

GEOMETRIC AND STATISTICAL MODELS FOR MULTI-OBJECT SHAPE ANALYSIS

Zhiyuan Liu

A dissertation submitted to the faculty at the University of North Carolina at Chapel Hill (UNC-CH) in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Computer Science in UNC-CH.

Chapel Hill
2022

Approved by:

Stephen M. Pizer

J.S. Marron

Martin A. Styner

James. N. Damon

Junier B. Oliva

©2022
Zhiyuan Liu
ALL RIGHTS RESERVED

ABSTRACT

ZHIYUAN LIU: Geometric and statistical models for multi-object shape analysis
(Under the direction of Stephen M. Pizer and J. S. Marron)

Shape analysis of multi-object complexes is important in many applications because it reveals additional information of interest over single-object shape analysis. For example, in medical applications where multiple structures in the human body often deform together, joint shape analysis of those interrelated structures facilitates robust and efficient algorithms. Specifically, shape correlation of functionally related structures allows us to understand the common underlying biological factors (e.g., disease). Also, beyond the within-object shape relations, between-object shape relations provide additional understanding of multi-object complexes.

Despite the need of multi-object shape analysis, this field has been challenged by many issues. For instance, shape variation is often coupled with pose and size variation between objects. Moreover, within-object shape variation is often coupled with between-object shape variation. These issues have prevented us from sufficiently understanding multi-object complexes.

To address the issues, this dissertation proposes geometric and statistical methods for joint analysis of multi-object complexes. In particular, I base my research on skeletal representations (i.e., s-reps) that are designed to provide intrinsic shape features with good correspondences. This dissertation improves the previous method fitting an s-rep to an object such that the fitted s-reps have desirable geometric and statistical properties. This improvement allows me to analyze intrinsic shape correlation between objects. To this end, this dissertation extends the existing statistical method to effectively extract joint shape variation, leading to a method called Non-EUclidean Joint and Individual Variation Explained (NEUJIVE). NEUJIVE shows notable robustness in analyzing multi-block non-Euclidean data with different variability. Last, to decouple within- and between-object shape variation, I develop non-branching linking structures for statistical analysis of between-object shape features. To capture geometric features that are insensitive to pose variation of multi-object complexes, this dissertation extends fitted local frames on s-reps to affine frames. The fitted local affine frames show special advantage because they free multi-object shape analysis from pre-alignment.

The driving problem of the proposed methods involves classifying and testing hypotheses on the shape of the hippocampus-caudate pairs between an autism group and a non-autism group. Also, this dissertation discusses other potential applications that can benefit from the proposed methods.

To my family and friends...

ACKNOWLEDGEMENTS

Complication is everywhere around us and within us. Both the Covid pandemic and wars bring terrible complication not only to the people in the center but also to the people far from it. During this complicated time, I come to the end of this dissertation. I am lucky that I have this opportunity to present my research. I am lucky because I met many people who helped me through difficult times and finally helped me to make this dissertation true.

I appreciate the great education and advice I had at UNC-CH. Especially, I would like to thank my principal advisors: Dr. Stephen M. Pizer (Department of Computer Science), Dr. J. S. Marron (Department of Statistics and Operations Research) and Dr. James N. Damon (Department of Mathematics). Moreover, I want to thank other members of my dissertation advisory committee: Dr. Martin A. Styner and Dr. Junier B. Oliver for their efforts on advising this dissertation.

In addition to the above advisors, there are other great people who are important to me and to this dissertation. I thank Dr. Z. Pei (from Duke University) for the encouragement when I lost my faith after the first oral exam. I thank Dr. S. Chang (Department of Radiation Oncology) for helping me to understand the purpose of life. I also thank Dr. S. Handa (Department of Public Policy) for offering me a summer research assistant position.

I thank my colleagues in my PhD program, including Dr. J. Schulz, Dr. J. Fishbaugh, Dr. J. Vicory, Dr. B. Paniagua and Dr. S. Jung. Moreover, I benefited from collaborations with my teammates, including M. Taheri, A. Sharma, A. Krishna, J. Hong, J. Prieto, J. Wang, M. Mostapha and others.

Most importantly, I thank all my family for their understanding. My parents are always supporting me. My wife brings me not only technical suggestions but also a lot of joy. As I said in the beginning, life is complex, filled with both happiness and sadness. It is the above people who make me feel that the past five years form one of the most valuable periods in my life.

TABLE OF CONTENTS

LIST OF TABLES	x
LIST OF FIGURES	xi
LIST OF ABBREVIATIONS	xv
1 Introduction	1
1.1 Problems and Challenges	1
1.1.1 Driving Problem: Classify Autism vs. Non-autism from the Shape of Multiple Subcortical Structures	1
1.1.2 Challenges in Anatomic Shape Analysis with Boundary Geometry	3
1.1.3 Challenges and Objectives in Analyzing Multi-Object Configurations (MOCs).....	4
1.1.4 Challenges from the Non-Euclidean Geometry of Shape Spaces.....	6
1.2 Proposed Methods	7
1.2.1 Improve S-reps Fitting Algorithm to Better Capture the Within-object Shape	8
1.2.2 Joint Analysis on MOCs using Non-Euclidean Statistics	9
1.2.3 Explicitly Modeling Between-object Shape with Non-branching Linking Structures.....	11
1.3 Thesis and Contributions	12
1.4 Overview of the Remaining Chapters.....	14
2 Background	15
2.1 Magnetic Resonance Imaging (MRI)	16
2.2 Materials: the Shape of Subcortical Structures from Infant Brain Images	17
2.3 Anatomic Shape Analysis	17
2.4 Differential Geometry in Shape Analysis	19
2.5 Theories and Methods of Skeletal Representations	20
2.5.1 Properties of the Blum Medial Axis	21
2.5.2 Shape Analysis as Functional Data Analysis	22

