Akash Krishna

Incorporating the Geometric Relationship of Adjacent Objects in Multi-Object Shape Analysis

Abstract:

Modeling the shared boundary region between adjacent 3D objects can provide useful information regarding the geometric relationship of objects when performing multi-object shape analysis. Our analysis goes about modeling the shared boundary region of 3D objects with a 2D s-rep. An s-rep is able to capture shape features like width, boundary locations, boundary normals, and object curvature. A 2D s-rep is fit to the shared boundary region of adjacent objects by mapping the shared boundary region onto the medial surface of an ellipsoid. Mapping onto a flat surface allows for the creation of the 2D s-rep, which we can map back to the curved shared boundary region. Incorporating the 2D s-rep of the shared boundary region along with the 3D s-reps of the adjacent objects allows for a more detailed multi-object representation to utilize for shape analysis. We conducted an experiment with computer generated data of pairs of deformed ellipsoids that were either stretched, bent, or shifted and analyzed how well our representation is able to classify a deformed pair of ellipsoids from an undeformed pair. Our study found that a representation that included a 2D s-rep of the shared boundary region increased the classification accuracy for classifying bent and shifted ellipsoids, but decreased the accuracy for stretched ellipsoids.

1. Introduction:

When attempting to perform classification or hypothesis testing for disorders that affect multiple structures in the human body, it is preferable to study the joint shape of multiple structures as opposed to studying individual structures one at a time. For example, many neurodevelopmental processes affect multiple structures in the brain. Individuals with autism have differently shaped and sized subcortical structures than individuals without autism. Additionally, multi-object shape analysis is utilized in the automatic segmentation of multiple objects from medical images, e.g., for radiation therapy

Figure 1: subcortical region in human brain

treatment planning. Understanding the shape prior of the whole region rather than an individual object could be expected to improve segmentation accuracy.
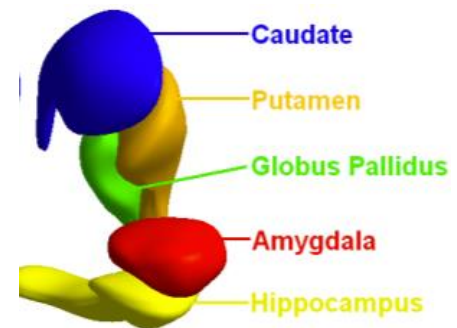
In the paper, our[1] approach to performing multi-object shape analysis involves not only modeling individual objects with 3-D s-reps, but also modeling the shared boundary region between adjacent objects with a 2-D s-rep. By including the model of the shared boundary region, we will be able to obtain shape statistics not only on the shape of the objects, but also on the geometric relationship of those adjacent objects. We have chosen s-reps as our means of using an object

representation because s-reps are able to capture a richer set of features compared to alternative methods like PDMs (Hong et al., 2016).

In section 2 entitled "s-reps," I will go deeper into what an s-rep is for a single object in 3D and how it is fitted to a particular anatomical structure. In section 3 entitled "Statistics," I will discuss the statistical techniques utilized for preparing the data and for conducting hypothesis testing and classification. Following the statistics section, in Section 4, "Multi-Object Shape Analysis", I will cover our method of getting a representation of a multi-object region. This will include our way of obtaining a geometric representation of a multi-object region in the human body. Section 5, entitled "2D s-rep," will cover my method of fitting an s-rep to a 2D object on a flat or curved surface. Section 6, entitled "Data Generation," will discuss how the data for our statistics was generated. Section 7, "Results and Analysis," will display our classification accuracy for conducting classification on our computer-generated test cases. Section 8, "Future Work," will cover our future applications of the research and possible improvements to the work.

