
Influence of Implementation Parameters on Registration of MR and SPECT Brain Images by Maximization of Mutual Information

Yang-Ming Zhu, PhD; and Steven M. Cochoff, MS

Nuclear Medicine Division, Philips Medical Systems, Cleveland, Ohio

Mutual-information maximization is one of the most popular algorithms for automatic image registration. However, many implementation issues have not been evaluated in a single, coherent context. **Methods:** Twenty-one registrations between MR and SPECT brain images (8 patients) were achieved by mutual-information maximization with different implementation strategies. The results of a popular strategy were chosen as the standard. All other results were compared with the standard, and the statistics of misregistrations were computed. The registration speed, accuracy, precision, and success rate were assessed. **Results:** Compared with trilinear interpolation, nearest-neighbor interpolation slightly sped the registration process, but with a lower success rate. The number of bins used to estimate the probability density function (pdf) affects the speed and robustness. Using fewer bins yielded a less robust registration. Adaptively changing the number of bins increased the registration speed and robustness. Simplex optimization increased the registration speed considerably, with a slightly degraded success rate. Simplex optimization with adaptive bin strategy improved the success rate and further decreased the registration time. Multiresolution optimization yielded a better success rate, with little effect on the accuracy and precision of registration. An increase in the number of resolution levels increased the success rate. Multisampling optimization also improved the success rate, but the results were less accurate and precise than those obtained with multiresolution optimization, with an increase in the number of levels decreasing the performance. Segmentation affected the registration speed and success rate. Because segmentation is problem specific, the effects were not conclusive. **Conclusion:** Different implementation strategies considerably affect the performance of automatic image registration by mutual-information maximization. On the basis of the experimental findings, we suggest that the best implementation strategy would include trilinear interpolation, adaptive change of the number of bins when estimating pdf, and exploitation of a simplex optimization algorithm with a multiresolution scheme.

Key Words: image registration; mutual information; MRI; SPECT

J Nucl Med 2002; 43:160–166

Received May 10, 2001; revision accepted Oct. 15, 2001.
For correspondence or reprints contact: Yang-Ming Zhu, PhD, Nuclear Medicine Division, Philips Medical Systems, 595 Miner Rd., Cleveland, OH 44143.
E-mail: yang-ming.zhu@cle.philips.com

Image registration and fusion of medical images are increasingly useful in research and patient care (1–3). Accurately associating data points on different images is required for diagnosis, treatment planning, and detection of changes. Many registration methodologies have been developed, reviewed (4), and evaluated (5–7).

From a practical point of view, automatic registration is desirable because its results are objective and because quantitative analysis is possible. Among the various automatic registration methods, maximization of mutual information of voxel intensities (8,9) is one of the most popular for multi- as well as single-modality image registration. The method is robust and accurate and does not require segmentation and preprocessing (5–7).

Mutual information is an information-theoretic concept. It quantifies the similarity between 2 random variables A and B and is defined by:

$$I(A, B) = \sum \sum p_{AB}(a, b) \log \frac{p_{AB}(a, b)}{p_A(a)p_B(b)},$$

where $p_{AB}(a, b)$ is the joint probability density function (pdf) and $p_A(a)$ and $p_B(b)$ are the marginal pdf's. In the context of image registration, random variable A is the voxel value in the reference image and random variable B is the voxel value in the floating image. A general flowchart of the implementation of mutual-information registration is given in Figure 1, where I_1 and I_2 are reference and floating images respectively. When the registration parameters are applied to the floating image, interpolation is generally needed such that the voxels are mapped to grid points. To compute the mutual information, one needs to estimate the pdf. To decide if a registration is optimal and, if it is not, how to update the registration parameters, one needs to use some optimization scheme.

Many implementation issues are involved in mutual-information image registration, such as interpolation methods, estimation of marginal and joint pdf's, optimization algorithms, multiresolution and multisampling schemes, and segmentation options. Although the reports (8–13) dealing with mutual-information registration discuss some aspects of these issues, not all of these parameters have been