2.5.3	Generalized Skeletal Representations with Good Correspondences in a Population	24
2.5.4	Algorithms of Fitting Skeletal Representations	25
2.5.5	Summarize the Basics of Skeletal Representations.....	26
2.6	Interior Geometry of Objects	26
2.6.1	Skeletal Coordinates of an Interior Point	26
2.6.2	Local Orthonormal Reference Frames	27
2.7	Statistics in Euclidean Space	27
2.7.1	Geometry of Euclidean Space and Subspaces	28
2.7.2	Correlation Analysis of Multi-block Euclidean Data	29
2.7.3	Traditional Methods of Correlation Analysis of Multi-block Euclidean Data	30
2.7.4	Joint Analysis of Two-block Euclidean Data via AJIVE	31
2.8	Non-Euclidean Data Analysis	34
2.8.1	Non-Euclidean Statistics in Shape Analysis	34
2.8.2	Principal Geodesic Analysis (PGA).....	35
2.8.3	Principal Nested Spheres (PNS)	36
2.9	Supervised Binary Classification	37
2.9.1	Statistical Goal	38
2.9.2	Support Vector Machine vs. Distance Weighted Discrimination	38
2.9.3	Evaluation Metrics	39
3	Fitting S-reps to Non-branching Objects	41
3.1	Introduction.....	41
3.2	What is an S-rep?	42
3.3	Automatic Initialization Methodology	44
3.4	Interpolation	47
3.4.1	Interpolate Skeletal Sheets and Attached Spokes	48
3.4.2	Interpolate Spokes in a Crest Region	49
3.5	Refinement	50
3.5.1	Relax Blum Conditions	50
3.5.2	Refinement Objectives for Interior Spokes	51

3.5.2.1	Penalize Boundary Distances	51
3.5.2.2	Penalize Non-orthogonality to the Boundary	52
3.5.2.3	Penalize Self-overlaps within an Object	52
3.5.3	Refine Edge Spokes	53
3.5.4	Implementation.....	54
3.6	Experimental Analysis	54
3.7	Conclusions and Discussion	56
4	Non-Euclidean Joint and Individual Variation Explained for Multi-object Complexes.....	58
4.1	Introduction.....	58
4.2	Classify ASD vs. non-ASD with Subcortical Shape	59
4.3	S-rep Implied Boundary Points	61
4.4	Euclideanize Skeletally Implied Boundary Points	63
4.5	Extract Joint Shape Variation in Two-object Complexes	63
4.5.1	Alignment in PNS Score Space.....	64
4.5.2	Ranks in NEUJIVE.....	65
4.6	Experiments	66
4.6.1	Joint Components of Two-block Non-Euclidean Data.....	66
4.6.2	Joint Shape Analysis in Classification	69
4.6.3	Hypothesis Testing on the Hippocampus-Caudate Data	72
4.6.4	Classification.....	74
4.6.5	Shape Differences Between Groups in the Joint Variation Subspace	76
4.7	Broader Impact	77
4.8	Conclusions and Discussion	78
5	Correlation of Within- and Between-object Features	80
5.1	Introduction.....	80
5.2	CoWBO Features: Correlation of Within- and Between-Object Features	82
5.3	Modified Linking Flow Captures Between-object Shape Features	83
5.3.1	Damon-Gasparovic's Linking Structures	83
5.3.2	Modified Linking Flow	85

5.3.3	Link Vectors and Between-object Shape Features	87
5.4	Local Affine Frames on Skeletal Representations	87
5.4.1	Local Affine Frames	88
5.4.2	Within- and Between-object Shape Features in Local Affine Frames	89
5.5	Experiments	89
5.5.1	Visualization of Between-object Shape Features from Modified Linking Structures	90
5.5.2	Classify Autism vs. Non-autism with Two-object Complexes	91
5.5.3	Local Affine Frames Capture Differentiating Local Geometry	92
5.5.4	Comparing with Various Feature Spaces	93
5.6	Conclusions and Discussion	95
5.6.1	Geometric Models	96
5.6.2	Statistical Analysis	96
6	Conclusions & Discussion	97
6.1	Contributions and Thesis	97
6.2	Limitations and Future Work	100
6.2.1	Fitting S-reps to Objects of More Complex Shape	100
6.2.2	Benefit from Interior Geometric and Appearance Features of Objects	101
6.2.3	Geometric and Statistical Correlation for Multi-object Shape Analysis	102
6.2.4	Other Applications	103
	BIBLIOGRAPHY	104

LIST OF TABLES

4.1	Test performance using different features in classification	72
4.2	Hypothesis testing of ASD vs. non-ASD with different joint features	75
5.1	Average test AUCs from using affine frames.	93
5.2	Average test AUCs with features from two-object complexes	93
6.1	Test AUCs/SEMs using different interpolation methods.....	102