2. S-reps:

An s-rep for a 3D or 2D object (Pizer et al., 2019) is composed of a skeleton and a collection of top and bottom spoke vector pairs on the skeleton that fill the inside of the object. An s-rep is a particularly useful representation for shape analysis because it is able to capture geometric properties like boundary normals and object width. Another useful feature of s-reps is that they are all going to be in correspondence because they are all formed from a common reference ellipsoid and richly depend on the object shape. Correspondence is especially important when attempting to produce statistics on a training population. Studies have further corroborated this idea that s-reps provide a better object representation to produce statistics. A study in 2016 by Hong et. al fit both PDMs and s-reps to hippocampi and performed classification on the resulting object representations. The ROC curves shown in Figure 2 indicate that Euclideanized s-reps were able to classify those with schizophrenia more accurately than methods using either Euclideanized or non-Euclideanized PDMs.
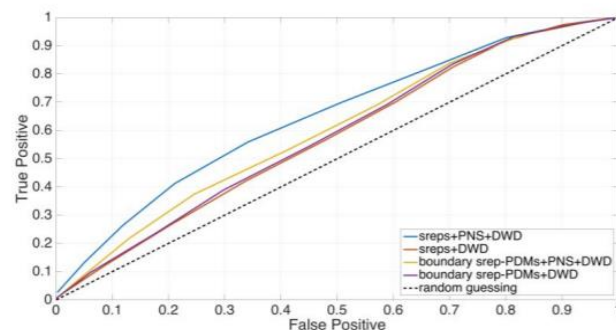

Figure 2: ROC curves displaying effectiveness of utilizing s-reps

Now, what is the process of fitting an s-rep to a target object in 3D represented by its boundary? The initialization portion of the fitting involves using a curvature flow to smooth the target object into that of an ellipsoid, calculating the s-rep for the given ellipsoid, and applying the inverse of the curvature flow along with an interpolation technique to deform the s-rep to fit the boundary of the target object.

Akash Krishna

The second stage of the s-rep fitting is the refines the s-rep. This stage optimizes the positional and orientational fit of the s-rep. Following the refinement step, the s-rep fitting process is complete. Figure 3 provides visuals of the s-rep fitting process.

Once an s-rep is fit to a given 3-D object, the features of the fitted s-rep are used for shape analysis. These features include: spoke length, spoke direction, and skeletal point locations.

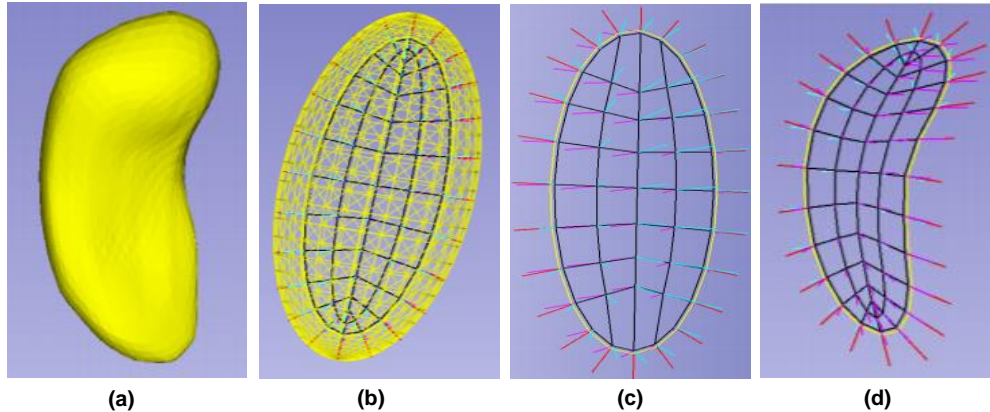
(a)   (b)   (c)   (d)

Figure 3:

(a): Mesh of the target object.
(b): Near ellipsoid shape after applying mean curvature flow.
(c): Computed s-rep of best fitting ellipsoid
(d): Final refined s-rep for target object

3. Statistics:

Before performing statistics, it is vital to Euclideanize and commensurate the features in order to improve classification performance. So, it is helpful to utilize principal nested spheres on direction vectors and apply the log to features with only positive values in order to Euclideanize them. Additionally, it is necessary to commensurate the data, and this can be done by normalizing the data by subtracting by the mean and dividing by the standard deviation.

Once you have Euclideanized and commensurated the features, there are various statistical algorithms to use to perform different tasks. One possibility, for classification, is to train an algorithm that discriminates the data between two classes. Distance Weighted Discrimination is a good method to utilize for this task. It is superior to that of support vector machines because support vector machines are prone to non-robustness due to data piling issues, especially for high dimensional data. DWD avoids this issue by using all sample points to determine the feature space direction that separates the classes.

