

An automated image-registration technique based on multiple structure matching

Paula L. Petti

Department of Radiation Oncology, University of California, San Francisco, San Francisco, California 94143-0226

Marc L. Kessler

University of Michigan Medical Center, Ann Arbor, Michigan 48109

Terri Fleming and Samuel Pitluck

Lawrence Berkeley Laboratory, Engineering Division, Berkeley, California 94720

(Received 27 August 1994; accepted for publication 24 June 1994)

A new image-registration technique that matches multiple structures on complementary imaging data sets (e.g., CT and MRI) has been developed and tested with both phantom and patient data. The algorithm assumes a rigid-body transformation and is suitable for correlating structures within the cranium or at the skull base. The basic premise of the new technique is that an optimum transformation is achieved when the relative volume lying outside of the intersection between a structure and its transformed counterpart is a minimum. This relative volume is calculated numerically using a random sampling approach, and a binary searching algorithm was used to step through the nine-dimensional parameter space consisting of three rotation angles, three scaling factors and three components of a translation vector. For the nine tests using phantom data, the automated structure-matching technique was able to predict the correct rotation angles to within $\pm 1^\circ$. The expected clinical performance of the new technique was assessed by comparing results obtained with the new method to those obtained using other techniques for 12 patients who were treated with charged particles at Lawrence Berkeley Laboratory (LBL) and who had image-registration studies performed as part of their treatment plan. For 9 of the 12 patients considered, the new structure-matching technique produced a significantly better registration than the older methods, as measured by the resultant average relative volume lying outside of the intersection between any structure and its transformed counterpart. For the other three patients, results were not significantly different for the new structure-matching method and the older techniques.

I. INTRODUCTION

Several image-registration techniques have been proposed that allow volumes-of-interest to be transformed from one imaging modality to another.¹⁻¹⁰ Registration between magnetic resonance imaging (MRI) and positron emission tomography (PET) provides a correlation between anatomy and function. Registration between computed tomography (CT) and MRI is a valuable tool in the implementation of conformal radiotherapy and radiosurgery. For example, precise dose localization in charged-particle radiotherapy requires both precise target definition, which may be achieved using MRI, and accurate dose calculations and compensator designs, which are derived from the tissue-density information provided by CT.^{11,12}

Most image-registration techniques rely on matching anatomical structures or surfaces²⁻⁹ or use externally placed markers (points or lines)²⁻⁴ to determine the registration. Recently, however, Woods *et al.*¹⁰ described a technique for MRI-PET registration that correlates pixel intensities instead of anatomical or fiducial landmarks. Although the images must be edited to exclude nonbrain tissue prior to the registration, this technique is automated and accurate to within 2 to 3 mm.

The point and line-matching methods require that external markers be placed in predetermined locations on the patient prior to both imaging studies. This is not always convenient or even possible, especially when the imaging studies are performed at different institutions. However, the advantage

of these techniques is that they are potentially very accurate. For example, Kessler *et al.* reported accuracies of between 1 and 2 mm using these methods.²

Structure-matching techniques fall into several categories. One class of methods, described by Kessler *et al.*,^{2,3} Pelizzari *et al.*,⁵ and Chen *et al.*,⁷ minimize the distance between corresponding points on the skull surface. These investigators reported accuracies of between 1 and 2 mm for this technique. Recently, van Herk and Kooy⁹ have proposed a new structure-matching algorithm using chamfer matching that is both accurate and fast. Another class of structure-matching image-registration techniques relies on the geometric properties of solid objects,^{6,8,13} for example, the property that any three-dimensional solid body has an inherent set of three principal axes, to determine a linear transformation between two imaging studies.

A final type of structure-matching techniques uses an interactive approach whereby several structures are aligned simultaneously.² The transformation parameters are adjusted manually while viewing results on a graphics display device. Although this technique relies to a large extent of the subjective judgement of the user, it has been applied successfully in the clinic.² The use of several structures to perform the registration may account for the success of this approach. In this paper, we describe a new image-registration method that was motivated, in part, by the success of the interactive technique. Like the interactive method, the new method correlates several structures simultaneously. However, the new




FIG. 1. The XOR-structure matching algorithm minimizes the volume outside of the intersection between a structure and its transformed counterpart.

technique requires no user intervention once the structures have been outlined on both image data sets. The new technique is evaluated with both phantom and patient studies. The phantom studies provide an objective assessment of the performance of the algorithm, whereas the patient studies demonstrate the performance of the algorithm in realistic clinical applications. For the patient studies, results obtained using the new technique are compared to those which were obtained as part of the treatment planning process using either the point-matching or interactive image-registration techniques.

