Automatic Registration for Multiple Sclerosis Change Detection

G.J. Ettinger[1,2] W.E.L. Grimson[1] T. Lozano-Pérez[1] W.M. Wells III[1,3]
S.J. White[2] R. Kikinis[3]

Abstract

We are developing an automated 3D change detection system which accurately registers medical imagery (e.g., MRI or CT) of the same patient from different times for diagnosing pathologies, monitoring treatment, and tracking tissue changes. The system employs a combination of energy-minimization registration techniques to achieve accurate and robust alignment of 3D data sets. The bases for the registration are 3D surfaces extracted from the 3D imagery. Resultant structural changes in the data are identified by using an adaptive segmentation technique to automatically determine tissue morphology. The novel contributions of this work are its end-to-end automation of the change detection process and its high accuracy in monitoring and highlighting such physiological changes. We have applied this system to a multiple sclerosis study in which each patient had been imaged over 20 times for the purpose of tracking lesion evolution. This report describes preliminary registration performance analysis using this data.

1 Change Detection Problem

A growing use of clinical imagery is the identification of medically-significant tissue changes over time. In studying conventional single images, medical professionals can locate possibly anomalous structures based on their knowledge of anatomy and pathology. But by *comparing current images against baseline or previous images*, clinicians may also be able to estimate the rate of change in the progression of a disease or as a reaction to a treatment. These improvements should facilitate earlier and more accurate diagnosis. The availability of relatively high resolution 3D raster data sets from diagnostic scanners and the growing emphasis on prevention and early detection of disease combine to underscore the need for accurate change detection technology.

[1] Artificial Intelligence Lab, MIT, Cambridge MA
[2] TASC, Reading MA
[3] Dept. of Radiology, Brigham and Womens Hospital, Harvard Medical School, Boston MA

The main issues with solving such change detection problems are

1. accurate *registration* of the imagery over time, and
2. reproducible *segmentation* of the imagery for highlighting changes.

In order to achieve accurate registration we use a series of matching techniques for first coarsely aligning the two data sets and then refining the match using energy-minimization techniques over both interpolated data and finely sampled data. In order to achieve data segmentation which is not only accurate but also stable over time, we use an adaptive Expectation-Maximization technique which accounts for local sensor gain variations. A key to our approach is automating these steps to achieve an effective end-to-end change detection tool. In other words, raw images from two time points are provided as input and the system automatically generates structural tissue changes as output. Automating these processes results in a beneficial tool by:

- minimizing the cost of performing change detection procedures,
- ensuring analysis of the complete data set (generally a large volume of data),
- minimizing error introduced by manual intervention, and
- increasing the efficiency of change detection evaluations.

We have applied our change detection system to the problem of registering magnetic resonance imagery (MRI) of multiple sclerosis (MS) patients over time for the purpose of tracking changes in the MS brain lesions. These lesions, possibly numbering in the hundreds, are small patches of diseased tissue in the brain's white matter. The size and structure of these lesions change over time and indicate the state and severity of the disease. Tracking changes in the lesions can thus provide critical information in understanding the progression of the disease and monitoring treatment.

0-8186-5802-9/94 $3.00

This change detection problem is a central issue in an NIH study being carried out by the Department of Radiology at Brigham and Women's Hospital, Boston MA, aimed at determining the optimal frequency for performing MR imaging of MS patients. The study covers patients exhibiting different stages of the disease, including chronic progressive, exacerbating-remitting, relapsing-remitting progressing, and inactive. The approximate MR imaging schedule is: weekly for the first eight weeks, bi-weekly for the next 16 weeks, and monthly for the next 26 weeks. Initial analysis of this data, without temporal registration, included collection of only global statistics over time, such as the changes in the total volume of MS lesions in the brain. But the accurate registration of the MRI data sets over time, described in this report, is facilitating tracking of the individual lesions and allowing study of local changes.

2 Change Detection Algorithm

Our change detection system consists of four basic steps, as depicted in Figure 1:

1. **Surface Point Extraction** – localization of surfaces of the same structure in both input images; we represent surfaces as sets of 3D points lying on the surface.

2. **3D Registration** – compute a transformation that maps points from one data set into the coordinate frame of the other data set; we are currently solving for a rigid (six degree of freedom) transformation.

3. **Tissue Class Segmentation** – classify voxels in the 3D imagery according to their tissue type, such as grey matter, white matter, cerebrospinal fluid, and lesions.

4. **Structural Change Localization** – using the transformation computed by the 3D Registration, transform the segmented images to the same coordinate frame and identify changes in the tissue structure via 4D (space and time) connected component analysis [19], image differencing, local statistics, or direct visualization.