LIST OF FIGURES

1.1	The axial (left), sagittal (middle) and coronal (right) view of a T1-weighted brain image. The blue regions are the caudate nuclei, while the red are the hippocampi.	2
1.2	Illustration of problematic representations of boundary geometry. Top Left: a partially corrupted boundary mesh resulting from segmentation of a caudate. The corruption region is highlighted by the red circle. Bottom left: a non-corrupted shape of a caudate. Right: Selected correspondence points from SPHARM-PDM on two hippocampi. The two red dots pointed to with arrows form a pair of correspondent points.	3
1.3	An illustration of radial flow level surfaces of a hippocampus from different viewpoints. Top row (a-e): individual level surfaces at different radial distances from the skeleton. Bottom row (f-g): the combination of all level surfaces. The white surface is the boundary. Copied from (Pizer et al., 2021) with permission.	5
1.4	2D radial flow level surfaces implied by a skeletal model. Left: a boundary of a 2D object. Middle: a discrete skeletal model fit to the boundary, wherein the vectors from the skeletal points to the boundary are called spokes. The blue and the red are, respectively, two segments of the boundary related to the two sides of the skeletal model. Right: the level surfaces associated with the discrete skeletal model. The dots are points on the spokes.	5
1.5	An illustration of the problems in the previous research. In the left figure with dark background, the implied boundary results from Hong’s method having kinks (marked by the blue box). In contrast, the middle figure with dark background results from an s-rep fitted with the proposed method. The kinks are resolved as highlighted by the red box. In the rightmost figure, the fitted s-rep of an hippocampus shows dense sampling around the center of the object, as compared to the ends of the object.	9
1.6	Flowchart of fitting an s-rep to an object with a close boundary.	10
2.1	Left: a 2D object with non-branching topology except for a boundary bump induced by noise. The noise induces an extra branch on the skeleton (blue). Right: a fitted skeleton to the object that obtains the noise-free object’s branching topology. This skeletal model is associated with a slightly modified (dotted) boundary.	22
2.2	Illustration of different generic properties of distance functions $f_p(x)$ as the interior point p varies. (a) In an ellipse three interior points (three crosses in boxes) and a boundary point (blue square). As the boundary point moves along the ellipse boundary counter-clockwise, the distances between the boundary point and the three interior points are shown on the vertical axes in (b-d). (b) The distance function corresponding to the interior point in the top black box. (c) The distance function corresponding to the medial point in the red box. The dashed horizontal line indicates that the two local minima are identical for medial points. (d) The distance function corresponding to the interior point in the bottom black box. The differences of the two local minima between (b) and (d) are discussed in the text.	23
2.3	Obtain PNS scores by fitting nested spheres to n samples, starting from the initial hypersphere \mathbb{S}^{n-1} . Blue dots illustrate pre-shapes on hyperspheres. Each subsphere A_{n-l-1} is parametrized by its center and radius (i.e., v_l and r_l) and is associated with a score vector of length n . The concatenation of the $n - 1$ score vectors is taken as the Euclidean representations of the n samples.	36
2.4	The blue curve is an ROC curve resulting from a classification problem. The gray curve is the ROC curve from random guess with AUC 0.5.	40

3.1	Left: a 3D s-rep for a hippocampus. The grid connects skeletal points. The yellow curve is the skeletal edge. The vectors proceeding from the skeletal points are the spokes. Middle: the 3D implied boundary of this s-rep. Right: the surface mesh of this hippocampus being fitted.	43
3.2	An oblique view of an ellipsoid’s skeletal representation, where the skeletal sheet is discretized as a grid. The interior black grid is the discretized colocated skeletal sheets M_{\pm} . Each grid point on M_{\pm} is associated with a spoke, pointing to either the north (magenta) or south (cyan) side of the ellipsoid boundary. The yellow curve is the edge M_0 , from which red spokes point to the crest curve of the ellipsoid.	44
3.3	Left: Input a 3D shape of a hippocampus represented as a surface mesh. Middle: The overlay of the target shape (red) and a near-ellipsoidal shape after mean curvature flow (black). Right: the best fitting ellipsoid (yellow transparent mesh) to the black mesh in the middle figure and its discrete s-rep (consisting of cyan and magenta spokes, a discrete skeletal mesh and an edge curve).	47
3.4	(a-c) The result of the initial fit to a hippocampus from different perspectives. The gray transparent mesh is the input target surface mesh. The black grid in the middle of the object is a discrete representation of the skeletal sheet. The cyan and magenta spokes indicate spokes pointing to different sides. The red spokes are edge spokes. (d) The white spokes are the interpolated northern spokes of the initial fit. (e) The refined s-rep.	47
3.5	Interpolating spokes’ directions using the interior coordinates defined in (Pizer et al., 2021). Left: the interpolation method first interpolates spokes’ directions along $\widehat{\nabla}\tau_1$ directions. These interpolated directions are shown as red arrows. The base points of the arrows are interpolated via a Hermite patch as described in text. Right: The interpolation along $\widehat{\nabla}\theta$ direction. The interpolated spokes directions are shown as blue arrows.	48
3.6	Interpolating boundary points for spokes in the crest region after interpolating spokes on the smooth skeletal sheet. The white line segments are the spokes of the input s-rep. Left: Interpolation of crest points by TPS. The ends of primary edge spokes are the input of TPS, which outputs the interpolated crest points. Right: Interpolation of boundary points around the crest. I interpolate boundary point along each spline to form spokes around the crest region. See the text for more details.	49
3.7	Fit of an s-rep to a caudate nucleus in the dataset described in section 2.2. Top: the initial fit. Bottom: the refined s-rep. In the refined s-rep, the skeletal points of edge spokes (red line segments) are relocated according to curvatures on the boundary.	53
3.8	(a): The implied boundary of the initial fit to a hippocampus. (b): The heat map of unsigned distances measured locally between the initial implied and target boundaries. (c): The implied boundary of the refined skeletal representation. (d): The local distances (in mm) from the target boundary to the implied boundary of the refined skeletal representation.	55
3.9	The average and standard deviation of critical parameters	55
3.10	The histogram of the curvedness on the initial and refined implied boundaries and target boundaries of a typical example. I apply the logarithm on the count of boundary point plus 1, yielding the y-value. The curvedness resulting from equation (3.15) is the x-value.	56
3.11	The conventional mean curvature flow develops singularities (around the circled region) for a mandible. Left: the input surface mesh of the mandible. Right: the result from a few iterations of the flow.	57

