

Learning an atlas from unlabeled point-sets

Haili Chui

Medical Imaging Group
R2 Tech.
Los Altos, CA

Anand Rangarajan

Department of CISE
Univ. of Florida
Gainesville, FL

Abstract

One of the key challenges in deformable shape modeling is the problem of estimating a meaningful average or mean shape from a set of unlabeled shapes. We present a new joint clustering and matching algorithm that is capable of computing such a mean shape from multiple shape samples which are represented by unlabeled point-sets. An iterative bootstrap process is used wherein multiple shape sample point-sets are non-rigidly deformed to the emerging mean shape, with subsequent estimation of the mean shape based on these non-rigid alignments. The process is entirely symmetric with no bias toward any of the original shape sample point-sets. We believe that this method can be especially useful for creating atlases of various shapes present in medical images. We have applied the method to create a mean shape from nine hand-segmented 2D corpus callosum data sets.

1 Introduction

The study of deformable shapes has recently been a very active area in medical imaging. Most of the effort has been primarily focused on understanding deformable shapes using a statistical approach. Analysis of deformable shapes has in turn helped create automated segmentation tools. A statistical understanding of shapes, their representations and deformations is therefore vitally important in many segmentation tasks. Statistical shape analysis using active shape models [4] is a representative example in this category. After the shapes of a certain single structure but from multiple subjects are extracted, much information can be gained by identifying, measuring and characterizing the shapes. Through such analysis, it is hoped that the subtle differences between different groups or populations, e.g. a normal group vs a diseased group, can be discovered and made available to aid clinical diagnosis. Recent work in brain image analysis [5, 10] can be seen as examples in this category.

1.1 Shape representations

Deformable shapes can have different representations. Curves and surfaces of the shape boundary can obviously be used [8, 9]. Despite being an intuitively natural choice for shape representation, curves and surfaces are somewhat difficult to directly use in statistical analysis. Using the intrinsic parameterization is difficult when the goal is atlas creation. Using one of many extrinsic parameterizations typically leads to a mixed representation defined on point locations and spline coefficients. Statistical shape analysis in this representation requires learning a density function defined on locations *and* coefficients.

Other than curves or surfaces, another natural choice is to use point location information. A point-set representation has the main advantage that statistical shape analysis can be well formulated [4] in that space. Not surprisingly, many recent statistical methodologies for shape analysis [1, 4, 6] use the point-set representation. Our work also focuses on using points to study deformable shapes.

1.2 Statistical shape analysis with unknown correspondence

Given a set of deformable shapes represented by point-sets, basic statistical measures such as the mean and the covariance enable us to quantify the set of shapes at the very least up to second-order statistics. The mean point-set is usually a placeholder for a normal representative shape. The covariance information, usually of the form of a high dimensional covariance matrix, further provides us with information about how the shapes can deform and vary from the mean shape (in a second-order statistical sense). By examining the covariance matrix's eigen-vectors and eigen-values, we can also observe the dominant variations or deformations present in the shape group [4]. If higher-order statistical information is desired, recent techniques such as Independent Component Analysis (ICA) can be pressed into service.

The primary technical challenge in using point-set repre-

representations of shapes is the *correspondence* problem. Computing a meaningful mean point-set (and then the covariance matrix) from multiple shape point-sets is only possible when the correspondences between all the shape point-sets are known. Automated correspondence estimation in shape point-sets is a non trivial task because of several reasons.

Typically, correspondences can be estimated once the point-sets are properly aligned with appropriate spatial transformations. An adequate transformation would obviously be a deformation, i.e., a non-rigid spatial mapping, since the objects at hand are deformable. Solving for non-rigid deformations between point-sets with unknown correspondence is a hard problem [3]. In fact, many current methods only attempt to solve for rigid transformations, e.g. affine transformations, for the alignment [6]. The correspondences resulting from these methods are, therefore, only rough approximations since they ignore the non-rigid aspect of the transformation. However, when dealing with point-sets, there is another aspect of the atlas problem, which we believe is more fundamental, and has largely been neglected (in the point-set literature).