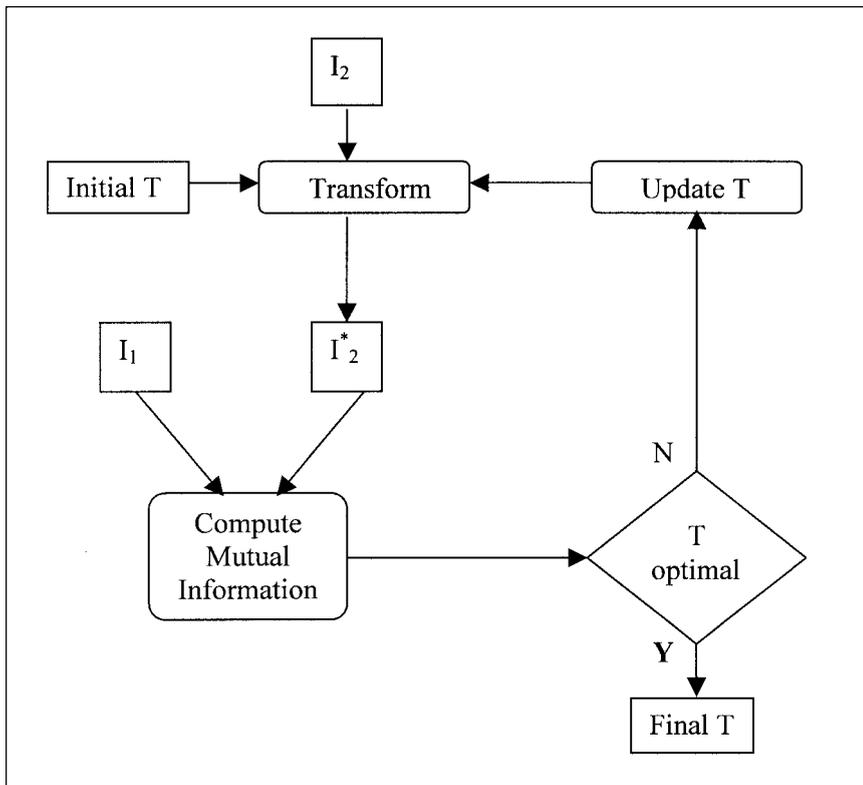


FIGURE 1. Flow chart of image registration by mutual-information maximization. I_1 = reference image; I_2 = floating image; I_2^* = transformed floating image; N = no; T = registration transformation; Y = yes.

evaluated in a single, coherent context. This study evaluated the use of different implementation strategies on 21 MR and SPECT brain image registrations, and we determined the best implementation strategy.

MATERIALS AND METHODS

Images

This study involved 8 patients who underwent brain MRI and nuclear transmission and emission scanning. The pertinent information on image files is in Table 1. The emission images using ^{201}Tl were excluded from this study because of their poor readability. Thus, 21 image pairs were available to be registered. All nuclear images were registered to their MRI counterparts.

Image Registration

All image pairs were registered by mutual-information maximization. We did not try to hand-register the images first. The source code of commercial software (Image Volume Registration; Philips Medical Systems, Cleveland, OH) was used as the base and was modified to incorporate different implementation strategies. The computer used was an AlphaStation (Digital Equipment Corp., Maynard, MA) with 192 MB of memory. The operating system was UNIX (version 4.0D; Digital).

All registration transformations were restricted to the rigid-body type. The registration parameter is a 6-dimensional vector, $(\theta_x, \theta_y, \theta_z, t_x, t_y, t_z)$, where $\theta_x, \theta_y,$ and θ_z are rotation angles in degrees around the x -, y -, and z -axes, respectively, and $t_x, t_y,$ and t_z are translation offsets in millimeters along the x -, y -, and z -axes, respectively.

Interpolation

When the registration transform was applied to the nuclear image, the original grid in the nuclear images was transformed to

a generalized nongrid coordinate system. The registration process needed to access the voxel values at grids of the transformed image. Thus, interpolation had to be used.

There are different interpolation methods. Those most popular are nearest neighbor and trilinear. Maes et al. (9) reported a trilinear partial-volume distribution method that is not an interpolation per se. In this study, we used nearest-neighbor and trilinear interpolations and compared their impact on registration.

Histogram

To compute the mutual information, one must have the marginal and joint pdf's of voxel values and their pairs. Two methods of estimating those pdf's have been published. Wells et al. (8) used a Parzen windowing method, which is time consuming. Maes et al. (9) used a histogram method, which is popular and used by many others (8–13). Because the latter approach is quick and easy to implement, we adopted it for this study.

To estimate the histograms, one maps the voxel values to an array or a sequence of bins. The number of bins is an important parameter. In our study, we used the same number of bins, n , for both images when estimating the marginal pdf's and used $n \times n$ when estimating the joint pdf. Values tested for n were 16, 32, and 64. The voxel values were linearly mapped to bins. Nonlinear mapping was not tried.

