

Multi-Object Analysis of Volume, Pose, and Shape using Statistical Discrimination

Kevin Gorcowski, Martin Styner, Ja-Yeon Jeong, J. S. Marron, Joseph Piven, Heather Cody Hazlett, Stephen M. Pizer, Guido Gerig

Abstract—One goal of statistical shape analysis is the discrimination between two populations of objects. In this paper, we present results of discriminant analysis on multi-object sets, a problem not yet sufficiently explored. Our choice of discriminant method is the distance weighted discriminant (DWD) because of its generalization ability in high dimensional, low sample size settings. Using an unbiased, soft discrimination score we can associate a statistical hypothesis test with the discrimination results. Furthermore, via the average discriminating axis, we can visualize the differences between the populations.

We explored the effectiveness of different choices of features as input to the discriminant analysis. In a multi-object setting, it may not be solely shape that is significant in discriminating between populations. The relative pose of objects to each other may also be of interest, which implies that the alignment procedure preceding the shape analysis must be selected carefully. We have chosen to compare the discrimination results using features of volume, pose, shape, and the combination of pose and shape. Our shape features are determined via the medial object representation m-rep.

We demonstrate our methods in an application to a longitudinal pediatric autism study with object sets of 10 subcortical brain structures in a population of 70 samples. The results show that in this study volume performed better than pure shape or pose in discriminating between diagnosis groups, although all give statistically significant results. Pure pose information excluding scale, as well as the combination of pose and shape is non-discriminating suggesting that the non-scale pose information is mostly noise. The medial radius or thickness, a pose-independent shape feature specific to medial descriptions, performs equally well as volume. This study illustrates that local thickness information is likely of higher relevance in shape analysis of multi-objects sets than the relative pose information.

I. INTRODUCTION

Statistical shape modeling and analysis [1], [2], [3] is emerging as an important tool for understanding anatomical structures from medical images. Clinical applications favor a statistical shape modeling of multi-object sets rather than one of single structures outside of their multi-object context. Neuroimaging studies of mental illness and neurological disease, for example, are interested in describing group differences and changes due to neurodevelopment or neurodegeneration. These processes most likely affect multiple structures rather than a single one. A joint analysis of the structures, therefore, should reveal more than studying them individually. Applications of multi-object analysis include segmentation and studying group differences. Litvin et al. [4], for example, have proposed methodology for building a multi-object shape prior with application in 2D curve evolution segmentation. Fillard et al. [5] modeled full brain variation through computation of

sulcal curves. In this manuscript, we will focus on studying group differences in neuroimaging studies using discrimination analysis.

A fundamental difficulty in statistical shape modeling is the relatively small sample size, typically in the range of 20 to 50 samples in neuroimaging studies, compared to a high dimensional feature space, commonly one to several orders of magnitude larger than the sample size. Given that we are describing the shape of several structures instead of a single one, the dimension of our feature space tends to be even higher. This difficulty must be considered when choosing among different methods for discrimination analysis [6]. We favor the distance weighted discrimination (DWD) [7], which is similar to Support Vector Machines (SVM), but it suffers less from data piling problems in high dimensional low sample size (HDLSS) settings. Previous work in discriminating single anatomical objects has been done by Golland et al. [8] using distance transforms for shape features and SVM to discriminate populations. Yuschkevich et al. [9] also used SVM to discriminate 2D m-reps of corpora collosa.

Another context-specific choice is the type of features to use as input to the shape analysis. Most neurological studies focus solely on volume for the sake of simplicity [10], [11], [12], [13], [14], [15]. However, Styner et al. [16], [17] have shown that for particular applications, the shape of an object can be more useful in discriminating populations than volume. In a multi-object setting, there may be an additional feature of interest: the relative pose of objects with respect to each other. A statistical description of multi-object pose variability was introduced in Bossa et al. [18]. Since multi-object analysis of subcortical structures is novel, we have chosen to evaluate and compare several different features, namely volume, pose, shape, and the combination of pose and shape.

Several different geometric shape representations have been used to model anatomy, such as landmarks [19], dense collection of boundary points [20], or harmonic coefficients [21], [22]. Unlike the above explicit description, Tsai et al. [23] and Yang et al. [24] propose an implicit statistical object modeling by level-sets with its inherent difficulties of topology preservation. Another shape-based approach focuses on the analysis of spatial deformation maps [25], [26], [27], [28]. In this work, we employed explicit deformable shape modeling with a sampled medial mesh representation called m-rep, introduced by Pizer et al [29]. Styner et al. [30] have systematically compared the use of boundary and medial representations for the analysis of subcortical structures.

The work in this paper could be applied well to other shape

descriptions, but we chose a medial description for several reasons. First, it gives a more intuitive representation of the interior of the object. The radius, which describes the distance from the medial axis to the boundary, serves as a localized measure related to the object’s volume. This is particularly interesting for neuroimaging work because of the widespread use of volume data. Bouix et al. [31] studied hippocampi using the radius function defined on a flattened 2D medial sheet. Medial representations are also advantageous when attempting to describe certain nonlinear shape deformations such as bending and twisting [32]. Simple boundary representations have difficulties to account for this type of variability. The sampled m-rep description is also relatively compact when compared to other shape representations. We can describe 10 subcortical structures using 210 medial atoms for a total of 1890 features. While this is much higher than the number of data samples we typically have, it is less than the spherical harmonic representation that we have also computed and which uses about 10,000 features. The results of the study presented will show that the choice of a medial description was crucial to find relevant shape differences.

In summary, this paper presents results of discriminant analysis on sets of inter-related 3D objects. We choose the distance weighted discriminant (DWD) method and feature sets of volume, pose, and shape. The latter is given by the sampled medial m-rep shape representation. Driving application is a longitudinal pediatric neuroimaging study to show the early growth pattern of different populations.

II. METHODS

In this section, we first discuss the methodology of the different features we use in our discrimination analysis. These are the m-rep shape features and the local pose change features. We then summarize the method of distance-weighted discrimination, along with the transformation of our raw data before it is input to the DWD. Finally, we explain our method for building an unbiased estimator of untrained samples’ classification using DWD.

A. M-rep Shape Description

The m-rep shape description for a 3-D object consists of a sheet of medial atoms, each of which is defined by a position, radius, and two unit-length normal vectors to the boundary (spokes). The radius represents the distance from the atom position to the corresponding point on the boundary of the object along the two normal vectors. The medial atom, seen in Fig. 1, is defined as $m = \{\mathbf{p}, r, \mathbf{U}^{+1}, \mathbf{U}^{-1}\} \in \mathcal{M}$, with $\mathcal{M} = \mathbb{R}^3 \times \mathbb{R}^+ \times \mathcal{S}^2 \times \mathcal{S}^2$.