1.3 Joint clustering and matching

The problem lies in the fact that the shape sample point-sets may not be *consistent* in the first place, i.e., points in each set may not be positioned at corresponding locations. Due to the vagaries of feature extraction, points in one set may not have exactly counterparts in another point-set. Without taking into account this feature extraction and sampling issue, the resulting correspondence achieved by any alignment method (rigid or non-rigid) is doomed from the start.

We face a dilemma here. On the one hand, we need to compare the point-sets, possibly through alignment, to each other in order establish correspondence. On the other hand, the alignment won't be entirely successful unless the sample point-sets have already satisfied the consistency criterion. The dilemma clearly shows that any process of using fixed sample point-sets to do the alignment is flawed. The problem is graphically depicted in Figure 1.

Finally, we also encounter the bias problem in atlas creation. Since we have more than two sample point-sets to align, a question that arises is: how do we align all the point-sets in a symmetric manner so that there is no bias toward any particular point-set? One way to circumvent this problem is to define a mean point-set and then align every sample point-set to the mean. We see this as an opportunity for mutual improvement instead of a fixed unidirectional process. As explained later, our approach consists of mutual refinement of the mean shape and the non-rigid deformations from the point-sets to the emerging mean.

After surveying the fundamental problems in atlas creation, we summarize our approach. A joint clustering and




Figure 1: The consistency problem. From left to right: i) the original point-set with densely distributed points; ii, iii) two subsampled point-sets; iv) the overlay of the subsampled point-sets. Note the difference between the two subsampled point-sets in the overlay.

matching approach is used to overcome the consistency problem. The matching is performed on *movable* cluster centers and not on the original points. The cluster centers are also used to estimate and re-estimate the atlas or mean shape through non-rigid matching and averaging. The estimated variables comprise (i) the cluster centers, (ii) the atlas (mean shape), (iii) the non-rigid transformations from (to) the cluster centers to (from) the atlas. We have designed a joint clustering and matching algorithm to carry out this iterative process. Embedded within a *deterministic annealing* scheme, the algorithm gradually refines all three sets of variables until convergence.

2 Review

After discussing the underlying motivations and rationale behind our method, we take a step back and briefly summarize *some* representative previous research work in this area. We also focus on the differences between these methods and ours.

The work presented in [8] is a representative method using an intrinsic curve parameterization to analyze deformable shapes. It is very close to previous efforts in [9] and [5] where intrinsic curve properties, like the curvature and the arc length, are used to align curves. In [8], the alignment is done in a pairwise manner. To somewhat reduce the bias introduced by this non-symmetrical procedure, the sample curves take turns in being the reference curve (atlas). The authors considered the symmetric alignment of multiple sample shapes to an emerging mean shape to be “intractable.” We are going to argue the opposite view in this paper. As with other methods using intrinsic curve or surface representations, further statistical analysis on these representations are much more difficult than the point representation *but the payoff may be greater* due to the use of intrinsic higher-order representations.

The active shape model proposed in [4] utilized points to represent deformable shapes. Their work successfully

initiated the research efforts in building point distribution models to understand deformable shapes [4, 11]. To be able to build such a model, a set of training data, which are a group of shape samples with known correspondences, are required. The training data are normally acquired through a more or less manual landmarking process where a person goes through all the samples and attempts, to the best of his/her knowledge, to mark corresponding points on each sample. It is a rather tedious process and the accuracy is limited. Since the derived point shape distribution is only going to be as good as the training data, it is highly desirable to automate the training data generation process and to make it more accurate. Different methods have been proposed to attack this problem. Rigid alignment [6] can be performed to get a rough estimate of the correspondences.

Bookstein pioneered the usage of non-rigid spatial mapping functions - specifically *thin-plate* splines, to analyze deformable shapes [1]. Since this method is landmark-based, it avoids the correspondence problem but suffers from the typical problems besetting landmark methods. Note that landmark-based methods largely overcome the consistency problem since the placement of corresponding points is driven by the visual perception of experts and not by an automated algorithm.