Multiresolution and multisampling optimization was used to avoid local minima. Suppose the image size is $16 \times 16 \times 16$. There then are, at most, $16 \times 16 \times 16$ gray-value pairs (when all voxels overlap). If 64×64 bins are used to estimate the joint pdf, then on average, there is only 1 voxel pair per bin. As a result, the statistical error in the joint pdf estimation would render the result meaningless. Thus, a small number of bins is desirable. As the following results show, however, a small number of bins can yield less robust registration, implying that a large number of bins is

TABLE 1
Image File Description

Patient	Modality	Image	Dimension	Voxel (mm ³)	Emission isotope	Comments
A	MR	Sag	256 ² × 128	1.0 ² × 1.5	—	M, 29 y old
	Tx	Axi	64 ² × 29	7.12 ³	^{99m} Tc	
	E	Axi	64 ² × 29	7.12 ³	^{99m} Tc	
B	MR	Sag	256 ² × 128	1.0 ² × 1.5	—	M, 21 y old
	Tx	Axi	64 ² × 24	7.12 ³	²⁰¹ Tl	
	E	Axi	64 ² × 24	7.12 ³	²⁰¹ Tl	Not used
	Tx	Axi	64 ² × 24	7.12 ³	^{99m} Tc	
C	MR	Sag	256 ² × 128	1.0 ² × 1.5	—	F, 70 y old
	Tx	Axi	64 ² × 22	7.12 ³	^{99m} Tc	
	E	Axi	64 ² × 29	7.12 ³	^{99m} Tc	
D	MR	Axi	192 × 256 × 24	1.17 ² × 6.0	—	M, 45 y old
	Tx	Axi	64 ² × 24	6.23 ² × 7.12	²⁰¹ Tl	
	E	Axi	64 ² × 34	6.23 ² × 7.12	²⁰¹ Tl	Not used
	Tx	Axi	64 ² × 31	6.23 ² × 7.12	^{99m} Tc	
	E	Axi	64 ² × 31	6.23 ² × 7.12	^{99m} Tc	
E	MR	Sag	180 ² × 120	1.41 ² × 1.5	—	M, 40 y old
	Tx	Axi	64 ² × 32	6.23 ² × 7.12	²⁰¹ Tl	
	E	Axi	64 ² × 32	6.23 ² × 7.12	²⁰¹ Tl	Not used
	Tx	Axi	64 ² × 26	6.23 ² × 7.12	^{99m} Tc	
	E	Axi	64 ² × 26	6.23 ² × 7.12	^{99m} Tc	
F	MR	Axi	256 ² × 124	0.94 ² × 1.0	—	M, 51 y old Before Diamox After Diamox
	E	Axi	64 ² × 47	3.50 ³	^{99m} Tc	
	E	Axi	64 ² × 45	3.50 ³	^{99m} Tc	
G	MR	Sag	192 × 256 × 120	1.41 ² × 1.5	—	M, 60 y old
	Tx	Axi	64 ² × 49	3.50 ³	²⁰¹ Tl	
	E	Axi	64 ² × 49	3.50 ³	²⁰¹ Tl	Not used
	Tx	Axi	64 ² × 53	3.50 ³	^{99m} Tc	
	E	Axi	64 ² × 53	3.50 ³	^{99m} Tc	
H	MR	Sag	192 × 256 × 120	1.41 ² × 1.5	—	M, 13 y old
	Tx	Axi	64 ² × 32	6.23 ² × 7.12	²⁰¹ Tl	
	E	Axi	64 ² × 32	6.23 ² × 7.12	²⁰¹ Tl	Not used
	Tx	Axi	64 ² × 28	6.23 ² × 7.12	^{99m} Tc	
	E	Axi	64 ² × 28	6.23 ² × 7.12	^{99m} Tc	

Tx = transmission scan; Sag = sagittal slices; Axi = axial slices; E = emission scan.

All SPECT scanners were PRISM 3000XP (Philips Medical Systems, Cleveland, OH). All MR scanners were Outlook (Philips Medical Systems).

required. To solve this paradox, we adaptively changed the number of bins. The number of bins was heuristically set to the value of the resolution; that is, if the resolution is 32, the number of bins is 32 (and 32 × 32 for the joint pdf). We did not try to optimize the number of bins at each resolution.