We are exploring methods for automating all four of these steps. Our primary focus is on automating the registration process since that step generally depends most on global calculation consistency (i.e., small errors are often magnified) and can be difficult to perform manually with high accuracy. Semi-automated techniques can often be used for the other steps with less efficient but often adequate results. Errors in surface point extraction can often be removed as outliers or canceled out in the registration process. Segmentation, while a difficult problem in general, can be less critical in change detection problems since imaging parameters can be tuned to highlight specific tissues of interest and changes can also be located with the raw imagery. The change localization step is usually a straightforward visualization or analysis operation which is not difficult to automate.

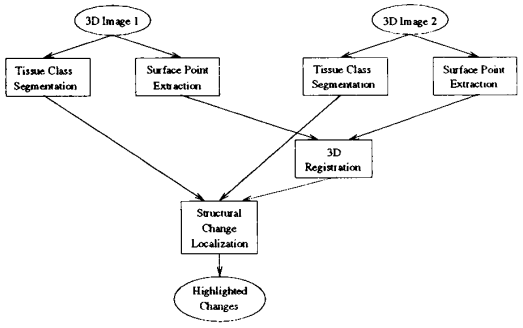

Figure 1: **Change detection algorithm.**

2.1 Surface Point Extraction

For surface point extraction we are currently pursuing techniques for localizing positionally stable surfaces such as skin or intra-cranial cavities (ICC). In some cases these surfaces can be extracted automatically via straightforward thresholding, connected components, and edge detection techniques. Skin surface, for example, can often be extracted using these simple operations. For the MS application described in this paper we used the intra-cranial cavity (the brain plus the surrounding cerebrospinal fluid) as our reference surface since it was generally completely imaged and, by virtue of lying against the inside of the cranium, exhibits little motion [4]. Skin points, on the other hand, can be incomplete and exhibit some movement from one time point to the next.

The surface data we used for the MS study was existing ICC boundaries extracted at Brigham and Women's Hospital using a semi-automated technique that is part of their routine 3D analysis of MR imagery. This technique consists of constructing a non-parametric intensity classifier for tissues that are internal and external to the ICC and then applying morphological filtering techniques to refine the ICC mask

[4] The patients in this MS study were imaged at roughly the same pose in the scanner, minimizing brain motion across scans. The imaging geometry, though, cannot be duplicated accurately enough at different times to avoid the need for registration.

[3, 5, 14]. Double echo MR imagery was used for the MS study at a resolution of $0.9375mm$ x $0.9375mm$ x $3.0mm$. We are also exploring a fully automated surface extraction technique using a variant of snakes to dynamically fit a surface to the ICC boundary [13].

2.2 3D Registration

The central part of our problem is to register two data sets to one another. Our initial registration approach searches for the best *rigid body transformation.* Of course, some medical registration problems will require more flexible registration, but for the head registration described here (as well as other applications), rigid transformations appear to suffice.

The inputs to the registration process consist of two data sets, represented as sets of 3D points, each in its own coordinate system. The points are assumed to lie on the same structural surface, although the coverages of the points do not need to exactly overlap and outliers may be present. Our problem is to determine a transformation that will map one data set into the other in a consistent manner.

We match the two data sets using the steps outlined in Figure 2 and described below.

2.2.1 Initial Match

If there is limited overlap in the coverage of the two data sets (i.e., if there are not both relatively complete models of the anatomy), then we need to get a rough initial alignment of the two data sets. In previous work, we have accomplished this by sampling a small number of widely spaced points from one of the data sets, and then using Interpretation Tree Search [7] to match those sampled points to data points in the other data set. This method basically searches over all possible ways of matching small sets of data features from the two data sets. For each pairing of data features from the two data sets, the method tests whether the pairwise distances between points are roughly the same. If all such tests are valid, the match is kept, and we compute the coordinate frame transformation that maps the data points from the first set into their corresponding points in the second set. These transformations form a set of hypotheses. Note that due to the sampling of the data, the actual corresponding points may not exist, so these hypothesized transformations are at best approximations to the actual transformation.

For the problem with which we are concerned here, we typically have complete data sets from the same (or nearly the same) portion of the anatomy. In this case, we can use a simpler method to get a rough initial