II. METHODS

A. Structure-matching algorithm: Overview

The basic premise of the new structure-matching algorithm is that an optimal transformation is achieved when the volume lying outside of the intersection between a structure and its transformed counterpart is a minimum. The volume contained within the intersection between the two structures is, thus, maximized, as illustrated in Fig. 1. Since the volume lying outside of the intersection represents the volume in one structure or the other, but not both, it is described as the exclusive-OR (XOR) volume. R_{XOR} is defined as the ratio of the XOR volume to the volume of the structure in question. After each transformation, the algorithm calculates the number of structures for which there is a significant improvement (i.e., a reduction) in the XOR ratio, R_{XOR} . The number of structures for which there is a significant worsening (i.e., an increase) in R_{XOR} is also calculated. If the number of structures for which R_{XOR} improves significantly exceeds the number of structures for which R_{XOR} is significantly worsened, then the current set of parameter values defining the transformation is deemed superior to the previous set.

The program is written in VAX-11 Fortran, and the investigations reported in this paper were executed on a Microvax 3500 workstation (Digital Equipment Corporation, Maynard, MA). The code has also been implemented on a DEC 3000 AXP model 500/500s workstation (Digital Equipment Corporation, Maynard, MA). The structure outlines are drawn using the Lawrence Berkeley Laboratory (LBL) treatment-planning system.¹⁴ Once the transformation is determined with the new algorithm, structures are transformed and sliced in planes defined by the slice indices of the second study using the structure-mapping code described by Kessler *et al.*^{2,3}

B. XOR ratio calculation

Relevant structures are outlined on successive slices through both sets of images either manually or using a threshold-edge detection algorithm.¹⁵ Polygonally tiled surfaces are then created to represent these structures.¹⁶ For any given structure and its transformed counterpart, R_{XOR} is determined numerically by generating a sufficiently large number of random points lying within a bounding box containing the two intersecting volumes. For each structure,

$$R_{\text{XOR}} = V_{\text{XOR}} / V \approx N_{\text{XOR}} / N, \quad (1)$$

where V_{XOR} is the XOR volume and V represents the structure volume. N_{XOR}/N is the ratio of the number of random points falling within the structure or its transformed counterpart, but not both, to the number falling within the original structure. To calculate this ratio, each transformed structure is sliced at levels that correspond to the slice indices of the second study. Random points are then distributed inside the bounding box on each slice. N_{XOR} equals the sum, over all slices, of the points lying inside the XOR region. N is calculated similarly. The number of random points used per slice per structure is defined by the user. The algorithm used to determine whether a point is inside or outside a given structure outline counts the number of times a horizontal line drawn through the point in question crosses the outline. If the number of crossings is odd, the point is inside, if the number is even, the point is outside. This standard technique works for both concave and convex polygons.

There is an inherent uncertainty in these volume calculations due to the approximate nature of the random-point method. The standard deviation in the XOR ratio defined by Eq. (1) for a given structure is

$$\sigma_{R_{\text{XOR}}} = R_{\text{XOR}} \sqrt{\frac{1}{N_{\text{XOR}}} + \frac{1}{N} - \frac{2}{N_0}}, \quad (2)$$

where N_0 represents the number of random points used in the calculation for the structure in question, and all of the other quantities are as defined for Eq. (1). A derivation of this equation is given in the Appendix. The optimization procedure takes this uncertainty into account as described in the next section.

For both the phantom studies and the patient studies, 400 random points per structure outline were used on each CT slice. The sensitivity of the results to the number of random points used in these calculations is discussed in Sec. IV.

C. Optimization procedure

Nine parameters describe the linear transformation between the two imaging data sets: three rotation angles, three scaling factors and three components of a translation vector. The scale factors describe potential linear scale distortions in the image data (particularly, in the MRI data), and are optimized only if the user requests it. In principle, these scale factors can be determined *a priori* by proper calibration of the MRI unit, and the image data can be corrected before the registration process. If the user requests no scale factor optimization, the values of all three scaling parameters are fixed at unity.

The new structure-matching image-registration algorithm uses a binary searching technique to step through the parameter space. The step size is initially set to a large value. For a given parameter, p , the algorithm begins by calculating the XOR ratios for $p=p_0$, where p_0 is the initial guess for this parameter, for $p=p_0+s$ and for $p=p_0-s$, where s represents the initial step size. Of these three parameter values, the one that maximizes $(S_{\text{improve}}-S_{\text{worse}})$, where S_{improve} is the number of structures for which the XOR ratio is *significantly* reduced, and S_{worse} is the number of structures for which the XOR ratio increases *significantly*, is then chosen as the new optimal value for p . The step size, s , is then reduced by a factor of 2, and the search proceeds until a user-defined number of binary steps, typically between three and five, have been taken. The number of binary steps required to optimize a given parameter is determined by the smallest practical increment in that parameter. For example, variations in the rotation angles of less than about 1° generally produce indiscernible changes in the transformed structures. Thus, if 1° is the smallest angular increment of interest, and, for example, the initial step size is set to 16° , five binary steps are required in the minimization process.