To obtain m-reps describing subcortical structures, we started with binary image segmentations from well-trained experts using semi-automated procedures.¹ We also needed a template m-rep that would be deformed to fit the binary image. We constructed these template medial models using the modeling scheme developed by Styner et al. [32] to

determine the minimum sampling required for each model. Given a binary segmentation and a model, the template model is deformed through an optimization process such that the model best fits the image without becoming too irregular in its geometry [33]. This process is applied individually to each of the 10 anatomical objects using the Pablo tool [34], while the correspondence across samples is implicitly established by the deformation process on the template model. Fig. 3 shows the medial atoms for a set of objects (a) and the implied surfaces (b).

B. Alignment and Pose

Alignment preceding shape analysis is well defined for shape analysis of single objects. In a multi-object setting, however, it must be decided how to remove unimportant shape variability through alignment. We call aligning the object set as a whole, where transformations are applied to all objects jointly, a global alignment. As seen in Fig. 2a and b, after this global alignment there are still local pose differences among the individual objects. In our case, we assumed these single object pose differences were important because they represent the inter-object changes within the multi-object set. Therefore, after the global alignment, we perform a second step referred to as the local alignment. In this step, we take the globally aligned object sets and align objects individually as would be done in a single object setting. It is these local pose changes that we include as part of the overall variability of the objects. The resulting m-reps after the local alignment are what we refer to as pure shape and can be seen in Fig. 2c. For the purposes of this paper, the global alignment included translation and rotation. This accounted for any pose differences between the original images. The local alignment included translation, rotation, and scale to remove all remaining pose. When we use the local pose changes as features for discriminant analysis, we have an 8-dimensional vector consisting of three elements for the translation, four for the orientation (stored as a quaternion), and one for the scale. After both global and local alignments have been completed, the final m-reps are in their mean pose position and are used as the pure shape features.

To align m-reps, we use a variation of the standard Procrustes method [35]. In a normal Procrustes alignment on a set of boundary points, the sum-of-squared distances between corresponding points is minimized. The standard Euclidean distance serves as the metric. For our purposes, we instead minimize the geodesic distance between m-reps because primitives do not lie in a Euclidean space. The geodesic distance $d(m_a, m_b)$ between two medial atoms m_a and m_b equals

$$\|\mathbf{p}_b - \mathbf{p}_a, \frac{r_b}{r_a}, \mathbf{R}(\mathbf{U}_b^{+1}) \cdot \mathbf{U}_a^{+1}, \mathbf{R}(\mathbf{U}_b^{-1}) \cdot \mathbf{U}_a^{-1}\| \quad (1)$$

where $\mathbf{R}(x)$ is the rotation of x to (1,0,0). For more details, see [36]. The distance between two m-reps is then the sum of geodesic distances between their corresponding atoms.

C. Distance Weighted Discrimination

Discriminant analysis is concerned with finding the direction which best separates two populations. An optimization

¹See <http://www.psychiatry.unc.edu/autismresearch/mri/roiprotocols.htm> for a detailed description of protocols and reliability results.

must be performed that somehow maximizes the distance between the discriminating hyperplane and the data points while separating the two classes. It is formulated in a general way as follows (see Fig. 4): given points x_i , class indicators $y_i \in \{+1, -1\}$, and w the normal to the separating hyperplane, the distance or residual, r , from the points to the hyperplane is

$$r_i = y_i(x_i'w + \beta) \quad (2)$$

where β determines the position of the separating hyperplane. One of the popular methods of discriminant analysis is Support Vector Machines (SVM). It attempts to maximize the minimum r_i . The main problem with this method is it tends to use only a small subset of the population, those closest to the separating plane, to completely define the discriminating axis. In high-dimensional feature space, it is manifested in the problem of “data piling” (see Fig. 5) where most of the samples from the same population group, when projected onto the normal of the discriminating axis, end up very close to each other. This leads to poor generalization performance when tested on new samples that were not included in the calculation of the discriminating axis: it is too specific to the samples from which it was computed.

Distance weighted discrimination is a method similar to SVM, but uses all sample points in the calculation of the discriminating axis.² It attempts to minimize the sum of the reciprocals of r_i . Through this, each point’s contribution to the calculation is weighted inversely proportional to the distance from that point to the opposite population. In this way, the DWD achieves a higher robustness when presented with new, untrained samples. This advantage is heightened further in the context of high dimensional feature spaces with low sample sizes where it is best to use all information available from the low number of samples.

D. Transformation of Raw Input Data

The m-rep shape description as well as the pose features contain rotational elements that are not part of a Euclidean space. This can lead to reduced performance of methods such as DWD that attempt to find a linear discriminant. Likewise, combining features with different units into one long feature vector can bias results towards features with larger variance. Finally, our data samples have unequal gender distributions within the two populations which could create a statistical bias. We must first account for each of these issues before running DWD analysis.

While the application of DWD to nonlinear features may give a reasonable solution, we found through experimentation that the linearized form of the m-rep features gives a better discrimination result (see Fig. 6). To obtain a linear instance of our curvilinear m-rep and pose features, we project them into the tangent space at the geodesic mean point [36]. This involves taking the log map of each of the non-Euclidean features. For the pose rotation, the log map of a unit-length

quaternion $q = (w, \mathbf{v})$ is defined as

$$\log q = \frac{\theta}{\sin(\theta/2)} \cdot \mathbf{v} \quad , \quad \theta = 2\arccos(w) \quad (3)$$

For the m-rep normal directions $\mathbf{U} = (x, y, z)$, the spherical log map is

$$\log \mathbf{U} = \left(x \cdot \frac{\theta}{\sin(\theta)}, y \cdot \frac{\theta}{\sin(\theta)} \right) \quad , \quad \theta = \arccos(z) \quad (4)$$

For the pose scale and m-rep radius factors it is just the logarithm function.

To concatenate features of different units, we first must make them commensurate to avoid unwanted bias. For our purposes, we have chosen to normalize each feature by subtracting the mean and dividing by the standard deviation. This makes the weighting of points equal among separate features in the DWD calculation. So for each feature, the input to the DWD procedure is of the form

$$Y = \frac{X - \bar{X}}{\sigma} \forall X \in \mathbb{R}^k \quad (5)$$

$$Y = \frac{\log(X - \bar{X})}{\sigma} \forall X \notin \mathbb{R}^k \quad (6)$$

The mean \bar{X} , however, is computed for each gender to avoid gender differences to become a confounding factor in the statistical analysis. By subtracting a gender-specific mean, we eliminate any disproportion in the gender sampling within our two populations. To build a gender-specific mean, we start with the lowest level of subcategories within our data and compute the mean of the samples in each. We then compute the mean of the subcategory means which are of the same gender. This gives us a gender-specific mean. Fig. 7 shows the process with subcategories according to the three criteria of gender, group, and time.