4.1	The intuitive illustration of extracting the joint components with NEUJIVE. (a) Input data live in non-Euclidean pre-shape spaces. (b) Construct the shared variation subspaces (\mathbb{R}^r represented by the yellow and green arrows) within the score spaces (\mathbb{R}^n shown as the yellow and green planes). (c) Project the data to obtain the joint components, shown as brown dots. Here, r denotes the dimension of the joint variation subspace. (d) Map the joint components (shown as the brown dots) of X_k back to the pre-shape space via ϕ_k^{-1} for each k	60
4.2	Fit an s-rep to a hippocampus and derive implied boundary from the s-rep. (a): SPHARM-PDM representation of the segmented boundary of a hippocampus. (b): A near ellipsoid (black transparent mesh) that is diffeomorphic to the hippocampal boundary. The ellipsoidal s-rep is discretized, yielding skeletal points (black grid points) and discrete spokes (colored line segments). (c): The discrete s-rep fit to the hippocampus. (d): The implied boundary derived from the hippocampus s-rep.	62
4.3	Transform s-rep implied boundary points (red points) of a caudate nucleus (top row) and those of a hippocampus (bottom row) into pre-shapes that are represented by transformed points (blue points) via eq. (2.15). Despite the translation and scaling effects, the geometry of the boundary is preserved. All these points are presented on smooth the implied boundary shown as transparent volumes underneath.	64
4.4	Left column shows simulated two blocks of points on \mathbb{S}^2 . The data tend to lie along a small circle with joint location on the circle indexed by the colors of points. The middle column shows the joint structure (magenta points) estimated by AJIVE in \mathbb{R}^3 . This AJIVE joint variation follows a linear pattern. The right column shows the joint structure (magenta points) estimated by NEUJIVE. This NEUJIVE joint structure captures the actual joint variation pattern in the two blocks.	67
4.5	Joint and individual structures from NEUJIVE on two correlated non-Euclidean blocks. The first row shows the simulation where the joint structures are designed as circular variables along Small Circles (SC) in the two blocks (see eq. (4.5)). Moreover, the two blocks share the Equal Noise Level (ENL) in which ε_k share the same standard deviation. The second row shows a different simulation study. In this simulation, the joint component is designed to be a Great Circle (GC) in block 1. But this component follows a small circle in block 2. Moreover, the two blocks have Different Noise Levels (DNL). On each sphere, the black dots are simulated data. The magenta dots are joint structures from each algorithm. The blue dots are individual structures from NEUJIVE and residual structures from AJIVE, respectively.	69
4.6	Root Mean Square Error (RMSE) between the actual joint structures and NEUJIVE estimated joint structures as increasing the noise level. Left panel: Simulated data as increasing the standard deviation $\sigma(\varepsilon_k)$ of noise from 0.01 to 0.05 and to 0.1. The two circular blocks are of different radii, i.e., $a_1 = 0.25$ while $a_2 = 0.35$. Right Panel: RMSE of the NEUJIVE joint structures of the two blocks as increasing standard deviation of noise. The red and the blue curve are, respectively, the RMSE of the block 1 and the block 2.	70
4.7	Joint analysis of two blocks of shape data formed by 2D landmarks from male gorillas' skulls. The left column shows the two blocks: block 1 (top) and block 2 (bottom). Each block has two groups (shown as the blue circles and the red crosses). I detail the generation of the two blocks in the text. The right two columns show the heat maps of the PNS matrices (the top row), the joint structures (the second row), the individual structures (the third row) and the residuals (the last row). The color represents the entry values of each matrix. See the text for detailed interpretation.	71

4.8	The cumulated Explained Variance Ratio (EVR) of the first 35 eigenmodes vs. the number of eigenmodes. The left and the middle panel show the EVR of shape features from single objects – the hippocampus and the caudate, respectively. The right panel shows the EVR of concatenated shape features of the hippocampus-caudate pairs. Best viewed in color images.	73
4.9	The results from the permutation test of concatenated PNS blocks (left) and the test of NEUJIVE joint components (right). The green dots on the bottom of each figure represent the test Mean Differences (MDs) from the permutations. The black curve represents the estimated density function of those MDs.	74
4.10	Results from the classification using joint structures. (a) The AUCs in classifying ASD vs. non-ASD with NEUJIVE joint features vs. the chosen initial ranks. (b) The histograms of test AUCs using joint structures of the hippocampus from different methods. (c) The average test AUCs from various methods using the joint structures of hippocampi (left part) and caudate nuclei (right part).	76
4.11	The point-wise differences (in millimeters) between a reconstructed ASD complex and a reconstructed non-ASD complex. The reconstructed ASD and non-ASD complexes have more difference in the red regions than in the blue regions.	77
5.1	Simulated links between two 2D objects. Compared to Damon-Gasparovic’s linking structure (top), the modified linking structure (bottom) changes the link directions to avoid self-linking, as highlighted in the yellow boxes.	84
5.2	Fitted local affine frames on the skeleton of a hippocampus. The skeleton is shown as the orange surface. At each skeletal point, there is an affine frame consisting of a $\nabla\tau_1$, $\nabla\theta$ and normal vector (see the zoomed picture above). The vectors of $\nabla\tau_1$ and $\nabla\theta$ are not necessarily orthonormal vectors.	89
5.3	Linking features of 5 caudate-hippocampus pairs: (a)-(e). The transparent surfaces are the caudate (above) and the hippocampus (below). Top row: spoke extensions (in mm) from a boundary to the external linking surface. The heat maps on the skeletons show the values of spoke extensions. Bottom row: link directions of link vectors at discrete skeletal points, shown as the arrows in magenta. The skeletons are shown as blue surfaces.	90
5.4	Left: Hypothesis testing using the correlated shape features of SPHARM-PDMs of the hippocampus and caudate. Middle: Hypothesis testing using the correlated shape features of the hippocampal s-rep features and the link features. Right: Average test AUCs over varying training data size. The error bar at each point shows the standard error of the mean.	92
5.5	Various shape features of a two-object complex consisting of a caudate (top row) and a hippocampus (bottom row). The left column shows boundary landmarks from SPHARM-PDM (Styner et al., 2006). The middle column shows spoke features from (Liu et al., 2021) – including coordinates of skeletal points, unit directions of spokes and lengths of spokes. The right column shows deformation field from an ellipsoidal skeleton to the skeleton of a target object (the caudate or the hippocampus). The arrows τ_1 and θ indicate the deformation in two directions.	94
6.1	Top: a caudate with partially corrupted boundary (highlighted by the red circle) from SPHARM-PDM (Styner et al., 2006). Bottom: The level surfaces of the radial flow of an s-rep of a caudate. These level surfaces describe the interior geometry of the object.	97
6.2	Deforming a mandible shape using conformalized mean curvature flow.	101

