

Registration of Cortical Anatomical Structures via Robust 3D Point Matching

Haili Chui, James Rambo, James Duncan, Robert Schultz and Anand Rangarajan

Departments of Diagnostic Radiology, Electrical Engineering
and Yale Child Study Center
Yale University, New Haven, CT 06520

Abstract. Inter-subject non-rigid registration of cortical anatomical structures as seen in MR is a challenging problem. The variability of the sulcal and gyral patterns across patients makes the task of registration especially difficult regardless of whether voxel- or feature-based techniques are used. In this paper, we present an approach to matching sulcal point features interactively extracted by neuroanatomical experts. The robust point matching (RPM) algorithm is used to find the optimal affine transformations for matching sulcal points. A 3D linearly interpolated non-rigid warping is then generated for the original image volume. We present quantitative and visual comparisons between Talairach, mutual information-based volumetric matching and RPM on five subjects' MR images.

1 Introduction

The recent development of brain imaging technologies (PET, MRI, fMRI) has provided rich information on the human brain. A potentially fruitful emerging area of research is human brain mapping [25] which requires a comprehensive statistical analysis of brain structure and function across diverse populations and different imaging modalities. A major requirement in brain mapping is that the imaging data from different subjects and modalities have to be placed in a common reference frame. Recent efforts have focused on using anatomical MR as the basis for such registration.

Inter-subject anatomical registration is a difficult task due to the complexity and variability of brain structures. This is most obvious in the cortical regions. The folding of the cortical surfaces—the sulci and gyri—vary dramatically from person to person and, in some cases [25] are not even always present in each subject. However, the folding pattern is not completely arbitrary. In fact, the sulci often serve as important cortical landmarks. Furthermore, many cortical areas have been associated with critical brain functionalities (vision, language, motor control etc.) with the sulci often representing important functional boundaries. Cortical registration, despite its enormous difficulty, is hence highly desirable as a basis for further statistical quantitative analysis.

Our approach is based on matching feature points representing the sulcal structures. The points were obtained using a tool [17] which allows a neuroanatomy expert to interactively trace sulci on a 3D skull-stripped MRI brain volume. In contrast to just choosing a few landmarks, the tool allows us to represent sulci using hundreds of 3D points. Also, major sulci can be identified and easily labeled.

We then match two sets of labeled sulci (extracted from two subjects' MRI) using a *robust point matching* (RPM) algorithm [20]. The method first determines the best global 3D affine transformation that brings the two sets of sulci into register. Then piecewise affine transformations are solved for each sulcus to further refine the registration. Afterward, a linearly weighted 3D volumetric warping is generated from the piecewise affine mapping.

RPM has been previously developed and used for 2D rigid alignment [20] and 2D affine warping [11]. For the first time, we have developed the technique for 3D affine and piecewise affine warping and applied it to real 3D sulcal features. Embedded within a *deterministic annealing* scheme, RPM allows us to jointly estimate the spatial mapping (affine, piecewise affine) and the point-to-point sulcal correspondences. Moreover, some sulcal structures in one subject may not have corresponding homologies in the other. RPM is able to reject a fraction of such non-homologies as *outliers*. Unlike other methods of point feature registration, RPM returns a one-to-one correspondence between sulcal points. Except for the extraction of the sulci, the whole process is done automatically and the registration and warping of one pair of brains only takes a few minutes.

2 Review

There are two principal approaches to non-rigid brain registration: voxel-based methods and feature-based methods.

Voxel-based approaches try to find the optimal transformation such that a local image intensity similarity measure is maximized. Most methods in this class allow highly complex transformations which are normally proportional to the size of the volume. Elastic media models, viscous fluid models [4] or local smoothness models [6] are introduced as constraints to guide the non-rigid spatial mapping. From these efforts, the need for non-rigid transformations is by now quite clear. Note, however that these algorithms are driven by local voxel intensities. Each voxel is treated equally without taking advantage of higher level geometric information (such as the sulcal and gyral patterns used here). In these methods, further anatomic validation is necessary to ensure that homologous sulci are indeed matched. Aware of this lack, landmarks were used as an initial step in [4]. As [5] also pointed out in their recent work, the voxel intensity approach worked well for deep subcortical structures, but sometimes had difficulty aligning sulci and gyri. To correct this, in their recent work, [5] used a chamfer distance measure [3] on sulcal points and combined it with their former voxel-based matching (ANIMAL) framework. All of these efforts are attempting

to incorporate neuroanatomic geometric feature expertise into their registration engines.