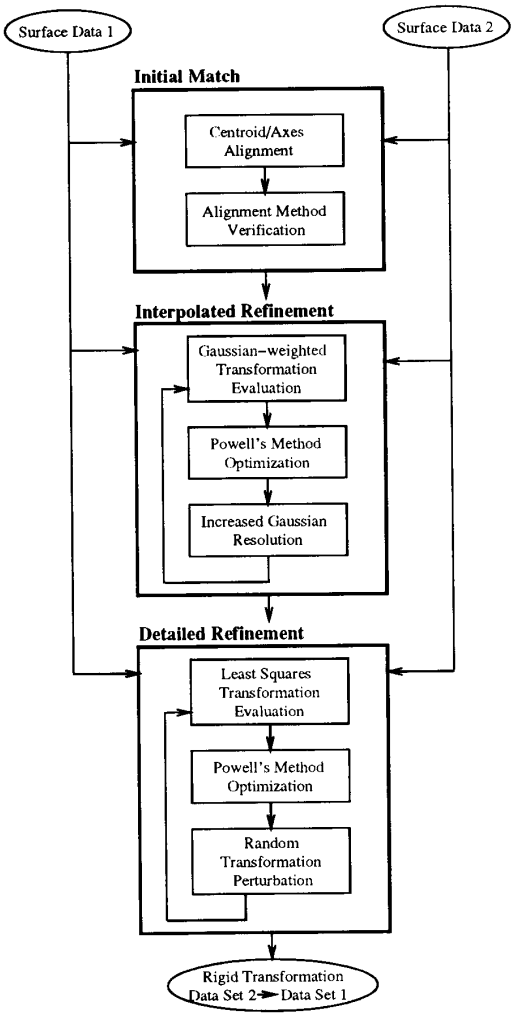

Figure 2: **3D registration algorithm.**

alignment. In particular, we can use the eigenvectors of the inertia matrix of each of the data sets. We can then translate the second data set so that its centroid aligns with the centroid of the first data set, and we can rotate the second data set so that its eigenvectors align with the corresponding eigenvectors of the first data set in all possible combinations. Note that due to a sign ambiguity in the direction of each of theses axes we need to keep a total of 144 different hypothesized transformations at this stage.

We use the Alignment Method [11] to filter these hypotheses. For each hypothesis, we transform all the points of one data set by the hypothesized transformation, and verify that the fraction of the transformed points that do not have a corresponding point from

the other data set, within some predefined distance, is less than some predefined bound. We discard those hypotheses that do not satisfy this verification. We can similarly use the RMS evaluation described in Section 2.2.3 to evaluate each hypothesis and select the best one(s).

2.2.2 Interpolated Refinement

For each verified hypothesis, we perform an initial refinement aimed at guiding the registration in the general direction of the global error minimum. To perform this refinement we evaluate the current pose by summing, for all transformed points (from data set 2), a term that is itself a sum of the distances from the transformed point to all nearby reference surface points (data set 1), where the distance is weighted by a Gaussian distribution [26]. This Gaussian weighted distribution is a method for roughly interpolating between the sampled reference points to estimate the nearest point on the underlying surface to the transformed data point. More precisely, if ℓ_i is a vector representing a data point, m_j is a vector representing a reference point, and $\mathcal{T}$ is a coordinate frame transformation, then the evaluation function for a particular pose (or transformation) is $E_1(\mathcal{T}) = -\sum_i \sum_j e^{-\frac{|\mathcal{T}\ell_i - m_j|^2}{2\sigma^2}}$. This objective function is similar to the posterior marginal pose estimation (PMPE) method used in [26]. One can visualize this objective function as if we placed a Gaussian distribution of some spread σ at each reference surface point, then summed the contributions from each such distribution at each point in the volume. Then the contribution of each transformed data point towards the evaluation function is simply the summed value at that point. Because of its formulation, the objective function is generally quite smooth, and thus facilitates "pulling in" solutions from moderately removed locations in parameter space. This evaluation function is iteratively minimized using Powell's method [23]. The result is an estimate for the pose of the second data set in the coordinate frame of the first data set.

We execute this refinement and evaluation process using a multiresolution set of Gaussians. Initially, a broad based Gaussian is used to allow influence over large areas, resulting in a coarse initial alignment, but one which can be reached from a wide range of starting positions. Subsequently, more narrowly tuned Gaussian distributions can be used to refine the pose, while focusing on only nearby data points to derive the pose.

2.2.3 Detailed Refinement

Based on the resulting pose of the interpolated refinement, we repeat the pose evaluation process using a rectified least squares distance measure. Each pose is evaluated by measuring the distance from each transformed data point to the nearest reference surface point, (with a cutoff at some predefined maximum distance to guard against outliers or missing data). The pose evaluation is the sum of the squared distances of each point. Powell's method is again used to find the least-squares pose solution. Here the evaluation function is $E_2(\mathcal{T}) = \sum_i \min\{d^2_{\max}, \min_j |\mathcal{T}\ell_i - m_j|^2\}$ where $d_{\max}$ is some preset maximum distance. This objective function is essentially the same as the maximum a posteriori (MAP) model matching scheme of [26]. It acts much like a robust chamfer matching scheme, similar to that used by [12]. The expectation is that this second objective function is more accurate locally, since it is composed of saturated quadratic forms, but it is also prone to getting stuck in local minima.