The work in [6] also uses 2D points to learn shape statistics. The method is quite similar to the active shape model method [4] except that more attention has been paid to the training data (or shape sample point-sets) generation process. The shape samples are first represented as curves and aligned with a rigid transformation. One curve is chosen as the reference and points are uniformly placed on that curve to form a reference point-set. To determine the location of corresponding points on the other curves, the extrinsic curvature information is compared. Points are allowed to slide along the curves until their curvatures are in better agreement with each other. This shape learning method improves the consistency of the shape samples under the assumption that the curvature information can be reliably computed for the shapes at hand. Since the extrinsic curvature is a rigid invariant, only a rigid mapping can be used. The process is not symmetric.

The correspondence problem, the shape sample consistency problem and the mean shape estimation problem are all inter-related to each other. Rather than treat each of them as a separate problem, we propose to regard them as three interlocking steps in a more general framework wherein we can simultaneously achieve meaningful answers for all three problems. We also believe that directly solving for non-rigid transformations, rather than using indirect information criteria or relying on curvature information, has many advantages in terms of efficiency and accuracy. By directly modeling the deformation, the results will be intuitively easier to evaluate. Using non-rigid transformations,

rather than rigid transformations is also one of the key aspects in improving the estimation of correspondence.

3 Methodology

In this section, we discuss the details of the proposed joint clustering and matching method.

The three inter-locking steps in our atlas estimation framework are (i) clustering, (ii) non-rigid mapping estimation and (iii) atlas (mean shape) estimation. The clustering step allows the cluster centers to be repositioned (to "slide" along the original points) so that they have a better chance to be consistent. The overall Euclidean distance between the cluster centers and the original point-set is the criterion for shape representation. The cluster center sets are associated with the emerging mean shape using non-rigid deformations. Since there is only one mean, this also provides an indirect connection between the sample cluster center sets. A spline deformation energy is the criterion for shape matching. With consistent sample cluster center sets positioned at corresponding locations, and with multiple non-rigid transformations accounting for the shape differences between the samples and the emerging mean, we apply the transformations and recompute the atlas (mean shape) by averaging over the warped cluster center sets. The average Euclidean distance of all warped cluster center sets to the mean shape is the criterion for atlas estimation.

We begin with a total of P shape samples represented by P point-sets, $X^1, X^2, \dots, X^p, \dots, X^P$. Each point-set X^p consists of N^p points $\{x_i^p, i = 1, 2, \dots, N^p\}$. Ultimately, we would like to compute a meaningful mean shape point-set Z given all the sample sets $\{X^p\}$. Note that the number of points in each sample set N^p may be different. Furthermore, the sample points may not be at the corresponding locations in each shape sample. To improve the consistency between the shape sample point-sets, we would like to refine their points' locations. For this purpose, we introduce a new variable V^p , for each sample point-set X^p , to represent such refined adjusted locations. We call the new set V^p as the "cluster center" set, which consists of K cluster centers $\{v_a^p, a = 1, 2, \dots, K\}$. These cluster centers are estimated by running a clustering process over the original sample points $\{x_i^p\}$. Since the cluster centers are being actively estimated, their locations can be constantly adjusted, while they would still always provide a faithful representation of the original shape sample.

A good way to understand this process is through the following example: if the shape sample of interest looks like a curved structure, the sample points would then all fall on this "imaginary curve". After clustering, the cluster centers from the sample points will also be constrained to roughly lie on the imaginary curve because otherwise they won't be a good shape representation of the original sample points.

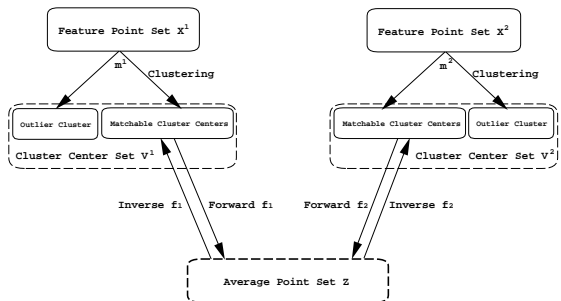


Figure 2: The super clustering-matching algorithm. Each original point-set (X^1 and X^2) is being clustered down to a set of cluster centers ($X^1 \rightarrow V^1$ and $X^2 \rightarrow V^2$). Each cluster set has an outlier cluster (v_{K+1}^1 in V^1 , v_{K+1}^2 in V^2) to account for the possible spurious points in each point-set. The rest of the cluster centers ($\{v_a^1, a = 1, 2, \dots, K\}$ and $\{v_a^2, a = 1, 2, \dots, K\}$) are being matched to the average point-set Z .