The uncertainty in the volume calculation due to the approximate nature of the random point method is taken into account by defining a significant improvement in the XOR ratio for a given structure in terms of the following inequality:

$$R_{\text{XOR}}^{\text{new}} < R_{\text{XOR}}^{\text{old}} - f\sigma_{R_{\text{XOR}}}^{\text{old}} \quad (3)$$

where $R_{\text{XOR}}^{\text{new}}$ and $R_{\text{XOR}}^{\text{old}}$ represent values of the XOR ratio for the new and old parameter sets, and $\sigma_{R_{\text{XOR}}}^{\text{old}}$ is the standard deviation in $R_{\text{XOR}}^{\text{old}}$ [Eq. (2)]. The multiplicative factor f was set to 1 when optimizing the rotation angles and to 2 when adjusting the scaling parameters. These values for f were determined empirically. A *significant worsening* in the XOR ratio for a given structure is defined by a similar relation:

$$R_{\text{XOR}}^{\text{new}} > R_{\text{XOR}}^{\text{old}} + f\sigma_{R_{\text{XOR}}}^{\text{old}} \quad (4)$$

The optimization procedure also taken into account the fact that, due to the arbitrary choice of the center of rotation, the rotation and scaling parameters are coupled with the translation parameters. The center of rotation is chosen to be the center of the entire set of structures on the first study, and, since this center is not, in general, equal to the true center of rotation, each change in the rotation angles or scaling parameters during the minimization process must be accompanied by an adjustment in the translation vector. This adjustment is determined by aligning the centers of a user-defined "alignment structure" after each rotation or scaling operation. To accurately determine the center of the alignment structure, this structure must be unambiguously identifiable on both imaging studies and completely contained in both studies. After optimizing the rotation angles and scaling factors, the algorithm verifies the coordinates of the translation vector by checking whether or not a small variation in these coordinates further minimizes the XOR ratios for any of the structures.




Fig. 2. The rotation phantom used for the phantom studies. (Reprinted from M. L. Kessler, S. Pitluck, P. L. Petti, and J. R. Castro. "Integration of Multimodality Imaging Data for Radiotherapy Treatment Planning." *Int. J. Radiat. Oncol. Biol. Phys.* **21**, 1653-1667 (1991), with kind permission from Pergamon, Headington Hill Hall, Oxford OX3 0BW, UK.)

If the user requests it, the algorithm reoptimizes the rotation angles and scaling factors a second or third time, using the parameter values resulting from the previous optimization as initial values. Making multiple passes through the parameter space was found to improve the results in some cases. At least two passes through the parameter space were made for both the phantom and patient studies described below.

D. Phantom studies

Data derived from the "rotation phantom" (Fig. 2) developed by Kessler *et al.*,^{2,3} and used in that investigation to evaluate other image-registration techniques, were also used as benchmark data in this investigation to test the XOR-structure-matching algorithm. This phantom is constructed from a human skull attached to a rotating Lucite base. CT scans were taken with the phantom rotated about the X axis (i.e., the axis orthogonal to the sagittal plane) by 5° , 10° , and 15° . The CT slice spacing was 0.3 cm, and the pixel size was 0.078 cm for each of these studies. The XOR-structure-matching algorithm was used to determine the transformations between the phantom in its unrotated and rotated states. To further test the algorithm, the structures defined on the unrotated study were rotated, via software utility programs, about the Y and Z axes (i.e., the axes orthogonal to the coronal and axial planes) by 5° , 10° , and 15° , and the XOR-structure-matching algorithm was also used to predict these transformations.

For all of these calculations, the frontal sinus served as the alignment structure, and the structures used to determine the transformation included the inner skull, the right and left maxillary sinuses and the frontal sinus. The left and right

maxillary sinuses were considered as a single structure, so that, in effect, three distinct structures were used to determine the transformation. The initial step size for the rotation angles was 8° in all cases. Because all of these transformations were from CT to CT, the scale parameters were not allowed to vary in these investigations, i.e., all scale factors were set equal to 1.0.