E. Unbiased Classification using Leave-Many-Out Experiments

To test the performance of the DWD, we conducted a leave-many-out, cross-validation experiment. We first divided our data samples into a training set and a testing set. The discriminating axis was computed using the training set. Each sample from the test set was then projected onto the DWD axis where the resulting one-dimensional projected value serves as the classification score (hence known as the DWD score). The DWD method produces both a discriminating axis and a threshold β . The threshold value is the amount by which the training data, after projected onto the DWD, must be shifted such that zero becomes the best dividing point between populations. Therefore, given a DWD axis w and a test sample feature vector x , the DWD score becomes

$$s = x'w + \beta \quad (7)$$

The discrete classification into one of the diagnosis groups is then simply defined as the sign of the DWD score.

In order to make the training set unbiased, we used the following strategy for selecting training samples (see Algorithm 1). We would alternately choose a random autism or

²See http://www.stat.unc.edu/faculty/marron/marron_software.html for a sample implementation.

a random control sample. With this sample from one group, we chose the sample from the other group that was a best match according to age and gender to the subject from the first group. This gave us one sample from each group. Since our data is longitudinal, we always included each sample’s counterpart across time. From here, the process was repeated but starting with a random sample from the opposite group than in the previous iteration. After several iterations, we end up with a training set with an equal number of samples from each group.

After several runs of this experiment, all of the data samples were included in the testing set of at least a few runs. From the results of these experiments, we could then build an unbiased estimate of each sample’s classification. For each sample, we computed it’s mean DWD score over those runs of the experiment for which it was in the test set. In this way, we calculate a classification for a sample only when the discriminating axis was computed without any knowledge of that sample. The box plots in the following sections are of these unbiased mean DWD scores.

Algorithm 1 Training Set T Selection

```

 $T = \emptyset, size = 0, i = 0$ 
while  $size < n$  do
  if  $i \bmod 2 = 0$  then
     $s =$  random sample from autism group
     $t =$  closest matching sample to  $s$  from control group
  else
     $s =$  random sample from control group
     $t =$  closest matching sample to  $s$  from autism group
  end if
   $s' =$  counterpart of  $s$  across time
   $t' =$  counterpart of  $t$  across time
   $T = T \cup \{s, s', t, t'\}$ 
   $size = size + 4$ 
   $i = i + 1$ 
end while

```

III. RESULTS

In this section, we describe our data set and the results of the leave-many-out experiment. We have divided the results into five sections corresponding to the features used in the discriminant analysis: volume, pose, shape, shape and pose combined, and local width represented by m-rep radii. We then finish with some visualizations of the discriminating differences between the populations.

A. Motivation and Clinical Data

The driving clinical problem of this research is the need for a joint analysis of the set of subcortical brain structures, over and above that of individual structures. The image data used in this paper is taken from an ongoing clinical longitudinal pediatric autism study [37]. This study includes autistic subjects (AUT) and typically developing, healthy controls (CONT) with baseline at age 2 and follow-up at age 4. For the results shown here, we have selected 23 subjects from the autism

group and 10 from the control group. For all of the autism subjects and 6 of the 10 controls, we have successful scans at age 2 and age 4. For the other 4 controls, we paired an age 2 scan of one subject with an age 4 scan of another unrelated subject. We also have 4 additional control age 2 scans that have no matching age 4 scan. This gives us a total of 70 samples: 46 autism and 24 control subjects.

In the classifier experiments, we always included all details of a selected training subject. Thus, the four unpaired control samples were always left out of the training set. From our specific data, we chose a training set, in the manner described above, consisting of 32 out of the 70 available samples. Thus, it included 16 samples from the control group and 16 from the autism group. The remaining 38 samples served as the test set. The experiment was then run 100 times. The number of runs was chosen heuristically such that each of the 70 samples was included in the test set for at least a few runs; the minimum number of runs in the test set for any sample was 4. From these, we could calculate unbiased mean DWD scores.

B. Volume

Because of its prevalence in neuroimaging studies, we first assessed the ability of object volumes to discriminate between the autism and control groups. The volumes were computed from the implied surface boundary of the m-reps. The 10 subcortical structures gave us a 10-dimensional feature space to serve as the basis for the discriminating axis. We computed the mean DWD score for each sample over the runs in which that sample was in the test set. This gives us an unbiased average classification score for each sample. Fig. 8 shows that there is a clear difference between the median score of the autism group and control group. This difference is statistically significant with $p < 0.001$. As another measure of the discrimination performance of the volume, Table I shows the average percentage (74%) of the 38 test samples that were correctly classified over the 100 runs of the leave-many-out experiment. Volume percent change from autism to control of the mean volumes of each structure is listed in Table III. All structures show a larger volume in the autism group.

C. Pose

The next step was to explore the significance of local pose changes. For each sample, these features totaled 70: three for translation, three for rotation, and one for uniform scale across 10 objects. The raw features were transformed as described in section II-D, thus reducing the quaternion representing the rotation to a three-dimensional vector. The significance in the volume discrimination led us to believe that the pose, which includes uniform object scale factors, would also show significance. This was the case for the mean DWD scores with $p = 0.01$ (Fig. 9). However, the test sample classification accuracy was considerably lower than volume as Table I illustrates. There also were 14 individual runs in which the classification accuracy was at or below 50%, a result that would be outperformed by a random coin flip. The translation and rotation components of the pose seemed to be adding mostly noise and instability to the DWD calculation because

the same experiment run with only the scale factors gave an average classification rate of 70% and $p < 0.001$ as opposed to 54% and $p = 0.1$ using the translations and rotations. From these results, we conclude that the pose does include some relevant information for discrimination but it is likely in the uniform scale factors.

D. Shape

Figure 10 shows the results of using only the m-rep shape features for the DWD calculations after objects have been individually aligned and scaled. Unlike with volume and pose, the mean DWD scores for the test samples were not significantly different ($p = 0.92$, Fig. 10). The classification accuracy of shape was less than volume and pose with an average correctness rate of 53%. Both of these results signify that there did not exist strong shape differences between the two population groups. However, as seen with the translations and rotations in pose, particular subsets of features which consist mostly of noise can be detrimental to discrimination. Restricting the analysis to a subset of shape features is discussed below.