So the curve-like shape structure will always be maintained. However, this doesn't prevent the cluster centers to have the extra freedom of being able to "slide" along the "imaginary curve" so that they can be better positioned to satisfy the consistency criteria.

The cluster centers are indirectly computed by first estimating an intermediate variable set $\{M^p\}$ called the *membership* matrix. This is determined between each of the original sample point-set X^p and each of the cluster center set V^p . The entries of M^p are set up such that, when a sample point x_i^p is considered to contribute to a cluster center point v_a^p , m_{ai}^p is closer to one otherwise it is closer to zero. Once M^p is known, computing each cluster center v_a^p can be as simple as averaging over x_i^p with weighting factors from the m_{ai}^p .

The newly estimated cluster center sets all have K clusters. To ensure that they are consistent with each other, we perform the clustering process along with a matching process. The matching process tries to align each and every cluster center set V^p with the emerging mean set Z by solving for a non-rigid transformation f^p between them. The mean point set Z then provides the link between different cluster center point sets to achieve the consistency.

Armed with the knowledge of adjusted shape representations via the cluster center sets $\{V^p\}$, and the non-rigid transformations $\{f^p\}$ for the alignment, we can then easily compute the mean point-set by averaging over the cluster center sets after the non-rigid transformations are applied to each one.

The whole setup is demonstrated in Figure 2. For simplicity, it is illustrated with only two sample point-sets ($P = 2$). Extending the setup to incorporate multiple point-sets is obvious since there is no bias toward either point-set

in Figure 2.

We now specify the overall energy function for joint clustering and matching. We intend to minimize an energy function E with 3 grouped sets of variables: the group consisting of $\{M^p\}$ and $\{V^p\}$, the set of transformations $\{f^p\}$ and the mean set Z .

$$E(Z, \{V^p, M^p, f^p, g^p\}) = \sum_p E^p(Z, V^p, M^p, f^p, g^p), \quad (1)$$

where

$$\begin{aligned} E^p(Z, V^p, M^p, f^p, g^p) = & \sum_{i=1}^{N_p} \sum_{a=1}^{K+1} m_{ai}^p \|x_i^p - v_a^p\|^2 \\ & + \sum_{a=1}^K \|z_a - f^p(v_a^p)\|^2 + \sum_{a=1}^K \|g^p(z_a) - v_a^p\|^2 \\ & + \lambda T \|Lf^p\|^2 + \lambda T \|Lg^p\|^2 + T \sum_{i=1}^{N_p} \sum_{a=1}^K m_{ai}^p \log m_{ai}^p \quad (2) \end{aligned}$$

where the membership matrix entries $m_{ai}^p \in [0, 1]$ and satisfy the constraint $\sum_{a=1}^{K+1} m_{ai}^p = 1$. We briefly go over the individual terms in this energy function and explain some of the new variables. The first term $\sum_{ai} m_{ai}^p \|x_i^p - v_a^p\|^2$ measures the average distance between the sample points and the cluster centers. Minimization of this term will essentially force all the cluster centers to be as close as possible to the original sample points, thus maintaining the original shape. The second and third terms $\sum_{ap} \|z_a - f^p(v_a^p)\|^2$, $\sum_{ap} \|g^p(z_a) - v_a^p\|^2$, model the deformations (both the forward f^p , which warps the cluster center sets V^p to Z , and the reverse g^p , which does the opposite) to align the cluster centers with the mean. We think it is possible to force the deformation to be consistent by requiring g^p to be $(f^p)^{-1}$ as in [2] but have not done so here. The fourth term $\|Lf^p\|^2 + \|Lg^p\|^2$ measures the amount of distortion or bending introduced to the space by the non-rigid warps f^p and g^p . By penalizing this bending measure, the algorithm will effectively place the mean shape points to be somewhat similar to all shape samples, since otherwise it will require a larger amount of bending to warp any of the shape samples onto the mean. Since all cluster center sets share a common mean point set Z , consistency between the cluster centers is also achieved. The fifth term $T \sum_{i=1}^{N_p} \sum_{a=1}^K m_{ai}^p \log m_{ai}^p$ arises from deterministic annealing [3]. The parameter T , termed the temperature, is used to control the fuzziness of the clustering membership matrices $\{M^p\}$: higher the temperature, greater the fuzziness. The $x \log x$ form of the barrier function effectively leads to the formation of Gaussian clusters. The temperature T can be interpreted as a common gaussian mixture variance parameter [3, 12]. However, it is manually controlled as opposed to being automatically adjusted. The fuzziness of the membership matrices is gradu-