E. Patient studies

The XOR-structure-matching algorithm was also applied to 12 representative patient cases. All of these patients were treated at LBL, and image-registration techniques (either the point-matching² or interactive-matching method,²) had been used to help identify the tumor volume during the treatment-planning process. In all cases, the transformation was from MRI to CT. Visual examination of the image data for the patients considered in this study indicated that, in most cases, the rotations involved were more complex than for the phantom studies. That is, two or more of the rotation angles were nonzero. The slice separation for the MRI studies ranged from 0.5 to 0.75 cm, while the slice separation for the CT data was between 0.3 and 0.5 cm. Pixel sizes for the CT and MRI data ranged from 0.078 to 0.1 cm. For all of these patient studies, the vitreous body of the right eye served as the alignment structure, and test structures used in the registration included some or all of the following structures depending on which structures could be unambiguously identified in each case: the inner skull, the right and left eyes, the right and left maxillary sinuses, the right and left remus of mandible, and the frontal sinus. As with the phantom studies, the initial step size for the rotation angles was 8° . The MRI data for the patients used in this study came from several different institutions, and it was unclear whether or not these data had been corrected for possible linear scale distortions. Thus the scale factors were optimized in these studies as well as the rotation angles, and the initial step size for the scaling parameters was 0.08.

III. RESULTS

A. Phantom studies

Results of the phantom studies are summarized in Table I. For these phantom studies, the XOR-structure-matching algorithm was able to predict the correct rotation values to within $\pm 1^\circ$. The root-mean-square (rms) deviation between the actual and predicted angles reported in Table I is 0.54° . This level of accuracy is comparable to that reported by Kessler *et al.*² for the image-registration techniques they evaluate.

Figures 3(a) and 3(b) illustrate the results for the 10° rotation about the X axis. Figure 3(a) shows the transformed maxillary sinuses and inner skull on an inferior level in the phantom, while Fig. 3(b) shows a more superior level containing the transformed frontal sinus and inner skull. In Fig. 3(b) there is little noticeable difference between the transformed structure outlines and the actual structures. Figure 3(a) shows a small discrepancy in the anterior edge of the inner skull outline. This is likely to be a result of the approximate way in which the structures are represented, that is, as

TABLE I. Phantom study results. The rotation angles determined with the XOR-structure-matching technique are compared to the actual rotation angles (in degrees). These calculations utilized 400 random points per CT slice per structure. The rms deviation between the actual and predicted angles is also listed at the end of the table.


Actual rotation $\theta_x, \theta_y, \theta_z$	Calculated rotation $\theta_x, \theta_y, \theta_z$
5,0,0	6,0,0
10,0,0	11,0,0
15,0,0	14,0,0
0,5,0	1,5,0
0,10,0	1,11,0
0,15,0	0,14,0
0,0,5	0,0,5
0,0,10	0,0,10
0,0,15	0,1,15
rms deviation:	0.54

polygonally tiled surfaces whose accuracy depends on both the slice separation and the number of points used to define a structure on each slice.

B. Patient studies

One cannot evaluate the patient studies in terms of the transformation parameters predicted by the XOR-structure-matching algorithm because, in contrast with the phantom examples cited above, the true values for the transformation parameters in the patient studies are unknown. Although informative, visual examination of the transformed MRI structures superimposed on the CT anatomy does not, by itself, constitute an adequate evaluation of the algorithm. To quantify this visual assessment, XOR ratios for the transformations determined by the XOR-structure-matching algorithm were compared with those determined by the image-registration method used as part of the treatment-planning process. For both the XOR-structure-matching algorithm and the older image-registration methods, structures defined on the patient's MRI scan were mapped to the corresponding CT study using the optimal transformation determined by each technique. XOR volumes were then calculated for each structure pair and for each mapping using the random point method. Some of these results are shown in Table II for two representative patients. Patient 1 was chosen to represent a case where the XOR-structure-matching algorithm produced results similar to the image-registration method used as part of the treatment-planning process (in this case, the interactive-matching method), while patient 5 represents a case where the XOR-structure-matching algorithm produced significantly better results. The standard deviation in the XOR ratios is also listed in Table II. According to Eq. (2), the standard deviation in R_{XOR} depends on the number of random points used in the calculation. Using 400 random points for each CT level within a structure, the standard deviation in R_{XOR} for the individual structures ranged from 2% to 8%.