LIST OF ABBREVIATIONS

AJIVE	Angle-based Joint and Individual Variation Explained
ASD	Autism Spectrum Disorder
AUC	Area Under Curve
CCA	Canonical Correlation Analysis
CoWBO	Correlation of Within- and Between-Object
CPNS	Composite Principal Nested Spheres Analysis
DiProPerm	Direction-Projection-Permutation
d-s-rep	discrete skeletal representation
DWD	Distance Weighted Discrimination
fMRI	functional Magnetic Resonance Imaging
MOC	Multi-Objecxt Configuration
NEUJIVE	Non-EUclidean Joint and Individual Variation Explained
PAA	Principal Angle Analysis
PCA	Principal Component Analysis
PDMs	Point Distribution Models
PGA	Principal Geodesic Analysis
PNS	Principal Nested Spheres
RMSE	Root of Mean Square Error
SPHARM-PDM	SPHERical HARMonic Point Distribution Model
SVD	Singular Value Decomposition
TPS	Thin Plate Spline

CHAPTER 1: INTRODUCTION

Shape correlation and between-object relations are critical ingredients in multi-object complexes. Joint analysis of multi-object complexes allows us to obtain important shape correlation, which leads to efficient and robust algorithms (e.g., segmentation, classification). Moreover, multi-object shape analysis leverages geometric information within and between objects, producing additional insights about the population, as compared to the analysis on single objects.

This dissertation focuses on medical image analysis, aiming at shape variation involving pairs of nearby structures in the human brain. Since structures in the human brain can often functionally and spatially interrelate, the joint analysis of these related structures can produce rich information about biology. To this end, Section 1.1 discusses the driving problem and the challenges to be addressed. Further, Section 1.2 introduces the geometric and statistical methods that address those problems. Specifically, I propose a) a method of fitting skeletal representations to single non-branching objects¹, b) a statistical method for extracting shape correlation from multi-object complexes and c) geometric methods to capture between-object relations. In Section 1.3 I summarize the contributions and the thesis of this dissertation. Finally, in Section 1.4 I give an overview of the following chapters.

1.1 Problems and Challenges

1.1.1 Driving Problem: Classify Autism vs. Non-autism from the Shape of Multiple Subcortical Structures

Autism Spectrum Disorder (ASD) is a childhood-onset neurodevelopmental disorder that affects about 1.4% of the world's population (Baird et al., 2006; CDC, 2016). Symptoms of ASD include social deficits, communication difficulties, repetitive behaviors, sensory difficulties and unusual abilities (Deserno et al., 2018). To date, there is no effective biological diagnostic test that detects ASD. However, brain functional images of ASD versus non-ASD children may have implications for diagnosis. These implications can lead to early diagnosis as well as early intervention.

While today the evidence on cortical structures has been particularly indicative of ASD (Van Rooij et al., 2017), it has been found that ASD symptoms are also associated with dysfunctional subcortical structures (JL et al., 2012; Hitti and Siegelbaum, 2014; Nicolson et al., 2006). In particular, the hippocampus and the caudate, the two major subcortical structures of concern in this dissertation, have been extensively implicated in the development of ASD.

The hippocampus (shown as the red region in fig. 1.1) plays an important role in memory and social behaviors. Previous studies in animals (Ferguson et al., 2000) and humans (Corkin, 2002; Murphy et al., 2012) suggest that deficits

¹In this dissertation, a non-branching object is also referred to a slabular object, whose length is greater than its breadth that is greater than its width.

in the hippocampus could lead to severe problems in social memory. Such problems presumably cause behaviors similar to autistic symptoms. Also, previous research in functional analysis of the human brain (Eilam-Stock et al., 2016) suggests that the caudate (shown as the blue regions fig. 1.1) is one of the major subcortical structures that facilitate desired movements. Deficits in the caudate are associated with unwanted and/or repeating movements, which are often observed on autistic subjects.




Figure 1.1: The axial (left), sagittal (middle) and coronal (right) view of a T1-weighted brain image. The blue regions are the caudate nuclei, while the red are the hippocampi.

Inspired by the association of ASD symptoms and subcortical functions, previous research (Nicolson et al., 2006; Murphy et al., 2012; Richards et al., 2020; Katuwal et al., 2015) has investigated the association between the development of ASD and the morphological changes of the hippocampus and of the caudate. Unfortunately, the findings from various research projects are inconsistent. This inconsistency can be due to the wide range of severity of ASD and the wide variety of sub-types of ASD, and it can also be due to complex confounding factors such as age and sex. For example, Groen et al. (Groen et al., 2010) measured the hippocampus volume in a small group of adolescents (12 through 18 years) including 23 ASD and 29 non-ASD cases. The researchers found that the left hippocampus is significantly enlarged in the ASD group compared with the control group. In contrast, Eilam-Stock et al. (Eilam-Stock et al., 2016) find a volume decrease of the hippocampus in a group of ASD cases with age range of 6-65 years.

Similarly, the results are also inconsistent from the research on the shape of the caudate in the development of ASD. Sears et al. (Sears et al., 1999) measure the volume of various subcortical structures by integrating areas inside contours in successive image slices of respective structures. They find an increased volume of the caudate in subjects with ASD. Especially, Hollander et al. (Hollander et al., 2005) find that the right caudate volume is significantly larger in autistic group than in control group. Yet, Wolff et al. (Wolff et al., 2013) find that the association between the caudate volumes and repetitive behaviors is different across various sub-types of ASD (including fragile X and idiopathic autism).

The geometric and statistical models in the above research cannot decouple (1) the shape variation due to the development of ASD with (2) the shape variation due to demographic differences. Specifically, volumes are used as the

major shape measurements that are insensitive to local morphology of objects. Therefore, it is necessary to develop geometric and statistical models for comprehensive analysis. Importantly, the geometric and statistical models should allow for joint analysis of multiple objects (i.e., the hippocampus and the caudate). In particular, a desirable geometric model should provide a) good positional and orientation correspondence across a population, b) rich and expressive local geometric information and c) robustness against noise. In the following, I discuss the challenges in obtaining such a desirable geometric model. Moreover, assuming a desirable geometric model, I discuss the challenges in analyzing multi-object shape with existing statistical methods.

1.1.2 Challenges in Anatomic Shape Analysis with Boundary Geometry

Boundary representations give straightforward understandings of the shape. Given a segmented 3D structure from an MRI image, its boundary can be represented in a discrete form as a triangular mesh. Such a representation can capture localized boundary geometry. However, this local boundary geometry can be problematic in anatomic shape analysis, as discussed below.