Feature-based methods, as the name implies, capitalize on the information from different identifiable brain structures. Features which represent important brain structures are extracted. The features run the gamut of landmark points [2], curves [22] or surfaces [24, 8]. Subsequently, these methods attempt to solve for the correspondence and transformation between the features. The spatial transformations resulting from feature matching are then propagated to the whole volume. Underlying the philosophy of feature matching is that homologous features always provide an effective anchor for registration. However, enthusiasm for these methods is usually tempered not only by the difficulty of feature extraction but also by the difficulty of simultaneously determining the correspondences or homologies and the spatial mapping. The first problem—feature extraction—usually calls for some residual manual intervention while the second problem—automated matching—involves the computationally demanding task of determining the correspondences and the spatial mapping. As our method basically belongs to this category, we discuss previous methods in some detail and compare them to ours.

Bookstein [2] pioneered the usage of landmark points for registration and shape analysis. The thin-plate-spline is used as the spatial mapping between the two landmark point sets to generate an elastic transformation in which the bending energy is minimized. Since this method basically relies on a few landmark points, the accuracy of their locations is essential. The homologies between all landmark points is deemed known (in advance). In contrast, in our approach, the correspondences and the spatial mapping are co-determined from hundreds of sulcal feature points. In addition, the anatomical variability between subjects can create many outliers, i.e., sulcal points which do not match. Since we are using hundreds of points to represent the structural information, it is statistically much more robust and the noise or point “jitter” which may be caused by various sources such as the tracing process or sampling error, should not significantly affect the final result.

In [24], 3D active surfaces are used to extract the surfaces of lateral ventricle and outer cortex which are developmentally fundamental for the brain. An initial surface is first constructed from some fiducial points and is then relaxed towards the edges until a final balance is reached between the edge attraction force and the surface smoothness measure. To better represent the deep cortical structures (sulci), parametric mesh surfaces are also interactively extracted. A point-to-point mapping between the two *surfaces* is then calculated and a linearly weighted 3D volumetric warping is generated. [8] has a similar framework where surface curvature maps at different scales are used for different brain structures. More consideration is given to the inhomogeneity within the brain. A more sophisticated elasticity model makes the algorithm more flexible at the ventricles and more powerful to account for some abnormal cases where, for example, tumors are involved. Both methods emphasize the importance of sulcal alignment and not surprisingly, the validation in [24] has shown that anatom-

ically homogeneous points can be accurately aligned. As our method is based on matching cortical structures, it is quite similar to both of these approaches. However, we use point-sets as a representation for the sulci rather than surfaces. The major sulci are labeled which imposes strong constraints on the matching. Moreover, the non-rigid matching of 3D surfaces (parameterized by surface normals for example) is a difficult problem. The parameterization of cortical structures as point-sets allows us to easily utilize Procrustes methods of shape analysis [12, 18] (by equating the atlas with the Procrustes mean). Eigenanalysis of the error covariance matrix (around the Procrustes mean atlas) also yields valuable information regarding the dominant modes of deformation present in a population [7].

We have presented a detailed review of competing approaches to solving point correspondence problems elsewhere [20]. Here we briefly discuss chamfer distances [3, 5] and the iterated closest point (ICP) matching algorithm [1]. The chamfer distance has been used in cortical registration by [22] and [5]. The main problem with the chamfer distance is that it uses a brittle nearest neighbor measure to assign correspondence. Nearest neighbor methods used in chamfer matching and ICP are problematic in the vicinity of outliers since they generate local minima [19]. Unlike the chamfer matching in [5] where a distance image is calculated from the Euclidean distance from each voxel to its nearest sulcal point feature, we directly use the sulcal point feature locations for the matching. Finally, we should mention the work presented in [13] where a maximum clique approach is taken to matching relational sulcal representations. Maximum cliques is a very difficult NP-complete problem [10] which in this case increases the likelihood of getting stuck in local minima. Also, it is difficult to explicitly model non-rigid spatial mappings in the maximum clique approach [13]. Consequently, the “engine” that does the work has to be pure sulcal correspondences making the problem more difficult.

3 Robust Sulcal Matching

3.1 Softassign and Deterministic Annealing

There are two important factors that make RPM different from other point matching methods. These two factors mostly account for RPM’s robustness, which proved to be well suited for matching of the complex sulcal patterns.