E. Shape and Pose

Finally, the high classification accuracy of the volume and pose features compared with shape led us to combine the latter two. This gave us the most complete description of the variability of the multi-object complex. The differences between the mean DWD scores were not significant ($p = 0.21$). Once again, the shape features seem to represent mostly noise since combining them with pose produced a nonsignificant result while pose alone was significant. The classification accuracy (58%) was slightly higher than shape but lower than pose features alone as seen in Table I.

F. Shape and Scale: Local Width via M-rep Radii

One of the ongoing questions with a discretely sampled m-rep shape description is the issue of correspondence. Currently, the correspondence is implicitly defined in the process of fitting a template model to each binary image. However, it is not clear that features such as the spokes and atom positions are sufficiently stable to give a meaningful correspondence across samples. Out of this arises the question of whether statistics on these features, which we used as part of our shape analysis above, could be tighter with a better correspondence.

To explore this question, we chose to run our same analysis on only the radius feature of each medial atom without any alignment. This is because the radius is less sensitive to a noisy correspondence than some of the other medial atom features; a small change in the atom position or spoke directions will generally not cause a large change in the radius. It is also invariant to any object-level translation and rotation, while including both shape and scale information in the form of local width. This makes it a more intrinsic measure than the other medial atom features. Compared to the other features in Table I, the mean classification accuracy over 100 runs of the leave-many-out experiment is best with 76% when using

only the atom radii. Fig. 12 shows the mean DWD scores using only the radii. While the radii give the best classification rate of any feature we studied, it is difficult to say outright that they are the best features to use. First, it may be our specific application that lends itself to a local measure that is related to volume like the radii. It is known, and evidenced above, that there are volume differences between autistic and typically developing brains. Also, Fig. 12 shows that the overlap between the populations is not drastically smaller than with volume. However, the improvement in classification accuracy from using all the medial atom features to only the radii suggests that there is a certain amount of noise in the other features which ends up being correlated to the detriment of the DWD calculation. At the same time, witnessing the local radii slightly outperform volume as a discriminating feature suggests additional discrimination power of shape versus volume and supports the choice of a medial shape description.

G. Evaluation of Bias

To verify that the mean classification scores were unbiased, we ran our same experiments using random, normally distributed input data. We used the same random number seeding and the exact same training and testing sets. The random data was generated with the same mean, variance, and dimension as our actual shape data. The p-value of the mean DWD scores for this case was 0.22 and the average classification accuracy was 49%.

H. Discrimination by Age

To complete our analysis, we ran the same experiments as above, with precisely the same gender-corrected data, but discriminating according to age instead of diagnosis. Through this we wanted to explore the performance of the discrimination in the context of presumably more separate groups. And indeed, as seen in Table II, most of the different features have an easier time discriminating by age than by diagnosis. Especially noteworthy is the pose, which does marginally better than the volume, m-rep radii, and scale factors (see Figs. 13 and 14). It is important to remember that the latter two are not a total brain volume or scale, which should be easily separated after correcting for gender, but not for the individual structures. Our previous work [38] had shown that the first principal mode of our pooled data (both time points) aligned very closely with time. The deformation along this mode showed a large global scaling effect, with the subcortical structures moving radially outward. Here, this change across time is represented in the local translation alignment parameters, which helps to explain why the entire local pose (translation, rotation, and scale) performs slightly better than scale or volume alone. Our conclusion is that since pose does well at discriminating by age, removing it for the analysis of diagnosis groups (and witnessing it not aid in the discrimination when included) leaves us with discriminating differences that are not heavily influenced by having pooled data across time. However, a direct method of correcting for age, such as that done for gender, would be preferable and will be explored in the future.

I. Visualization

To visualize the changes in shape along the DWD direction, we start with the mean m-rep of the autism group. Then, we deform the autism mean m-rep along the unit-length DWD which points toward the control group. The distance along the DWD direction by which the autism mean is deformed is defined as the distance between the mean of each group’s projections onto the DWD line. The final m-rep which has been deformed this full distance is then used to represent the control group.

For robustness, we chose to use the mean DWD direction over all the runs instead of using a single run from the leave-many-out experiment. We also used all 70 samples to compute the distance between the projected group means, which gives us the distance to deform along the DWD direction, and we used all 46 autism samples for the autism mean m-rep. Fig. 15a shows colormaps of surface distances between the two multi-object sets representing each of the two diagnosis groups using shape only. The measurement is the distance when starting from the autism group and deforming towards the control group. Red, green, and blue coloring denote inward, zero, and outward deformations respectively. In this image, we see that the amygdala and hippocampus undergo strong shape changes between the groups relative to the other three structures. There is a distinct inward deformation of the hippocampus tail (seen in far right of Fig. 15a) as well as an outward change in the midsection. A large portion of each amygdala also presents a difference in shape.

Given that the m-rep radii performed best as a discriminating feature, we also wanted to visualize these differences. We used the same procedure as above to obtain a representative set of radii for each diagnosis group. Fig. 15b shows the surface distance between m-reps of each group with only the radii modified. The strongest individual radii changes appear to be in the hippocampus and the caudate. To assess the overall radii differences between groups, we calculated $\Delta r = \log(r_c/r_a)$ for each corresponding atom with r_c the radius of the control group and r_a the radius of the autism group. Fig. 16 shows spheres plotted at the mean atom positions between the two groups with size proportional to Δr . The color denotes the sign of Δr , with red being a decrease in radius from autism to control and blue an increase. We see a clear decrease in local widths when deforming from autism to control at almost all positions across both structures. The hippocampus shows mainly a large decreases in the posterior head and body section. These radii decreases are supported by Table III which lists the percent change from autism to control of the mean volumes of each structure at ages 2 and 4. All structures show a larger volume in the autism group.

IV. CONCLUSION

This research presents work in progress towards shape analysis and group discrimination of multi-object complexes. Traditionally, shape analysis is mostly concerned with representation and statistical analysis of single objects, mostly following a well developed mathematical framework that proposes linear alignment and subsequent statistical analysis

of corresponding features. Statistical analysis of complexes of interrelated shapes has not yet been sufficiently studied and is not just to be seen as a straightforward extension of methodologies for single shape analysis.

In a multi-object setting, the alignment step has to be reconsidered. Linear alignment of a population of sets of objects will remove global translation, rotation and scale, but will not account for relative object pose variability. A joint analysis of only globally aligned sets of shapes will therefore include these residual pose differences into the statistical shape model. Here, we discuss and explore various options for global and local alignment of sets of shapes. We propose an initial global alignment with rotation and translation to map each dataset into a common coordinate frame. This step is followed by a local alignment of each object individually, but the alignment parameters translation, rotation and scale are kept as pose parameter vectors. Shape analysis of the joint set of objects will therefore use *pure* shape features not affected by any residual pose differences. Features are mapped into Riemannian symmetric space, the appropriate choice for features that include rotational frames and positive reals. Features used here include medial atoms and the vectors of the joint pose parameters of the multi-object complexes. It is then straightforward to chose pose, shape, or pose and shape as features for statistical group discrimination.