4. Multi-Object Analysis:

There are many circumstances in which performing analysis of multiple structures jointly could be advantageous over studying structures individually. For example, it is found that neurodevelopmental processes affect multiple structures in the brain. Thus, when attempting to use shape analysis to classify whether an individual has some neurodevelopmental disorder like autism, it is valuable to study multiple objects jointly. The image on the right shows how Alzheimer's affects entire
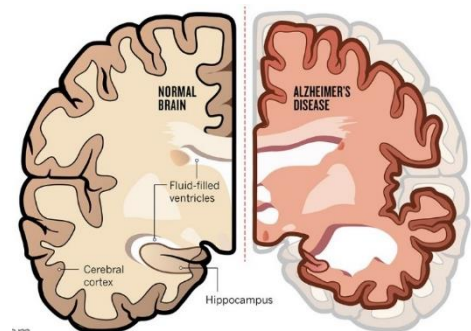

Figure 4: Image comparing the brain of a healthy individual vs. brain of individual with Alzheimer's disease

regions in the brain. Early analysis of the joint shape of multiple structures in the brain can serve to be useful in the early diagnosis of Alzheimer's disease.

In order to perform multi-object analysis, it is still necessary to fit the individual objects within the multi-object system with s-reps. It is desirable to ensure that the 3D s-reps of adjacent objects are consistent with each other. This consistency of s-reps is to make sure that the implied boundary of the two objects do not interpenetrate or pull away from each other. The s-reps are made consistent by ensuring collinearity of each pair of s-reps' spokes that share a boundary point on the shared boundary region.

Once the s-reps are fit to the objects in the multi-object region, we can concatenate features of all the s-reps and perform statistics on the joint representation of the entire region. Additionally, there are more features we can include in this joint representation. Our method includes information on the geometric relationship between adjacent objects. We propose fitting a 2D s-rep to the shared boundary region of adjacent objects and including its features in the joint representation of multi-object regions. Including the geometric relationship of adjacent objects can serve to provide additional features to hopefully improve classification accuracy.

5. 2D s-rep

Similar to the motivation of fitting s-reps to target objects in 3D, a 2D s-rep provides a geometric representation that optimizes first order and second order fit of the target object in 2D. The process of fitting a 2D s-rep to a 2D target object is similar to the same process in 3D. The 2D s-rep fitting process involves two stages. The first stage involves a curve evolution to smooth the boundary into that of an ellipse, calculating the 2D s-rep for the ellipse, and applying the inverse of the curvature flow to deform the s-rep to fit the 2D target object. The second stage refines the s-rep to optimize the positional and orientational fit of the s-rep.

2D s-reps can be fit to objects on flat spaces or curved spaces. I was the primary developer of the code of fitting 2D s-reps to 2D objects, and I will go over the procedure for both fitting a 2D s-rep in the following sections.

Flat s-reps:

Given a target object as seen in Figure 5, we may desire to fit a 2D s-rep to the object to perform statistical shape analysis. We begin be using a curvature flow to smooth the boundary of the object until it is near the shape of an ellipse. Each iteration of the curvature flow works by going through all the discrete points on the objects boundary, and fitting a circle through the current points and its two adjacent points. From this circle, we can obtain a curvature value and move the point along the direction to the normal of the curve by the curvature value. Once each point has been moved in the
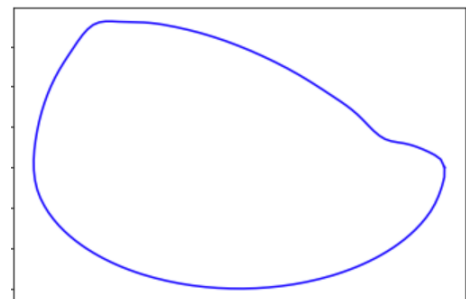

Figure 5: Boundary of 2D target object

direction of its normal, an iteration of the curvature flow is completed. It is important to keep track of all the local diffeomorphisms after each iteration of the curvature flow in order to be able to apply the inverse of this process later in the s-rep fitting. After every iteration of the curvature flow, the best fitting ellipse to the points along the boundary of the object is computed.