The first is the *softassign* technique. Let’s suppose we have two point sets $\{X_i, i = 1, 2, \dots, N_1\}$ and $\{Y_j, j = 1, 2, \dots, N_2\}$, where N_1 and N_2 are the numbers of points in each set. ($X_i = (1, X_i^1, X_i^2, X_i^3)^T$. We are using homogeneous coordinates with a 4x4 affine spatial mapping so that the whole transformation could be simply written as AX_i .) The point matching problem is then equivalent to solving the following optimization problem:

$$\min_{M,A} E(M, A) = \min_{M,A} \sum_{i=1}^{N_1} \sum_{j=1}^{N_2} M_{ij} \|X_i - AY_j\|^2 - \alpha \sum_{i=1}^{N_1} \sum_{j=1}^{N_2} M_{ij} \quad (1)$$

subject to: $\sum_{i=1}^{N_1+1} M_{ij} = 1, \forall j \in \{1, \dots, N_2\}, \sum_{j=1}^{N_2+1} M_{ij} = 1, \forall i \in \{1, \dots, N_1\}$, where $M_{ij} \in \{0, 1\}$. A represents the set of transformation parameters we are trying to solve. M is the *binary* correspondence matrix [20, 11] with an extra row and an extra column introduced to account for outliers. The second term in (1) controls the degree of robustness. Greater the value of α , less points are rejected as outliers and vice-versa.

Obviously, the transformation parameters, represented by A , belong to the set of continuous variables; on the other hand, the correspondence matrix M is binary. The softassign technique provides a way to solve the optimization problem with two such variables of different natures. Instead of forcing M_{ij} to be binary, we relax it to be continuous in the interval $[0, 1]$, but with the row and column sum constraints still intact. In addition to being just a numerical technique, it also give us a new way of treating correspondence. Now, one point does not necessarily just correspond to only one other point; it could have multiple memberships with all others with one membership being much larger than the rest. This property is clearly desirable if you have one point in one set lying in between two points in the other set. It does not have to choose immediately which one it belongs to but instead keeps a degree of “fuzziness” while preferring the closest one a little bit more. This also suggests that during the registration process when the transformation is optimized gradually, the correspondence memberships would change continuously and gradually as well without jumping around in the space of permutation matrices (and outliers). In more formal terms, making the correspondences fuzzy smoothes the energy function ridding it of poor local minima [19]. The fuzzy correspondence matrix still has to satisfy the row and column constraints. It turns out that the Sinkhorn balancing procedure of alternating row and column normalizations is an ideal vehicle to satisfy the row and column constraints [20]. The softassign essentially keeps all correspondences positive and then uses Sinkhorn’s theorem to ensure that all rows and columns sum to one (except for the outlier row and column).

Another classic point matching method is the ICP algorithm [1, 9]. ICP uses a nearest neighbor heuristic to set binary correspondences. The algorithm iterates between the spatial mapping and the nearest neighbor correspondences until convergence. As in the chamfer distance [3], the brittleness of the nearest neighbor measure can create local minima in many cases [19]. Some efforts have been made to improve ICP’s robustness by including an adaptive thresholding [9]. Also, there is no guarantee that ICP will return one-to-one correspondences. While correspondence does not have to be a pre-requisite for registration, it does play a more significant role in the creation of probabilistic atlases [15]; the atlas formation step requires averages and covariance matrices to be computed over all the corresponding points in a training set. We expect the one-to-one correspondence returned by RPM to play a significant role in the formation of probabilistic atlases.

Deterministic Annealing [27, 28] is the other important technique used in RPM, which is a good companion to softassign. It is closely related to simulated annealing except that all operations are deterministic. The temperature

parameter T in deterministic annealing specifies the degree of fuzziness of the correspondence matrix—higher the temperature, greater the fuzziness. At each temperature, the initial condition from the previous temperature is used and a straightforward deterministic descent is performed on the energy function. The process is repeated at lower and lower temperatures until M becomes almost binary. The method is more robust than classical gradient methods in that more configurations are allowed at higher temperature, and this makes the energy function smoother and less vulnerable to local minima. At very low temperatures, RPM is very similar to ICP with the added benefit of one-to-one correspondence.