In order to avoid such local minima, we randomly perturb the solution and repeat the least squares refinement. The observation is that while the above method always gets very close to the best solution, it can get trapped into local minima in the minimization of E_2. We thus continue this perturbation and refinement process, keeping the new pose if its associated RMS error is better than our current best. We terminate this process when the number of such trials that have passed since the RMS value was last improved becomes larger than some threshold. The final result is a pose, and a measure of the residual deviation of the fit to the reference surface.

We collect such solutions for each verified hypothesis, and rank order them by smallest RMS measure. The result is a highly accurate transformation of one data set into the coordinate frame of the other data set.

2.3 Related Registration Algorithms

Several other groups have reported registration methods similar to ours. Of particular interest are four such approaches. First, Pelizzari and colleagues [18, 20, 21, 22] have developed a method that matches retrospective data sets, such as MRI or CT or PET, to one another. Similar to our approach, this work uses a least squares minimization of distances between data sets, although their system uses a different distance function for minimization, in particular using for each data point, the distance from that data point to the model surface, as measured along a ray from the

data point to the centroid of the model. Typical reported RMS errors are on the order of $3-5mms$. As opposed to our system, however, this approach does require some operator intervention to set a decent initial starting position, which our system does not. It also apparently requires some operator intervention to steer the system towards the correct solution, suggesting that local minima are a potential problem. Our system avoids this difficulty by randomly perturbing near final solutions to find better nearby minima.

A second related approach is that of Lavallee, Szeliski, and colleagues [2, 16, 17, 25]. This method also does a least-squares minimization of a distance function to match data sets. Here, the distance is weighted by an estimate of the inverse variance of the noise in the measurements, and a Levenberg-Marquardt method is used to find the minimum. It appears that at present the method requires a reasonable initial starting position, though the authors observe that sampling over the view sphere could remove this restriction. Once a first estimate of the solution is found, points with large errors are removed (they are considered to be outliers) and the minimization process is repeated to refine the pose. It is unclear whether the removal of outliers is sufficient to keep the method from getting trapped into local minima.

A third approach is that of Ayache, Gueziec, and colleagues [1, 8, 9, 10] who perform automatic rigid registration of 3D surfaces by matching ridge lines which track points of maximum curvature along the surface. The ridge lines are characterized by five intrinsic parameters (curvature, torsion, maximum curvature of the surface, angle between curve and surface normal, and angle between curve tangent and direction of maximum curvature) which are used to hash the ridge points of the model data set into a five-dimensional hash table. During matching, the hash table is used to efficiently find model ridge points which are similar to ridge points from a second data set. A rigid transform is computed for each pair of matched ridge points and the results are collected in a six-dimensional transform accumulator. The most accessed cell in the accumulator is selected as the best transform. While this technique works well, it requires relatively dense data for its curvature computation. Our technique also works well for sparse data.

Finally, the work of Bajcsy et al [15] uses moments of inertia to align two data sets, similar to our initial match, followed by an elastic matching between the aligned data sets to handle small remaining variations. While we use eigenvectors of the inertia matrix as a first stage in our method, the multi-stage minimization of deviations between the aligned data sets provides an alternative method for finding the global minimum. We have found that using the axes to align the data sets can still result in large deviations, primarily because the two data sets may not actually sample exactly the same portions of the anatomy. Because of this, there may be a need for alternative methods for obtaining initial alignments, so that the minimization step is close enough to the global minimum to guarantee convergence. Our method has the ability to use initial alignment methods based on interpretation tree search, which do not suffer the same drawbacks.

2.4 Tissue Class Segmentation

For change detection applications we are usually interested in localizing and quantifying changes in particular tissues. Attempting to identify changes directly from the raw (MR) imagery is often problematic due to global and local sensor gain artifacts and magnetic susceptibility variations. By performing tissue segmentation we can systematically treat these problems, focus on the tissues of interest, and also provide improved context for change detection by introducing spatial knowledge of related anatomical structures.

Currently, semi-automatic statistical segmentation methods [3] are routinely used to segment MR images. These methods are based upon single-voxel intensity measurements, and the quality of the results is often limited by the intensity variability that is usually present even in the best MR images. Because of this problem, manual outlining on the multi-slice image data sets is frequently employed as the expensive and tedious alternative, when high accuracy is needed.

We have recently developed a robust statistical method that is used to automatically segment MR brain images of MS patients [27]. It uses statistical knowledge of tissue properties and gain inhomogeneities to correct the gain artifacts of MRI. The Expectation-Maximization algorithm [4] is used to iterate two components to convergence: tissue classification and gain field estimation. The result is a powerful new method for estimating tissue class and RF gain. This method is effectively immune to the intensity inhomogeneities that are usually present in such images, and is significantly more powerful in this regard than other published methods.