The best fitting ellipse is found by performing an Eigenanalysis on the the second-moment matrix about the center of mass of the points (Stojmenovic and Nayak, 2017). Once we have the best fitting ellipse, we calculate sum of the distances between the points on the boundary of the target object and their closest points on the boundary of the ellipse. Once the distances between the points are small enough, we can end the curvature flow. The target object following the curvature flow and its best fitting ellipse are shown in Figure 6.
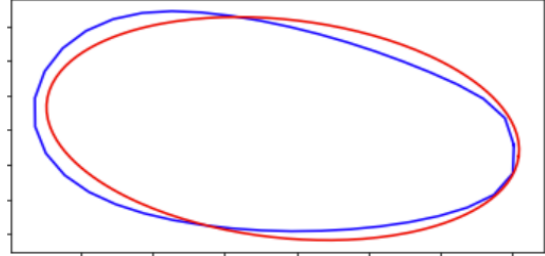
Figure 6: (blue): target object after curvature flow (red): best fitting ellipse

After the curvature flow, we calculate the s-rep for the ellipse. The s-rep for the ellipse is found by first sampling skeletal points along the major axis of the ellipse. The first and last skeletal points are found by moving along the major axis from both of the vertices of the ellipse towards the other vertex by the radius of curvature value. The rest of the skeletal points are found by sampling $m_1\cos(\theta)$ where m_1 is the length of the skeleton and θ is a constant value. After the skeletal points are sampled, we find the boundary points of the spokes by calculating the nearest point on the boundary of the ellipse from the skeletal point. Each skeletal point will have an up and down spoke, each of which will go from the skeletal point to the nearest point on the ellipse on its repective side. The calculated s-rep for a given ellipse is shown in Figure 7.
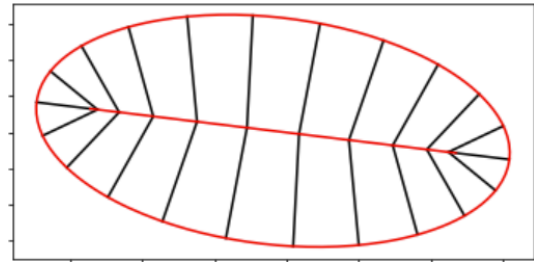
Figure 7: Best fitting ellipse with computed s-rep. (black): spokes of s-rep (red line): the red line is the spine of the s-rep

Following the calculation of the s-rep for the ellipse, we deform the s-rep to fit the boundary of the near-ellipse shape of the last diffeomorphism of the curvature evolution. The transformation of the skeletal and boundary points of the s-rep is conducted via Thin Plate Splines where the points on the ellipse are the source points and the points on the near ellipse are the target points. Now, we conduct an iterative process of deforming the s-rep into fitting the original target object's boundary. Using the local diffeomorphisms of the curvature flow, we have the K^{th} diffeomorphism providing the source points and the K-1th diffeomorphism providing the target points, and we continue to transform the s-rep. Once we are back the original boundary of the target object, we have obtained an initial s-rep fit for the target object.
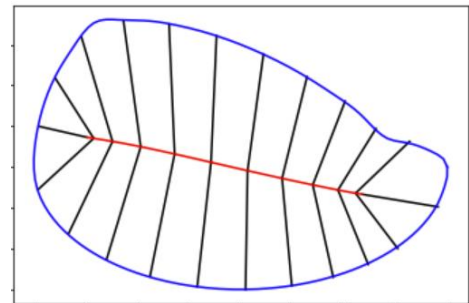
Figure 8: Initialized s-rep for target object

The previous steps were part of the Initialization portion of fitting an s-rep to a 2D target object's boundary. Following the Initialization, we refine the fit of the s-rep to capture first and second order properties of the object boundary. This stage utilizes an optimization technique to obtain an s-rep that minimizes deviations from Damon's medial conditions (Damon 2003).
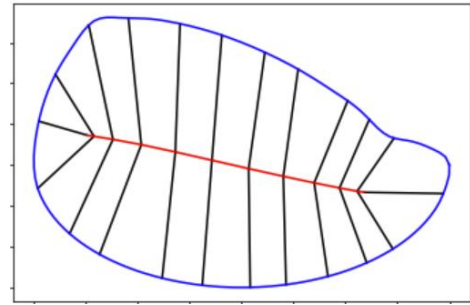

Figure 9: Refined s-rep for target object

The optimization function involves three penalties. The first penalty is the sum of the square distances from the spoke boundary points to the target object boundary, the second penalty penalizes non-orthogonality of the spokes to the target boundary, and the last penalty penalizes the crossing of spokes. Once we have optimized the fit to the target object, we have successfully obtained a flat 2D s-rep.