First and foremost, the boundary of an anatomic structure is often problematic due to noise in images and/or errors in segmentation. The top left figure in fig. 1.2 gives an example caudate in which the tail of the caudate is seriously corrupted, compared to a smooth caudate boundary that is shown in the bottom left. Such corruption can be hardly smoothed in a post-processing of the boundary. Thus, the local geometry around the corrupted tail is useless in statistical analysis.




Figure 1.2: Illustration of problematic representations of boundary geometry. Top Left: a partially corrupted boundary mesh resulting from segmentation of a caudate. The corruption region is highlighted by the red circle. Bottom left: a non-corrupted shape of a caudate. Right: Selected correspondence points from SPHARM-PDM on two hippocampi. The two red dots pointed to with arrows form a pair of correspondent points.

Second, it is difficult to define shape correspondences on 3D boundary surfaces. Shape correspondence is one of the most fundamental elements in many shape analysis tasks, e.g., shape registration, retrieval, understanding, etc. (see (Van Kaick et al., 2011) for a good review). Previous research (Cates et al., 2007) proposes establishing good correspondences across 3D anatomic objects based on landmark-based boundary models, also known as Point Distribution Models (PDMs). In that work the ordering of the boundary points implies the correspondence across samples. To be specific, consider a set of n single objects drawn from a population of objects. Each boundary point on a shape is embedded in a d -dimensional ambient Euclidean space. In most relevant work, $d \in \{2, 3\}$. A PDM of an object is constructed by sampling m d -dimensional points on the boundary. Each of the points is called a *correspondent point* (Cates et al., 2006). Dryden and Mardia (Dryden and Mardia, 2016) refer to the set of the m points as the *configuration*.

Typically, the correspondence among these landmarks are based on mathematical equivalence. Researchers expect the variation of the correspondent points can shed light on local shape variation. However, the mathematical correspondences often deviate from anatomic correspondences. An example is shown in the right panel of fig. 1.2. The dots (in both blue and red) on the two hippocampi result from SPHARM-PDM (Styner et al., 2006), and the red dots are mathematical correspondent points. Because these mathematical correspondent points are placed at inconsistent anatomic locations, their positional variation among samples is not helpful in revealing the geometric variation of the tail of the hippocampus.

The above two issues are the consequence of ignoring interior geometry and high order geometry of an object. On the one hand, the interior geometry can provide robust shape descriptions of an object, focusing on how an object is smoothly filled. In this regard, the boundary is understood as the bound of a set of smooth level surfaces (see fig. 1.3) that fill the object's interior. On the other hand, the high order geometry is sensitive to regions of high curvatures and thus allows recognition of vertices, edges and other salient anatomic landmarks.

Using both interior geometry and high order geometry of objects, skeletal models have been developed to provide shape features with good correspondence. A skeletal model consists of a skeletal sheet and a mapping from the skeleton to the boundary. As shown in figs. 1.3 and 1.4, a skeletal model is associated with a set of non-intersecting surfaces that fills the interior of an object. Moreover, the skeletal model can capture salient anatomic landmarks that have extreme curvatures because the process of developing a skeletal model utilizes high order geometry of an object. To establish good correspondence, skeletal models are consistently discretized across a population; see the details in Section 1.2.1.

1.1.3 Challenges and Objectives in Analyzing Multi-Object Configurations (MOCs)

Structures in the human body, especially the brain structures, often functionally interrelate and interact with each other. As such, the development of ASD can cause simultaneous morphological changes of multiple structures. Thus, it is necessary to analyze MOCs not only focusing on individual shape variation but also on the shape correlation. The




Figure 1.3: An illustration of radial flow level surfaces of a hippocampus from different viewpoints. Top row (a-e): individual level surfaces at different radial distances from the skeleton. Bottom row (f-g): the combination of all level surfaces. The white surface is the boundary. Copied from (Pizer et al., 2021) with permission.




Figure 1.4: 2D radial flow level surfaces implied by a skeletal model. Left: a boundary of a 2D object. Middle: a discrete skeletal model fit to the boundary, wherein the vectors from the skeletal points to the boundary are called spokes. The blue and the red are, respectively, two segments of the boundary related to the two sides of the skeletal model. Right: the level surfaces associated with the discrete skeletal model. The dots are points on the spokes.

individual variation accounts for the shape variation respective to each single object, while the shape correlation can provide insights related to common underlying biological factors (i.e., ASD).

In this dissertation, each MOC is formed by a fixed number of functionally correlated or spatially nearby objects, i.e., the hippocampus and the caudate. In this setup, I discuss two major challenges in analyzing MOCs below.

First, it is challenging to build geometric models for MOCs. A simple way to represent an MOC is to concatenate the single shape models since there exist a variety of well-developed single shape models. The concatenated shape models can separately correspond across MOCs. However, it has been shown that such analysis can be biased by components of different variability in MOCs. For example, Gollmer et al. (Gollmer et al., 2012) concatenated single-object models in an automatic segmentation task. They found that the simple concatenation of the shape models (of the liver and the

spleen) leads to poorer performance than separate segmentation of the two organs. Because the between-object pose variation dominates the variation of MOCs, the researchers proposed decoupling pose variation with the concatenated shape models. As a consequence, the segmentation performance is improved by a large margin.

Though removing the between-object pose can increase the robustness of multi-object shape analysis, it is necessary to understand the geometric relations between an object of interest and its nearby objects; namely, the between-object shape is important in anatomic shape analysis. Nonetheless, it is challenging to define geometric correspondences for the between-object shape across a population because a small change of single-objects' poses can cause a dramatic change of between-object relations (see e.g., (Gasparovic, 2012)). For instance, two adjacent objects can have inconsistent abutting regions across a population of two-object complexes. These inconsistent abutting regions can imply largely different geometric relations across the population.