3.2 The Spatial Mapping—3D Affine Transformations

With the above background regarding softassign and deterministic annealing in place, it is reasonably straightforward to develop the method for a 3D affine spatial mapping. The complete form of the energy function is:

$$\begin{aligned}
& \min_{M,A} \max_{\mu \nu} E(A, M) \\
&= \min_{M,A} \max_{\mu \nu} \left\{ \sum_{i,j}^{N_1, N_2} M_{ij} \|X_i - (A + I)Y_j\|^2 + \lambda \text{trace}(A^T A) - \alpha \sum_{i,j}^{N_1, N_2} M_{ij} \right. \\
&\quad + \sum_i^{N_1} \mu_i \left(\sum_j^{N_2+1} M_{ij} - 1 \right) + \sum_j^{N_2} \nu_j \left(\sum_i^{N_1+1} M_{ij} - 1 \right) \\
&\quad \left. + T \sum_{i,j}^{N_1, N_2} M_{ij} (\log M_{ij} - 1) \right\} \tag{2}
\end{aligned}$$

Even though there are six terms, only the first two will be directly involved when we are going to solve for the transformation A (actually $A + I$ where I is the identity transformation). The transformation A is now in 3D. The first term is the error measure. Assume for the moment that the correspondence M is known. The second term is the regularization on A . Basically we are assuming that the affine transformation should be close to identity. The degree of deviation from identity depends on λ . Typically, we begin with a high value of λ and quickly decrease it, with the consideration being that at first the correspondences are still far from the right answer and the transformation should not be too committed. Though this may add some complexity to the algorithm, we have found it worthwhile for two reasons. The first is that the algorithm does not seem to be very sensitive for slightly different choices of λ annealing schedules, i.e., as long as the starting value is high enough so that the transformations are not too large in the beginning and the final value small enough so that the transformations won't always be forced too close to identity. The second reason stems from an observation that because of the extra constraint we put on the transformation, we could then choose not to use that robustness term $-\alpha \sum_{i,j}^{N_1, N_2} M_{ij}$ at all. Actually, in all our experiments α was set to 0.

With M held fixed, the energy function w.r.t. A is:

$$E_{\text{affine}}(A)|_M = \sum_{i,j}^{N_1, N_2} M_{ij} \|X_i - (A + I)Y_j\|^2 + \lambda \text{trace}(A^T A) \quad (3)$$

which is a standard least squares problem for the matrix A . By taking the derivative $\frac{\partial E_{\text{affine}}}{\partial A} = 0$, we can get the closed-form solution for A .

$$A = \left[\sum_{i,j}^{N_1, N_2} M_{ij} (X_i Y_j^T - Y_j Y_j^T) \right] \cdot \left[\sum_{i,j}^{N_1, N_2} M_{ij} Y_j Y_j^T + \lambda I \right]^{-1} = P \cdot Q \quad (4)$$

We will briefly describe the solution for the correspondence mainly for the sake of completion. The fourth and fifth terms are the row and column constraints expressed via Lagrange parameters. The Sinkhorn algorithm within the softassign process will automatically satisfy these constraints so we do not need to explicitly solve for the Lagrange parameters μ_i and ν_j [20]. The sixth term is an entropy term which can also be regarded as a barrier function [14]. Solving for M_{ij} (keeping the Lagrange parameters μ_i and ν_j fixed), we get:

$$M_{ij} = e^{-\frac{\|X_i - (A+I)Y_j\|^2 - \alpha - \mu_i - \nu_j}{T}} \quad (5)$$

Having specified both the spatial mapping in (4) and the correspondences in (5), we summarize the algorithm in the following pseudo-code.

The Robust Point Matching (RPM) Algorithm

Initialize M, T, A, λ

Begin A: Deterministic Annealing. Do A until $T < T_{\text{final}}$

Begin B: Softassign and Relaxation. Do B until M converges or # of iterations $> I_0$

$$Q_{ij} \leftarrow \|X_i - (A + I)Y_j\|^2 - \alpha$$

$$M_{ij} \leftarrow \exp\left(-\frac{Q_{ij}}{T}\right)$$

Begin C: Sinkhorn. Do C until M converges or # of iterations $> I_1$

$$M_{ij} \leftarrow \frac{M_{ij}}{\sum_{j=1}^{N_2+1} M_{ij}} \quad (\text{row normalization})$$

$$M_{ij} \leftarrow \frac{M_{ij}}{\sum_{i=1}^{N_1+1} M_{ij}} \quad (\text{column normalization})$$