As a convenient means of summarizing the data for each of the 12 patients considered in this study, the XOR ratio was averaged over all of the structures used to determine the



(a)



(b)

FIG. 3. Phantom study results: Transformed structures are superimposed on the CT scan of the unrotated phantom for the 10° rotation about the X axis: (a) an inferior cut showing the inner skull and maxillary sinus outlines; and (b) a more superior cut illustrating the inner skull outline as well as the frontal sinus.

transformation. These average values are listed in column 3 of Table III for the image-registration method used during the treatment-planning process and in column 4 of the same table for the XOR-structure-matching algorithm. The uncertainty in each of the values in column 4 was estimated by averaging $\sigma_{R_{XOR}}^2$ over all of the structures used to determine the transformation. That is,

$$s = \sqrt{\frac{\sum \sigma_{R_{XOR}}^2}{N_S}}, \quad (5)$$

where N_S represents the number of structures used to determine the transformation. For 9 of the 12 patients considered in this study, the average XOR-ratio obtained for the XOR-structure-matching technique was less than that obtained using the older image-registration techniques by more than two

TABLE II. Comparison of the XOR-structure-matching algorithm and the interactive-matching image-registration technique described by Kessler *et al.*² for two representative patients. The structures are coded as follows: ISK denotes the inner skull; LRM and RRM denote the left and right remus of mandible; LMS and RMS denote the left and right maxillary sinuses; and LEY and REY denote the left and right eyes.

Patient number	Structure	$R_{XOR} \pm \sigma_{R_{XOR}}$ algorithm used during treatment planning	$R_{XOR} \pm \sigma_{R_{XOR}}$ XOR- structure- matching algorithm
1	ISK	0.133 ± 0.005	0.131 ± 0.005
	LRM	0.58 ± 0.03	0.59 ± 0.03
	RRM	0.68 ± 0.03	0.59 ± 0.03
	RMS	0.70 ± 0.03	0.63 ± 0.03
	REY	0.28 ± 0.01	0.31 ± 0.01
	LEY	0.41 ± 0.02	0.38 ± 0.02
5	ISK	0.270 ± 0.006	0.303 ± 0.007
	LRM	0.92 ± 0.05	0.85 ± 0.03
	RRM	0.97 ± 0.05	0.73 ± 0.04
	RMS	0.56 ± 0.02	0.24 ± 0.01
	REY	0.68 ± 0.02	0.26 ± 0.01
	LEY	0.38 ± 0.05	0.27 ± 0.01

standard deviations. For the other three patients, namely patients 1, 7, and 9, the average value of R_{XOR} obtained with the new method was within two standard deviations of that obtained with either the point-matching or interactive-matching techniques.

The R_{XOR} values listed in Tables II and III are significantly greater than zero because, for all of the 12 patients considered in this investigation, some of the structures were

TABLE III. Summary of patient study results. The second column indicates whether the treatment-planning image registration was performed with the interactive structure-matching technique (designated by "I") or the point-matching technique (designated by "P"). The fourth column lists the average value of the XOR ratio for all of the structures used to determine the transformation with the XOR-structure-matching algorithm along with the estimated standard deviation in this value, and the third column gives this average ratio for the treatment-planning image registration. The standard deviation, s , is defined in the text.

Patient number	algorithm used during TMT planning	$R_{XOR} \pm s$ algorithm used during planning	$R_{XOR} \pm s$ XOR- structure- matching algorithm
1	I	0.46 ± 0.02	0.44 ± 0.02
2	I	0.53 ± 0.04	0.40 ± 0.03
3	I	0.56 ± 0.04	0.47 ± 0.03
4	I	0.54 ± 0.04	0.44 ± 0.03
5	I	0.63 ± 0.03	0.45 ± 0.03
6	P	0.58 ± 0.04	0.32 ± 0.02
7	P	0.44 ± 0.03	0.44 ± 0.03
8	I	0.46 ± 0.03	0.38 ± 0.03
9	P	0.39 ± 0.02	0.36 ± 0.03
10	I	0.71 ± 0.04	0.52 ± 0.03
11	P	0.58 ± 0.03	0.30 ± 0.02
12	I	0.50 ± 0.02	0.43 ± 0.02






FIG. 4. Patient study results: Transformed MRI structures are superimposed on the CT scan for patient 1: (a) an inferior cut showing the inner skull outline, the right maxillary sinus and the right and left ramus of mandible; (b) a more superior cut illustrating the transformed eyes, the right optic nerve, and the inner skull.

FIG. 5. Patient study results: Transformed MRI structures are superimposed on the CT scan for patient 5: (a) an inferior cut showing the inner skull (brain stem), the right maxillary sinus and the right and left ramus of mandible; (b) a more superior level showing the eyes, inner skull and tumor volume (nasopharyngeal carcinoma).

only partially imaged on either MRI or CT or both. For example, for patients 1, 2, 7, 9, 10, 11, and 12, the inner skull was fully imaged on the CT scan, whereas, the superior portion of the skull was not imaged on MRI. If only partial volumes are imaged, the algorithm operates only on the portion of the volume that is defined. In these cases, the XOR ratio does not approach zero, but rather a value reflecting the relative amounts of the structure imaged in each study. Aligning the centers of the alignment structure (the only structure that must be fully contained on both imaging studies) serves to place the other structures in approximately the correct position.