Curved S-reps:

Our method of fitting an s-rep to a 2D target object on a curved surface involves mapping that curved surface to a flat region, fitting a flat 2D s-rep to the mapped boundary, and mapping that s-rep back to the curved space of the target object. In context of our goal, we are looking to fit an s-rep to the shared boundary region between adjacent objects. In order to map the shared boundary region of the 3D target objects to a flat surface, we begin by mapping the shared boundary region onto the skeleton of the 3D s-rep of the primary object. The process of mapping a point on the boundary of an object onto its skeleton involves interpolating the spokes on the skeleton and finding the skeletal point of the spoke whose boundary point matches the target point on the boundary. Once the shared boundary region is mapped onto the skeleton, we sample the (U,V) coordinates along the

Figure 10: Points on the shared boundary region (black) get mapped to points onto the object's skeleton (white).

boundary of this mapped region. The (U,V) coordinates are found by utilizing the skeletal grid as a coordinate system. The mapped shared boundary region on the skeleton is shown in Figure 11.
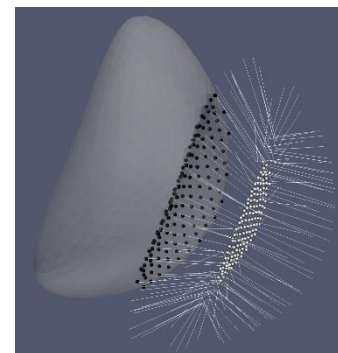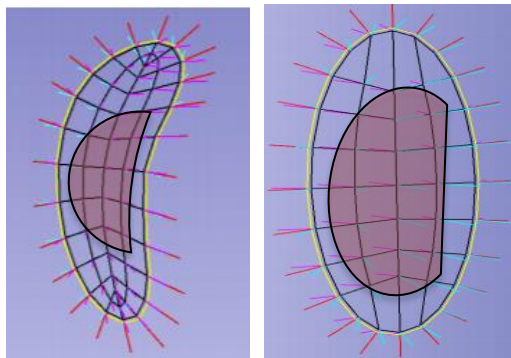

Figure 11: (left): shared boundary region mapped onto an adjacent object's skeleton. (right): shared boundary region mapped onto the adjacent object's best fit ellipse's skeleton

Once we have mapped the shared boundary region onto an object's skeleton and found the (U,V) coordinates, we now able to map this region to a flat surface. During the s-rep initialization step for the 3D object whose skeleton we mapped the shared boundary region onto, we would have found the best fitting ellipsoid for that 3D object and calculated the s-rep for the ellipsoid. The skeleton of the best fitting ellipsoid is flat, and we can map the region from the 3D object's skeleton onto the skeleton of the ellipsoid. We map this region by mapping points from their (U,V) coordinate on the 3D object's skeleton to their

respective (U,V) coordinate on the ellipsoid's skeleton. We now have a 2D boundary on a flat surface that we can use to fit an s-rep.To fit an s-rep for the region on the flat skeleton of an ellipsoid, we utilize the s-rep fitting process for flat surfaces as detailed in the previous section.