End C

$$A \leftarrow \left[\sum_{i,j}^{N_1, N_2} M_{ij} (X_i Y_j^T - Y_j Y_j^T) \right] \cdot \left[\sum_{i,j}^{N_1, N_2} M_{ij} Y_j Y_j^T + \lambda I \right]^{-1}$$

End B

$$T \leftarrow T * T_{\text{anneal-rate}}$$

$$\lambda \leftarrow \lambda * \lambda_{\text{anneal-rate}}$$

End A

3.3 Global/Piecewise Affine Registration and Warping

Given two brains' sulcal point-sets, the registration is done in two steps. The first step finds the global affine transformation to account for the overall translation, orientation, scale and skew. After that, we further allow each sulcus to move locally to refine the alignment by solving for a piecewise affine transformation for each of them. To make sure that the sulcus only does *local* adjustment, the regularization is increased compared to the first step so that only small transformations are allowed.

We then tried to propagate the transformations found for the sulcal points to the whole 3D volume to generate a 3D warping. A weighted linear combination of all the sulci's piecewise affine transformations is calculated based on the shortest distance between a voxel and each sulcus. More specifically, we have a total number of N sulci with each of them (n th) denoted by a set of points, $\{X_l^{(n)}, l = 1, 2, \dots\}$ and a set of affine transformations $A^{(n)}, n = 1, 2, \dots, N$. For the current voxel Y_{ijk} (other than the sulcal points locations, where the transformation is unknown and need to be calculated), the shortest distance to the n th sulcus is found, $d_{ijk}^{(n)} = \min_l \|Y_{ijk} - X_l^{(n)}\|, n = 1, 2, \dots, N$. A set of weights is then defines as:

$$w_{ijk}^{(n)} = \frac{\frac{1}{d_{ijk}^{(n)}}}{\sum_{n=1}^N \frac{1}{d_{ijk}^{(n)}}} \quad (6)$$

and the final voxel transformation is the weighted summation of all $A^{(n)}$:

$$A_{ijk} = \sum_{n=1}^N w_{ijk}^{(n)} \cdot A^{(n)} \quad (7)$$

This is done for each voxel to warp the entire volume.

4 Experiments and Results:

4.1 Interactive Sulcus Tracing

The sulcus tracing was done on an SGI graphics platform [17] with a ray-casting technique that allows drawing in 3D space by projecting 2D coordinates of the tracing onto the exposed cortical surface. A screenshot of the tool is shown on the left in Fig. 1. The inter-hemispheric fissure and 10 other major sulci (superior frontal, central, post-central, Sylvian and superior temporal on both hemispheres) were extracted as point features. A sulcal point-set extracted from one subject is shown on the right in Fig. 1.

4.2 RPM Applied to Sulcal Point Sets

The original sulcal point-sets normally contain around 3,000 points each. The point-set is first sub-sampled to have around 300 points by taking every tenth point. The original MRI volume's size is 106(X) x 75(Y) x 85(Z, slices). With that in mind, it is reasonable to assume that the average distances between points before registration should be in the range of 10 - 100. We set our starting temperature to be roughly in the same scale. After registration, we would expect the average distance between corresponding points to be within a few voxels (say, 1 - 10). Our final temperature should be slightly smaller.




Fig. 1. Left: A screenshot of the sulcus tracing tool with some traced sulci on the 3D MR brain volume. Right: Sulci extracted and displayed as point-sets.

From these considerations, we set the RPM annealing parameters to be the following: $T_{init} = 50$, $T_{final} = 1$, $T_{anneal-rate} = 0.95$. The regularization parameter λ is set to force A to be small at first. We use the value of $\lambda_{init} = \max_{ij}[P_{ij}]$ (P , as defined in (4)) and decrease it by $\lambda_{anneal-rate} = 0.8$ at the end of every temperature iteration. As mentioned above, the idea is that the regularization should prevent the affine to be over determined by the initial fuzzy correspondence at first; once the algorithm starts moving towards the right correspondence, which usually happens within the first few iterations, the regularization should be relaxed by decreasing λ faster than the temperature. Actually, we observed that normally any annealing rate between 0.7 and 0.9 works quite well.

Fig. 2 shows one example of RPM in action. The circles and crosses stand for two sets of sulcal points and the gray links indicate the most significant correspondences ($(M_{ij} > \frac{1}{N_1}$ or $\frac{1}{N_2})$) between the two point-sets at that moment. The first one is taken in the middle of the registration procedure in which clearly you can see that correspondence is still “fuzzy”. The second one is taken towards the end of the process and the correspondence is close to binary so fewer links are seen.

