A Comparison of Manual and Automatic Methods for Registering Scans of the Head

Torre D. Zuk and M. Stella Atkins*

Abstract—**This article compares four widely utilized, yet fundamentally different, approaches for registering medical scans of the head. Comparisons are made on the basis of method, accuracy, robustness, computer requirements, and usability. This examination is intended to provide a means for determining an appropriate method for any given application. These approaches are: 1) an iterative method based on the repeated manual selection of 1–2 corresponding points, 2) an approach using the manual selection of 9–15 corresponding points, 3) an automatic surface matching method, and 4) an automatic approach based on voxel similarity. The methods are tested bothon simulated data to provide a gold standard of accuracy, and on real data. All registrations are performed in the same visualization environment created for multipurpose image processing. Simulated data tests provided mean transformation errors and time requirements for the different methods, as well as the displacement errors for a set of anatomical landmarks. These results show all of the methods provide good accuracy when the data is not highly distorted and has a large amount of overlap. From the tests using real data both transformations and time requirements are tabulated for comparison. All of the techniques successfully aligned the real data with the exception of surface matching, which failed on the PET-MRI. Each method exhibits strengths and weaknesses that should be understood in order to utilize the most appropriate technique for a given problem. Based on our examination, the voxel-similarity approach proved in general to be the method of choice.**

I. Introduction

THE registration of medical scans is the process of finding correct alignment or the proper spatial relation of one medical scan with reference to another. This correlation can provide important insights into each of the scans. For example, many time series data are registered to be used for clinical evaluation of the progress of disease and/or treatment [1], [2]; other clinical uses arise in functional neuroimaging studies [3], image guided neurosurgery planning [4], and radiotherapy treatment planning [5].

Although a plethora of different methods for achieving registration have been proposed, objective comparisons between them are scarce. Most often the comparisons focus on the accuracy of the techniques, but do not probe deeply into the reasons behind the results. Accuracy is highly dependent on the particulars of the data such as dimension, resolution, and noise etc., so it is difficult to generalize the accuracy of a method for all the different types of scans that are in clinical use. For this reason this paper compares several approaches not only on the basis of accuracy, but also on method, robustness, usability, and computer requirements.

Four different approaches for retrospective registration are examined, viz: an iterative manual method based on the repeated selection of 1–2 corresponding points; a direct method using the manual selection of 9–15 corresponding points; a fully automatic surface matching approach; and a fully automatic voxel-similarity-based method. These four methods represent a cross sampling of many available techniques for rigid body transformations; other techniques such as multiscale matching [6] are less widely used so are not considered here.

To facilitate this comparison all four different methods are implemented in the same visualization environment called WiT.[1] The results obtained on a variety of brain scans are analyzed to provide a relative evaluation of their strengths and weaknesses.

The registration procedure can be broken down into three separate processes: segmentation, matching, and verification. These processes are often three sequential phases, but not necessarily disjoint. Numerous ways to perform the different processes exist, each of which may be suited to one particular type of registration [e.g., positron emission tomography (PET) to magnetic resonance imaging (MRI)]. Decomposition of the registration process into three components aids in the comparison. The segmentation process is the isolation of features which can be directly compared.[2] The matching process takes one set of segmented features from scan A and one set from scan B and computes the optimal transformation to map B's features (dynamic) onto A's (static). The matching is usually formulated as an optimization, expressed as a minimization, of the comparison function which measures the proximity and differences between the two sets of features. The comparison function is dependent on the segmented features and may not be monotonic. This complicates the matching process by creating local minima which are difficult

Manuscript received May 30, 1995; revised June 27, 1996. This work was supported by the Natural Science and Engineering Council of Canada under Grant 8020. The Associate Editor responsible for coordinating the review of this paper and recommending its publication was C. R. Meyer. *Asterisk indicates corresponding author.*

T. D. Zuk is with the School of Computer Science, Simon Fraser University, Burnaby, B.C. V5A 1S6 Canada. He is also with the Multiple Sclerosis Research Group at the University of British Columbia Hospital, Vancouver, B.C. V6T 1Z4 Canada.

*M. S. Atkins is with the School of Computer Science, Simon Fraser University, Burnaby, B.C. V5A 1S6 Canada. She is also with the Multiple Sclerosis Research Group at the University of British Columbia Hospital, Vancouver, B.C. V6T 1Z4 Canada.(e-mail: stella@cs.sfu.ca).

Publisher Item Identifier S 0278-0062(96)07296-5.

[1] A product of Logical Vision Ltd., Suite 265, 4299 Canada Way, Burnaby, BC.

[2] The spatial (image) domain was chosen for all the registrations, although the segmentation itself does not necessarily have to occur in the spatial domain.

0278–0062/96$05.00 © 1996 IEEE

for a computer program to distinguish from the global or "true" minimum. Unless a method can be shown to be either extremely robust or reliably detect failure the final verification process must be achieved through a visual check.

1) Transformations: The transformation needed to correctly register one scan to another can be arbitrarily complex and can be characterized in many ways. The transformation may be global and have a domain of the entire volume or locally affect sub-regions of the volume differently. Global changes are commonly due to different patient positioning, while local ones are usually more complex and can be caused by object deformation, joint movement, or distortions. Local transformations may be handled by cutting the data into smaller regions, each of which can then be handled individually, or with complex nonlinear transformations that affect each region differently. A major difficulty with local transformations is the avoidance of discontinuities that may appear on the boundaries between local regions.

In comparing the four methods only global rigid body motion is assumed, as this model has often been utilized for registrations of the head [7]–[9]. It is considered best to independently correct for as many other variations as possible. When more complex distortions are suspected, as for example when registering computed tomography (CT) and magnetic resonance (MR) images of the head, where different surface contours may be extracted due to the sensitivity of MR to the skin and CT to the bone, the rigid body solution should be used as a first approximation, as explained below. Thus during gross alignment computation is not wasted adjusting numerous other parameters that may only result in an incorrect solution (such as the optimization being trapped in a local minimum).

Six parameters are required for the rigid body transformation: three for translation (labeled x, y for intraslice dimensions and z for interslice) and three for rotation (about these same axis). The spatial separation of voxels is obtained from the scanner configuration, and these are not modified by scaling. The transformation can be wholly represented by a 4×4 ($M_{B \Rightarrow A}$). This matrix maps the dynamic scan B onto the static scan A and taking the inverse of the matrix reverses the roles of the scans ($M_{A \Rightarrow B}$). Because the computed solutions are not exact, the final transformations obtained when optimizing the comparison function of $f(B \Rightarrow A)$ and $f(A \Rightarrow B)$ are not necessarily the inverse of each other ($M_{f(B \Rightarrow A)} \neq M^{-1}_{f(A \Rightarrow B)}$). This difference in solutions certainly arises when registering different modalities and/or spatial resolutions. Because the efficiency and accuracy of the algorithms are not commutative, the choice of dynamic scan plays an important role. However transformations more complex than global rigid body motion may be difficult to invert and thereby determine the dynamic scan.

2) Rationale for Choosing Rigid Body Transformations: Rigid body motion maintains spatial relationships within a scan, so the registration process will not corrupt any of this information; this is especially important for quantitative methods where areas or volumes are calculated. However, some quantization error will be introduced by interpolation (resampling) during the transformation stage. This error can be made negligably small by use of sinc interpolation functions [10]. However, the simpler trilinear interpolation is used here for consistency of measurements, as it is difficult to use a sinc interpolation function on anisotropic voxels [11].

With simulations, using from four to 50 landmarks pairs with added three-dimensional (3-D) Gaussian noise, Timmens [12] compares singular value decomposition (SVD) which calculates a rigid body transformation only, with the direct linear method which calculates an affine transformation, for solving strictly rigid body motion problems. In summary, he shows that although both methods attain very similar root-mean-square (rms) error in matching noisy landmarks pairs, the SVD method provides a more accurate registration (as based on noise-less landmarks present throughout the data sets). The optimization process may find small components of skew and scaling due to noise, even when no such components actually exist. In this way the accuracy of the registration may be significantly reduced. Therefore transformational components should only be used when those variations can be shown to exist.

Turkington *et al.* have performed accuracy comparisons [13], [14] which suggest that voxel spacing should not be modified (i.e., scaled) if it is accurately known. When scaling is allowed, only 1% scaling has been found necessary for the best-fit in matching PET and MR brain images [14].

When accurate *a priori* voxel spacing is known a scaling parameter's deviation from unity may indicate poor segmentation and so the data should be resegmented in order to improve correspondences [15]. To correct for errors in voxel spacing an extremely accurate segmentation must be utilized or it will result in a translational "shift" of other interior features. The "shift" created by erroneous scaling will be either toward or away from the centre of the segmented feature. Segmentation of the same features in two modalities with different effective point spread functions will result in one feature's surface being slightly larger than the other. With the lower effective resolutions of PET and SPECT this difference may create significant scaling errors if the voxel sizes are modified in the registration process.