After computing flat s-rep, we need to map this s-rep back onto the curved surface of the shared boundary region. This involves obtaining the (U,V) coordinates of the skeletal and boundary points of the flat s-rep. We also sample points along the spokes of the s-rep and obtain the coordinates of those sampled points. Using the (U,V) coordinates we map the s-rep onto the skeleton of the 3D object. We then map the s-rep onto the shared boundary region by finding the boundary points of the spokes of the 3D object s-rep whose skeletal points share a point with the 2D s-rep. This completes the procedure of fitting an s-rep to a shared boundary region on a curved surface.
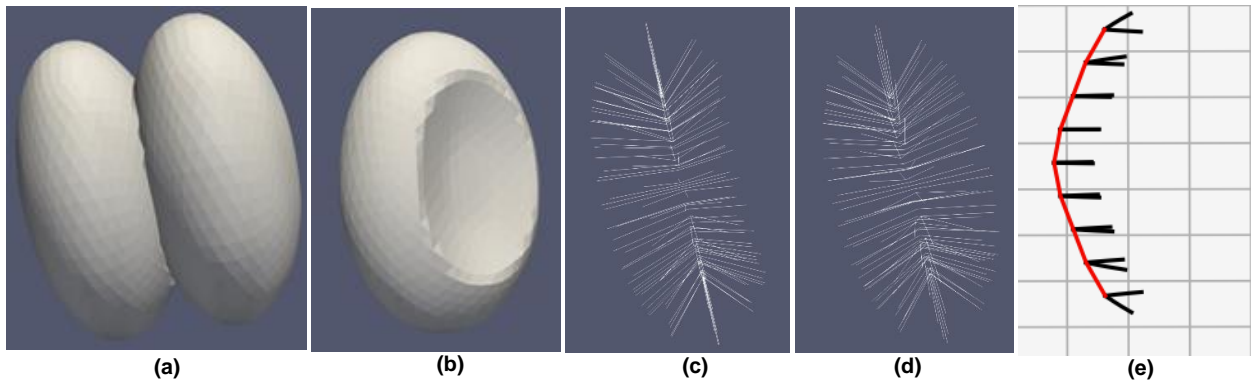

| (a) | (b) | (c) | (d) | (e) |

Figure 12: (a): Pair of adjacent deformed ellipsoids. (b): Indented region displays the shared boundary region on one of the ellipsoids (c): s-rep of left-most object. (d): s-rep of right-most object. (e): 2D s-rep of shared boundary region

6. Data Generation:

I have generated adjacent deformed ellipsoids to test the effectiveness of including an s-rep of the shared boundary region when performing classification tasks. I have a control group of pairs of adjacent ellipsoids with minimal bending and minimal streching. I have also generated test cases where the adjacent ellipsoids have both bent by the same amount, each lengthened by different amounts, or shifted in respect to the other. I attempt to see how well we are able to classify the control group ellipsoids from the pairs of ellipsoids with each category of deformation. In particular, we seek to compare how incorporating consistent s-reps and including the s-rep of the shared boundary region affects the classification accuracy.
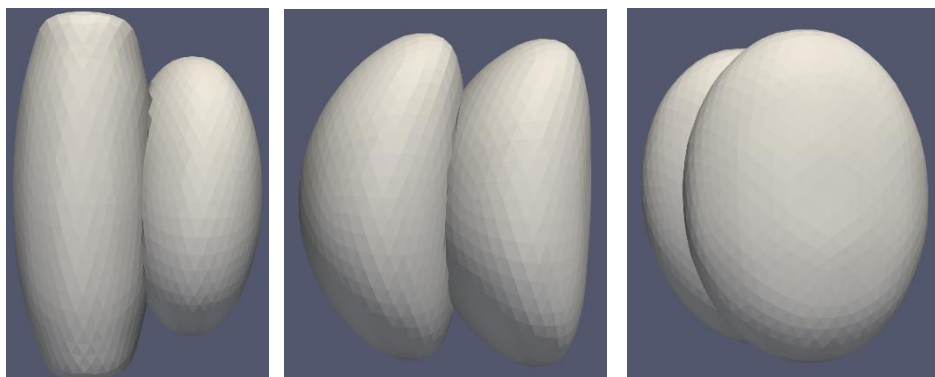

Figure 13: Pairs of generated deformed ellipsoids used for classification (left): Pair of ellipsoids where one is stretched. (middle): Pair of ellipsoids with bending. (right): Pair of ellipsoids where one has shifted

Akash Krishna

For the control group of ellipsoid pairs, each ellipsoid has axes lengths of 16, 9, and 4 and have their centers at point (0, 0, 1.5) and (0, 0, -1.5). The control group has minimal amount of stretching, bending, and shifting determined by a normal probability distribution. Stretching is done by lengthening the major axis by a factor. Bending is done by subtracting the z-component of a point by a bend factor times the x-component squared: $[z = z - (b*x^2)]$ where b is the bend factor. Shifting is done by moving the ellipse along the y-axis. The normal distributions for the amount of stretch factor, bend factor, and shift distance for each of the four classes is given in Table 1 below.