Figures 4 and 5 show XOR-structure-matching results for patients 1 and 5 (the same patients referred to in Table II). In Fig. 4(a) (patient 1), transformed MRI structures are super-

imposed on an inferior cut through the CT data. These structures include the right maxillary sinus, the inner skull and the left and right ramus of mandible. Figure 4(b) is a cut through a more superior level that includes sections of the transformed eyes, the right optic nerve as well as the inner skull. In general, the transformed MRI structures correspond well with the CT anatomy. There is some discrepancy in the posterior border of the inner skull and with the anterior edge of the maxillary sinus. This may be due, in part, to uncertainty in identifying bony anatomy on MRI.

Figures 5(a) and 5(b) show similar results for patient 5. The transformed tumor volume is also shown since the ultimate goal of image registration was, in this case, to transform the tumor (a nasopharyngeal carcinoma) and critical structure volumes from MRI to CT. The transformed tumor

volume agrees well with the abnormal region on CT. The inner skull, right maxillary sinus, and right and left eyes were used to calculate the transformation, and sections of these transformed structures are shown superimposed on the CT anatomy. There is some discrepancy with the lateral skull outlines at the superior level, but overall, the agreement between the transformed MRI structures and the CT anatomy is good.

IV. DISCUSSION

A. Choosing a function or quantity to optimize

The outcome of any optimization problem depends critically on the choice of a function or quantity to optimize. In this study, the registration between two image data sets was optimized by maximizing $S_{\text{improve}} - S_{\text{worse}}$, where S_{improve} represents the number of structures for which the XOR ratio is significantly reduced, and S_{worse} represents the number of structures for which the XOR ratio increases significantly. Other possibilities for optimizing the overlap between corresponding sets of structures include:

- (1) maximizing S_{improve} alone;
- (2) maximizing $S_{\text{improve}} - S_{\text{worse}} - S_{\text{unchanged}}$ where $S_{\text{unchanged}}$ represents the number of structures for which the XOR ratio does not change significantly;
- (3) maximizing $S_{\text{improve}} - S_{\text{worse}}$ without taking into account the uncertainty in R_{XOR} , i.e., with $f=0$ in Eqs. (3) and (4);
- (4) minimizing the average XOR ratio taken over all structures, i.e., minimizing $(1/N_S)\sum_{i=1}^{N_S} R_{\text{XOR}}$, where N_S is the number of structures used to do the registration.

The tests using phantom data described in Secs. II D and III A were repeated with each of these alternative optimization strategies. In each of these tests, the initial conditions were exactly the same as for the calculations described in Sec. II D. In particular, the same structures were used to perform the registration, and 400 random points per CT slice per structure were employed. None of the alternative strategies performed as well as maximizing $S_{\text{improve}} - S_{\text{worse}}$. The rms deviation between the actual and predicted angles was 1.39° for test 1, 1.40° for test 2, 1.00° for test 3, and 1.07° for test 4. This compares to a value of 0.54° obtained by maximizing $S_{\text{improve}} - S_{\text{worse}}$. The results of the first two tests suggest that some level of penalty is required for choosing an incorrect parameter value, and that subtracting S_{worse} is an adequate penalty. The results of the third and fourth tests suggest that it is important to take the uncertainty in R_{XOR} into account in the optimization procedure, and that using R_{XOR} values for the individual structures is a better strategy than averaging R_{XOR} over all structures.

B. Calculation time

It is advantageous to use as few random points as possible when calculating R_{XOR} to reduce the calculation time required for the optimization. To assess the implications of using fewer random points in the optimization, the phantom studies were repeated using 200 random points per CT section through each structure. In nearly all of the phantom

studies, using 200 as opposed to 400 random points led to less accurate results. In most cases, the predicted angles were within $\pm 2^\circ$ of the actual angles, as compared to within $\pm 1^\circ$ when 400 points were used. The rms deviation between the actual and predicted angles was 1.22° using 200 random points.

Another consideration in improving the speed of the calculation while maintaining a sufficiently high level of accuracy is the manner in which points are sampled. Random sampling techniques have been compared to sampling points on a regular Cartesian grid in two recent publications^{17,18} for problems related to treatment-plan evaluation. In particular, Lu and Chin¹⁸ assert that to calculate dose-volume histograms, using a regular Cartesian grid is significantly more efficient than using randomly selected points. Whether or not this is true for XOR-structure-matching image-registration technique will be the subject of future investigations.