It is important not to forget that what is being registered is based only on what has been segmented. This inevitably leads to better correspondences or smaller errors at the segmented features, and also to the conclusion that the comparison function used for the optimization may be a poor indicator of the overall registration accuracy. Hemler [16] shows in a cadaver study that a surface registrations' comparison function may be a poor indicator of overall accuracy.

Stereotactic frames (fiducials) may allow very high precision registration throughout the scan volume because of the following.

1) The markers can be designed to be segmented consistently.
2) In general, by segmenting features spaced further apart one can attain more accurate estimates of the rotation parameters.
3) The least constrained parameter will tend to dominate the error in the solution, and so the markers are constructed to constrain all the rigid body transformation parameters.

However, the fiducial approach can suffer from the problem that the region of interest (e.g., the brain) may not remain stable relative to the markers. If similar precision is to be attained with retrospective approaches then all these aspects must be addressed.

II. Materials and Methods

A. Iterative Manual Matching of Corresponding Points

Iterative manual matching can be performed in various ways. Using the Interactive Data Language[3] (IDL) Pietrzyk has created a method [17] that utilizes the repeated selection of two-dimensional (2-D) corresponding points. The points are different for each modality and for each scan; in general, for both this manual method and the following "direct matching" manual method, the points are chosen at the edges and extrema of anatomical features such as the bridge of the nose and the ends of the interhemispheric fissure. The user chooses one slice from each of three orthogonal views (one view along the acquired slice planes, the two others generated by resampling on perpendicular planes), and the dynamic scan is resampled to provide those same spatially located slice planes, resulting in three sets of corresponding views. The transformation can be adjusted by either selecting corresponding points for one of the views, or by manually adjusting a transformation parameter directly. All of the views of the dynamic scan are then updated and the user can continue the process until satisfied with the result. Without formal accuracy analysis Pietrzyk expects the results to be better than the half-maximum width of the line spread function of the device with lower spatial resolution. This expectation is based on a user's ability to detect variations larger than this.

Kapouleas [18] uses another approach for MRI-PET: he first identifies the interhemispheric fissure plane in both scans and uses them to solve for the two rotation parameters and the translation parameter perpendicular to the plane. The user then interactively adjusts the two translational parameters lying on the interhemispheric plane, and the rotational parameter for an axis through the sagittal view. By overlaying the result of some form of MRI edge–detection on the PET images the user can then more easily see misalignment. Kapouleas does mention that when the fissure deviates from a plane it may be necessary to go back and resolve for the fissure plane after making other parameter adjustments. A similar approach has recently been described by Ge [19].

Kessler [20], in a presentation of four different methods, describes an iterative approach utilizing the manual manipulation of transformation parameters by keyboard only. In this method the user is simultaneously provided with 12 slices of the static data set with the segmentation of the dynamic scan (contours) overlaid on those slices. The parameters are then manipulated to best match up the contours with the static scan's information. Kessler estimates that with CT-MRI an accuracy acceptable for radiotherapy treatment planning ($\pm$ 2.0 mm) can be obtained.

[3] A product of Research Systems Inc., 2995 Wilderness Place, Boulder, CO.

The iterative approach described here is based on the one created by Pietrzyk [17]. This method was chosen as its presentation of all three dimensions provides the user with more of the information needed to perform the process, and at the same time provides flexibility in the choice of correspondences. Using a commercial programming environment WiT, most of the options Pietrzyk provides in his IDL program have been duplicated. The implementation provides three orthogonal views (axial, sagittal, coronal) of the two datasets to the user. Contours can be extracted from the dynamic and static scans and overlaid on the opposing scan to visually reveal misalignment. The user may then select one of the views to manipulate or one of the parameters. If the user chooses a view they then select either one or two pairs of corresponding points. This provides the translation and/or rotation needed to align the scans in the two dimensions shown in any one view. The choosing of point pairs will usually have to be performed more than once per view as the transformation on the two dimensions shown in a single view affect the match obtained in the other views. This process should be repeated in varying views until the user is satisfied with the overall match. The decomposition of a 3-D problem into that of repeatedly solving 2-D problems is often performed to reduce problem complexity. Thus the user understands the process much more easily and the method inherently provides validation as the user continues until satisfied with the result.

B. Direct Matching of Corresponding Points

With this approach a set of corresponding points are manually selected in each scan. The manual selection of corresponding points is a common approach in practical use, some examples are Hill [21], Henri [22], Maguire [23], and Schiers [24]. Foundational work in this area is reviewed by van den Elsen [25] and by Gerlot [26] in their surveys.

This method is called "direct matching" in that the point selection is performed once and then a solution is obtained, as opposed to the repeated manual interaction common to the approaches discussed in the preceding section. Given the set of corresponding points a transformation must then be calculated which maps them onto each other. This calculation is based on the type of transformation to be utilized and may be either an approximation or interpolation of the points.

For rigid body transformations only three noncollinear corresponding points are necessary to determine the parameters,[4] but it has been found beneficial to greatly overdetermine the system to achieve the best approximation. Various methods are available to solve for the transformation parameters. Timmens [12] showed with simulations that when using from four to 50 pairs of landmarks, with added 3-D Gaussian noise, the errors with both SVD and the Direct Linear Approach can be reduced by using more point pairs, but after 25 the improvement is insignificant. Hill [21] has shown in tests using from four to

[4] Except if the points form the vertices of an isosceles triangle, where a degenerate case can occur potentially causing misregistration due to a "flipping" of the image across any symmetric bisector of the triangle. We used matched landmark points in the same order on each image, to prevent this degenerate case; furthermore, we always use more than three points.

26 points with random error that mean displacement error is proportional to $\frac{1}{\sqrt{N}}$ with N being the number of points.

The implementation here utilizes the Numerical Recipes [27] Levenberg–Marquadt method for nonlinear model-fitting adapted to 3-D rigid body motion by Menke [28]. This routine was utilized for the real-time tracking of patient movement during a single scanning session. The well defined nature of the rigid body problem makes this specific technique very effective for calculating a fast and accurate solution.

After the solution is found, the misfit on each pair of corresponding points can be determined. Then the pair with the largest misfit can be excluded if the value of the misfit exceeds a specified threshold. The threshold should be large enough to only exclude clearly incorrect correspondences (i.e., mistaken features), as excluding other outliers will decrease the accuracy gained through the averaging of more points. The optimization can then be performed again to obtain the best match using the remaining points. This may be repeated as many times as necessary to exclude all incorrect correspondences.

C. Surfaces Based

When surfaces are the most consistent and easily segmented features across both data sets, one surface can be transformed (the dynamic) and compared to the other one (the static). This type of matching is based on the idea, as Pelizzari [29] describes it, of fitting a "hat" on the "head." Some other surface matching implementations are described in [1], [30], and [31]. Collignon [9] compares a number of different surface matching methods and reviews the basic components of each. He also attempts a quality constrained cost analysis, but admits difficulty in finding common ground for a good comparison. In summary he finds that surface based approaches are a viable alternative to point and voxel-similarity-based methods, and expects that a hybrid approach will probably yield the best results. An application of surface registration with PET-MRI is described by Bidaut [32], [14]; Levin [33] also provides case studies with CT-MRI and MRI-PET registrations.

The surface registration technique used in this analysis uses a threshold value calculated as the maximum mean voxel intensity of all the slices plus a fraction of the standard deviation of that same slice. What fraction of the standard deviation is to be added must be interactively determined for a given data type. Everything above this threshold is outlined, and this provides the contours that form the surfaces. This segmentation technique has provided results usable for registration [34].

The comparison function for two segmented surfaces is often the root-mean-square distance from a set of points on one surface to the closest points on the other. To provide this type of comparison function, a distance transform [35], [36] is used here. It is well suited for registering multiple scans in a time series back to a baseline. The precomputation of the comparison function is performed once based on the static baseline surface, and then can be used for all subsequent scans. Generating an accurate segmentation is the critical part of this approach. If the surfaces are not highly corresponden the optimization may be lead astray unless other constraints are added.

The implementation of surface matching is most similar to that of Jiang's [37]. An earlier version of this approach is described in [34], but to increase efficiency the optimization algorithm of Powell's direction set [38] taken from "*Numerical Recipes in C*" [27] is used. To provide matching at multiple scales, hierarchical subsampling at different rates is used. To exclude outliers included by the automatic segmentation, thresholds limit the maximum distance any single point can contribute to the comparison function value. To detect local minima multiple starting points have been added to the minimization process.