Normal Distribution of Deformations (mean, std)			
	Stretching Distribution	Bending Distribution	Shifting Distribution
Control	(0, 0.03)	(0, 0.015)	(0, 0.25)
Stretch	(0.1, 0.03)	(0, 0.015)	(0, 0.25)
Bend	(0, 0.03)	(0.035, 0.015)	(0, 0.25)
Shift	(0, 0.03)	(0, 0.015)	(0.7, 0.25)

Table 1: Displays the mean and the standard deviations of the normal distributions of deformations applied to each class of ellipsoid pairs. The stretching value is the factor that the major axis is stretched. The shift value is the amount shifted in the y-direction

For each of the four classes of pairs of ellipsoids, I have generated 30 test cases for each class and have fit two 3D s-reps for the adjacent objects and a 2D s-rep for the shared boundary region for each test case. The representation for each pair of objects consists of the two 3D s-reps for each individual object and a 2D s-rep for the shared boundary region. The features utilized from the 3D s-rep are the skeletal point locations, the spoke lengths, and the spoke directions. The features utilized in the 2D s-rep are the skeletal point locations, and the spoke points along evenly sampled distances along the curved spokes. The features from the 3D s-reps and the 2D s-reps are concatenated together and make up the feature set for a given sample. After obtaining the features for each of the samples, I euclideanized the features and used DWD to see the classification accuracy of the different representations.

Results and Analysis:

I tested the classification accuracy using different representations of the deformed ellipsoid pair. The first representation was simply the two 3D s-reps of the individual objects, the second representation was the two 3D s-reps after being made consistent with one another, and the third representation was the pair of consistent s-reps and the 2D s-rep of the shared boundary region. The classification accuracy for the three representations are displayed in Table 2.

Classification Accuracy			
	Stretch	Bend	Shift
3D s-reps	88%	84%	92%
consistent 3D s-reps	88%	85%	92%
consistent 3D s-reps + 2D s-rep	89%	85%	94%

Table 2: Displays classification accuracy for different representations of the adjacent ellipsoid pairs.

Akash Krishna

For the first row, using information about the two 3D s-reps only, including their relative positions, already a fair level of classification is achieved. The other two rows use sources of additional information characterizing the shape relationships of the two objects other than their relative positions.

The results in the second row do not clearly convey the effect of making the s-reps of the adjacent objects consistent by modifying spokes along the shared boundary to be collinear with each other. Using the consistent s-reps representation either led to no change or a slight improvement in classification accuracy. The consistency mapping requires interpolations which may lower the information about the shape of the object. The interpolations and the consistency mapping effects may counteract each other in a way that is poorly understood.

In all of the deformation classes adding the 2D s-rep appears maintain or improve classification accuracy over the modified consistent s-reps alone. Additionally, the combination of consistent s-reps with the 2D shared boundary s-rep always led to a higher classification accuracy compared to the unmodified s-reps.

There are many possible explanations for the wide range of outcomes of the classification accuracy. Firstly, there were only 45 samples for the different classes, so the lack of samples may not have led to accurate results.

For classifying the streched ellipsoids, the accuracy remained the same after utilizing consistent 3D s-reps over the typical 3D s-rep. We believe this is because of there was little information to be gained from the typical s-rep representation as the classification accuracy was already high. The accuracy goes up slightly after including the 2D s-rep of the shared boundary to the representation with the consistent s-reps. This is potentially because the 2D s-rep offered information on the shape of the shared boundary region not offered by the consistent s-rep like curvature and width.

For classifying bent ellipsoids from the control group, there was a small accuracy increase from the regular 3D s-rep to the consistent s-rep. The accuracy did not further increase after including the 2D s-rep of the shared boundary region. This may have occurred because the consistent s-reps already provide some insight into the shape of the shared boundary region, and the additional features provided by the 2D s-rep were not useful in discriminating the two classes.