Our results, given the clinical pediatric imaging application, show that pose features do not give statistically significant discrimination. Shape features also did not show significance, except when isolating a particular feature of the m-rep shape description, namely the radius measure of local width. This feature combines the locality of a shape description with the scale information known to be discriminating in our particular application. Using this local size parameter results in better generalized discrimination performance.

Although sampled medial representations use a lower number of features than densely sampled surfaces, we still face the HDLSS problem (high dimensionality low sample size). This problem is even more pronounced with the analysis of object sets, resulting in a feature space dimensionality which is magnitudes larger than the number of samples. In typical applications similar to the one described here, two populations of 25 samples are each represented by 2000 features to provide a sufficiently detailed representation for 10 3-D objects. For classification, we applied the distance-weighted discrimination (DWD) method, which is a variant of support vector machine discrimination but is designed to be robust for HDLSS data analysis problems. Unbiased statistical analysis by repeated leave-many-out experiments finally results in classification rates and significance values (p-values).

The driving application is a pediatric autism study with autistic and typically developing children imaged at 2 and 4 years of age. We focus on a joint analysis of five left and right subcortical structures represented as sampled medial representations after model fitting. Please note the discrepancy of relatively low classification rates in the presences of highly significant population differences. This might possibly be explained by the nature of the underlying clinical problem. Morphologic phenotypes in neurodevelopmental disorders are

often reflected by only subtle differences and increased heterogeneity. In the future, we will explore multi-variate classification by selection of a best-separating subspace rather than a single axis. Further, we will have to develop a technique to explore the covariance structure of sets of shapes in order to explain their interrelationship. This will help clinicians to explore links between morphological changes and underlying biological processes.

V. ACKNOWLEDGEMENTS

This research is supported by the NIH NIBIB grant P01 EB002779, the NIH Conte Center MH064065, and the UNC Neurodevelopmental Research Core NDRC, subcore Neuroimaging. The MRI images of infants, caudate images and expert manual segmentations are funded by NIH RO1 MH61696 and NIMH MH64580.

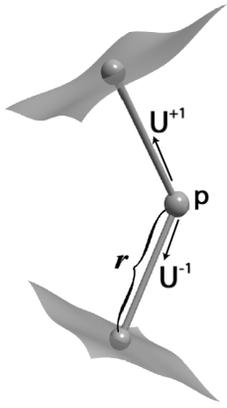


Fig. 1.

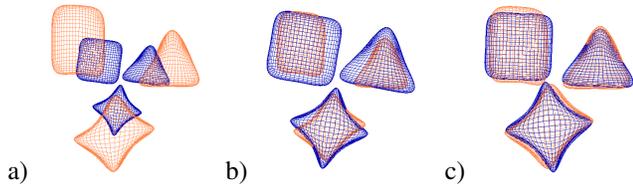


Fig. 2.

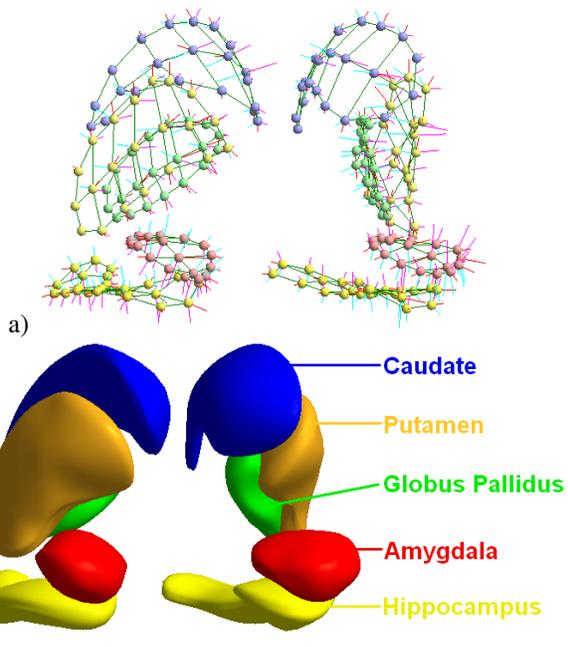


Fig. 3.

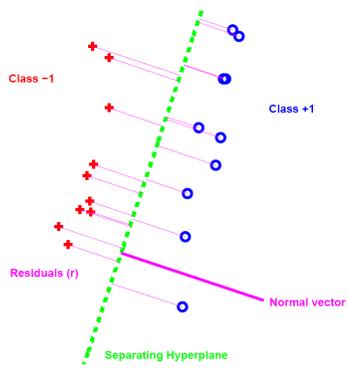


Fig. 4.

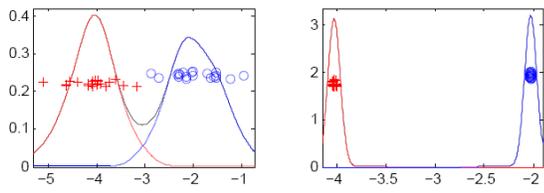


Fig. 5.

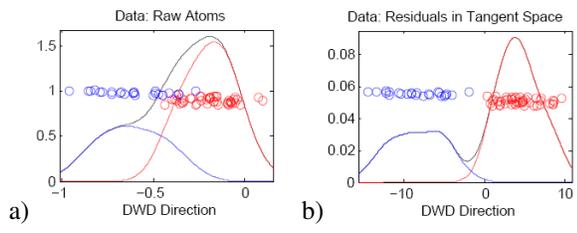


Fig. 6.

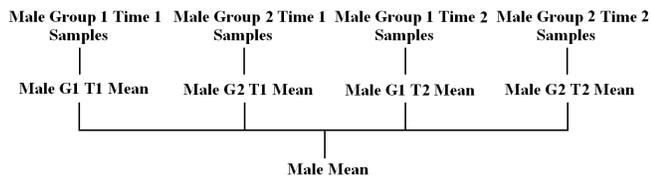


Fig. 7.

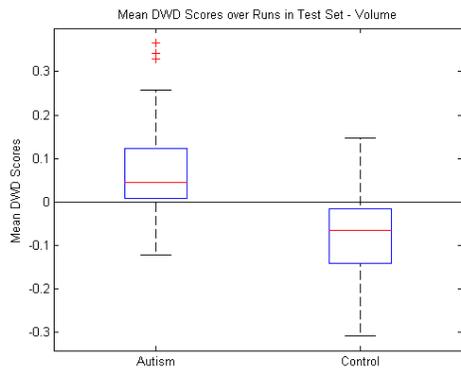


Fig. 8.