D. Voxel Similarity Based

The implementation is based on Woods' method [8], [39]. With this method each intensity value (0–255 for 8-b data) is assumed to represent a different group or "segmentation" of tissues. For any transformation a single intensity (e.g., 120) in one scan is mapped onto a distribution of intensities (e.g., 100–135) in the other. The matching process is based on maximizing the uniformity, or minimizing the standard deviation of these distributions. Hill and Studholme [10], [40] have proposed a variation using the 3rd order moment over a restricted range of intensities; this is done with the aim of registering PET-MRI without masking out everything in the MRI but the brain, as required by Woods' approach. From their study of how intensities of one scan get mapped to another across different modalities, they determined that by maximizing the skew (any odd moment about the mean] they can achieve registration. They also introduce better interpolation, based on the 3-D sinc function, with the goal of extremely high accuracy. Gerlot-Chiron [26] has formalized the voxel-similarity comparison criteria with his multimodality work. Collignon [41] takes the voxel-comparison process even further by segmenting all voxels during the registration process in order to more accurately compare voxels of the same tissue type. He also presents a summary [42] of the various voxel-similarity- and correlation-based methods and introduces another comparison function based on entropy.

In the implementation here, the second-order moments (squared standard deviations) of the previously described intensity distributions are minimized. This is done symmetrically for A's intensities mapping onto B's and B's mapping onto A's. These deviations are based only on the intensity distributions and are not normalized as Woods, but are similarly weighted by the number of voxels contributing to each. This comparison function is minimized using Powell's Direction Set algorithm, so that each intensity is mapped as closely as possible to a single intensity in the other. This assumes that each intensity represents only a single tissue type, and the degree to which this holds determines the robustness of our comparison function. The minimization can operate using a subset of the original intensity ranges, possibly one over which the aforementioned assumption holds. With hierarchical levels an increasing fraction of all the voxels are used as well as

different intensity ranges. As with the surface based method, multiple starting points are utilized to deal with local minima.

III. Results

The results of the different methods are provided for various problem types, both simulated and real. Visual verification is not included in the tabulated time requirements for each of the methods because it requires highly variable amounts of time. An appropriate amount of user time could be added to each of the method's shown time requirements, with the exception of the iterative manual method which inherently provides a form of verification. While this is not relevant to the simulations it may have significant implications for the real cases in which the exact solution cannot be determined. For all the studies described in this paper, x refers to the patient right-left direction, y refers to the anterior-posterior direction, and z refers to the axial direction in the scanner (inferior to posterior).

A. Simulations

Through the use of simulations very precise values of accuracy for the different methods can be calculated. By resampling an arbitrary scan after a given rigid body transformation, represented by the matrix M_{sim}, a simulated scan is obtained for which the registering transformation back to the original is known, i.e., M_{sim}^{-1}. If M_{sim} is multiplied with the solution found by one of the methods M_{reg}, the error transformation M_{err} is obtained.

The rotational errors about each axis and the translational errors can be seen directly by decoupling the rotational and translational components of M_{sim}^{-1} and M_{reg} about the same rotational origin. Note that the translational errors are only for the rotational origin, as the actual displacement errors vary throughout the volume. In order to see how the actual displacement error varies over the scan the error transformation M_{err} can be used to transform a user selected set of 3-D landmarks. From this the exact positional errors for any given set of points in the scan can be obtained, as is shown later for an MRI simulation.

Current "automatic" methods often require manual intervention during the segmentation and/or matching processes. They may then go on to require great lengths of time to perform a global search for the transformation parameters. From a practical point of view, these two processes are somewhat contradictory in nature. If manual intervention is required then simple interaction may also be used to calculate good initial approximations of the registration parameters, thus eliminating the large time costs of a global search. The two automatic implementations described here required no user interaction at all.

If an approach is to be completely automatic then a global search strategy must be used, but no methods to date have been shown to be robust enough to be completely automatic over a large variety of transformations and data. Notwithstanding, when the problem is constrained we make the claim that completely automatic methods are feasible. This occurs in many cases when the patient can be positioned in the scanner in a "standard" position and a large amount of overlap between scans can be obtained. When this occurs the range of rotations and translations that may be required for registration is substantially reduced, and it is in these cases that the use of a fully automatic approach is most promising. This issue is also addressed later in the discussion section on "Robustness."

TABLE I
PET Simulation Transformations

Simulation no.	Rotation (degrees) x	y	z	Translation (mm) x	y	z
1	−8.0	9.0	−3.0	−9.0	−11.0	9.0
2	3.0	8.0	−10.0	11.0	9.0	8.0
3	−10.0	−2.0	10.0	9.0	−11.0	−8.0
4	10.0	10.0	10.0	−11.0	−11.0	−8.0

The range of rigid body transformations over which to test an automatic method is subject to the ideas of local and global search. When more than one rotation angle becomes large (> 10 degrees) the complexity of the search usually increases dramatically. This is because the comparison function's parametric gradients may lead away from the global minimum, and also results from the function's nonlinear dependence on the rotation parameters. Therefore a local optimization method should be tested thoroughly to determine the "capture range" over which it converges to the correct solution with a high probability of success. This parameter range can then be exploited to determine an efficient global optimization strategy. With this in mind the simulation tests described here fall into the two categories; where local search is sufficient, and where global search is needed.

The results of the four described methods are presented for a set of simulation transformations on a PET scan and a MRI scan. In order to "ground" these simulations, fiducial based registration has been mimicked by artificially introducing eight fiducial markers into the scans. These fiducials are created as two sets of four coplanar blobs located just outside the epidermis. The markers provide a variation of the direct manual method in which a few truly corresponding points can be easily segmented. This adds a fifth registration method for these simulations. The simulations are intended only to illustrate the different problem scopes. They are not meant to indicate the relative accuracy of the different methods, as that is highly dependent on implementation details and the specifics of the data.

B. PET Simulations

For the following simulations an ^{18}FDG PET scan with 31 axial slices having 128×128 voxels is used. The voxel separation is 2.0 mm in x and y (intraslice) and 3.375 mm in z (interslice). No interpolation is utilized to create cubic voxels and so the data remains anisotropic throughout the process. The transformation values used for the PET simulations are presented in Table I. After resampling the 8-b data with trilinear interpolation the intensity range of the scan was compressed from the original range 0–255 down to 0–200. Another complicating factor was that the simulated scan was resampled into a scan volume the identical size to the original; this effectively clipped everything that was transformed outside the

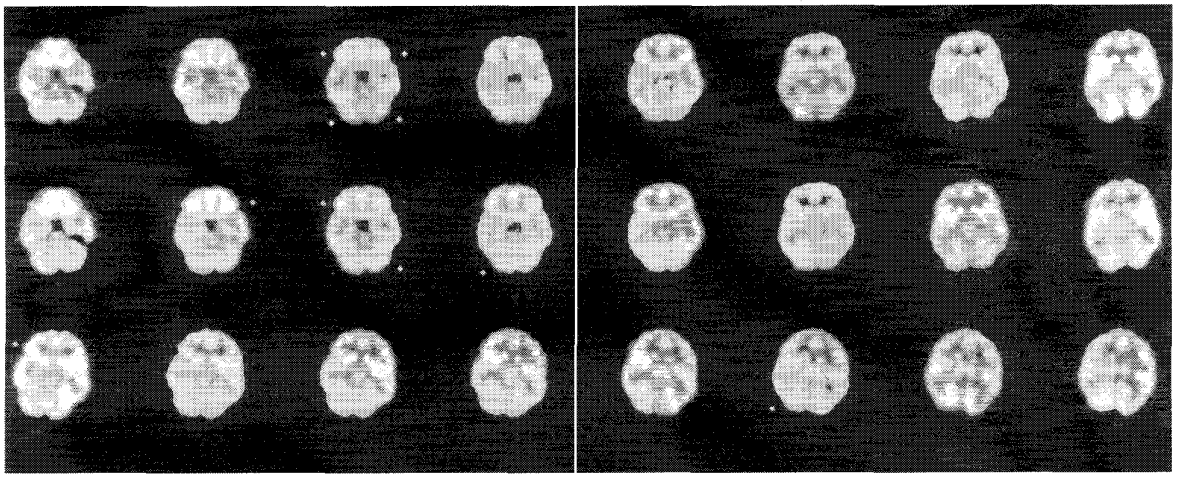

Fig. 1. PET simulation. From the top to bottom row: Consecutive slices of the static scan; results of registration using the direct manual approach, and the same slices in the simulated scan. Note that the four fiducials in column 3 of the first row are all visible in the registered slice (in the second row). Residual error, due to interpolation and misregistration, is visible in the surrounding columns of row 2. The fiducials in row 3 (the artificially rotated and translated scan) appear in different slices due to the simulation rotation and translation.

TABLE II
PET SIMULATION MEAN RIGID BODY TRANSFORMATION ERRORS

Method	Rotation (degrees)						Translation (mm)					
	x	σ_x	y	σ_y	z	σ_z	x	σ_x	y	σ_y	z	σ_z
Iter. Man.	0.42	0.25	0.59	0.42	0.67	0.40	0.83	0.23	0.78	0.52	0.56	.09
Dir. Man.	0.58	0.35	0.47	0.44	0.76	0.29	0.17	0.03	0.21	0.17	0.33	0.25
Fid. Mark.	0.30	0.19	0.38	0.10	0.21	0.16	0.36	0.38	0.59	0.43	0.49	0.40
Surface	0.43	0.28	0.23	0.12	0.24	0.01	0.09	0.07	0.12	0.13	0.27	0.23
Voxel Simil.	0.25	0.16	0.21	0.19	0.13	0.01	0.07	0.06	0.13	0.11	0.30	0.02

Voxel spacing is 2.00mm in x & y and 3.37[illegible]mm in z

original bounding volume. The intensity range compression, resampling errors, and truncation allowed the creation of more realistic or phantom-like problems.