This method may be applied to multiple tissue classes which have tissue intensity models represented by arbitrary non-parametric probability distributions. In addition, multiple channels of intensity information are used – in the MS study, these are proton-density and T2 weighted images. This method successfully adapts to intensity variations that are due to spatial inhomogeneities in the sensitivity of the imager, in-

tensity variations associated with different dates of acquisition (due to equipment changes) , and those due to the varying appearance of tissues among patients. Within our MS project, all of these sources of image intensity variation are accommodated with the same tissue intensity models, thus no per-scan or per-patient training is required.

As a further improvement, methods that correct for spatial distortions due to magnetic susceptibility differences between different tissue types citebinfordl can be incorporated to improve the accuracy of the segmentation.

2.5 Structural Change Localization

Once we have registered the two data sets, we transform one of the data sets into the coordinate system of the other. In order to accurately view changes we then resection the data to create new 2D slices of the transformed data that correspond to the slice positions of the reference data set. The reslicing can be applied to the raw data, in which case we employ a simple trilinear interpolation to resample the new slices, or to the segmented data, in which case we employ a nearest value sampling method. In either case, we perform an image difference[5] to detect positive and negative changes.

3 Multiple Sclerosis Study Results

We ran the registration system on 21 MRI data sets of the same patient collected over a period of one year. One of the data sets was selected as the standard and the other 20 were registered to it in order to generate a fixed coordinate frame in which to evaluate changes. Pertinent parameters for these runs are:

- The resolution of all the data sets was $0.9375mm$ x $0.9375mm$ x $3.0mm$.
- Values summed into the RMS evaluations are cut off at a maximum of $25mm^2$.
- Subsampling of the data sets for registration resulted in an average of 7060 points in the reference data set (factor of 4 subsampling) and 1415 points in the transformed data set (factor of 20 subsampling).

The registration results for one data set pair are shown in Figure 3 which overlays the transformed ICC surface points from one data set onto the surface of the other ICC. Most of the points appear to have been registered well. The principal error sources arise from the brain stem, where one scan included fewer slices, and from the top of the ICC, where the tangency of the surface to the slicing plane leads to partial voluming artifacts and the 2D sampling of surface points in these reduced cross sectional areas leads to sparser data. The RMS error for this run was $1.96mm$, with a median residual distance error of $1.57mm$. These values are close to the expected limits as dictated by the sampled resolution of the data.

Figure 4 shows the results of image differencing two slices at the same position in the reference and transformed data sets. The intensity images are resliced using trilinear interpolation while the segmented images are resliced using nearest values. Note that the main source of change inside the brain is due to MS lesion growth.

For the 20 test runs the average final RMS error was $1.92mm$. The average RMS after initial alignment was $2.38mm$. For these runs the gaussian interpolated refinement was not performed since the relatively close initial alignment and the high density of the data resulted in highly accurate registrations with just the detailed RMS minimization refinement[6]. Although the initial inertia axis alignment provided close enough starting points for registration refinement, it was generally not sufficient as a final result.

4 Registration Issues

The registration algorithm outlined in this report is designed to achieve automatic accurate alignment of rigid 3D data sets via the novel integration of segmentation, matching, and optimization techniques. The resultant system can be used not only for diagnosis via change detection, but also for enhanced reality visualization, frameless stereotaxy, and image-guided surgery [6]. In addition to brain studies, potential domains of application (with extensions to flexible registration) include orthopedics, mammography, craniofacial surgery, and minimally-invasive surgeries.

In designing our registration algorithm we focused on achieving robust performance, as dictated by the following principles:

- **Stability in the presense of input data errors.**
  Our registration technique combines matching contributions across all the available data, but limits the impact of data deviations by placing a limit on any one data point's contribution to the

[5] If differencing the raw images, their intensities are first normalized.

[6] Our experience with other data sets, in which the data is much more sparse, indicates that such interpolated refinement can be a powerful tool in avoiding local minima.