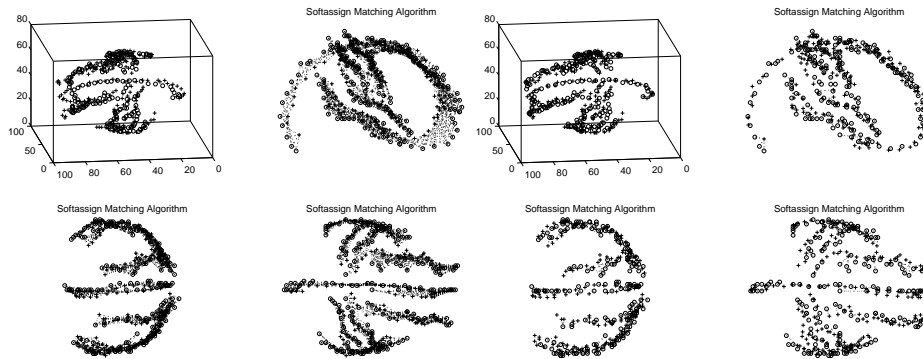


Fig. 2. Demonstration of the robust point matching process. Left four: 3D point sets and their three 2D projections in the middle of the matching. 3D point sets are shown as circles and crosses. Their most significant correspondences are shown as dotted lines. Right four: Towards the end of the matching.

4.3 A Comparison between Talairach, Voxel-based matching and RPM.

We applied RPM to five sulcal point-sets and compared it with two other methods which also use affine (and piecewise-affine) transformations for brain registration.

As we mentioned in the review section, we suspected that the voxel-based methods' performance would not be as satisfying as feature-based methods for sulcal alignment. To test this, we compare RPM with a voxel-based affine matching method [23] which maximizes the mutual information between the two volumes.

By defining a common coordinate system, the Talairach method is a standard technique for brain alignment. A piecewise affine transformation is applied to 12 rectangular regions of brains defined by landmark points of anterior and posterior commissures and extrema of the cortex. We used the Talairach program available as part of the MEDx package (from Sensor Systems Inc.) to align 5 brains. Sulcal points were traced on the resulting brain volumes.

The volumes were then matched by the voxel-based method described in [23] and the resulting spatial mapping was applied to the sulcal points. RPM was separately run on the sulcal point-sets, and both the result from a simple global affine transformation and piecewise affine transformations are shown in Fig. 3. Since we register every brain to the first one, after registration, the minimum distance from each sulcal point in the current brain to the first is calculated. The mean and variance of such minimum distances for each sulcus is calculated for quantitative comparison. The results are shown in Fig. 4.




Fig. 3. Sulcal points alignment. The alignment results of five brains' sulcal point-sets on the left side of the brain are shown together. Denser, closely packed distributions of sulcal points suggest that they are better aligned. We clearly see the improvement of RPM over both Talairach and the voxel-based approach, especially for RPM with the piecewise affine mapping.

The above comparison of Talairach with RPM and voxel-based approaches shows that RPM can significantly improve upon Talairach in most cases even though it may have less degrees of freedom. The voxel-based method's performance is mixed; it gives bigger error for 5 of the 11 sulci. The significant improvement from the global affine transformations by allowing piecewise transformations confirmed the belief of the importance of non-rigid transformations.

4.4 3D Warping and Comparison

The three dimensional warping of the brain volumes is calculated from the transformations found for the sulci as described above. The insufficiency of the Talairach

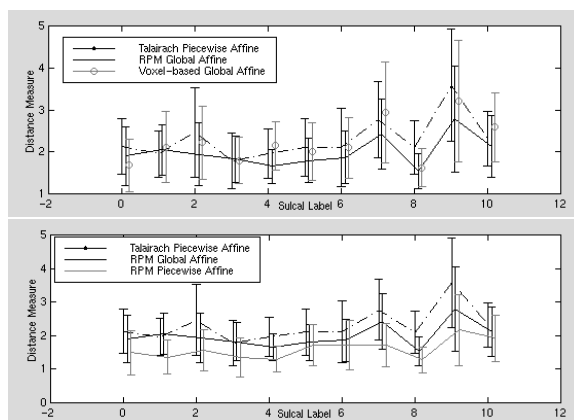


Fig. 4. The minimum distance measure (as described in text) for each sulcus (sulcal label: 0 - interhemispheric fissure; 1,2 - central sulcus; 3,4 - Sylvian fissure; 5,6 - superior temporal sulcus; 7,8 - post-central sulcus; 9,10 - superior frontal sulcus.) Top figure shows the comparison between Talairach (dashed line), voxel-based method (light dotted line) and RPM with a global affine transformation (solid line). Bottom figure shows the first two again with results of RPM with piecewise affine transformation (light line).