For the direct manual approach 15 points were selected and the two worst outliers excluded. Due to the low resolution of the scans there was difficulty locating corresponding landmarks. As specified in the previous section, points were generally chosen at the extrema of edges in the images. Fig. 1 shows some of the data and the results of the direct manual approach on simulation number 4. The markers provide a clearly visible indicator of the error, but they were only utilized when performing the fiducial based registrations. Two of the fiducial registrations were performed with seven corresponding fiducial points and the other two with all eight. The two cases that used only seven fiducial points were a result of one fiducial marker being clipped in the process of creating the simulated scan (i.e., it was outside the bounds of the new resampled volume). The global search strategy of multiple starting points was utilized for the surface matching method. This was necessary as with a single starting point the algorithm may get stuck in an obviously incorrect local minimum. However, the voxel-similarity comparison function was robust enough that a single starting point sufficed.

Table II shows the mean parameter error and the standard deviations of the errors over the four simulations for each of the different methods.

The standard deviations of the errors are sufficiently small that four simulations were considered to give an adequate

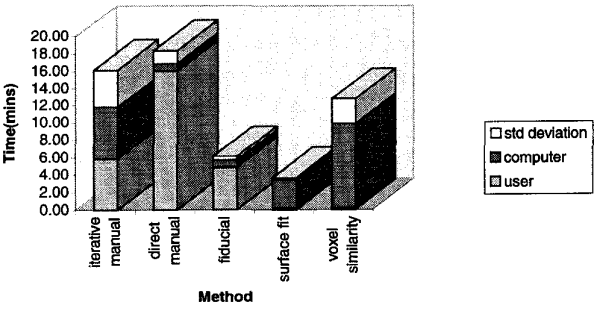

Fig. 2. PET simulation mean registration times. Note that the standard error for each method is shown above the total time for each method.

indication of the algorithms' general performance. The mean time required for each method is illustrated in Fig. 2, broken down into user interaction and computer time.

C. MRI Simulations

Proton density (PD) and T2-weighted MRI scans obtained simultaneously in a double-echo sequence provide a registered pair of scans with variations typical to certain different modalities where resolution is comparable, such as for CT to MRI registration which is used for surgical and radiation therapy planning. By resampling one of the two scans after a transformation, a situation which has many of the difficulties of the true multimodality registration problem is obtained, and

TABLE III
MRI SIMULATION TRANSFORMATIONS

Simulation	Rotation (degrees)			Translation (mm)		
no.	x	y	z	x	y	z
1	17.0	−13.0	2.0	−8.0	7.0	9.0
2	−5.0	17.0	−18.0	−6.0	8.0	8.0
3	20.0	−3.0	15.0	7.0	−7.0	−9.0
4	20.0	20.0	20.0	−7.0	−7.0	−12.0

yet the exact transformation between them is known. The MRI simulations were performed using a double-echo PD and T2 scan of 22 slices with 256×256 voxels per slice. The voxel separation is 0.820 mm in x and y and 5.00 mm in z (with no gap).

The transformation values used for the MRI simulations are listed in Table III. As with the PET simulations the manual direct method utilized 15 corresponding point pairs and the two worst outliers were excluded one at a time. For the fiducial registration three simulations utilized six corresponding points and the other seven. The larger rotational values with the MRI simulations necessitated the use of a global search strategy for both of the automatic methods. This, along with the increased amount of data, increased the time required for automatic registration.

Table IV illustrates the mean parameter error and the standard deviation of the errors over the four simulations for each of the different methods. These errors are discussed in detail under the Accuracy heading in Section IV.

The mean time required for each method is graphically illustrated in Fig. 3.

Table V shows the actual displacement errors for a selected set of anatomical landmarks based on simulation number 4. These points are chosen just to illustrate typical displacement errors over a range of positions. These data should not be used to measure the accuracy of the registration methods as the different methods will match different regions of the brain more accurately. Rather, these data should be used to see how the different methods do achieve different accuracies for selected points; for example, the surface fitting method has the smallest error of all the methods for a voxel located at the aqueduct and at the red nucleus, but not for all the points.

Fig. 4 illustrates some of the data and the results using the voxel-similarity-based method on the same simulation.

D. Case Studies

The uses for registration may be broadly grouped into two categories.

Subtractive: The aim is to reveal subtle variations or differences more accurately. These studies are usually uni-modal.

Additive: In which complementary information is combined to form a better understanding of the whole. These studies are usually multimodal.

Examples of subtractive applications are the registration and subtraction of Pre- and Post-contrast scans, and the cine of a single slice through time after registering a time series. Common additive uses are the matching of structure to function, or radiotherapy treatment planning using registered MRI-CT. A typical case study of each type is detailed in the following two sections.

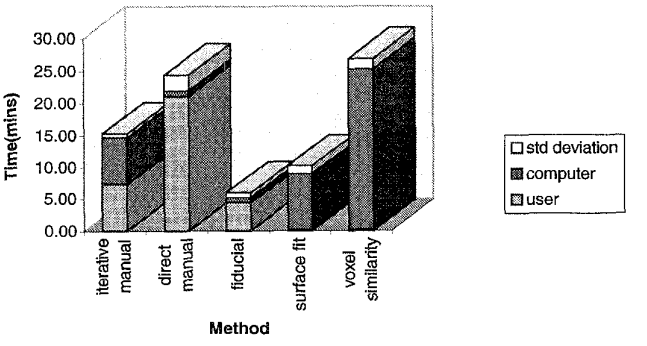

Fig. 3. MRI Simulation mean registration times. Note that the standard error for each method is shown above the total time for each method.

1) MRI Time Series: Using the four methods, a time series of two T1-weighted MRI scans taken four years apart was used and the second scan was registered to the first (baseline) scan. Each slice is 256×256 with spacing of 0.937 mm. There are 2.50-mm gaps between the eighteen 5.00 mm thick slices, creating a voxel spacing of 7.50 mm in z. The methods were utilized as previously described for the MRI simulations. Transformations found by each of the methods and the time required are presented in Table VI.

A subset of the slices and one set of results are given in Fig. 5. Even with perfect registration, variations between the baseline and resampled data would be present due to physical changes over time, volume averaging, different regions in the gaps, and interpolation errors. Taking all of these factors into account the results still provide comparable slices.

2) MRI-PET: Registration of a T1-weighted MRI scan to a PET Fluorodopa scan was performed. The MRI scan's 256×256 voxels have intraslice spacing of 0.780 mm, and 22 slices 7.00 mm thick. The PET scan represented a mean scan calculated by averaging five immediately consecutive PET scans after being registered with Woods' software [39]. This was done to provide a much higher signal to noise ratio, and was performed before we obtained the data.[5] The 31 slices are separated by 3.375 mm in z, and the 128×128 intraslice voxels have a separation of 2.608 mm in x and y. The results of the different registration methods are shown in Table VII.

The iterative manual approach proceeded in the standard fashion with the exception of the visual testing of the user. Because the functional data does not necessarily correspond with the structural, it was important that the PET data not protrude beyond the overlaid MRI boundaries, while variations within them were acceptable. With the direct manual approach only nine corresponding points could be selected and, due to this low number, no outliers were excluded.[6] For the surface matching technique, the PET surface that was segmented was a poor approximation of the edge between epidermis and

[5] This is standard procedure for dynamic PET scans, although the scans are sometimes not registered because patient motion is minimized through use of a head mask.

[6] Again, points were chosen from edges and extrema.

TABLE IV
MRI SIMULATION MEAN RIGID BODY TRANSFORMATION ERRORS

Method	Rotation (degrees)						Translation (mm)					
	x	σ_x	y	σ_y	z	σ_z	x	σ_x	y	σ_y	z	σ_z
Iter. Man.	0.59	0.22	0.96	0.32	0.29	0.14	0.61	0.35	0.37	0.14	1.05	0.32
Dir. Man.	0.87	0.38	0.42	0.25	0.26	0.17	0.15	0.09	0.19	0.13	1.06	0.57
Fid. Mark.	0.37	0.18	0.13	0.10	0.08	0.05	0.24	0.03	0.28	0.06	0.55	0.24
Surface	0.26	0.07	0.22	0.13	0.08	0.05	0.07	0.05	0.09	0.05	0.25	0.14
Vox. Simil.	0.39	0.32	0.24	0.07	0.10	0.06	0.06	0.04	0.13	0.10	0.89	0.10

Voxel spacing is 0.820mm in x & y and 5.00mm in z

air surrounding the head. This led to the clearly incorrect solution found by this automatic approach. Without a PET transmission scan to yield a cleaner segmentation there was little to no chance of a good surface-based solution. In order to keep the voxel-similarity method fully automatic no masking was performed on either scan. The assumption of a single intensity representing a single tissue type is violated by the values present outside the brain with the same values as those within. This reduces the effectiveness of the voxel-similarity comparison function, but with functional to anatomical registration the assumption has already been violated to a large extent. Automatic segmentation methods such as Ardekani's intradural space detection [43] should in the future replace the manual masking that is commonly utilized to improve registration results. The data is exemplified with the results of the manual iterative method in Fig. 6.