ally reduced by slowly annealing T in a pre-defined linear scheme. Note that the effective regularization parameter λT is a linear function of T . This allows us to focus on first estimating the rigid transformation parameters and later on the non-rigid transformations.

Having specified this objective function with three groups of unknown variables, we conduct a grouped coordinate descent procedure to achieve the energy minimization. Essentially, an alternating algorithm is used to cycle between the updates of each of the three groups of variables.

The update steps can be summarized as the following:

(i) Update the cluster center sets: We first estimate the membership matrices $\{M^p\}$

$$m_{ai}^x = \frac{q_{ai}^x}{\sum_{a=1}^{K+1} q_{ai}^x}, \forall a, i \quad (3)$$

where

$$q_{ai}^x = e^{-\frac{\|x_i - v_a\|^2}{T}}, \text{ for } a \in \{1, \dots, K\}, \text{ and } i \in \{1, \dots, N\}. \quad (4)$$

Then we compute the cluster centers,

$$v_a^p = \frac{1}{2} \left(\frac{\sum_{i=1}^{N_p} m_{ai}^p x_i^p}{\sum_{i=1}^{N_p} m_{ai}^p} + g^p(z_a) \right) \text{ for } \forall a, p. \quad (5)$$

Note that the cluster centers are determined both by the original sample points $\{x_i^p\}$ as well as the common mean Z . As we mentioned before, the cluster center sets are linked to each other through the mean set Z which enables them to achieve consistency. The update equation for $\{M^p\}$ is exact whereas the update equation for $\{V^p\}$ is *approximate*. We have neglected the term $\sum_{a=1}^K \|z_a - f^p(v_a^p)\|^2$ in computing $\{V^p\}$. This is because we believe that the ‘‘inverse’’ term $\sum_{a=1}^K \|g^p(z_a) - v_a^p\|^2$ is a sufficient constraint for estimation. A closed-form solution for $\{V^p\}$ cannot be obtained if the term $\sum_{a=1}^K \|z_a - f^p(v_a^p)\|^2$ is taken into account.

(ii) Update the atlas (mean shape) point-set:

$$z_a = \frac{1}{P} \sum_{p=1}^P f^p(v_a^p) \text{ for } a = 1, 2, \dots, K. \quad (6)$$

This update equation for Z is approximate since we have neglected the term $\sum_{a=1}^K \|g^p(z_a) - v_a^p\|^2$ in computing the mean Z . To get an initial estimate of Z , we initialize the algorithm by setting $\{f^p, g^p\}$ to be identity transformations. This is because we believe that the forward term $\sum_{a=1}^K \|z_a - f^p(v_a^p)\|^2$ is a sufficient constraint for estimation. A closed-form solution for Z cannot be obtained if the term $\sum_{a=1}^K \|g^p(z_a) - v_a^p\|^2$ is taken into account. To get an initial estimate of Z , we initialize the algorithm by setting $\{f^p, g^p\}$ to be identity transformations.