alignment for the sulcal structure is clearly seen in Fig. 5. Even though our warping strategy based on piecewise affine transformations is quite simple, the results show further improvement upon the global affine transformation.

Two subjects' brains as well as some of their major sulci are shown in Fig. 4. The variability of the sulci can be appreciated from this figure.




Fig. 5. Two brain volumes after Talairach alignment with their sulci are shown here. Left to right: Both sides of one brain A, both sides of another brain B.

By displaying the reference brain's sulci on the other brain volumes warped using the Talairach technique, global affine transformation from RPM and piecewise affine transformations from RPM, we can see the improvement of sulcal structure alignment by our method. We should also note that the better alignment is accompanied by increased degree of brain deformation. These are shown in Fig. 5, 6 and 7.




Fig. 6. One brain's (A) sulci overlay on the other's (B) volume after Talairach alignment for both side. Clear mismatch can be seen at the temporal lobe area near the Sylvian fissure and the superior temporal sulcus.




Fig. 7. Same except that the volume shown here is deformed by the global affine transformation found by RPM. An improvement over the previous Talairach result can be seen. Locally around the Sylvian fissure area, a small misalignment is seen in the left figure.

5 Discussion and Conclusion

Our simulation and experiments with real data indicate that sulcal point matching is a fast, robust and accurate tool for the registration of cortical anatomical structures. We now mention several enhancements that could further improve our point feature-based non-rigid registration. First, statistical shape models can be computed using the correspondence information returned by RPM. From these models, more meaningful deformation modes based on principal components can be constructed. Also, an arc length-based ordering of the points (akin to curves) can be imposed. This would have the effect of radically reducing the correspondence search. Finally, using a mixture model [21], we can extend the matching algorithm to the problem of matching a labeled sulcal atlas to an unlabeled or partially labeled sulcal point-set. This would allow us to automatically label the sulci extracted from a new brain image. We have reported preliminary work on matching labeled point-sets to unlabeled features (though not sulcal point features) elsewhere [16].




Fig. 8. Same as above except that the volume shown is deformed by the 3D warping based on piecewise affine transformations from RPM. The better global alignment seen from the figure above is maintained and some further local refinement could also be seen (again, notice the area around Sylvian fissure in the left figure). We could also see that the brain's shape has been noticeably deformed.

Since we are using point-sets, which are quite flexible—i.e., it does not matter if those points all lie on a curve, a surface or a more complicated geometrical object—deep cortical structures (for example representations like *sulcal ribbons* [26]) can be easily incorporated into our framework. Future work will focus on hierarchical (labeled and unlabeled) point-set representations of the cortical structures.

6 Acknowledgments

We thank Larry Staib for useful discussions. We thank Colin Studholme for providing us with the mutual information-based voxel matching algorithm and Larry Win for extensive help with neuroanatomy. A.R and H. C are partially supported by a grant from the Whitaker Foundation.