IV. DISCUSSION

The variety of results that may be obtained when using the different registration techniques has been demonstrated. Recently Strother [44] has quantitatively compared registrations based on Talairach space, surfaces, direct point matching, and voxel similarity, and states the voxel-similarity method provided the most accurate results on MRI-MRI, PET-PET, and MRI-PET registrations. Nevertheless, under the proper conditions each of the four methods examined can attain sub-voxel accuracy. With multiple options available to the user the determination of the most applicable approach for a given problem should be based on the specifics of the data and the purpose for which the results are to be utilized. Based on these results the major characteristics that can be used to make this decision are described in detail below.

A. Comparative Analysis

1) Computer Requirements: Time and memory requirements of each of the different methods can vary dramatically with specific implementation details, but they are bounded by certain factors. Computer times were measured on a SUN SPARCstation 2 to provide a commonly attainable processing speed. Tables VI and VII provide the time requirements for each method for the particular case study, but with the automatic methods the time required for the optimization process can be highly variable. Thus the mean required times in Figs. 2 and 3 are more representative of typical registration times, and the breakdown between user and computer time enables a prediction of future results. With the ever increasing speed of computers, the computer component of the graphs will shrink. Although direct manual point matching is the only method that is extremely dominated by user interaction time, its time requirements could possibly be greatly reduced when new interactive point selection methods become feasible with faster computers.

TABLE V
LANDMARK DISPLACEMENT ERRORS FOR MRI SIMULATION NO. 4

Method	3D Displacement (mm)			
	aqueduct	Red nucleus	Globus palidus	top of lateral ventricle
Iterative Manual	1.62	1.57	1.36	2.34
Direct Manual	0.90	0.67	0.61	0.29
Fiducial Markers	0.76	0.76	0.82	0.69
Surface	0.52	0.54	0.65	0.46
Voxel Similarity	0.80	0.72	0.59	0.77

The test cases were chosen to be within the overall capture range of the automatic methods, and so are somewhat biased. If the rotational components of the transformations were much larger than 25 degrees, these approaches would have been more likely to fail (although so would the manual direct approach). Adding more starting points can provide a solution to this, but at the cost of a large increase in time requirements. This demonstrates that the robustness of the automatic methods can often be increased at the cost of increased computer time. Therefore the time required is dependent on the level of robustness desired.

All of the methods have very similar computer memory requirements. Basically they work most efficiently when there is enough computer memory to hold both scan volumes at once. The direct manual method can cope with limited memory the most easily as only memory for two image slices is essential. The iterative approach only has to work on one scan at a time, and so can make do with half the memory requirements of that needed for the voxel-similarity method. Usually surface-based approaches require little more than the amount needed to hold one scan, but our implementation utilized a floating point distance transform, and so required four times the memory needed for an 8-b (0–255) scan. This illustrates another tradeoff that is often possible, that of increased memory usage for decreased time requirements. With current computer memory management, physical memory requirements can often be relaxed through the use of disk space (virtual memory), but at the cost of increased execution time.

2) Accuracy: From the results shown in the tables, all the methods provide near sub-voxel accuracy. Table IV shows that the largest voxel dimension (z in this case) dominates the

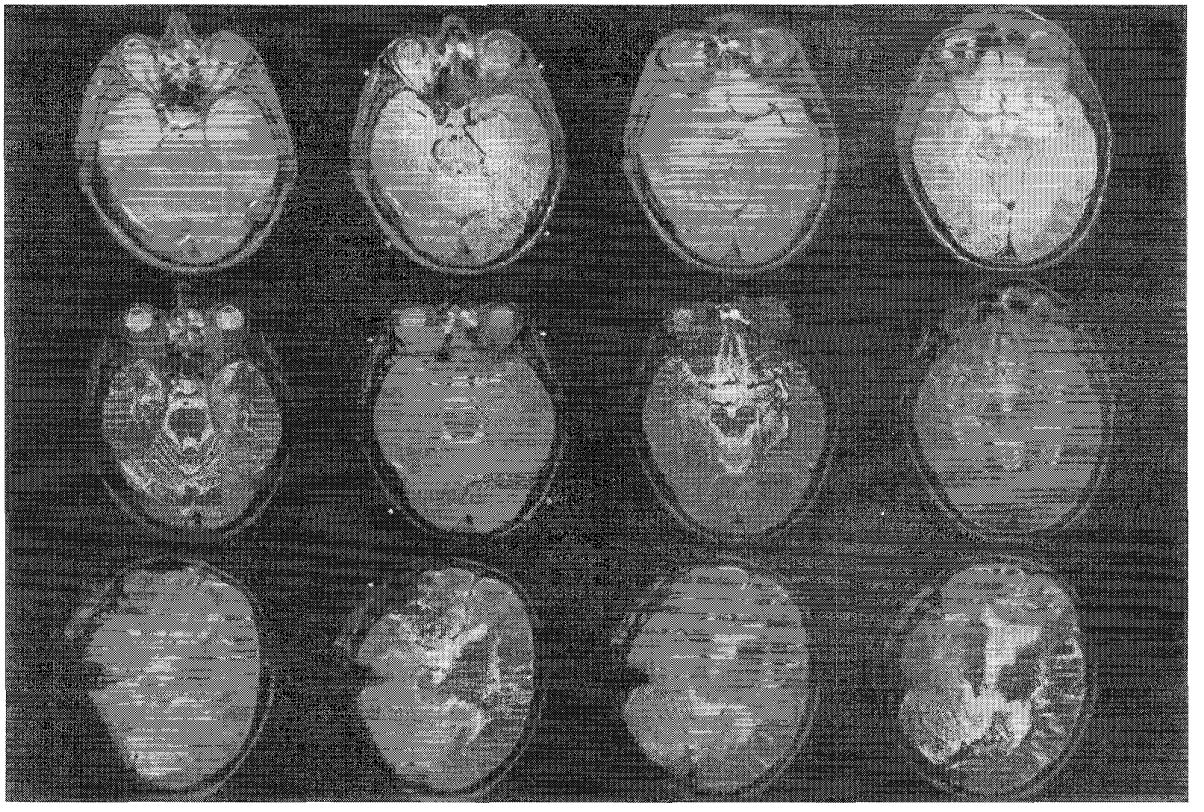

Fig. 4. MRI simulation. The top row illustrates slices of the static scan, the middle shows results of registration using the voxel-similarity method, and the bottom row the same slices in the simulated scan. Note that the four fiducials in the second column of the base scan have been faithfully registered in the second column of the second row.

TABLE VI
MRI TIME SERIES TRANSFORMATION RESULTS

Method	Rotation (degrees)			Translation (mm)			Time
	x	y	z	x	y	z	(min)
Iterative Manual	15.54	−0.67	0.18	−8.77	−2.83	0.20	8
Direct Manual	14.72	−0.80	0.68	−8.24	−3.50	−1.62	15
Surface	15.58	−1.28	0.32	−8.48	−2.26	−0.92	20
Voxel Similarity	15.77	−1.38	0.26	−8.44	−2.13	−1.26	12

Voxel spacing is 0.937mm in x & y and 7.50mm in z

translation error, and that the rotation error arises around the axes which lie in the plane of the largest voxel dimension (x and y in this case). Note that the test cases were also reasonably constrained in order for the automatic methods to work with a high probability of success. These results are confirmed by other researchers such as Pelizzari [15], who compares a surface-based approach with the direct matching of corresponding points when registering a PET ^{18}FDG, $H_2^{15}O$, and transmission scan to MRI. Their comparison found that the two automatic approaches gave comparable and accurate results as far as could be determined by identifying homologous points across modalities.

Tables II and IV illustrate that when specific landmarks cannot be segmented with a high degree of accuracy, then by using a greater number of points (iterative manual < direct manual < surface < voxel similarity) more accurate results can be obtained.[7] However, the variation in results was small enough that modifications on each method or different data could change the relative accuracies. When consistent features cannot be segmented, as with surfaces in the MRI-PET example, the number of points has very little impact.