Optimization

The optimal registration is found by an iterative process, and iterative optimization is an important part of a registration method. Different optimization algorithms have been used previously: Powell, simplex, steepest decent, conjugate gradient, quasi-Newton, Levenberg-Marquardt least squares, and simulated annealing, among others (8–14). Because Powell and simplex are the most frequently used algorithms in this context, we evaluated only these. In the simplex algorithm, a simplex in hyperspace is deformed in a well-defined manner to enclose the minima (14). The Powell algorithm, a direction-set method in multidimensions, loops on a set of directions and finds the minima along each direction. Brent's method in 1-dimensional search was used in our

implementation (14). The direction matrix was initialized to a unitary matrix, consisting of the vector ($\theta_x, \theta_y, \theta_z, t_x, t_y, t_z$). The order of these parameters was fixed, and the termination condition for both algorithms was the same (tolerance = 0.001).

Multiresolution and Subsampling

Optimization methods cannot guarantee that a global optimal value will be found. The value can easily be trapped to a local minimum. To find a true global optimal value, simulated annealing or genetic algorithms can be applied. The speed of these algorithms limits their application to 3-dimensional image registration. In practice, an approach using multiresolution and subsampling proves to be helpful (10).

In both multiresolution and subsampling, the idea is to register the coarse images first and then to use the result as the starting point for fine-image registration. For multiresolution, the coarse image is obtained by averaging the voxels in the sampled neighborhood. For subsampling, the coarse image is obtained simply by sampling the original input image.

The number of levels (i.e., the number of resolutions or subsamplings) in multiresolution and subsampling and the resolutions or sampling frequencies in each level can influence the registration performance. In this study, we fixed the sizes of images at $8 \times 8 \times 8$, $16 \times 16 \times 16$, $32 \times 32 \times 32$, or $64 \times 64 \times 64$ for both multiresolution and subsampling. When the number of levels was 4, images of all above-mentioned sizes were used, from coarse images to fine images. When the number of levels was 3, 2, or 1, only the latter 3 finer images, 2 finer images, or 1 finer image, respectively, was considered.

Analysis of Registration Results

For these retrospective registrations, the ground truths are unknown and we cannot assess the absolute accuracy of different implementation strategies. To assess the influence of implementation parameters on registration results and compare the registration performance under different conditions, we established an artificial standard. We defined a standard implementation as one that uses trilinear interpolation, a fixed histogram bin (64×64) to estimate the marginal as well as joint histogram, and the Powell optimization algorithm, with a multiresolution strategy (resolutions change from 8 to 16 to 32 to 64). This implementation is typical of those used by other researchers (except that we used more levels of resolution) (10).

The misregistration parameter was defined as the difference between the registration parameter and the standard. A study has shown that a trained clinician can detect differences from the registration parameter of 4° in the x - and y -rotation angles, 2° in the z -rotation angle, 2 mm in the x - and y -translations, and 3 mm in the z -translation (10). If all misregistration parameters were within these detection thresholds, the registration was regarded as a success. If any misregistration parameter was outside this detection threshold, the registration results were visually assessed, because if the standard registration parameter is on the borderline of a failed registration and the current registration is on the opposite borderline, the difference between them can be as large as twice the detection threshold. If the visual inspection identified the registration as being good, it was treated as a success.

The mean and SD of the misregistration parameters of successful registrations were then calculated. If the mean was close to zero, the result was regarded as accurate. If the SD was small, the result was regarded as precise. A robust implementation gave an accurate and precise registration with a high success rate.

The accuracy of registration can also be evaluated by the absolute position error of selected regions or points. This error is known to comprise the error in translation offsets and the error caused by rotation angles (5). Thus, we did not report the absolute position error for this study.

The average run time of each implementation was compared with the standard implementation. If the other implementation ran faster than the standard, we defined a speedup as:

$$\text{speedup} = \frac{\text{standard run time} - \text{run time}}{\text{run time}} \times 100\%;$$

otherwise, we defined a slowdown as:

$$\text{slowdown} = \frac{\text{run time} - \text{standard run time}}{\text{standard run time}} \times 100\%.$$

RESULTS

Standard Implementation

The standard implementation achieved 20 successful registrations (95%) as judged by visual inspection, with an average run time of 170 s. A typical registration is shown in Figure 2. The angle parameter ranged from -24° to 0° , and the translation parameter ranged from -12 to 54 mm. We replaced the failed registration with the registrations given by the adaptive bin scheme that was judged to be successful by visual inspection. As the following results show, the adaptive bin scheme gave registrations very close to those of the standard, in terms of the extent of the difference and the SD of the difference. If the failed registration had been a success, the registration results would have been close to those given by the adaptive bin scheme. Thus, this replacement did not have an adverse effect on our evaluations; that is, what was replaced was reliable.

