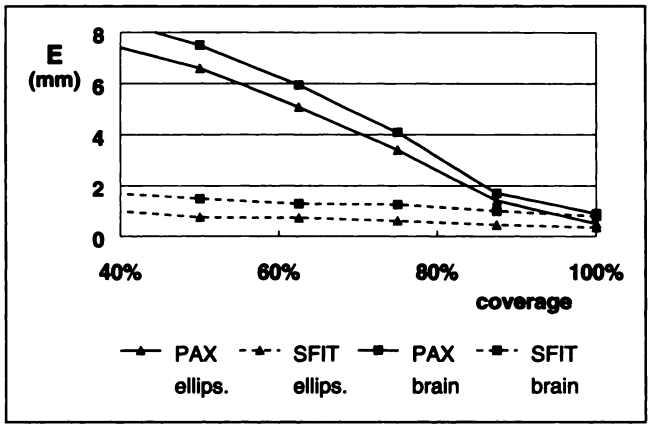

FIGURE 5. The effect of longitudinal brain coverage (%) on the accuracy of PAX and SFIT registration methods.

4.7 mm. The implementation based on the filled interior processed 225,112 pixels and yielded significantly better accuracy (translational error of 0.7 mm, rotational error of 1.1 mm and a mean error equal to 1.3 mm. Since a similar pattern was observed in other experiments, only the results based on filled interior sampling are reported below.

Applied to the same objects, the three variants of SFIT yielded essentially the same error, E = 0.3 mm. The two variants based on computing the scaling factors, however, required nearly three times as many iterations as the first variant. With one exception (see the results of global deformation below), the same behavior was found in other simulations. Unless noted otherwise, SFIT results are reported only for the first variant using nominal scaling factors.

Figure 5 shows the effect of varying the longitudinal brain coverage of the target scan while keeping the other scan parameters as described above. This experiment simulated scans acquired with PET cameras having a limited number of detector rings. The performance of PAX (most notably its rotational error, not shown in the figure) deteriorated for coverage <80%, while the SFIT method remained unaffected by the incomplete data. The errors in registering the brain and the ellipsoidal phantom were nearly equal. It should be noted that Figure 5 corresponds to the target scan that is centered on the object along the z-axis, i.e., approximately the same number of sections are missing from the top as from the bottom. A separate experiment, in which only the top sections or only the bottom sections are removed, confirmed a large error in PAX and no effect on SFIT.

Figure 6 illustrates the effect of misalignment between the source and target scan by a rotation about the x-axis. This experiment simulates the variability of the orientation of the scan plane. The angle ax of misalignment ranged from 0° to 90°. The misalignment of 90° simulated the registration of axial versus coronal sections. The longitudinal coverage was 100% for the source scan (18 sections) and 87% for the target scan (15 sections). The accuracy of SFIT method remains relatively stable, whereas the error of PAX method increases with misangulation. Similar results were obtained for scans misaligned by the y-axis rotation. For both methods, misalignment by translation and by z-axis (within image) rotation had no significant effect on accuracy.

In the next experiment we varied the number of sections with the coverage and misalignment kept as in the first experiment. Within the range 6–48 sections per scan, mean error was 0.3–1.7 mm for SFIT and 0.5–1.9 mm for PAX. Similar lack of effect was observed when we varied the section thickness (range 2–12 mm), the number of samples per contour (range 60–720) and the eccentricity of the object (range 1.1–1.3).

There are several practical situations when the assumption of rigidity is violated. For example, when registering CT and MR images of the head, different contours may be extracted due to the sensitivity of MR to the skin and CT to the bone. We have used ellipsoidal phantoms to model the deformation of a scan. When modeling *global* deformation, we changed the length of each semiaxis of the ellipsoid; for *local* deformation, only one octant (12.5% of the surface) was changed. The results (Fig. 7) suggest that the effect of local deformations is larger than the effect of global deformations. SFIT is nearly as sensitive as PAX to local deformations; it is significantly more sensitive to global deformations that exceed 3 mm. Results of modeling global deformation were remarkable, as they showed significantly better accuracy for the third variant of SFIT, i.e., the variant where the scaling factors are treated as search parameters, then replaced by the nominal values. In the 0–8 mm range, the error was 52% lower than error obtained using the first variant. There was no corresponding improvement for local deformation of contours.

The effect of random perturbation of contour points is minimal for the SFIT method; noise >1.5 mm (s.d.) begins to affect the accuracy of PAX (Fig. 8). One run of PAX registration took 3.6 sec, while the run time of SFIT was 3–11 min, depending on the number of iterations.