Specific errors vary throughout the image volume, as shown in Table V, but the results in Table IV indicate that the principle components of this error are z translation and x and y rotation. This is certainly due to the lower axial resolution (thick slices), and so this area of the scanner protocol must be addressed if highly accurate results are required. Thicker slices may also increase the difficulty in manually selecting corresponding points because the volume averaging will blur structure boundaries. The automatic methods do not make the difficult decisions about single point correspondences and so will not suffer to the same degree from the decreased resolution.

The surface-based approach's accuracy is largely determined by the accuracy of the segmentation and the degree to which the segmented region reflects the transformation on the region of interest. A voxel-similarity method's implicit segmentation avoids the first limitation, but cannot exclude distorted regions as easily as the surface technique. Both com-

[7] This statement must not be confused with a registration method such as the principal axis technique discussed in [45] which uses many points to compute the principal axes, but which then discards much of the information by computing an average for registration.

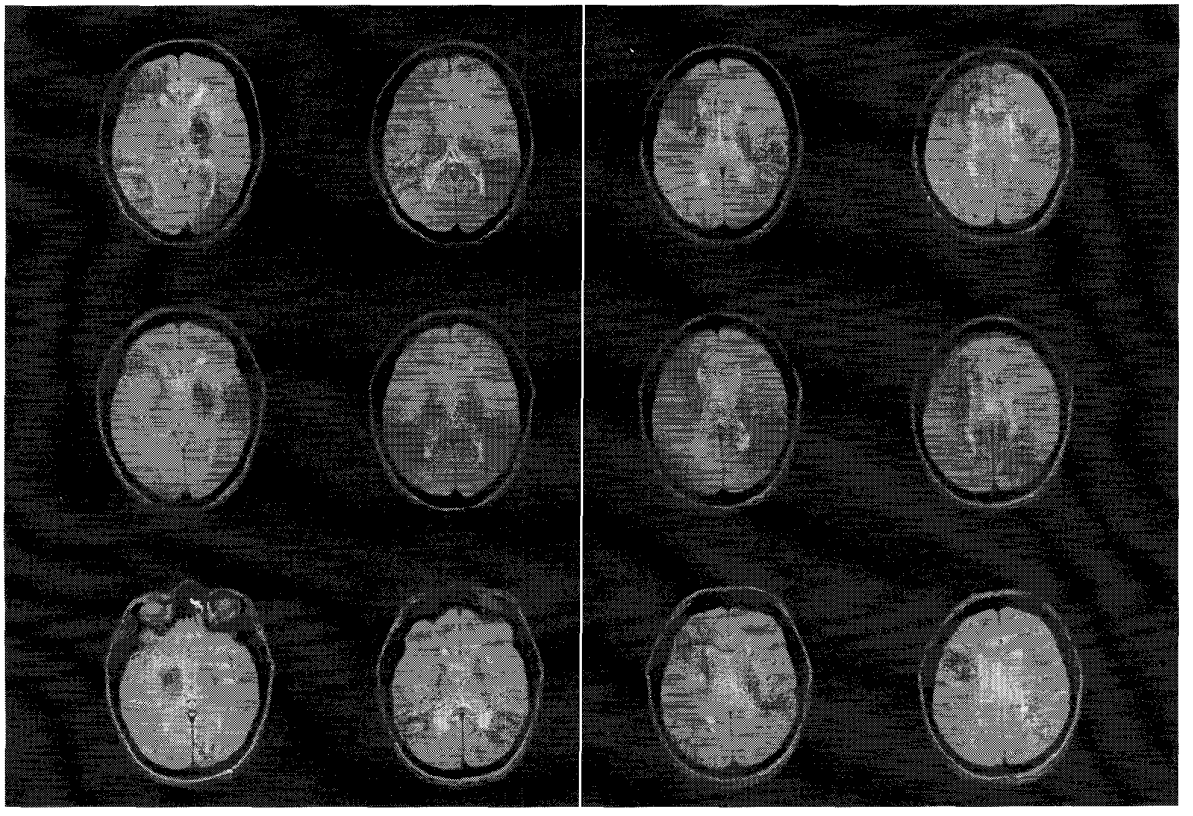

Fig. 5. MRI time series. From top to bottom the rows contain: the baseline scan, results of the surface-based registration, and original slices in the dynamic scan. Although the brain tissues have changed between the baseline scan and the dynamic scan, there is sufficient stability for an adequate registration. However, note the interpolation errors in the scalp in the last image in row 2.

TABLE VII
MRI-PET fluorodopa transformation results

Method	Rotation (degrees)			Translation (mm)			Time
	x	y	z	x	y	z	(min)
Iterative Manual	-17.80	-1.99	-3.51	-66.63	-77.38	17.74	28
Direct Manual	-17.45	0.07	-1.89	-68.62	-81.83	19.56	21
Surface	12.07	5.52	19.96	-64.02	-81.65	-2.54	6
Voxel Similarity	-16.64	-0.18	0.06	-66.67	-78.84	19.71	7

Voxel spacing for PET/MRI is 2.608/0.780mm in x & y and 3.375/7.00mm in z

parison functions for our automatic methods were based on nearest neighbor interpolation in-slice, and linear interpolation between the slices. Trilinear interpolation was used whenever resampling the data for viewing. The use of better interpolation functions for both the manual and automatic methods would provide more accurate solutions, with increased time requirements. For quantitative use of the automatically registered scans the increased time requirements of trilinear interpolation would certainly be justified. Whether a resampling filter based on the sinc function is worth the time is probably dependent on the application requirements (Cepstral filtering could provide further increases in accuracy when using the complex data from MRI [46]). The correction of other, more complex, distortions has not been addressed in this article, but these aspects are crucial in obtaining higher accuracy.

3) Robustness: Robustness refers to the probability of a method succeeding given new data and/or greater misalignment. The iterative manual approach is the most robust of the methods. The ability to interactively guide the registration process throughout allows the user's knowledge to be fully exploited. The direct manual method is quite robust with small rotations, but breaks down as the out of plane (x and y axis) rotations become larger. This is especially true with PET and SPECT images due to the lower resolution and lack of well-defined features, and so the user must face the more difficult problem of mental 3-D rotation for landmark identification (mental 3-D rotation requires cognitive processing while translation and in-plane rotation are visual cortex processes and thus relatively much simpler). The interactive method's provision to gradually align the scans allows the user to transcend this visualization problem.

The automatic methods can provide a high degree of robustness under well-constrained problems, such as in a time series study, where the patient is immobilised in a mask and every

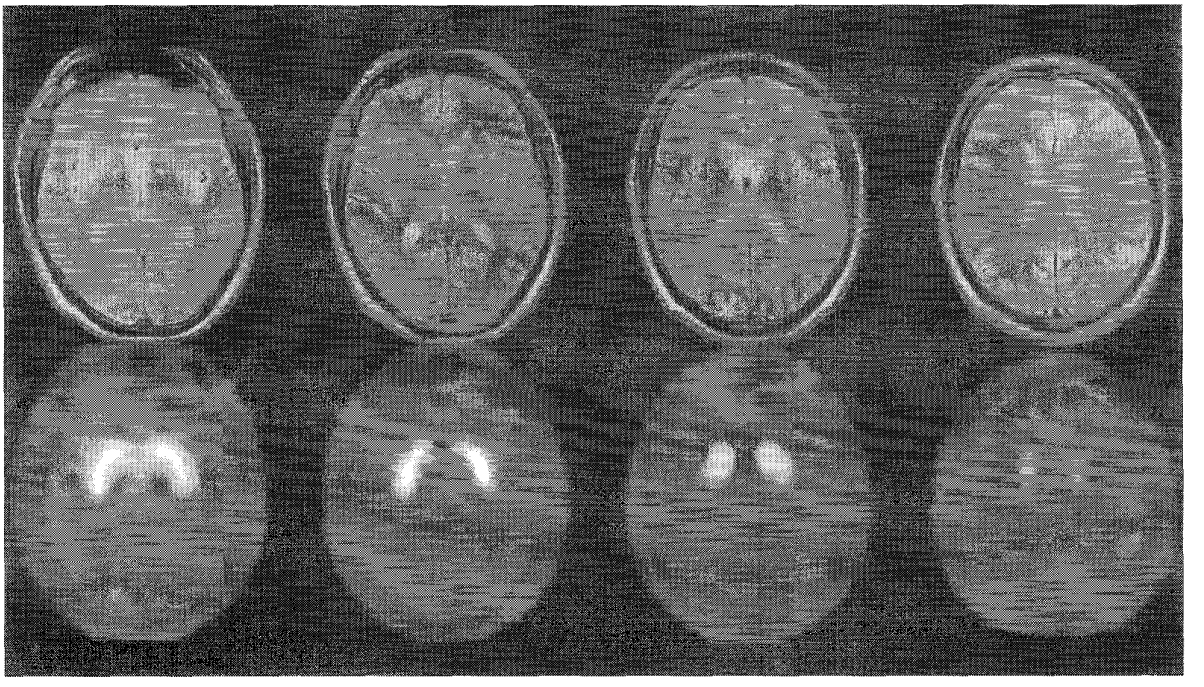

Fig. 6. MRI-PET registration. The top row displays consecutive slices of the MRI scan, and the bottom row shows the PET resampled to lie along those same slice planes, after registration with the manual iterative method.

effort is made to re-position the patient in a standard way. When facing new types of problems the automatic methods must be specially adapted or they will often find erroneous solutions. The surface-based approaches, by performing data reduction before the matching process, lead to a more efficient representation than the voxel-similarity approach. In doing so they run the risk of eliminating too much information, but may regain robustness by exploring more possible solutions in the same time period. The robustness is tightly linked to the consistency of the dynamic and static surface segmentations. If a specific structure's surfaces are believed to be the most consistent feature present, then the surface-based approach may lead to greater robustness. Once the segmentation is performed the robustness will depend on the surface's complexity and symmetry, as this determines the presence and number of local extrema in the parameter space.