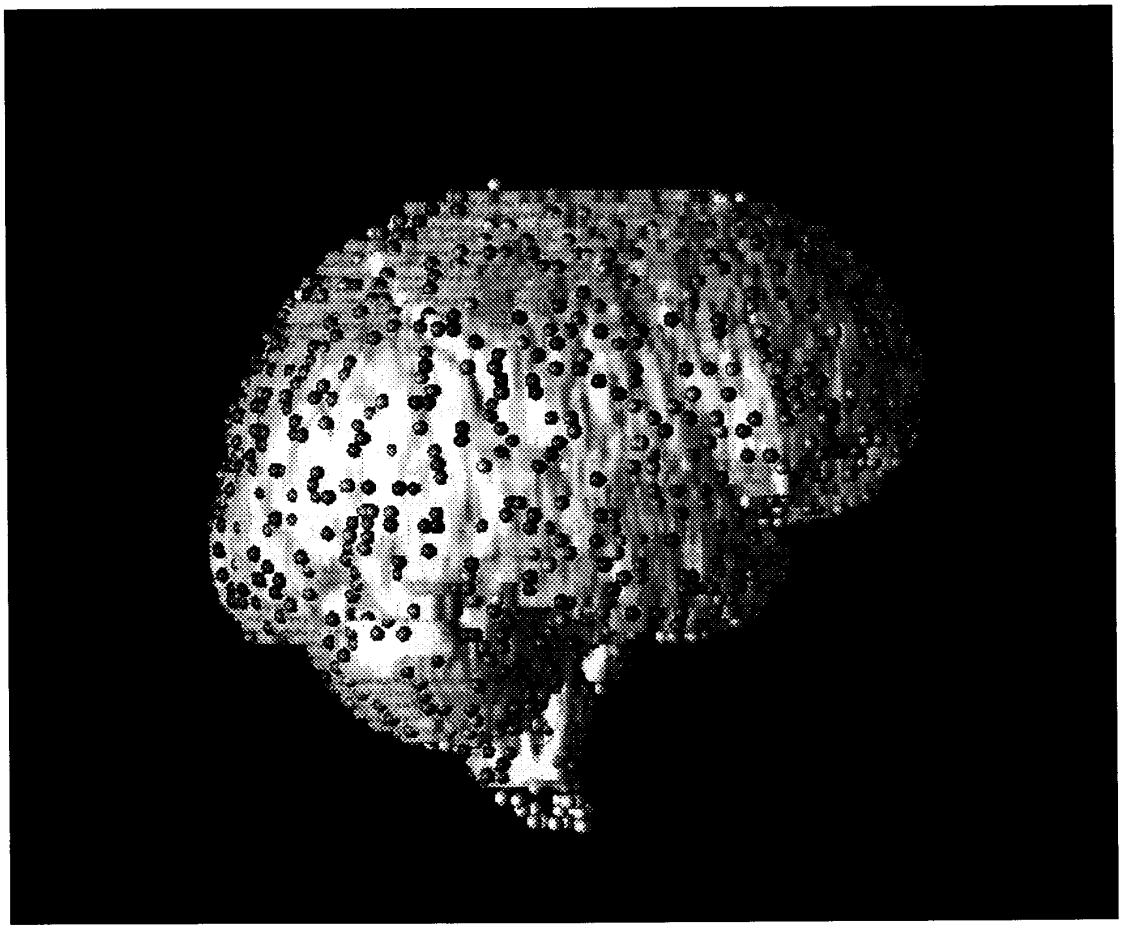

Figure 3: **Side view of registered points from one data set overlaid on surface of the other data set. Dot brightness varies from black ($0mm$ distance error) to white ($5mm+$ distance error).**

alignment evaluation. The resultant truncated least-squares approach is designed to achieve accurate registration in the presence of:

- surface extraction deviations; e.g., data segmentation errors, surface sampling artifacts,
- data outliers; e.g., clutter points and surface points present in one data set but not the other,
- imaging distortions.

- **Effective performance for both dense and sparse data.**
  Since our registration technique is based on aligning points to surfaces to solve for a consistent global transform, we are not dependent on local or dense data at any one point. Furthermore, by avoiding the issue of feature correspondence we avoid the need for dense data which is often needed for accurate feature extraction. In general we perform the registration by transforming the sparser smaller-coverage data set to the coordinate system of the denser broader-coverage data set.

  In applying our registration to dense data, such as the MS application discussed in this paper, we generally achieve efficient performance by randomly subsampling the data to achieve a uniform, lower resolution coverage across the data set. The main effect of increasing the subsampling factor, though, is to reduce the "depth" of the global minimum, resulting in a higher probability that