The simplest way to reduce the calculation time for an optimization procedure such as the one described in this paper is to use a faster computer. The average CPU time required for each of the phantom studies reported in Table I using a Microvax 3500 workstation was 3.3 h. However, using a DEC 3000 AXP model 500/500s workstation, these calculations took only between 6 and 7 min.

C. Structure definition

As with any structure-oriented image-registration method, some level of uncertainty is introduced by the manner in which structures are represented and by the ability of the user to identify structures unambiguously. In LBL's treatment-planning system, which was used in this investigation to define the structures, contours are represented as sets of contiguous points. The software fits a third-order polynomial between points entered by the user (or between points derived by slicing transformed structures) to obtain a set of contiguous points. This provides a reasonably accurate description of each structure outline, but has the disadvantage of increasing the calculation time required to determine whether a point is inside or outside a given contour. For large regularly-shaped structures, such as the inner skull, it is not necessary to use the large number of points supplied by the LBL software to represent the contour. These contours were, thus, thinned before performing the inside-outside tests. Contours for smaller more irregularly-shaped structures, such as the frontal sinus, were not thinned. Whether or not all structure contours can be thinned safely without loss of accuracy will be studied further in the future.

Regardless of how they are represented, some structures are inherently difficult to identify on MRI. In particular, it is often difficult to define bony structures, especially when they are adjacent to air cavities. This uncertainty may lead to errors in the registration. No systematic investigation has been undertaken to attempt to quantify these errors. Overall, the results presented in the previous section suggest that the new algorithm performs satisfactorily, despite uncertainties in identifying bony structures on MRI. Using several structures to perform the registration may mitigate the effects of small errors in contour definition for some of the structures. Furthermore, using a well-defined spherical object (e.g., the

globe of the eye) as the alignment structure allows the algorithm to predict the translation between the two studies reasonably accurately, possibly reducing the error introduced by any ambiguity in defining the other structures.

V. CONCLUSIONS

A new automated image-registration technique based on multiple structure matching has been proposed and evaluated using both phantom and patient data. In a collective sense, the algorithm minimizes the relative volume lying outside of each structure and its transformed counterpart, and uses a random sampling approach to calculate this relative volume. Due to the approximate nature of the random sampling method, there is an inherent uncertainty associated with the relative volume calculation, and phantom studies suggest that it is important to incorporate this uncertainty into the optimization procedure. Phantom studies further demonstrate that an accuracy of $\pm 1^\circ$ can be achieved with this method. The algorithm also performed well when realistic patient MRI and CT data were employed. Registrations obtained with the new technique were either significantly better than or as good as those obtained using two older techniques that have demonstrated success in the clinic.

ACKNOWLEDGEMENTS

The authors are grateful to Dr. Lynn Verhey for a careful reading of this manuscript and for several helpful suggestions. We also thank Dr. Randy Stone for providing an efficient algorithm for checking whether a point is inside a contour. This work was supported by NCI Grant No. CA19138 and DOE Contract No. DE-AC0376SF00098.

APPENDIX: DERIVATION OF EQ. (2)

For any structure, the XOR ratio is given by Eq. (1), i.e.,

$$R_{\text{XOR}} \approx N_{\text{XOR}}/N,$$

where N_{XOR} is the number of random points falling within the XOR volume, and N represents the number of points falling within the structure itself. The variance in R_{XOR} divided by R_{XOR}^2 is

$$\frac{\sigma_{R_{\text{XOR}}}^2}{R_{\text{XOR}}^2} = \frac{\sigma_{N_{\text{XOR}}}^2}{N_{\text{XOR}}^2} + \frac{\sigma_N^2}{N^2}. \quad (\text{A1})$$

N and N_{XOR} may be described by binomial distributions, and the variance in these quantities may be estimated as

$$\sigma_N^2 = N(1 - N/N_0) \quad (\text{A2})$$

and

$$\sigma_{N_{\text{XOR}}}^2 = N_{\text{XOR}} \left(1 - \frac{N_{\text{XOR}}}{N_0} \right), \quad (\text{A3})$$

where N_0 represents the number of random points used to determine the volume ratios. Substituting Eqs. (A2) and (A3) into Eq. (A1) gives Eq. (2), i.e.,

$$\sigma_{R_{\text{XOR}}} = R_{\text{XOR}} \sqrt{\frac{1}{N_{\text{XOR}}} + \frac{1}{N} - \frac{2}{N_0}}.$$