(iii) Update the non-rigid transformations:

$$f^p = \arg \min_{f^p} \left(\sum_{a=1}^K \|z_a - f^p(v_a)\|^2 + \lambda T \|L f^p\|^2 \right) \forall a \quad (7)$$

with a similar update for the ‘‘inverse’’ transformation g^p . Since this is equivalent to a landmark-based approach, we can solve for f^p and g^p in closed form. Our approach can therefore be viewed as an automated landmarking process. If diffeomorphisms are required for f^p , a single closed form solution is usually not possible. Instead, we would have to introduce a ‘‘time’’ parameter and construct an ODE to take each V^p onto Z as in [7]. The pseudo-code of the algorithm is as follows.

The Super Clustering-Matching Algorithm:

Initialize parameters $T = T_0$ and λ .

Initialize all membership matrices $\{M^p\}$ (e.g., use uniform matrix).

Initialize all transformation matrices $\{f_p\}$ (e.g., use identity transformation).

Initialize the average point-set Z (e.g., use the center of mass of all point-sets).

Begin A: Deterministic Annealing.

• **Begin B: Alternating Update.**

Step 1: Update cluster centers $\{V^p\}$ based on current $\{M^p\}$, $\{g^p\}$ and Z .

Step 2: Update average point-set Z based on cluster centers $\{V^p\}$ and $\{f_p\}$.

Step 3a: Update transformations $\{f^p, g^p\}$ based on current V^p and Z .

Step 3b: Update cluster memberships $\{M^p\}$ based on current $\{V^p\}$.

End B

Decrease $T \leftarrow T \cdot r$ until T_{final} is reached.

End A

Instead of having the flavor of essentially a gradient descent algorithm, which can be plagued by many local minima, our algorithm utilizes deterministic annealing (DA) with its proven benefit of robustness especially in clustering problems [12]. The working mechanism of deterministic annealing has been discussed extensively in many other papers [12, 3]. We would like to point out here that utilizing annealing enables us to compute the cluster centers in a coarse to fine fashion, and hence also the computation of the mean. The annealing is also applied on the regularization of the non-rigid transformation (we use λT instead of a constant λ). The annealing technique is the key to solving this difficult deformable shape averaging problem.

4 Results

After discussing the details of our algorithm and its implementation, we present a simple but yet demonstrative example of the joint clustering and matching algorithm for 2D atlas estimation. While our framework is obviously valid in 3D as well as in 2D, our initial results are for 2D shapes.

The structure that we are interested in is the corpus callosum as it appears in MR brain images. The corpus callosum structure normally has an elongated shape. Constructing an atlas for the corpus callosum and subsequently analyzing the individual shape variation from "normal" anatomy has been regarded as potentially valuable for the study of brain diseases such as agenesis of the corpus callosum (ACC), and fetal alcohol syndrome (FAS).

We manually extracted points on the outer contour of the corpus callosum from nine normal subjects (as shown in Figure 3). The resulting atlas (mean point-set) is shown in Figure 6. The cluster center sets, which are our newly refined shape samples, are shown in Figure 4. The correspondences and the transformations between the cluster centers and the mean are also shown in Figure 5. As we described earlier, all these results are computed simultaneously and automatically. This example clearly demonstrates that our joint clustering and matching algorithm can simultaneously align multiple shape sample point-sets and compute a meaningful atlas (mean point-set) while automatically refining the sample points' location to achieve better consistency.

5 Conclusion

There are two important parameters in our cost function which need to be estimated; the regularization parameter and the number of clusters. The annealing schedule, initial temperature and final temperature are not as critical since they can be readily set as in any other scale-space strategy. We wish to point out an interesting relationship between these two unknown parameters and a connection to entropy-based approaches. If the number of cluster centers is increased, then it is possible to better approximate the shape of a point-set. However, if the number of clusters exceeds a limit, its effectiveness decreases since the inherent feature "jitter" increases (resulting in increased cluster center variance). Increased cluster center variance negatively impacts on the deformation. If the regularization parameter is too high, then the deformation cannot compensate for the cluster center "jitter" resulting in a bad spline fit. If the regularization parameter is too low, then the spline happily follows the cluster center jitter and overfits the deformation. Obviously there is going to be a tradeoff regime for both parameters. A reasonable criterion for parameter estimation is the minimization of the overlap entropy of the original point-sets. It seems fairly straightforward to compute the overlap entropy (for different settings of the two parameters) and use the minimum entropy criterion to pick adequate values of the two parameters.