When registering PET and MRI brain images Turkington [14] has shown surface matching to be robust even when using partial surfaces. The voxel-similarity methods appear to add robustness by moving the segmentation problem into the heart of the matching process. At the same time segmentation issues must inevitably be faced, but it is hoped that higher robustness can be attained by using more of the information present in the scan. A limitation of the voxel-similarity method is that it is based on an overlap criterion, and these criterion (e.g., stochastic sign change, sum of absolute differences) are usually used as a correlation function over the entire image or volume space. This is because a correlation will result in numerous local extrema, and so the entire domain should be traversed in order to find the global extremum. This is most clear when the two objects being registered do not overlap at all and so the overlap-based comparison function provides no information on which direction to move. In comparison, a distance metric can extend over an infinite domain and so will generally point the algorithm in the right direction. However, when the two objects initially overlap to any great extent, as is the case with most scans of the head, the overlap criterion's gradient is likely to point toward the global extremum. The presence of multiple extrema with the overlap-based comparisons can often be managed through the use of multiresolution matching. At the lowest resolution the features are smoothed so that there are very few extrema. Then the largest extrema can more easily be found, and provide the starting point for the next higher resolution optimization.

All of the methods proved to be robust over the simulated test cases. However, with the diverse range of clinically acquired scans there is often little consideration for the registration process. This is most evident when there is very little physical overlap between the scans, and so automatic methods will have great difficulty coping with the increased number of local minima. By using a large number of starting points, it is sometimes still possible to obtain an accurate automatic registration. With the larger time factors involved in these multiple searches, it may not be worth waiting, especially when the results may still be incorrect.

4) Usability: The iterative manual approach is very simple in nature, and after the method has been used once, the user basically understands it completely. Because visual verification is part of the process, the user will almost always be satisfied with the solution. The direct approach is conceptually even simpler, but suffers from the mental rotation problem. Without an interactive way to transform the data the user

must make tough decisions (i.e., difficult to reverse) about correspondences. The results must then be verified separately, albeit the user should be satisfied with the solution if they had high confidence when selecting corresponding points. The simple process behind the manual methods would indicate high usability, but this rating does not reflect the expert knowledge needed to select corresponding points. Both manual methods require the user to complete a lengthy and tedious task, thus resulting in a major drop in usability when registrations must be performed frequently. These factors, along with the possibly larger ones of repeatability, interoperator variability, and the operator's time, justify the research toward automatic methods.

Conversely, the automatic methods falsely provide the appearance of high usability because the user does not have to perform as many actions (this may not be true with most other groups' automatic registration methods which require semi-automatic segmentation or masking, taking up to 30 minutes per scan [8]). Because the user should understand all the variables controlling the automatic method the usability drops substantially. These controls become more complex as the programs utilize more knowledge of the problem. For an automatic method to be robust it must usually be highly configured to the specific problem (e.g., ^{18}FDG PET-MRI of the head) and so when the problem deviates slightly from the norm the configuration must be modified, and it now requires expert knowledge of not only the scans themselves but also of the complex registration program. The voxel-similarity method usually requires no separate segmentation, and so it is more usable than the surface-based approach. The results of both automatic methods must be verified separately as the comparison function does not provide a sound means to judge the accuracy of the solution. When the automatic methods fail the user must decide whether to adjust the parameters to try again or to utilize some other method.

V. Conclusion

This paper agrees and expands on many of the aspects of Kessler's [20] presentation of point, line, surface, and interactive registration methods. The results do not indicate that one single method is clearly superior for all registrations of the head. Rather, all of the methods should be provided to allow flexibility in solving the different types of registration problems. When scans with a large amount of overlap frequently need to be registered we feel the voxel-similarity approach should be used. In "hiding" the segmentation process it gains usability, robustness, and has the potential for increased accuracy, while at the same time not having large computer requirements.

To attain the best results from each of the different methods not only their strengths, but their limitations must be well understood. To meet this end, when presenting new methods the focus should be on where the methods fail as well as succeed. Completely automatic methods have been shown to be feasible for some constrained problems. This is often true when a large number of scans are performed for a study, and so the proper protocol can be devised to aid in the registration process.

Recent work has often progressed toward hybrid approaches that combine the strengths of the different approaches. Collignon has described a method [30] that uses corresponding points to constrain the transformation parameters search and surfaces to minimize total error, resulting in greater robustness. With a similar goal, Hill constrains the parameters through the combination of surfaces and a form of knowledge representation [21]. These types of specialized solutions can provide accurate results when other methods fail, but they also become more dependent on assumptions which may only hold under certain circumstances. Another feasible hybrid approach would utilize interactive manual alignment to get a quick initial estimate and then automatically perform exhaustive matching over a greatly restricted range of parameters. Improved understanding of the problems and faster computers will certainly provide a means for increased accuracy and robustness in the near future.

Acknowledgment

The authors wish to thank A. Riddehough and D. Paty [MRI Group, University of British Columbia (UBC) Hospital], S. Morrison and B. Snow (PET Group, UBC Hospital), D. Hill (Image Processing Group, Guy's Hospital, United Kingdom) for providing data and information about it; M. Menke (TRIUMF) for his help in utilizing the Levenberg–Marquadt optimization routine; U. Pietrzyk (Max–Planck-Institut für Neurologische Forshung) and R. P. Woods (Division of Nuclear Medicine, UCLA) for making their registration software available; S. Jang, C. Romanzin, and K. Booth (Dept. of Computer Science, UBC) for information and computing resources; and A. Collignon (Laboratory for Medical Imaging Research, ESAT/MI2, Belgium) for his helpful dialogue.

References

[1] G. J. Ettinger, W. E. L. Grimson, T. Lozano-Pérez, W. M. Wells III, S. J. White, and R. Kikinis, "Automatic registration for multiple sclerosis change detection," in *IEEE Workshop on Biomedical Image Analysis*, Seattle, WA, USA, 1994, pp. 297–306.

[2] S. Wong, R. Knowlton, M. Chew, and H. Huang, "Integrating multidimensional imaging, multimodality registration, and multimedia database for epilepsy diagnosis," *SPIE—Med. Imag.*, p. 2431, 1995.

[3] J. Hajnal, N. Saeed, E. Soar, A. Oatridge, I. Young, and G. Bydder, "Detection of subtle brain changes using subvoxel registration and subtraction of serial MR images," *J. Comput. Assist. Tomogr.*, vol. 19, no. 5, pp. 677–691, 1995.

[4] B. Davey, R. Comeau, P. Munger, L. Pisani, D. Lacerte, A. Olivier, and T. Peters, "Multimodality interactive stereoscopic image-guided neurosurgery," *SPIE Visualization in Biomed. Computing*, vol. 2359, pp. 526–536, 1994.

[5] R. Wasserman, J. Rajapakse, and R. Acharya, "Multimodality medical imaging for radiotherapy treatment planning," *IEEE Workshop on Biomed. Image Anal.*, 1994, pp. 235–244.

[6] A. Liu, S. Pizer, D. Eberly, B. Morse, J. Rosenman, E. Chaney, E Bullitt, and V. Carrasco, "Volume registration using the 3D core," *SPIE Visualization in Biomed. Computing*, vol. 2359, pp. 217–226, 1994.

[7] P. A. van den Elsen, J. B. A. Maintz, E. D. Pol, and M. A. Viergever, "Image fusion using geometrical features," in *SPIE Visualization in Biomed. Computing*, vol. 1808, pp. 172–186, 1992.

[8] R. P. Woods, J. C. Mazziotta, and S. R. Cherry, "MRI—PET registration with automated algorithm," *J. Comput. Assist. Tomogr.*, vol. 17, no. 4, pp. 536–546, July/Aug. 1993.

[9] A. Collignon, D. Vandermeulen, P. Suetens, and G. Marchal. "Surface-based registration of 3D medical images," in *SPIE Med. Imag.: Image Processing*, vol. 1898, pp. 32–42, 1993.

[10] D. L. G. Hill, C. Studholme, and D. J. Hawkes. "Voxel similarity measures for automated image registration," in *SPIE Visualization in Biomed. Computing*, vol. 2359, pp. 205–216, 1994.