References

1. P. J. Besl and N. D. McKay. A method for registration of 3-D shapes. *IEEE Trans. Patt. Anal. Mach. Intell.*, 14(2):239–256, Feb. 1992.
2. F. L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Trans. Patt. Anal. Mach. Intell.*, 11(6):567–585, June 1989.
3. G. Borgefors. Hierarchical chamfer matching: a parametric edge matching algorithm. *IEEE Trans. Patt. Anal. Mach. Intell.*, 10(6):849–865, 1988.
4. G. Christensen, S. Joshi, and M. Miller. Volumetric transformation of brain anatomy. *IEEE Trans. Med. Imag.*, 16(6):864–877, 1997.
5. D. Collins, G. Goualher, and A. Evans. Non-linear cerebral registration with sulcal constraints. In W. Wells, A. Colchester, and S. Delp, editors, *Medical Image Computing and Computer-Assisted Intervention*, volume 1496 of *Lecture Notes in Computer Science*, pages 974–984. Springer, 1998.
6. D. Collins, C. Holmes, T. Peters, and A. Evans. Automatic 3D model-based neuro-anatomical segmentation. *Human Brain Mapping*, 3(3):190–208, 1995.
7. T. Cootes, C. Taylor, D. Cooper, and J. Graham. Active shape models: Their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.
8. C. Davatzikos. Spatial transformation and registration of brain images using elastically deformable models. *Computer Vision and Image Understanding: Special Issue on Medical Imaging*, 6(2):207–222, 1997.
9. J. Feldmar and N. Ayache. Rigid, affine and locally affine registration of free-form surfaces. *Intl. J. Computer Vision*, 18(2):99–119, May 1996.
10. M. R. Garey and D. S. Johnson. *Computers and intractability: a guide to the theory of NP-completeness*. W. H. Freeman, San Francisco, CA, 1979.
11. S. Gold, A. Rangarajan, C. P. Lu, S. Pappu, and E. Mjolsness. New algorithms for 2-D and 3-D point matching: pose estimation and correspondence. *Pattern Recognition*, 31(8):1019–1031, 1998.
12. C. Goodall. Procrustes methods in the statistical analysis of shape. *J. R. Statist. Soc. B*, 53(2):285–339, 1991.
13. G. Lohmann and D. von Cramon. Sulcal basin and sulcal strings as new concepts for describing the human cortical topography. In *Workshop on Biomedical Image Analysis*, pages 41–54. IEEE Press, June 1998.

14. D. Luenberger. *Linear and Nonlinear Programming*. Addison-Wesley, Reading, MA, 1984.
15. J. Mazziotta, A. Toga, A. Evans, P. Fox, and J. Lancaster. A probabilistic atlas of the human brain: theory and rationale for its development. *NeuroImage*, 2(2):89–101, 1995.
16. S. Pappu, S. Gold, and A. Rangarajan. A framework for non-rigid matching and correspondence. In D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, editors, *Advances in Neural Information Processing Systems 8*, pages 795–801. MIT Press, Cambridge, MA, 1996.
17. J. Rambo, X. Zeng, R. Schultz, L. Win, L. Staib, and J. Duncan. Platform for visualization and measurement of gray matter volume and surface area within discrete cortical regions from MR images. *NeuroImage*, 7(4):795, 1998.
18. A. Rangarajan, H. Chui, and F. Bookstein. The softassign Procrustes matching algorithm. In *Information Processing in Medical Imaging (IPMI '97)*, pages 29–42. Springer, 1997.
19. A. Rangarajan, H. Chui, and J. Duncan. Rigid point feature registration using mutual information. *Medical Image Analysis*, 1999. (accepted).
20. A. Rangarajan, H. Chui, E. Mjolsness, S. Pappu, L. Davachi, P. Goldman-Rakic, and J. Duncan. A robust point matching algorithm for autoradiograph alignment. *Medical Image Analysis*, 4(1):379–398, 1997.
21. R. A. Redner and H. F. Walker. Mixture densities, maximum likelihood and the EM algorithm. *SIAM Review*, 26(2):195–239, 1984.
22. S. Sandor and R. Leahy. Surface based labeling of cortical anatomy using a deformable atlas. *IEEE Trans. Med. Imag.*, 16(1):41–54, 1997.
23. C. Studholme, D. Hill, and D. Hawkes. An overlap invariant entropy measure of 3d medical image alignment. *Pattern Recognition*, 32:71–86, 1999.
24. P. Thompson and A. W. Toga. A surface-based technique for warping three-dimensional images of the brain. *IEEE Trans. Med. Imag.*, 5(4):402–417, August 1996.
25. A. Toga and J. Mazziotta. *Brain Mapping: The Methods*. Academic Press, 1996.
26. M. Vaillant, C. Davatzikos, and R. Bryan. Finding 3D parametric representations of the deep cortical folds. In A. Amini, F. L. Bookstein, and D. Wilson, editors, *Proc. of the Workshop on Mathematical Methods in Biomedical Image Analysis*, pages 151–159. IEEE Computer Society Press, 1996.
27. A. L. Yuille. Generalized deformable models, statistical physics, and matching problems. *Neural Computation*, 2(1):1–24, 1990.
28. A. L. Yuille and J. J. Kosowsky. Statistical physics algorithms that converge. *Neural Computation*, 6(3):341–356, May 1994.