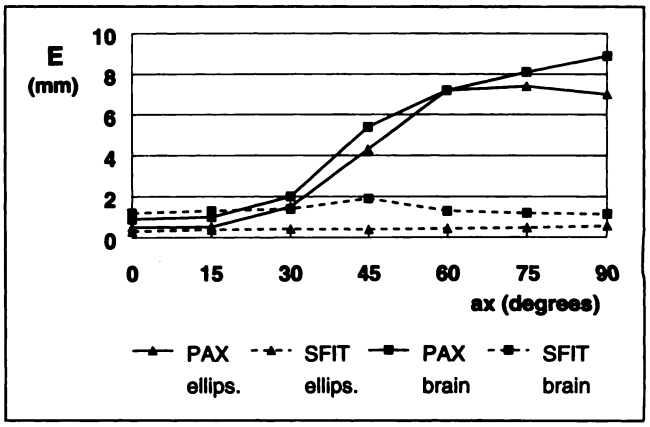

FIGURE 6. The effect of misalignment on accuracy. The two scans are misaligned by rotation ax about the x-axis. The coverage of the target scan is 87.5%. See text for description of other scan parameters.

DISCUSSION

We have identified parameters affecting the accuracy of two automatic three-dimensional registration techniques. Both methods were sensitive to the systematic deformation of contours. Moreover, the PAX method was found to be sensitive to incomplete scan coverage and to the x-axis and y-axis misangulation. The accuracy of both methods were relatively unaffected by varying the number of sections, section thickness and the number of contour points.

The ellipsoid, with symmetry along each axis, may seem to be an unrealistic test object. It should be noted, however, that the presence of symmetry makes the task of registration more difficult. Our experiments indicate increased error for registering "round" objects with eccentricity approaching 1.0. As seen in Figures 5 and 6, the accuracy of registering brain scans is closely matched by the accuracy of registering ellipsoidal objects using similar scan parameters.

Gamboa-Aldeco et al. registered multiple CT scans of the vertebrae with a radius of 25 mm (*4*). Using the PAX method and surface sampling, the registration error was 0.8° for a 20° misalignment about the y-axis and 1.6° for a 20° misalignment about the x-axis. While our results indicate larger errors, the discrepancy could be due to the difference in shape between the ellipsoidal object and the vertebrae, which have distinctive cross-sections along each plane. The trend of large rotational errors for incompletely scanned objects reported by Gamboa-Aldeco et al. (*4*) agrees with our results.

Alpert et al. describe improvement in the performance of the PAX method when the principal axes are derived from the volumes rather than the surfaces (*9*). They also demonstrate relatively insignificant translation errors and less than 2° z-axis rotation error for a section thickness as large as 20 mm. Our results confirm these findings, but indicate a significant sensitivity of PAX to the x-axis and y-axis misangulation, especially in the presence of random noise.

Pelizzari et al. used anthropomorphic head phantoms and fiducial markers to establish the accuracy of SFIT (*1*). The RMS error of registering PET images with CT was 1.4 mm; for PET/MR registration the error was 2.5 mm. Their study suggests that the slice thickness and pixel size are the key parameters affecting accuracy.

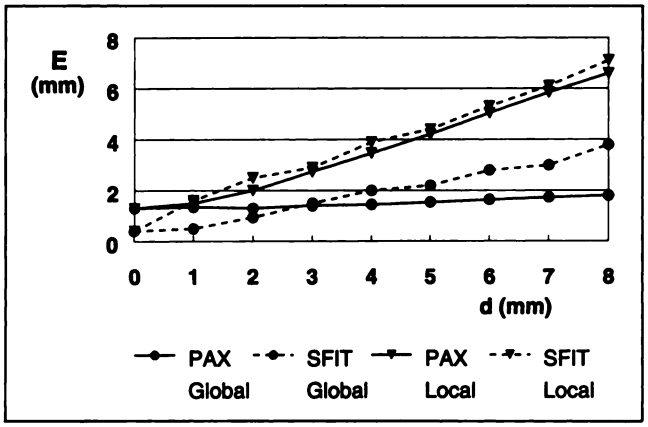

FIGURE 7. The effect of contour deformation on accuracy. The target scan is displaced by distance D while keeping the source scan fixed. In modeling global deformation, each semiaxis of the ellipsoid is changed. For local deformation, only one octant (12.5% of the surface) is changed.

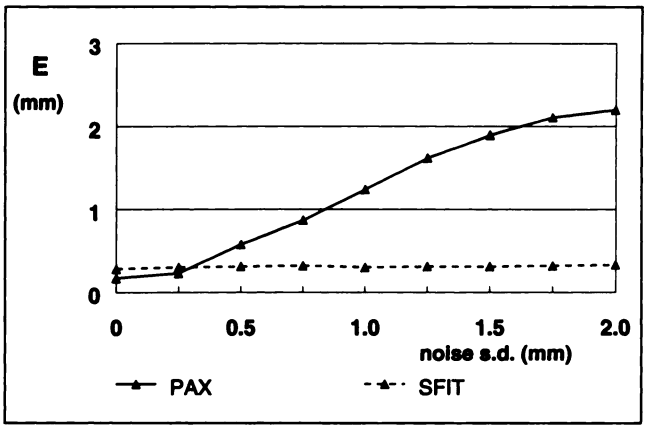

FIGURE 8. The effect of random deformation on accuracy. The coordinates of each contour point were perturbed by adding a normally distributed random error along the radial direction. See text for description of other scan parameters.