The registration achieved through mutual-information image registration has been reported to have a subvoxel accuracy (5). Visual inspection of the results of standard registration was satisfactory. The results of this standard implementation were also consistent when different initial registrations were introduced. Thus, we assume that the standard registration is close to the unknown truth.

Interpolation

To evaluate the effects of nearest-neighbor interpolation, we used that interpolation and kept all other implementation parameters the same as for the standard implementation. The results showed that this implementation failed to register 5 image pairs, with a success rate of only 76%. The mean and SD of misregistration parameters are shown in Table 2. The average run time for this implementation was 161 s, which amounts to a speedup of 6%.

Comparing the data in Table 2 with those in Tables 3–6, we find that the implementation using nearest-neighbor interpolation has a large SD. It is generally believed that nearest-neighbor interpolation can speed up the registration because less computation is involved in the interpolation. Our data do not support this belief. It is true that each iteration cycle takes less time if nearest-neighbor interpolation is used, but the registration time is determined by the number of iterations and the time per iteration. The implementation using nearest-neighbor interpolation is likely to take more iterations. Because the nearest-neighbor interpolation is insensitive up to 1 voxel, there was concern that subvoxel registration accuracy could not be achieved. Our data indicate that although this implementation yields a result with a large SD, the registration result, compared with that obtained using the standard implementation, still achieved subvoxel accuracy (the error from translation offsets was 1.24 mm on average).

Bin Size

In the standard implementation, we changed the number of bins to 16×16 and 32×32 . The number of bins was also changed adaptively. The statistics of that implementation are tabulated in Table 3.

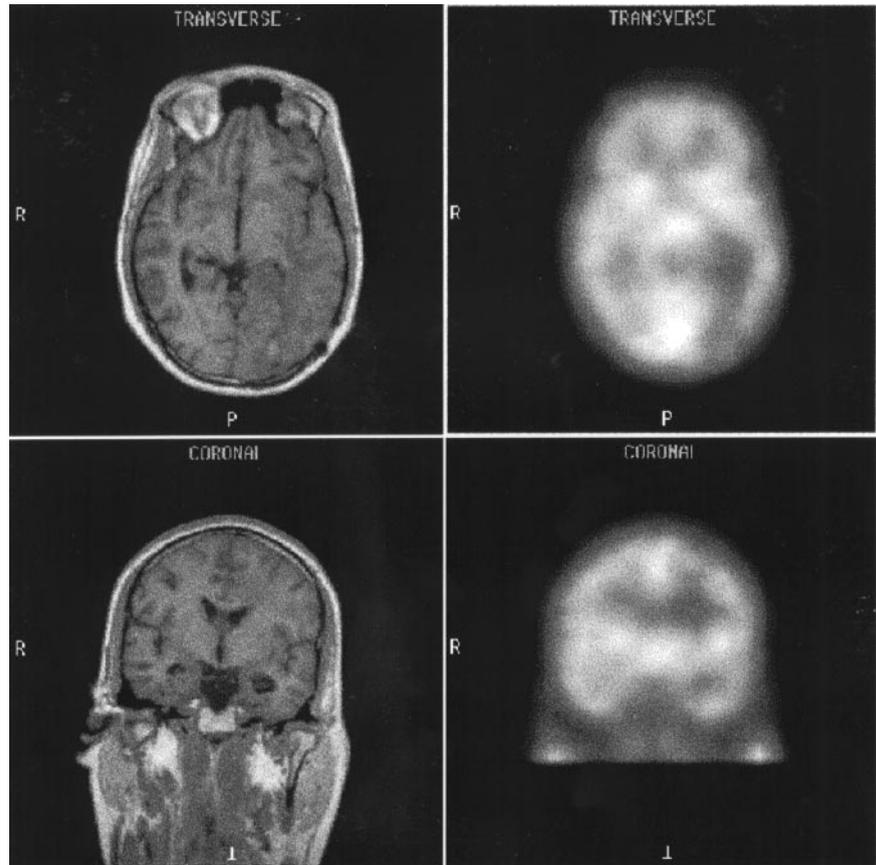


FIGURE 2. Typical result for registration of MR and SPECT images. Transverse and coronal views of MR images are displayed on left, and those of SPECT images are on right.