We have presented anecdotal evidence demonstrating that it is indeed possible to simultaneously estimate landmarks, deformations and an atlas from unlabeled point-sets. And we have argued, we hope convincingly, that correspon-

dence, sample consistency, deformations and atlases belong together and should be simultaneously estimated in order to avoid bias, landmark inconsistency and a host of other problems. Obviously a good validation effort is required in order to show the effectiveness of our approach.

Acknowledgments

This work is supported by NSF IIS 9906081.

References

- [1] 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.
- [2] G. Christensen. Consistent linear-elastic transformations for image matching. In *Proceedings of Information Processing in Medical Imaging—IPMI 99*, pages 224–237. Springer-Verlag, 1999.
- [3] H. Chui and A. Rangarajan. A new feature registration framework using mixture models. In *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis—MMBIA 2000*. IEEE Press, 2000.
- [4] 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.
- [5] 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.
- [6] N. Duta, A. K. Jain, and M.P. Dubuisson-Jolly. Learning 2D shape models. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 8–14, 1999.
- [7] S. Joshi and M. Miller. Landmark matching via large deformation diffeomorphisms. *IEEE Trans. Image Processing*, 9:1357–1370, 2000.
- [8] T. B. Sebastian, J. J. Crisco, P. N. Klein, and B. B. Kimia. Construction of 2D curve atlases. In *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis—MMBIA 2000*, pages 70–77, June 2000.
- [9] H. Tagare. Shape-based nonrigid correspondence with application to heart motion analysis. *IEEE Transactions on Medical Imaging*, 18(7):570–579, July 1999.
- [10] P. Thompson, D. MacDonald, M.S. Mega, C.J. Holmes, C.J. Evans, and A.W. Toga. Detection and mapping of abnormal brain structure with a probabilistic atlas of cortical surfaces. *Journal of Computer Assisted Tomography*, 21(4):567–581, 1997.
- [11] Y. Wang and L. H. Staib. Boundary finding with prior shape and smoothness models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(7):738–743, July 2000.
- [12] A. L. Yuille, P. Stolorz, and J. Utans. Statistical physics, mixtures of distributions, and the EM algorithm. *Neural Computation*, 6(2):334–340, March 1994.