Second, it is challenging to extract joint shape variation using existing statistical methods. To capture the joint shape variation among subcortical structures, Akhoundi-Asl et al. (Akhoundi-Asl and Soltanian-Zadeh, 2007) applied Principal Component Analysis (PCA) to the concatenated single shape representations of the putamen, thalamus and caudate. However, the joint shape variation is often overwhelmed by the individual shape variation and thus can hardly be perceived by PCA. Later, Schwarz et al. (Schwarz et al., 2010) propose constructing a joint shape space with correlated (eigen-)modes of shape variation. The experiments show that the target object is better segmented in the joint shape space that takes into account multi-object shape restrictions, as compared to using single-object segmentation methods. However, statistical modeling of joint shape variation is often challenged by high dimensionality of feature spaces of MOCs. Statistics in a high dimensional space require a huge amount of data so that the learned statistical model will not overfit data (see Chapter 2). In anatomic shape analysis the data size is often limited due to expensive acquisition of the shape data. Therefore, a good method for multi-object shape analysis should be data-efficient and reliable. As explained in the next section, understanding the geometry of data space can help to develop such an efficient algorithm.

1.1.4 Challenges from the Non-Euclidean Geometry of Shape Spaces

It is useful to understand the geometry of shape spaces in order to develop an effective and efficient solution for multi-object shape analysis. The literature (Kendall, 1989; Dryden and Mardia, 1998) has laid out the concept of shape spaces that provides a firm mathematical basis for statistical endeavors. Because the shape of an object is invariant under similarity transformations and because there exists no notion of “zero shape”, the shape space is often understood as a non-Euclidean space. For instance, PDMs of objects can be transformed to data on a unit hypersphere (Dryden and Mardia, 1998).

One of the most important goals in analyzing shape is to estimate the mean shape and to characterize the shape variation of an object around the mean. The estimation of both the mean and the variation relies on the distance metric between the shape data. Riemannian geometry generalizes the Euclidean distance metric to shape spaces (Le et al., 1993). This generalization allows us to measure distances between any two shape data using the generalized inner product. However, the generalization needs assumptions on the geometry of the data space. For example, Fletcher et al. (Fletcher et al., 2004b) proposed the generalized Fréchet mean, which minimizes the overall squared geodesic distances from the mean to samples in Hilbert spaces. Moreover, the researchers used geodesic curves to characterize a distribution of objects' shape around the mean shape, yielding a statistical method called Principal Geodesic Analysis (PGA). The underlying assumption in PGA is that the shape are distributed along some geodesics, with these geodesics all based at the Fréchet mean in the shape space.

With different assumptions on the shape space, other statistical approaches also showed strengths in learning shape distributions for a specific set of objects. Pizer et al. (Pizer et al., 2019) provide a comprehensive review on those non-Euclidean statistical methods. Among these methods, Huckemann et al. (Huckemann et al., 2010) proposed Geodesic PCA (GPCA) to generalize PCA to quotient spaces. They pointed out that the pre-determined mean that relies on projection onto Euclidean space may not be a good representation of the data due to the curvature of the shape space. As a result, the best fitting geodesic is not necessarily passing through any pre-determined mean. More recently, Jung et al. (Jung et al., 2012) found that non-geodesic curves can sometimes better represent the principal components than geodesics. Based on this observation, the researchers proposed Principal Nested Spheres (PNS) that can fit data with either non-geodesic curves or geodesics, depending on the distributions of the data. PNS assumes that the data live on a unit hypersphere. Thus, the distance metric can be obtained using spherical geometry. The Euclidean representations from PNS lead to good performance in statistical tasks as classification (Hong et al., 2016) and hypothesis testing (Schulz et al., 2016).

It is challenging to understand the geometry of the data space of multi-object data. It can be problematic to assume that the data space of MOCs is singularity-free and smooth by concatenating single-object shape models. As mentioned above, a small pose change of a single object can lead to a dramatic change of between-object relations. The analysis of MOCs should decouple the within-object shape with the between-object relations.

1.2 Proposed Methods

To address the challenges mentioned in Section 1.1, I develop geometric and statistical models for more effective, interpretable and robust multi-object shape analysis. Specifically, Section 1.2.1 introduces the improvement on fitting skeletal representations (i.e., s-reps) to single non-branching objects. Section 1.2.2 introduces the proposed statistical method that I call *Non-EUclidean Joint and Individual Variation Explained* (NEUJIVE). NEUJIVE aims to jointly

analyze MOCs from the perspectives of the individual shape variation and especially of the joint shape variation. The resulting joint shape variation is sensitive to common underlying biological factors (i.e., ASD) yet are insensitive to the different variability among multiple objects. Section 1.2.3 introduces geometric methods that capture between-object shape features with good correspondence. Specifically, I discuss non-branching linking structures that aim to produce smooth between-object shape features. These linking structures allow comprehensive analysis of an object including within-object shape features and its relations with nearby objects.

1.2.1 Improve S-reps Fitting Algorithm to Better Capture the Within-object Shape

A skeletal model of an object consists of a skeleton and connections from the skeleton to the boundary of an object. Such a model is promising because i) it can capture a rich set of geometric features within an object, ii) it can provide better correspondence across objects than boundary models and iii) it leads to intrinsic shape features that do not require pre-alignment.

Specifically, a skeletal model captures not only the boundary but also the interior geometry of an object. The connections from the skeleton to the boundary are formed as vectors, called *spokes*, defined on the skeleton (Pizer et al., 2019). Taking derivatives of spokes yields insights on how an object tapers toward its tail as well as how that object bends. Moreover, a smooth spoke field maps a smooth skeleton to level surfaces of some radial distance function (e.g., Damon’s radial distance (Damon, 2003)). Therefore, a smooth spoke field leads to smooth interior geometry of an object (see Chapter 2 for more details). In fig. 1.3, the smooth boundary (i.e., the level surface at $\tau = 1$) that is associated with a skeletal model is referred to as the *implied boundary*.

To develop robust skeletal representations with good correspondences for single objects, this dissertation improves an automatic method due to Hong (Hong, 2018) that fits skeletal representations (also called s-reps) to non-branching objects. By fitting s-reps to objects’ boundaries, the skeletal topology is consistent across a population despite boundary noise. Because this dissertation focuses on slabular objects such as the hippocampus and the caudate, the fitted s-reps should be sheet-like without branches. Compared with Hong’s algorithm, this dissertation overcomes three major issues, as follows.