When the number of bins was small, less computation was involved. However, this did not necessarily speed up registration, as the data in Table 3 reveal. When the bin was 16×16 , the registration slowed by 5%, from 170 to 178 s. When the bin was 32×32 , the speed improved by 11%. The success rate for the latter case was reduced. When the number of bins was adaptively changed, the speedup was 13%.

When the number of bins was smaller, the mean and SD of the misregistration parameters were larger, indicating that a smaller number of bins may degrade registration accuracy and precision. The adaptive scheme registered all image pairs better than did the standard implementation, and the results were accurate and precise.

Optimization Algorithm

The Powell direction-set method and the simplex downhill method were the optimization algorithms used for comparison, and the results are shown in Table 4. The result of simplex optimization with a fixed number of bins was close

to that of the standard but failed to register 2 image pairs. The speedup was 91%.

When we added the adaptive bin scheme to simplex optimization, all image pairs were successfully registered, and the registration speed was improved further (speedup, 110%). However, the mean and SD of $\delta\theta_x$ and δt_y increased. The misregistrations of the 2 cases that failed by the simplex-only method were large. Nevertheless, the registration quality was good by visual inspection.

Multiresolution and Subsampling Strategy

The implementation using different levels of resolution or sampling were studied, and the statistics of the misregistration parameters are reported in Table 5. For multiresolution, when the number of resolution levels was small, the success rate decreased and the processing time increased. However, accuracy and precision were practically independent of the number of levels. The registration speed was related to the number of levels. The registration slowdowns were 32%,

TABLE 2
Mean and SD of Misregistration Parameters of Implementation Using Nearest-Neighbor Interpolation

Misregistration parameter						Time (s)	Success rate (%)
$\delta\theta_x$ (deg)	$\delta\theta_y$ (deg)	$\delta\theta_z$ (deg)	δt_x (mm)	δt_y (mm)	δt_z (mm)		
0.66 ± 1.32	-0.25 ± 0.53	-0.33 ± 1.14	-0.14 ± 0.49	-0.29 ± 1.00	1.20 ± 1.26	161	76

TABLE 3
Mean and SD of Misregistration Parameters of Implementation Using Different Number of Bins

No. of bins	Misregistration parameter						Time (s)	Success rate (%)
	$\delta\theta_x$ (deg)	$\delta\theta_y$ (deg)	$\delta\theta_z$ (deg)	δt_x (mm)	δt_y (mm)	δt_z (mm)		
16	-0.33 ± 2.00	0.21 ± 0.99	0.00 ± 0.98	0.15 ± 0.47	0.50 ± 1.58	1.43 ± 1.52	178	95
32	0.22 ± 0.97	0.30 ± 0.81	-0.07 ± 0.57	0.23 ± 0.41	-0.17 ± 0.65	0.54 ± 0.80	153	90
AD	0.12 ± 0.65	0.07 ± 0.58	-0.23 ± 0.62	0.04 ± 0.35	-0.08 ± 0.58	0.24 ± 0.42	150	100

AD = adaptive bin.

TABLE 4
Mean and SD of Misregistration Parameters of Implementation Using Simplex Optimization

Al	Misregistration parameter						Time (s)	Success rate (%)
	$\delta\theta_x$ (deg)	$\delta\theta_y$ (deg)	$\delta\theta_z$ (deg)	δt_x (mm)	δt_y (mm)	δt_z (mm)		
SF	-0.23 ± 0.58	0.27 ± 0.71	-0.01 ± 0.70	0.10 ± 0.34	0.24 ± 0.35	-0.20 ± 0.46	89	90
SA	-1.02 ± 1.99	0.03 ± 0.75	-0.07 ± 0.67	-0.07 ± 0.38	0.62 ± 1.18	-0.09 ± 0.46	81	100

Al = algorithm; SF = simplex with fixed number of bins; SA = simplex with adaptive number of bins.