Classifying the shifted ellipsoids had the highest gain in percentage after inclusion of consistent s-reps and a 2D s-rep of the shared boundary region. Using a representation of consistent 3D s-reps over the regular 3D s-rep did not have an increase in accuracy. However, including the 2D s-rep along with the consistent s-rep further increases the accuracy by 2% over using just the individually fit s-reps. This is likely due to the shifting of the ellipsoids significantly changing the shape of the shared boundary region.

Akash Krishna

Summary of Contributions:

The code for this research can be found here: https://github.com/AkashK23/MultiObjectShapeAnalysis
My contribution to this code base includes the following:
1. Making 3D s-reps of adjacent objects consistent with each other.
2. Computing the shared boundary and projecting it onto the (curved) skeleton of one of the objects.
3. Mapping the projection of the shared boundary on the curved skeleton onto the flat skeleton of an ellipsoid.
4. Fitting 2D s-reps to 2D objects on a flat surface.
5. Mapping the 2D s-rep on the flat skeleton back onto the curved skeleton and thence back onto the shared boundary.
6. Integrating all of these modules into a pipeline that allows for the simulations and analyses described in the experiment.

All of this code is available to be integrated with the web-resident Slicer SALT shape analysis toolkit.

I also conducted this experiment. I generated 120 ellipsoid pairs, used my code to deterime the consistent 3D s-reps and the curved 2D s-rep, and tied in PNS and DWD code to determine the classification accuracy.

Future Applications:

There are many possible applications of this work. I plan on applying my code to classify real disorders. Many neurological disorders like Parkinson's, Autism, Alzheimer's, etc. affect the shape of structures in the brain. Using real-life examples will provide a clear indication of the usefullness of including information on the geometric relationship of adjacent objects.

The usecase of modeling the geometric relation of adjacent objects could also be useful in the automatic segmentation of anatomical structures from medical images. In many cases, there is not a contrast in pixel intensity of adjacent structures, so having prior information as to the geometric relationship of the objects can serve to be useful in accurately segmenting objects.

Desired Future Improvements:

There are also many possible improvements and modifications to be made to the work done. For one, our method of obtaining the 2D s-rep of the shared boundary region involved mapping the boundary region onto the skeleton of the undeformed ellipsoid, mapping it to a flat surface, fitting an s-rep, and mapping it back onto the shared boundary. It could be useful to see how mapping this shared boundary region onto the other ellipsoid's skeleton affects the classification accuracy or if including both 2D s-reps into our representation has any impact on the classification accuracy.

Additionally, there are current limitations in our approach of obtaining the 2D s-rep. Our approach does not account for cases where the shared boundary region maps onto the top and bottom side

of the deformed ellipsoid's skeleton. A possible way to account for this would be to utilize the spherical topology of the (U, V) representation of the 3D s-rep.

Acknowledgements:

Akash Krishna

Citations:

Damon, J., 2003. Smoothness and geometry of boundaries associated to skeletal structures I: sufficient conditions for smoothness. Annales de l'Institut Fourier 53, 1941–1985.

Gorczowski, K., Styner, M., Jeong, J., Marron, J., Piven, J., Hazlett, H., Pizer, S. and Gerig, G., 2007. Statistical Shape Analysis of Multi-Object Complexes.

Junpyo Hong,∗ , Jared Vicory , Jörn Schulz, Martin Styner, J.S. Marron, Stephen M. Pizer, 2016. Non-Euclidean classification of medically imaged objects via s-reps.

Liu, Z., Vicory, J., Pizer, S., Damon, J. and Hong, J., 2020. Fitting unbranching skeletal structures to objects. *Medical Image Analysis*.

Milos Stojmenovic and Amiya Nayak, 2017. Direct Ellipse Fitting and Measuring Based on Shape Boundaries

Pizer, S. and Marron, J., 2017. Object Statistics on Curved Manifolds. *Statistical Shape and Deformation Analysis.*,.

*Pizer, SM, *J Hong, *J Vicory, *Z Liu, *J.S Marron (*principal co-authors), and H-Y Choi, J Damon, S Jung, B Paniagua, J Schulz, A Sharma, L Tu, J Wang. "Object Shape Representation via Skeletal Models (s-reps) and Statistical Analysis," in *Riemannian Geometric Statistics in Medical Image Analysis,* (X Pennec, S Sommer, and T Fletcher, eds.): 233-272, Academic Press, 2019.