Our results confirm the feasibility of registration of multimodality scans with accuracy better than 2 mm using SFIT. However, we did not observe improvement in registration accuracy due to decreased slice thickness. The discrepancy could be due to the differences in implementation or to the difficulties in measuring registration error using fiducial markers.

In applications to nuclear medicine data, image registration is followed by three-dimensional reformatting of one modality along the transformed coordinates. Algorithms of various complexity have been proposed for the interpolating step (*11,12*). While the detailed discussion of these methods is beyond the scope of this report, it should be noted that the reslicing process introduces errors of its own. The factors affecting the interpolation error include section thickness, image resolution and the high frequency content of the image.

When registering functional and structural brain images using either SFIT or PAX, it is important to choose appropriate surfaces to be matched. Pelizzari et al. and Levin et al. have chosen external skull contours, easily extracted from CT and PET transmission scans (*1–3*). On the other hand, PET-to-PET coregistration by Levy et al. and SPECT-to-SPECT coregistration by Turkington et al. are based on the cortical brain surfaces (*6,13*). The use of bony surface makes coregistration insensitive to abnormalities in cortical tracer distribution, but it may introduce errors due to misalignment between transmission and emission images. In addition, transmission images are usually unavailable in SPECT and could be of poor quality in PET. Our analysis of local and global deformations (Fig. 7) can be helpful in selecting surfaces to be matched. For example, if the cortical boundary can be extracted from the functional scan with 4 mm accuracy and if there are no local deficits

in tracer distribution, the expected registration error is approximately 2 mm.

Our simulations suggest that SFIT is the method of choice for registering brain images. We have established that this method does not require user supervision and that there is no need to routinely include the scaling factors as search parameters. The major weakness of both methods is their sensitivity to local deformations in surface contours. In practical applications, however, the user of SFIT method can detect problematic cases by inspecting the residual distance D between surfaces.

ACKNOWLEDGMENTS

The authors thank Dr. C. Pelizzari from the University of Chicago for informative discussions. This work was supported by National Institute of Health grants 2RO1MH43965-04 and 3P3OAGOH051-0451. This work was also supported by NIMH grant 44833 and BNL contract DE-AC02-76CH00016.

REFERENCES

1. Pelizzari CA, Chen GTY, Spelbring DR, Weichselbaum RR, Chen C-T. Accurate three-dimensional registration of CT, PET, and/or MR images of the brain. *J Comput Assist Tomogr* 1989;13:20–26.
2. Levin DN, Pelizzari CA, Chen CTY, Chen C-T, Cooper MD. Retrospective geometric correlation of ME, CT, and PET images. *Radiology* 1988;169: 817–823.
3. Levin DN, Hu X, Tan KK, et al. The brain: integrated three-dimensional display of ME and PET images. *Radiology* 1989;172:783–789.
4. Gamboa-Aldeco A, Fellingham LL, Chen GTY. Correlation of 3D surfaces from multiple modalities in medical imaging. *Proc SPIE* 1986;626:467–473.
5. Belsole RJ, Hilbelink D, Llewellyn JA, Dale M, Stenzler S, Rayhack JM. Scaphoid orientation and location from computed, three-dimensional carpal models. *Orthop Clin North Am* 1986;17:505–510.
6. Levy AV, Brodie JD, Russel AG, Volkow ND, Laska E, Wolf AP. The metabolic centroid method for PET brain image analysis. *J Cereb Blood Flow Metab* 1989;9:388–397.
7. Vannier MW, Gayou DE. Automated registration of multimodality images. *Radiology* 1988;169:860–861.
8. Moshfeghi M, Rusinek H. Three-dimensional registration of multimodality images using the principal axes technique. *Phillips J Res* 1992;47:81–97.
9. Alpert NM, Bradshaw JF, Kennedy D, Correia JA. The principal axes transformation: a method for image registration. *J Nucl Med 1990;31:1717–1722.*
10. Siddon RL. Prism representation: a 3-D ray-tracing algorithm for radiotherapy applications. *Phys Med Biol* 1985;30:817–824.
11. Parker JA, Kenyon RV, Troxel DE. Comparison of interpolating methods for image resampling. *IEEE Trans Med Imag* 1988;2:31–39.
12. Goshtasby A, Turner DA, Ackerman LV. Matching of tomographic slices for interpolation. *IEEE Trans Med Imag* 1992;11:507–516.
13. Turkington TG, Jaszczak RJ, Greer KL, Coleman RE, Pelizzari CA. Correlation of SPECT images of a three-dimensional brain phantom using a surface fitting technique. *IEEE Trans Nucl Sci* 1992;39:1460–1463.