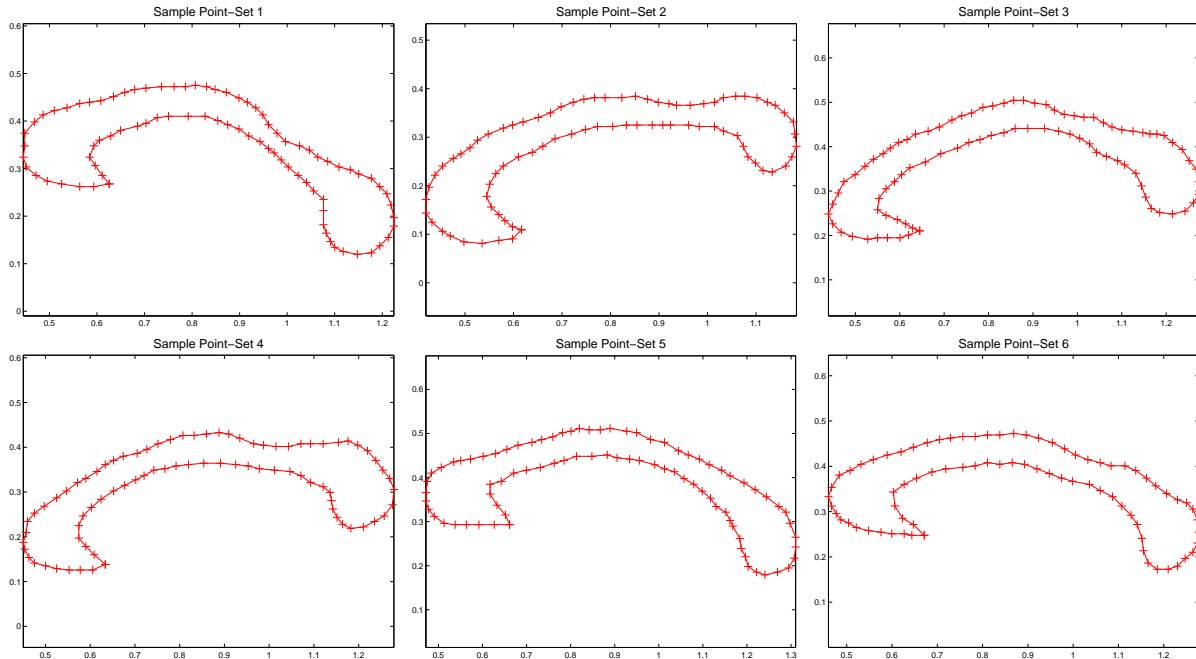


Figure 3: Six of the nine corpus callosum 2D shapes are shown. Care was taken to approximately extract 2D contours at the same orientation.




Figure 4: The estimated cluster centers for all nine corpus callosum shapes are shown. Note that the cluster centers for the most part lie on the shape outline. The cluster centers are in correspondence as well because they share a common index.

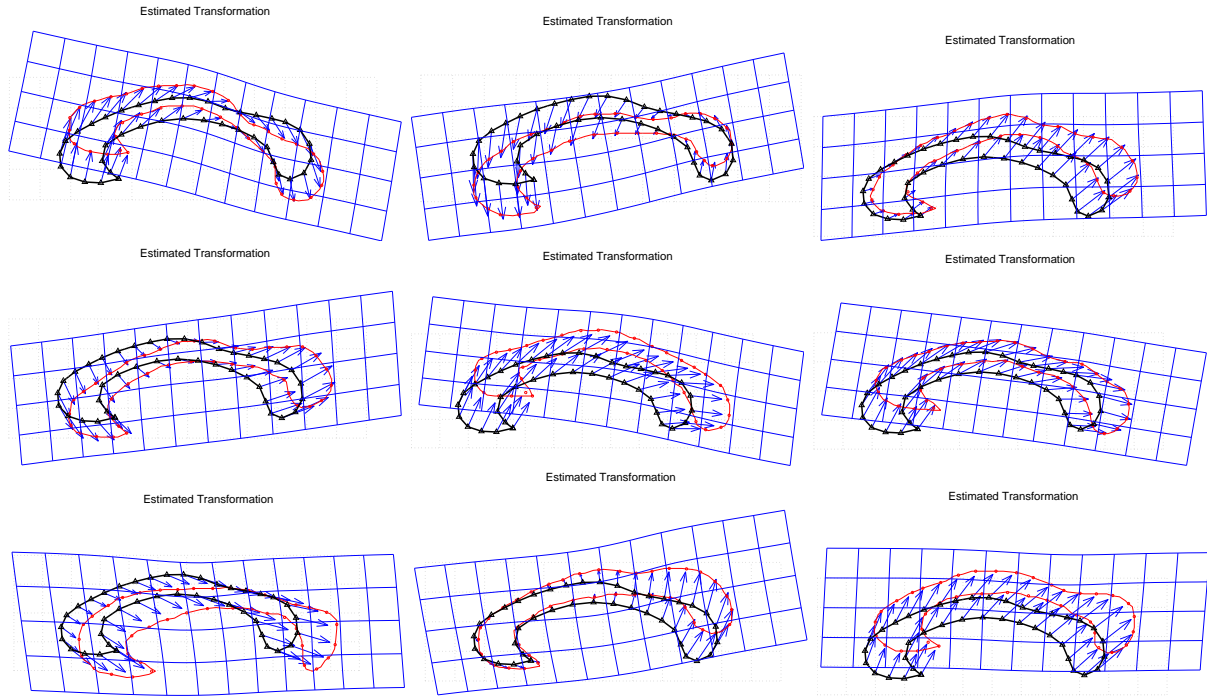


Figure 5: The deformations of each set of clusters onto the atlas are shown.




Figure 6: Left: The estimated atlas is shown superimposed over all the point-sets. Right: An atlas contour is traced and shown superimposed over all the original contours.