TABLE 5
Mean and SD of Misregistration Parameters of Implementation Using Different Levels of Resolution and Sampling

Level	Misregistration parameter						Time (s)	Success rate (%)
	$\delta\theta_x$ (deg)	$\delta\theta_y$ (deg)	$\delta\theta_z$ (deg)	δt_x (mm)	δt_y (mm)	δt_z (mm)		
R1	-0.36 ± 0.99	-0.04 ± 0.52	0.12 ± 0.69	-0.08 ± 0.27	0.26 ± 0.60	0.06 ± 0.26	225	62
R2	0.13 ± 0.54	0.32 ± 0.67	0.10 ± 0.54	0.12 ± 0.38	-0.06 ± 0.35	-0.07 ± 0.55	186	86
R3	0.26 ± 0.80	0.08 ± 0.55	0.03 ± 0.74	0.16 ± 0.39	-0.07 ± 0.48	0.11 ± 0.55	181	95
S1	0.06 ± 0.74	0.03 ± 0.72	-0.25 ± 0.81	0.59 ± 0.65	1.18 ± 0.46	-1.79 ± 1.26	216	86
S2	0.25 ± 0.94	0.38 ± 0.86	0.11 ± 0.68	0.71 ± 0.69	1.11 ± 0.80	-1.91 ± 1.18	197	95
S3	-0.08 ± 0.79	0.19 ± 0.67	0.01 ± 0.66	0.59 ± 0.63	1.32 ± 0.65	-1.91 ± 0.96	211	86
S4	0.49 ± 0.86	0.32 ± 0.90	0.18 ± 0.81	0.61 ± 0.71	0.90 ± 0.46	-2.09 ± 1.00	215	86

R1–R3 = multiresolution technique with 1, 2, or 3 resolution levels; S1–S4 = subsampling technique with 1, 2, 3, or 4 subsampling levels.

TABLE 6
Mean and SD of Misregistration Parameters of Implementation Using Segmentation

Segment	Misregistration parameter						Time (s)	Success rate (%)
	$\delta\theta_x$ (deg)	$\delta\theta_y$ (deg)	$\delta\theta_z$ (deg)	δt_x (mm)	δt_y (mm)	δt_z (mm)		
L2	-0.01 ± 0.93	0.29 ± 0.92	-0.04 ± 0.54	0.08 ± 0.60	0.01 ± 0.78	0.03 ± 0.49	187	95
L5	0.12 ± 1.11	-0.08 ± 0.70	0.04 ± 0.72	0.00 ± 0.51	0.08 ± 0.93	0.54 ± 0.58	165	100
L10	-0.45 ± 1.48	0.10 ± 0.85	0.05 ± 0.97	0.08 ± 0.49	0.49 ± 0.81	0.72 ± 1.28	201	81
H80	-0.19 ± 0.88	0.10 ± 0.67	0.01 ± 0.70	0.07 ± 0.33	0.12 ± 0.55	0.03 ± 0.42	178	95
H90	0.22 ± 0.97	0.38 ± 0.89	0.09 ± 0.56	0.17 ± 0.48	-0.31 ± 0.88	0.03 ± 0.52	178	95

L2, L5, and L10 = clamping of voxels whose values are <2%, <5%, and <10%, respectively, of highest voxel value; H80 and H90 = clamping of voxels whose values are >80% and >90%, respectively, of highest voxel value.

9%, and 6% for 1, 2, and 3 levels, respectively. The standard implementation (with 4 levels) was the fastest.

For subsampling, the success rate was maximized when the number of sampling frequencies (subsampling levels) was 2. Further increasing the number of sampling frequencies did not improve the success rate. In fact, the success rate decreased when the number of sampling frequencies was 3 and 4, probably because the interdependency among neighboring voxels decreased. The means and SDs of δt_y and δt_z were large, indicating that the subsampling strategy was not accurate and precise and thus was not ideal. The registration slowdowns were 27%, 16%, 24%, and 26% for 1, 2, 3, and 4 sampling frequencies, respectively. On average, the multiresolution approach was faster than the subsampling approach.

The multiresolution and subsampling registrations were not identical for 1 resolution level and 1 subsampling level. This occurs because the voxels in the high-resolution MR image are averaged in the multiresolution technique but are directly sampled in the subsampling technique. The nuclear image in either case was the same.

In the multiresolution registration, the neighboring voxels were averaged. Averaging was slower than pure sampling but helped the overall registration performance.

Segmentation Effects

A meaningful segmentation may depend on the image content and the properties of the voxel-value histograms (13). A fixed segmentation of this kind would not be applicable to all images. Nevertheless, our segmentation was based on the percentile of voxel values. If the voxel value was less than some percentage of the maximum voxel value (background noise, for example), it was clamped. If the voxel value was larger than some percentage of the maximum voxel value (hot spots in SPECT and PET, for instance), it was clamped. We clamped only 1 end, although low-end and high-end clamping can be applied simultaneously.