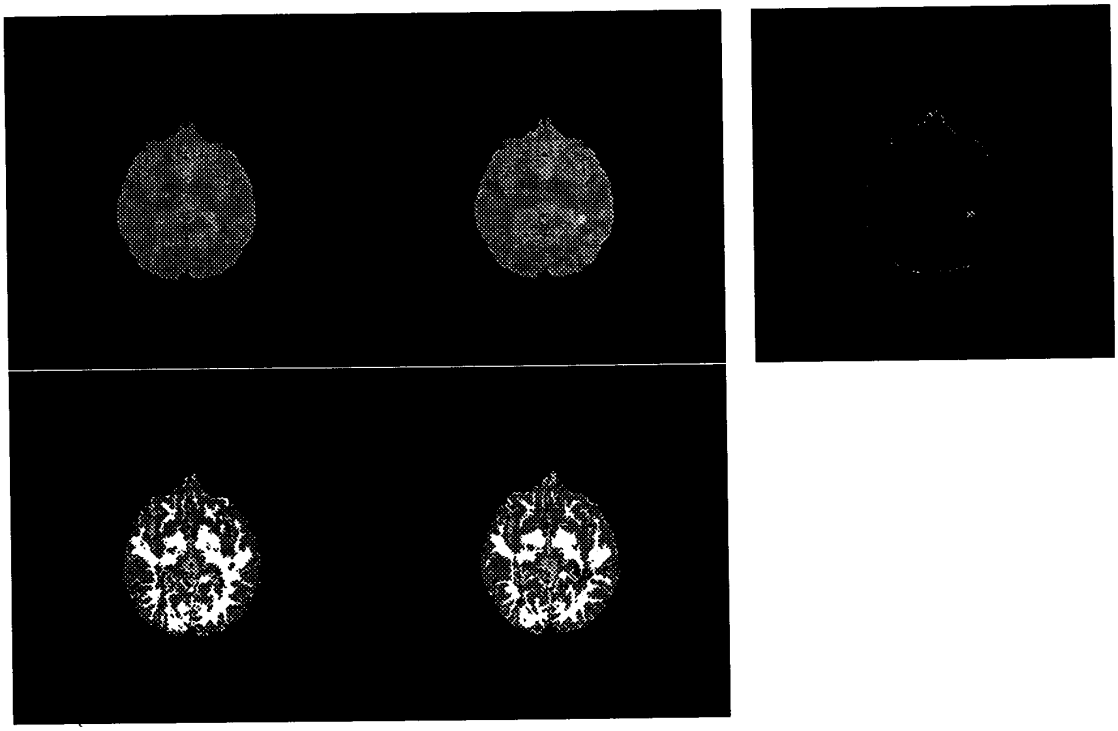

Figure 4: **Changes in the same slice position. First column is transformed data which was registered to second column. Imagery in second column was taken eight months later. Top row is normalized intensity imagery from one of the two MR data echos used in the study, along with the absolute difference between the two. Bottom row is the segmented imagery: dark grey = grey matter, white = white matter, light grey = CSF, and black = lesions.**

a sub-optimal alignment result will be computed. But as long as the subsampling factor is not increased to this point, the same final transform is generated for a coarser data set as for a higher resolution data set (although the lower resolution data will have a higher RMS error).

- **Minimum dependence on initial alignment.** By incorporating techniques to automatically generate initial alignments and then refine them using both interpolation and fine sampling techniques, we are able to register data sets independently of any input alignments. If known, such initial alignments can be exploited to accelerate the registration process, but are not required.

- **Avoidance of local minima.** Since we are using energy-minimization optimization techniques on a complex underlying evaluation function, a key issue is reaching the global minimum without getting trapped into local minima. We have incorporated two techniques to treat this problem: interpolation of the evaluation function and random perturbation of resultant transformations.

Acknowledgments

This report describes research supported in part by ARPA under ONR contract N00014-91-J-4038 and by the Medical Informatics training grant T 15 LM 07092 from the National Library of Medicine.

References

[1] Ayache, N., J.D. Boissonnat, L. Cohen, B. Geiger, J. Levy-Vehel, O. Monga, P. Sander, "Steps Toward the Automatic Interpretation of 3-D Images", In *3D Imaging in Medicine*, edited by H. Fuchs, K. Hohne, S. Pizer, NATO ASI Series, Springer-Verlag, 1990, pp. 107–120.

[2] Champleboux, G., S. Lavallee, R. Szeliski, L. Brunie., "From Accurate Range Imaging Sensor Calibration to Accurate Model-Based 3D Object Localization", *IEEE Conference Computer Vision and Pattern Recognition*, 1992, pp. 83–39.

[3] Cline, H.E., W.E. Lorensen, R. Kikinis, F. Jolesz, "Three-Dimensional Segmentation of MR Images of the Head Using Probability and Connectivity", *Journal Computer Assisted Tomography*, **14**(6), 1990, pp. 1037–1045.

[4] Dempster, A.P., N.M. Laird, and D.B. Rubin, "Maximum Likelihood from Incomplete Data via the EM Algorithm", *Journal Royal Statistical Society*, **39**, 1977, pp. 1–38.

[5] Gerig, G., W. Kuoni, R. Kikinis, O. Kübler, "Medical Imaging and Computer Vision: an Integrated Approach for Diagnosis and Planning", *Proceedings 11'th DAGM Symposium*, Hamburg FRG, Springer, 1989, pp. 425–443.

[6] Grimson, W.E.L., T. Lozano-Pérez, W.M. Wells III, G.J. Ettinger, S.J. White, R. Kikinis, "An Automatic Registration Method for Frameless Stereotaxy, Image Guided Surgery, and Enhanced Reality Visualization", *IEEE Conference Computer Vision and Pattern Recognition*, 1994.

[7] Grimson, W.E.L., *Object Recognition by Computer: The Role of Geometric Constraints*, MIT Press, Cambridge, 1990.