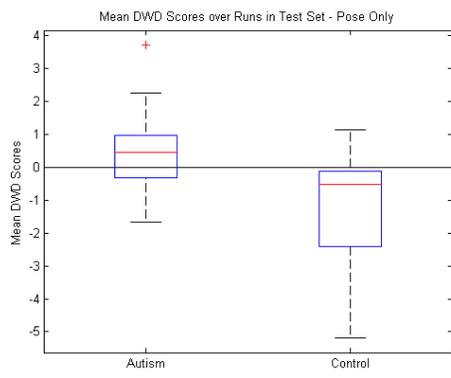


Fig. 9.

TABLE I
CLASSIFICATION ACCURACIES FOR TEST SAMPLES OVER 100 RUNS OF
LEAVE-MANY-OUT EXPERIMENT

Feature	Mean	Std. Dev.
Volume	74 %	± 8 %
Pose	61 %	± 8 %
Pose (Scale Only)	70 %	± 8 %
Shape	53 %	± 6 %
Shape and Pose	58 %	± 8 %
Shape and Scale (Radii Only)	76 %	± 7 %

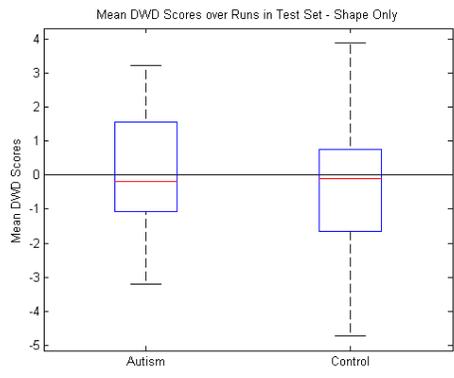


Fig. 10.

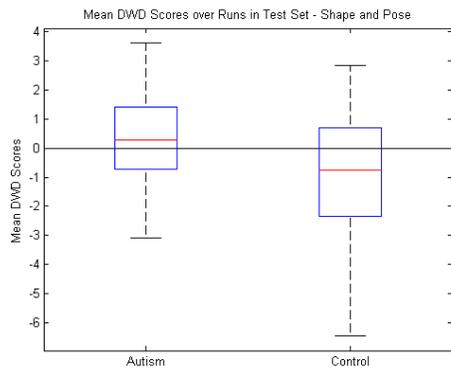


Fig. 11.

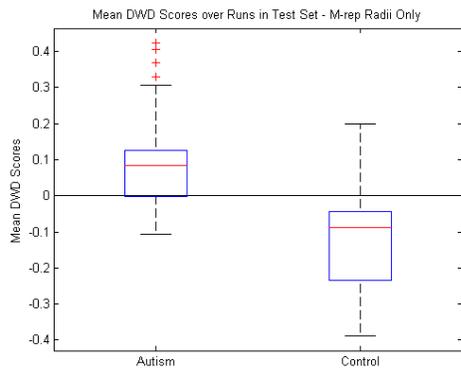


Fig. 12.

TABLE II
CLASSIFICATION ACCURACIES FOR DISCRIMINATION BY AGE

Feature	Mean	Std. Dev.
Volume	72 %	± 5 %
Pose	74 %	± 6 %
Pose (Scale Only)	69 %	± 6 %
Shape	60 %	± 7 %
Shape and Pose	73 %	± 6 %
Shape and Scale (Radii Only)	70 %	± 6 %

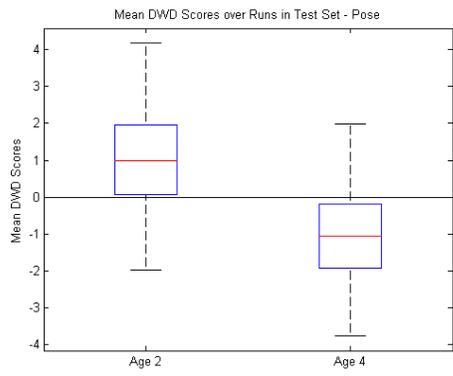


Fig. 13.

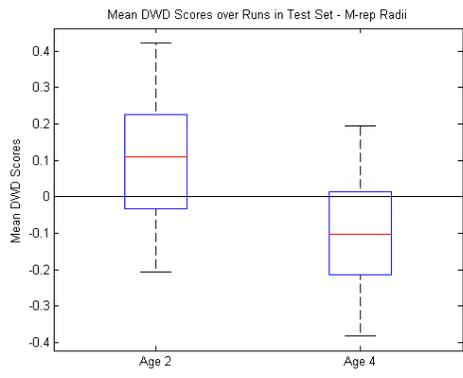


Fig. 14.

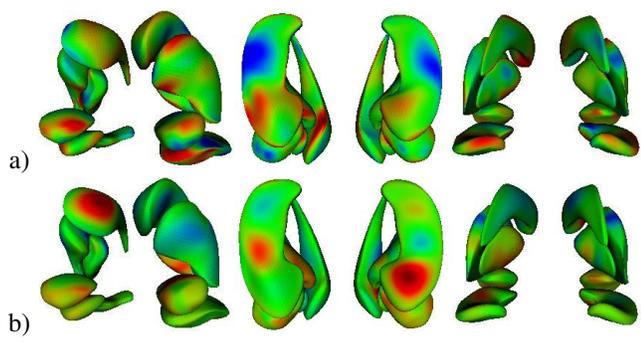


Fig. 15.

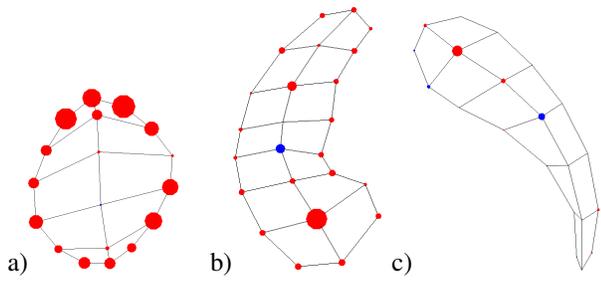


Fig. 16.