[11] J. Hajnal, N. Saeed, A. Oatridge, E. Williams, I. Young, and G. Bydder, "A registration and interpolation procedure for subvoxel matching of serially acquired MR images," *J. Comput. Assist. Tomogr.*, vol. 19, no. 2, pp. 289–296, 1995.

[12] G. R. Timmens, P. A. van den Elsen, F. J. R. Appelman, and M. A. Viergever, "Landmark-based registration of images: A comparison of algorithms," in *Int. Conf. on Volume Image Processing*, Utrecht, The Netherlands, June 1993, pp. 145–148.

[13] T. G. Turkington, R. J. Jaszczak, C. A. Pelizzari *et al.*, "Accuracy of registration of PET, SPECT and MR images of a brain phantom," *J. Nucl. Med.*, vol. 34, no. 9, pp. 1587–1594, 1993.

[14] T. G. Turkington, J. M. Hoffman, R. J. Jaszczak *et al.*, "Accuracy of surface fit registration for PET and MR brain images using full and incomplete brain surfaces," *J. Comput. Assist. Tomogr.*, vol. 19, no. 1, pp. 117–124, Jan./Feb. 1995.

[15] C. A. Pelizzari, A. C. Evans, P. Neelin, C. T. Chen, and S. Marrett, "Comparison of two methods for 3-D registration of PET and MRI images," *Annu. Int. Conf. IEEE Eng. in Med. Biol. Soc.*, 1991, vol. 13, no. 1, pp. 221–223.

[16] P. F. Hemler, S. Napel, T. S. Sumanaweera *et al.*, "Registration error quantification of a surface-based multimodality image fusion system," Computer Science Department, Stanford University, Stanford, CA, *Tech. Rep. KSL-94-66*, 1994.

[17] U. Pietrzyk, K. Herholz, and W. Heiss, "Three-dimensional alignment of functional and morphological tomograms," *J. Comput. Assist. Tomogr.*, vol. 14, no. 1, pp. 51–59, 1990.

[18] I. Kapouleas, A. Alavi, W. M. Alves, R. E. Gur, and D. W. Weiss, "Registration of three-dimensional MR and PET images of the human brain without markers," *Radiol.*, vol. 181, no. 3, pp. 731–739, 1991.

[19] Y. Ge, J. M. Fitzpatrick, J. R. Votaw *et al.*, "Retrospective registration of PET and MR brain images: An algorithm and its stereotactic validation," *J. Comput. Assist. Tomogr.*, vol. 18, no. 5, pp. 800–810, Sept./Oct. 1994.

[20] M. L. Kessler, S. Pitluck, P. Petti, and J. R. Castro, "Integration of multimodality imaging data for radiotherapy treatment planning," *Int. J. Radiation Oncology, Biology, Physics*, vol. 21, no. 6, pp. 1653–67, Nov. 1991.

[21] D. L. G. Hill and D. J. Hawkes, "Medical image registration using knowledge of adjacency of anatomical structures," *Image and Vision Computing*, vol. 12, pp. 173–178, 1994.

[22] C. J. Henri, A. Cukiert, D. L. Collins, A. Olivier, and T. M. Peters, "Toward frameless stereotaxy: Anatomical-vascular correlation and registration," in *SPIE Visualization in Biomed. Computing*, vol. 1808, pp. 214–224, 1992.

[23] G. Q. Maguire, Jr., M. E. Noz, H. Rusinek, J. Jaeger, E. L. Kramer, J. J. Sanger, and G. Smith, "Graphics applied to medical image registration," *IEEE Computer Graphics and Applications*, vol. 11, no. 2, pp. 20–28, Mar. 1991.

[24] C. Schiers, U. Tiede, and K. H. Höhne, "Interactive 3D registration of image volumes from different sources," in *International Symposium on Computer Assisted Radiology*, Berlin (West), Federal Republic of Germany, 1989, pp. 667–670.

[25] P. A. van den Elsen, E. J. D. Pol, and M. A. Viergever, "Medical image matching—A review with classification," *IEEE Eng. Med. Biol.*, vol. 12, no. 1, pp. 26–39, Mar. 1993.

[26] P. Gerlot-Chiron and Y. Bizais, "Registration of multimodality medical images using region overlap criterion," *Comput Vision Graphics Image Process: Graphical Models and Image Processing*, vol. 54, no. 5, pp. 396–406, Sept. 1992.

[27] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge, MA: Univ. of Cambridge, 1988.

[28] M. Menke, M. S. Atkins, and K. R. Buckley, "A software compensation method for videometrically detected motion during head PET scans," *IEEE Trans. Nuclear Sci.*, vol. 43, no. 2, pp. 310–317, Feb. 1996.

[29] C. A. Pelizzari, G. T. Y. Chen, D. R. Spelbring, R. R. Weichselbaum, and C. Chen, "Accurate three-dimensional registration of CT, PET, and/or MR images of the brain," *J. Comput. Assist. Tomogr.*, vol. 13, no. 1, pp. 20–26, Jan./Feb. 1989.

[30] A. Collignon, D. Vandermeulen, P. Suetens, and G. Mar.al, "Registration of 3D multi-modality medical images using surfaces and point landmarks," in *Int. Conf. Volume Image Processing*, Utrecht, The Netherlands, June 1993, pp. 27–30.

[31] M. van Herk and H. Kooy, "Three-dimensional CT-CT, CT-MR, and CT-SPECT correlation using chamfer matching," in *Int. Conf. Volume Image Processing*, Utrecht, The Netherlands, June 1993, pp. 25–26.

[32] L. M. Bidaut, "Accurate registration of various medical imaging modalities: An application to PET and MRI for anatomical localization and kinetic modeling," *IEEE Melecon*, pp. 1233–1237, 1991.

[33] D. N. Levin, C. A. Pelizzari, G. T. Y. Chen, C. Chen, and M. D. Cooper, "Retrospective geometric correlation of MR, CT, and PET images," *Radiol.*, vol. 169, pp. 817–823, 1988.

[34] T. Zuk, M. S. Atkins, and K. Booth, "Approaches to registration using 3-D surfaces," in *SPIE Med. Imag.: Image Processing*, vol. 2167, pp. 176–187, 1994.

[35] G. Borgefors, "Distance transformations in arbitrary dimensions," *Comput Vision Graphics Image Process*, vol. 27, pp. 321–345, 1984.

[36] G. Borgefors, "Hierarchical chamfer matching: A parametric edge matching algorithm," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 10, no. 6, pp. 849–865, Nov. 1988.

[37] H. Jiang and R. Robb, "Image registration of multimodality 3D medical images by chamfer matching," in *SPIE Biomed. Image Processing 3-D Microscopy*, vol. 1660-1635, pp. 356–366, 1992.

[38] M. J. D. Powell, "An efficient method for finding the minimum of a function of several variables without calculating derivatives," *Comput. J.*, vol. 7, pp. 155–162, 1964.

[39] R. P. Woods, S. R. Cherry, and J. C. Mazziotta, "Rapid automated algorithm for aligning and reslicing PET images," *J. Comput. Assist. Tomogr.*, vol. 16, no. 4, pp. 620–633, July/Aug. 1992.

[40] C. Studholme, D. L. G. Hill, and D. J. Hawkes, "Using voxel similarity as a measure of medical image registration," in *British Machine Vision Conference*, E. Hancock, Ed. BMVA Press, 1994, pp. 235–244.

[41] A. Collignon, D. Vandermeulen, and P. Suetens, "Automatic registration of 3D images of the brain based on fuzzy objects," in *SPIE Med. Imag.: Image Processing*, vol. 2167, pp. 162–175, 1994.

[42] A. Collignon, D. Vandermeulen, P. Suetens, and G. Mar.al, "3-D multimodality medical image registration using feature space clustering," to appear in *CVR Med'95*.

[43] B. A. Ardekani, M. Braun, I. Kanno, and B. F. Hutton, "Automatic detection of intradural spaces in MR images," *J. Comput. Assist. Tomogr.*, vol. 18, no. 6, pp. 963–969, Nov./Dec. 1994.

[44] S. C. Strother, J. R. Anderson, X. Xu, J. Liow, D. C. Bonar, and D. A. Rottenberg, "Quantitative comparisons of image registration techniques based on high-resolution MRI of the brain," *J. Comput. Assist. Tomogr.*, vol. 18, no. 6, pp. 954–962, Nov./Dec. 1994.

[45] H. Rusinek, W.-H Tsui, A. Levy, M. Noz, and M. Leon, "Principal axes and surface fitting methods for three-dimensional image registration," *J. Nuclear Med.*, vol. 34, pp. 2019–2024, 1993.

[46] E. Bandari, Q. S. Xing, and J. Little, "Visual echo registration of magnetic resonance images," in *AAAI Spring Symp. Applications of Computer Vision in Medical Image Processing*, Mar. 1994, pp. 38–41.