The statistics of misregistration parameters are shown in Table 6. The data reveal that clamping the low-end voxels had a more dramatic effect on the registration performance. For example, when we clamped voxels whose values were <10% of the highest voxel value, the success rate was reduced. For <5% and <10% low-end clamping, the SDs of some registration parameters increased, indicating that different parameters have different sensitivities to low-end clamping. Clamping the high-end voxels had almost no effect on the registration performance (if more voxels had been clamped, the performance would have been degraded). The difference in the influence of low-end and high-end voxels was caused by the difference in the population of low-end and high-end voxels.

CONCLUSION

The performance of mutual-information maximization for registration of MR and SPECT brain images has been evaluated using different implementation strategies, includ-

ing interpolation methods, histogram bin strategies, optimization algorithms, multiresolution or subsampling schemes, and segmentation options. Our results indicate that different implementation strategies have a disparate performance in terms of speed, accuracy, precision, and success rate. On the basis of this study, we suggest that the best implementation will use trilinear interpolation, adaptively changing the number of bins when estimating pdf, using the simplex optimization algorithm, and exploiting a multiresolution strategy with as many levels as possible. Segmentation can be combined with the registration process, which has to be problem specific. Although this recommendation is based on the study of MR and SPECT images, we hope it will also be applicable to the registration of images from other modalities.

ACKNOWLEDGMENTS

The authors thank Dr. Michael Hartshorne, who provided the image files used in this study. Thanks are also due to the anonymous referees for helpful and constructive comments and suggestions and to Tom Brendza for proofreading the manuscript.

REFERENCES

1. Vera P, Kaminska A, Cieuta C, et al. Use of subtraction ictal SPECT coregistration to MRI for optimizing the localization of seizure foci in children. *J Nucl Med.* 1999;40:786-792.
2. Brinkmann BH, O'Brien TJ, Aharon S, O'Connor MK, Mullan BP. Quantitative and clinical analysis of SPECT image registration for epilepsy studies. *J Nucl Med.* 1999;40:1098-1105.
3. Pfluger T, Vollmar C, Wismuller A, et al. Quantitative comparison of automatic and interactive methods for MRI-SPECT image registration of the brain based on 3-dimensional calculation of error. *J Nucl Med.* 2000;41:1823-1829.
4. Maintz JBA, Viergever MA. A survey of medical image registration. *Med Image Anal.* 1998;2:1-36.
5. West J, Fitzpatrick JM, Wang MY, et al. Comparison and evaluation of retrospective intermodality brain image registration techniques. *J Comput Assist Tomogr.* 1997;21:554-566.
6. Fitzpatrick JM, Hill DLG, Shyr Y, West J, Studholme C, Maurer CR. Visual assessment of the accuracy of retrospective registration of MR and CT images of the brain. *IEEE Trans Med Imaging.* 1998;17:571-585.
7. West J, Fitzpatrick JM, Wang MY, et al. Retrospective intermodality registration techniques for images of the head: surface-based versus volume-based. *IEEE Trans Med Imaging.* 1999;18:144-150.
8. Wells WM, Viola PV, Atsumi H, Nakajima S, Kikinis R. Multi-modal volume registration by maximization of mutual information. *Med Image Anal.* 1996;1:35-51.
9. Maes F, Collignon A, Vandermeulen D, Marchal G, Suetens P. Multimodality image registration by maximization of mutual information. *IEEE Trans Med Imaging.* 1997;16:187-198.
10. Studholme C, Hill DLG, Hawkes DJ. Automatic three-dimensional registration of magnetic resonance and positron emission tomography brain images by multiresolution optimization of voxel similarity measures. *Med Phys.* 1997;24:25-35.
11. Ritter N, Owens R, Cooper J, Eikelboom RH, van Saarloos PP. Registration of stereo and temporal images of the retina. *IEEE Trans Med Imaging.* 1999;18:404-418.
12. Maes F, Vandermeulen D, Suetens P. Comparative evaluation of multiresolution optimization strategies for multimodality image registration by maximization of mutual information. *Med Image Anal.* 1999;3:373-386.
13. Carrillo A, Duerk JL, Lewin JS, Wilson DL. Semiautomatic 3-D image registration as applied to interventional MRI liver cancer treatment. *IEEE Trans Med Imaging.* 2000;19:175-185.
14. Press WH, Teukolsky SA, Vetterling WT, Flannery BP. *Numerical Recipes in C: The Art of Scientific Computing.* 2nd ed. Cambridge, U.K.: Cambridge University Press; 1999:394-455.