TABLE III

VOLUME PERCENT CHANGE FROM AUTISM TO CONTROL

Object	$\% \Delta_2$	$\% \Delta_4$
Left Amygdala	-19.85	-15.41
Right Amygdala	-21.05	-18.19
Left Caudate	-14.38	-3.84
Right Caudate	-13.96	-5.49
Left Hippocampus	-10.92	-7.98
Right Hippocampus	-10.29	-6.69
Left Globus Pallidus	-10.19	-16.17
Right Globus Pallidus	-2.57	-8.27
Left Putamen	-6.43	-3.66
Right Putamen	-4.54	-2.24

$$\% \Delta_A = \frac{\overline{Vol}_{cont} - \overline{Vol}_{aut}}{\overline{Vol}_{cont}} \times 100, \text{ at age } A$$

Fig. 1. Medial atom: position (\mathbf{p}), radius (r), two normals to boundary (\mathbf{U})

Fig. 2. Multi-object alignment. a) Global translation and rotation. b) Global translation, rotation, and scale. c) Local translation, rotation, and scale after initial global translation and rotation.

Fig. 3. M-reps of a multi-object complex. a) Medial atoms. b) Implied boundary surfaces of medial description.

Fig. 4. Illustration of two-class discrimination with separating hyperplane and residuals. The samples marked with boxes would act as the support vectors in SVM, whereas all samples are included in DWD.

Fig. 5. Left: Projection onto normal of optimal separating hyperplane. Right: Projection onto normal of separating hyperplane which exhibits data piling.

Fig. 6. Separation of 70 multi-object m-reps into two populations given by DWD axis. a) Raw, nonlinear medial atom data. b) Atom data after projection into tangent space and subtraction of mean.

Fig. 7. Illustration of gender-specific mean calculation given subcategories of gender, group, and time. The same process is applied to obtain a female mean.

Fig. 8. Volume features: Box plot (median, 25 and 75 percentiles, min/max) of mean DWD scores of each group over those runs in which the samples were in the test set. Greater than zero classified as autism, less than zero classified as control. $p < 0.001$.

Fig. 9. Pose features: Box plot (median, 25 and 75 percentiles, min/max) of mean DWD scores of each group over those runs in which the samples were in the test set. Greater than zero classified as autism, less than zero classified as control. $p = 0.01$.

Fig. 10. Shape features: Box plot (median, 25 and 75 percentiles, min/max) of mean DWD scores of each group over those runs in which the samples were in the test set. Greater than zero classified as autism, less than zero classified as control. $p = 0.92$.

Fig. 11. Shape and pose features combined: Box plot (median, 25 and 75 percentiles, min/max) of mean DWD scores of each group over those runs in which the samples were in the test set. Greater than zero classified as autism, less than zero classified as control. $p = 0.21$.

Fig. 12. M-rep radii features: Box plot (median, 25 and 75 percentiles, min/max) of mean DWD scores of each group over those runs in which the samples were in the test set. Greater than zero classified as autism, less than zero classified as control. $p < 0.001$.

Fig. 13. Pose features (by age): Box plot (median, 25 and 75 percentiles, min/max) of mean DWD scores of each group over those runs in which the samples were in the test set. Greater than zero classified as age 2, less than zero classified as age 4.

Fig. 14. M-rep radii features (by age): Box plot (median, 25 and 75 percentiles, min/max) of mean DWD scores of each group over those runs in which the samples were in the test set. Greater than zero classified as age 2, less than zero classified as age 4.

Fig. 15. Colormap of surface distances from autism mean m-rep to deformed m-rep along DWD direction using a) shape only, b) m-rep radii only. Red, green, and blue are inward distance, zero distance, and outward distance respectively.

Fig. 16. Visualization of radii change from autism to control in a) right amygdala, b) right hippocampus, c) right caudate. Size of ball at an atom position is proportional in size to the log change in radius from control to autism. Red signifies that radii in autism are larger than in controls ($\log(r_c) - \log(r_a) < 0$), and blue the opposite.