- ¹P. A. van den Elsen, E. J. D. Pol, and M. A. Vieregger, "Medical Image Matching—A Review with Classification," *IEEE Eng. Med. Biol.* **12**, 26–39 (1993).
- ²M. L. Kessler, S. Pitluck, P. L. Petti, and J. R. Castro, "Integration of Multimodality Imaging Data for Radiotherapy Treatment Planning," *Int. J. Radiat. Oncol. Biol. Phys.* **21**, 1653–1667 (1991).
- ³M. L. Kessler, "Integration of Multimodality Imaging Data for Radiotherapy Treatment Planning," Doctoral thesis, University of California, Berkeley (1989).
- ⁴B. A. Fraass, D. L. McShan, R. F. Diaz, R. K. Ten Haken, A. Aisen, S. Gebarski, G. Glazer, and A. S. Lichter, "Integration of Magnetic Resonance Imaging into Radiation Therapy Treatment Planning: Technical Considerations," *Int. J. Radiat. Oncol. Biol. Phys.* **13**, 1897–1908 (1987).
- ⁵C. Pelizzari, G. T. Y. Chen, D. R. Spelbring, R. R. Weichselbaum, and C. T. Chen, "Accurate Three-Dimensional Registration of CT, PET and/or MR Images of the Brain," *J. Comput. Assisted Tomogr.* **13**, 20–26 (1989).
- ⁶E. J. Holupka and H. M. Kooy, "A Geometric Algorithm for Medical Image Correlations," *Med. Phys.* **19**, 433–438 (1992).
- ⁷G. T. Y. Chen, M. L. Kessler, and S. Pitluck, "Structure Transfer in Three-Dimensional Medical Imaging Studies," in *Proceedings of the National Computer Graphics Association*, Dallas, 171–177 (1985).
- ⁸A. Gamboa-Aldeco, L. L. Fellingham, and G. T. Y. Chen, "Correlation of 3-D Surfaces from Multiple Modalities in Medical Imaging," in *Proceedings SPIE*, edited by R. H. Schneider and S. J. Dwyer III (SPIE Press, Bellingham, WA, 1986), Vol. 626, pp. 467–473.
- ⁹M. van Herk and H. M. Kooy, "Automatic three-dimensional correlation of CT-CT, CT-MRI, and CT-SPECT using chamfer matching," *Med. Phys.* **21**, 1163–1178 (1994).
- ¹⁰R. P. Woods, J. C. Mazziotta, and S. R. Cherry, "MRI-PET Registration with Automated Algorithm," *J. Comput. Assisted Tomogr.* **17**, 536–547 (1993).
- ¹¹M. H. Phillips, M. L. Kessler, F. Chuang, K. A. Frankel, J. T. Lyman, J. I. Fabrikant, and R. P. Levy, "Image Correlation of MRI and CT in Treatment Planning of Intracranial Vascular Malformations," *Int. J. Radiat. Oncol. Biol. Phys.* **20**, 881–889 (1991).
- ¹²J. M. Collier, P. L. Petti, I. Daftari, T. Whitcomb, S. Pitluck, and M. L. Kessler, "The Clinical Use of Image Correlation (MRI to CT) for Tumors of the Head in Charged Particle Radiotherapy," *Med. Phys.* **17**, 552(A) (1990).
- ¹³H. R. Rusinek, W. H. Tsui, A. V. Levy, M. E. Noz, and M. J. de Leon, "Principal Axes and Surface Fitting Methods for Three-Dimensional Image Registration," *J. Nucl. Med.* **34**, 2019–2024 (1993).
- ¹⁴G. T. Y. Chen, R. P. Singh, J. R. Castro, J. T. Lyman, and J. M. Quivey, "Treatment Planning for Heavy Ion Radiotherapy," *Int. J. Radiat. Oncol. Phys.* **5**, 1809–1819 (1979).
- ¹⁵T. Pavlidis, *Algorithms for Graphics and Image Processing* (Computer Science, Rockville, MD, 1982).
- ¹⁶H. N. Christiansen and T. W. Sederberg, "Conversion of Complex Contour Line Definitions into Polygonal Element Mosaics," *Comput. Graphics* **12**, 187–192 (1981).
- ¹⁷A. Niemierko and M. Goitein, "Random Sampling for Evaluating Treatment Plans," *Med. Phys.* **17**, 753–762 (1990).
- ¹⁸X.-Q. Lu and L. M. Chin, "Sampling Techniques for Evaluation of Treatment Plans," *Med. Phys.* **20**, 151–161 (1993).