[8] Gueziec, A., N. Ayache, "Smoothing and Matching of 3-D Space Curves", *Proceedings Second European Conference on Computer Vision*, May 1992, pp 620–629.

[9] Gueziec, A., "Large Deformable Splines, Crest Lines and Matching", *INRIA TR 1782*, October 1992.

[10] Gueziec, A., N. Ayache, "New Developments on Geometric Hashing for Curve Matching", *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, June 1993, pp 703–704.

[11] Huttenlocher, D., S. Ullman, "Recognizing Solid Objects by Alignment with an Image", *International Journal Computer Vision*, **5**(2), 1992, pp. 195–212.

[12] Jiang, H., R.A. Robb, K.S. Holton, "A New Approach to 3-D Registration of Multimodality Medical Images by Surface Matching", *Visualization in Biomedical Computing–SPIE*, 1992, pp. 196–213.

[13] Kapur, T., "Extracting ICC Surface from an MR Image," AAAI 1994 Spring Symposium Series, *Applications of Computer Vision in Medical Image Processing*, March 1994.

[14] Kikinis, R., M. Shenton, F.A. Jolesz, G. Gerig, J. Martin, M. Anderson, D. Metcalf, C. Guttmann, R.W. McCarley, W. Lorensen, H. Cline, "Routine Quantitative Analysis of Brain and Cerebrospinal Fluid Spaces with MR Imaging", *Journal MRI*, **2**, 1992, pp. 619-629.

[15] Kovacic, S., J.C. Gee, W.S.L. Ching, M. Reivich, R. Bajcsy, "Three-Dimensional Registration of PET and CT Images", *Images of the Twenty-First Century–Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 1989, pp. 548–549.

[16] Lavallee, S. L. Brunie, B. Mazier, P. Cinquin, "Matching of Medical Images for Computed and Robot Assisted Surgery", *Proceedings 13th Annual International Conference Engineering in Medicine and Biology*, **13**(1), 1991, pp. 39–40.

[17] Lavallee, S., R. Szeliski, L. Brunie, "Matching 3D Smooth Surfaces with their 2D Projections using 3D Distance Maps", *SPIE – Geometric Methods in Computer Vision*, 1991, pp. 322–336.

[18] Levin, D.N., X. Hu, K.K. Tan, S. Galhotra, C.A. Pelizzari, G.T.Y. Chen, R.N. Beck, C.T. Chen, M.D. Cooper, J.F. Mullan, J. Hekmatpanah, J.P. Spier, "The Brain: Integrated Three-Dimensional Display of MR and PET Images", *Radiology*, **172**(3), 1989, pp. 783–789.

[19] Metcalf, D., R. Kikinis, C. Guttmann, L. Vaina, F. Jolesz, "4D Connected Component Labelling Applied to Quantitative Analysis of MS Lesion Temporal Development", *Proceedings IEEE EMBS Conference*, October 1992.

[20] Neiw, H.M., C.T. Chen, W.C. Lin, C.A. Pelizzari, "Automated Three-Dimensional Registration of Medical Images", *Medical Imaging V: Image Capture, Formatting, and Display: SPIE*, **1445**, 1991, pp. 259–264.

[21] Pelizzari, C.A., G.T.Y. Chen, D.R. Spelbring, R.R. Weichselbaum, C.T. Chen, "Accurate Three-Dimensional Registration of CT, PET, and/or MR Images of the Brain", *Journal of Computer Assisted Tomography*, **13**(1), 1989, pp. 20–26.

[22] Pelizzari, C.A. K.K. Tan, D.N. Levin, G.T.Y Chen, J. Balter, "Interactive 3D Patient-Image Registration", *Information Processing in Medical Imaging. 12th International Conference, IPMI '91 Proceedings*, 1991, pp. 132–141.

[23] Press, W.H., S.A. Teukolsky, S.T. Vetterling, B.P. Flannery, *Numerical Recipes in C, The Art of Scientific Computing, Second Edition*, Cambridge University Press, 1992.

[24] Sumanaweera, T.S., G.H. Glover, T.O. Binford, J.R. Adler, "MR Susceptibility Misregistration Correction", *IEEE Transactions Medical Imaging*, **12**, 1993, pp. 251–259.

[25] Szeliski, R., S. Lavallee, "Matching 3D Anatomical Surfaces with Non-Rigid Deformations using Octree-Splines", *SPIE Geometric Methods in Computer Vision II*, **2031**, 1993.

[26] Wells III, W. M., *Statistical Object Recognition*, Ph.D. Thesis, MIT, 1993. (MIT AI Lab TR 1398)

[27] Wells III, W.M., W.E.L. Grimson, R. Kikinis, F. Jolesz, "In-Vivo Intensity Correction and Segmentation of Magnetic Resonance Image Data, AAAI 1994 Spring Symposium Series, *Applications of Computer Vision in Medical Image Processing*, March 1994.