REFERENCES

- [1] D. Thomson, *On Growth and Form*. Cambridge University Press, second edition, 1942.
- [2] I. Dryden and K. Mardia, "Multivariate shape analysis," *Sankhya*, vol. 55, pp. 460–480, 1993.
- [3] C. G. Small, *The statistical theory of shape*. Springer, 1996.
- [4] A. Litvin and W. C. Karl, "Coupled shape distribution-based segmentation of multiple objects," in *Information Processing in Medical Imaging*, 2005, pp. 345–356.
- [5] P. Fillard, V. Arsigny, X. Pennec, K. M. Hayashi, P. M. Thompson, and N. Ayache, "Measuring brain variability by extrapolating sparse tensor fields measured on sulcal lines," *NeuroImage*, vol. 34, no. 2, pp. 639–650, January 2007.
- [6] I. Dryden and K. Mardia, *Statistical Shape Analysis*. Wiley, 1998.
- [7] J. Marron and M. Todd, "Distance weighted discrimination," Operations Research and Industrial Engineering, Cornell University, Technical Report 1339, 2002, available at <http://www.stat.unc.edu/postscript/papers/marron/HDD/DWD/DWD2.pdf>.
- [8] P. Golland, W. E. L. Grimson, M. E. Shenton, and R. Kikinis, "Detection and analysis of statistical differences in anatomical shape," *Medical Image Analysis*, vol. 9, pp. 69–86, 2005.
- [9] P. Yushkevich, S. M. Pizer, S. Joshi, and J. Marron, "Intuitive, localized analysis of shape variability," in *Information Processing in Medical Imaging*, 2001, pp. 402–408.
- [10] R. McCarley, C. Wible, M. Frumin, Y. Hirayasu, J. Levitt, I. Fischer, and M. Shenton, "Mri anatomy of schizophrenia," *Biological psychiatry*, 1999.
- [11] J. Levitt, R. McCarley, C. Dickey, M. Voglmaier, M. Niznikiewicz, L. Seidman, Y. Hirayasu, A. Ciszewski, R. Kikinis, F. Jolesz, and M. Shenton, "Mri study of caudate nucleus volume and its cognitive correlates in neuroleptic-naive patients with schizotypal personality disorder," *Am J Psychiatry*, vol. 159, no. 7, pp. 1190–1197, 2002.
- [12] R. S. E. Keefe, L. J. Seidman, B. K. Christensen, R. M. Hamer, T. Sharma, M. M. Sitskoorn, S. L. Rock, S. Woolson, M. Tohen, G. D. Tollefson, T. M. Sanger, and J. A. Lieberman, "Long-term neurocognitive effects of olanzapine or low-dose haloperidol in first-episode psychosis," *Biological psychiatry*, vol. 59, no. 2, pp. 97–105, 2006.
- [13] J. Juranek, P. A. Filipek, G. R. Berenji, C. Modahl, K. Osann, and M. A. Spence, "Association between amygdala volume and anxiety level: magnetic resonance imaging (mri) study in autistic children," *Journal of child neurology*, vol. 21, no. 12, pp. 1051–8, Dec. 2006.
- [14] M. Langen, S. Durston, W. G. Staal, S. J. M. C. Palman, and H. van Engeland, "Caudate nucleus is enlarged in high-functioning medication-naive subjects with autism," *Biol Psychiatry*, 2007.
- [15] J. Munson, G. Dawson, R. Abbott, S. Faja, S. J. Webb, S. D. Friedman, D. Shaw, A. Artru, and S. R. Dager, "Amygdalar volume and behavioral development in autism," *Archives of general psychiatry*, vol. 63, no. 6, pp. 686–93, June 2006.
- [16] M. Styner, A. Lieberman, R. K. McClure, D. R. Weinberger, D. W. Jones, and G. Gerig, "Morphometric analysis of lateral ventricles in schizophrenia and health controls regarding genetic and disease-specific factors," *Proc. of the National Academy of Sciences*, vol. 102, no. 13, pp. 4872–4877, March 2005.
- [17] G. Gerig, M. Styner, M. Shenton, and J. Lieberman, "Shape versus size: Improved understanding of the morphology of brain structures," in *MICCAI*, 2001, pp. 24–32.
- [18] M. N. Bossa and S. Olmos, "Statistical model of similarity transformations: Building a multi-object pose," in *2006 Conference on Computer Vision and Pattern Recognition Workshop (CVPRW'06)*, 2006, p. 59.
- [19] F. Bookstein, "Shape and the information in medical images: A decade of the morphometric synthesis," in *MMBIA*, 1996.
- [20] T. Cootes, C. Taylor, D. Cooper, and J. Graham, "Active shape models -their training and application," *Comp. Vis. Image Under.*, vol. 61, pp. 38–59, 1995.
- [21] A. Kelemen, G. Székely, and G. Gerig, "Elastic model-based segmentation of 3d neuroradiological data sets," in *IEEE Trans. Med. Imaging*, vol. 18, 1999, pp. 828–839.
- [22] L. Staib and J. Duncan, "Model-based Deformable Surface Finding for Medical Images," *IEEE Trans. Med. Imaging*, vol. 15, no. 5, pp. 1–12, 1996.
- [23] A. Tsai, A. Yezzi, W. Wells, C. Tempany, D. Tucker, A. Fan, E. Grimson, and A. Willsky, "Shape-based approach to curve evolution for segmentation of medical imagery," *IEEE Transactions on Medical Imaging*, vol. 22, no. 2, pp. 137–154, Feb 2003.
- [24] J. Yang, L. H. Staib, and J. S. Duncan, "Neighbor-constrained segmentation with level set based 3d deformable models," *IEEE Transactions on Medical Imaging*, vol. 23, no. 8, pp. 940–948, Aug 2004.
- [25] C. Davatzikos, M. Vaillant, S. Resnick, J. Prince, S. Letovsky, and R. Bryan, "A computerized method for morphological analysis of the corpus callosum," *J. of Comp. Assisted Tomography*, vol. 20, pp. 88–97, Jan./Feb 1996.
- [26] J. Csernansky, S. Joshi, L. Wang, J. Haller, M. Gado, J. Miller, U. Grenander, and M. Miller, "Hippocampal morphometry in schizophrenia via high dimensional brain mapping," *Proc. Natl. Acad. Sci. USA*, vol. 95, pp. 11406–11411, September 1998.
- [27] P. Thompson, M. Mega, and A. Toga, *Brain Mapping: The Disorders*. Academic Press, 2000, ch. Disease-Specific Brain Atlases.
- [28] P. Thompson, J. Giedd, R. Woods, D. MacDonald, A. Evans, and A. Toga, "Growth patterns in the developing brain detected by using continuum mechanical tensor maps," *Nature*, vol. 404, pp. 190–193, 2000.
- [29] S. Pizer, D. Fritsch, P. Yushkevich, V. Johnson, and E. Chaney, "Segmentation, registration, and measurement of shape variation via image object shape," in *IEEE Trans. Med. Imaging*, vol. 18, 1999, pp. 851–865.
- [30] M. Styner, J. A. Lieberman, D. Pantazis, and G. Gerig, "Boundary and medial shape analysis of the hippocampus in schizophrenia," in *MedIA*, 2004, pp. 197–203.
- [31] S. Bouix, J. C. Pruessner, D. L. Collins, and K. Siddiqi, "Hippocampal shape analysis using medial surfaces," *NeuroImage*, vol. 25, pp. 1077–1089, 2005.
- [32] M. Styner, G. Gerig, J. Lieberman, D. Jones, and D. Weinberger, "Statistical shape analysis of neuroanatomical structures based on medial models," *Medical Image Analysis (MEDIA)*, vol. 7, no. 3, pp. 207–220, Sept 2003.
- [33] S. Pizer, T. Fletcher, Y. Fridman, D. Fritsch, A. Gash, J. Glotzer, S. Joshi, A. Thall, G. Tracton, P. Yushkevich, and E. Chaney, "Deformable m-reps for 3d medical image segmentation," *International Journal of Computer Vision IJCV*, vol. 55, no. 2, pp. 85–106, 2003.
- [34] S. M. Pizer, P. T. Fletcher, S. Joshi, A. G. Gash, J. Stough, A. Thall, G. Tracton, and E. L. Chaney, "A method and software for segmentation of anatomic object ensembles by deformable m-reps," *Medical Physics Journal*, vol. 32, no. 5, pp. 1335–1345, May 2005.
- [35] C. Goodall, "Procrustes methods in the statistical analysis of shape," *Journal of the Royal Statistical Society*, vol. 53, no. 2, pp. 285–339, 1991.
- [36] P. Fletcher, C. Lu, S. Pizer, and S. Joshi, "Principal geodesic analysis for the study of nonlinear statistics of shape," in *IEEE Transactions on Medical Imaging*, vol. 23, 2004, pp. 995–1005.
- [37] H. Cody Hazlett, M. Poe, G. Gerig, R. Gimpel Smith, J. Provenzale, A. Ross, J. H. Gilmore, and J. Piven, "Magnetic resonance imaging and head circumference study of brain size in autism," *The Archives of General Psychiatry*, vol. 62, pp. 1366–1376, December 2005.
- [38] M. Styner, K. Gorcowski, T. Fletcher, J. Jeong, S. M. Pizer, and G. Gerig, "Multi-object statistics using principal geodesic analysis in a longitudinal pediatric study," in *Lecture Notes in Computer Science*, G.-Z. Yang, T. Jiang, D. Shen, L. Gu, and J. Yang, Eds., no. 4091. MIAR Conference, August 2006, pp. 1–8.