First, this dissertation fits s-reps with good geometric properties. Hong’s method essentially deforms an s-rep of an ellipsoid to fit an object through a reverse diffeomorphism of the boundary. Such a deformation only focuses on two level surfaces, i.e., the skeleton and the implied boundary. As a result, the deformation can result in non-smooth interior geometry, which can be observed on the implied boundary (see the left figure in fig. 1.5). This dissertation resolves this issue by taking into account the smoothness of interior geometry of an object.

Second, this dissertation utilizes higher-order geometry of an object in the fitting process to produce good correspondence across a population. In fitting s-reps, the proposed method not only encourages the boundary match

(i.e., the implied boundary is close to the segmented boundary) but also encourages the orthogonality between spokes and the boundary. By doing so, the spokes' directions can capture local shape measured by boundary curvature.

Third, this dissertation reparametrizes s-reps to capture geometric properties at regions of high curvature. The rightmost figure in fig. 1.5 shows a resulting s-rep from Hong's method, in which the skeletal points are contracted toward the center of the object (i.e., a hippocampus), ignoring geometric features around the ends of the object. To solve this issue, this dissertation utilizes ellipsoidal coordinates (see Chapter 3) to more densely sample the spokes around the ends of the object.




Figure 1.5: An illustration of the problems in the previous research. In the left figure with dark background, the implied boundary results from Hong's method having kinks (marked by the blue box). In contrast, the middle figure with dark background results from an s-rep fitted with the proposed method. The kinks are resolved as highlighted by the red box. In the rightmost figure, the fitted s-rep of an hippocampus shows dense sampling around the center of the object, as compared to the ends of the object.

To implement the above improvements, this dissertation develops a two-stage pipeline that fits s-reps with good geometric and statistical properties. Fig. 1.6 illustrates the pipeline. The initialization stage generates s-reps that have skeletons of consistent topology. On the skeleton I sample the skeletal points (and thus the spokes) using an ellipsoidal coordinate system, which allows dense spokes around the ends of the object. Moreover, because the refinement stage resolves crossing spokes within the object the fitted s-rep can produce smooth level surfaces of Damon's radial flow.

As shown in fig. 1.6, the refinement algorithm focuses not only on the geometric properties of the discrete spokes from the initialization stage but also on the geometric properties of interpolated spokes. To this end, this dissertation improves the interpolation method due to Vicory (Vicory, 2018) to obtain arbitrarily dense spokes from an s-rep, which is reparametrized using the ellipsoidal coordinates.

1.2.2 Joint Analysis on MOCs using Non-Euclidean Statistics

Both joint and individual shape variation are of interest in the analysis of MOCs. Joint shape variation captures jointly varying shape features that are affected by common factors (e.g., a disease), while individual shape variation explains significant shape variation specific to each object other than the joint shape variation. Despite the achievements

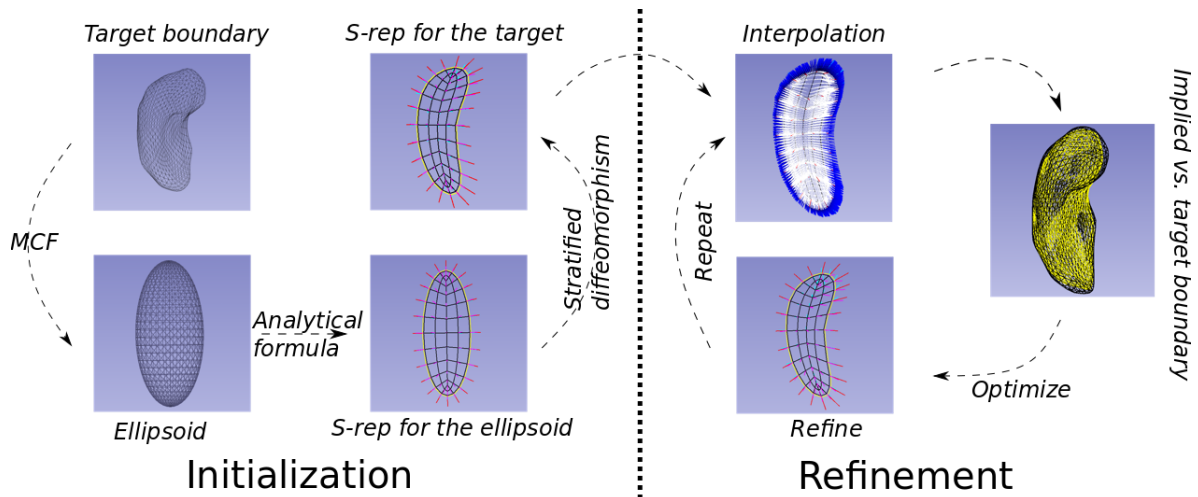


Figure 1.6: Flowchart of fitting an s-rep to an object with a close boundary.

in shape statistics and those in multi-block data analysis, there previously existed no effective method that extracts joint and individual shape variation from MOCs.

Because of the invariance properties of shape data under similarity transformations, shape statistics often utilize Riemannian geometry to efficiently analyze shape variation of single-object complexes, as discussed in Section 1.1.4. However, it is difficult to utilize the geometry of the data space for MOCs. Some previous methods represent an MOC by the concatenation of single shape representations, assuming the data space in question is smooth. For example, to obtain integrated shape information from MOCs, Pizer et al. (Pizer et al., 2003) presented one of the first multi-organ extensions of s-reps by computing statistics of all the organs taken together. Later, Jeong et al. (Jeong et al., 2006) incorporated inter-organ neighboring effects to the model. In this framework, the variation of an organ is decomposed into two parts: a) within-object variation that describes the variation of the organ itself and b) between-object variation that captures the between-object interactions. However, such an analysis based on concatenation of shape models can hardly be generalized to other datasets where the between-object shape overwhelmingly varies compared to the within-object shape in the population.

This dissertation regards multi-object shape analysis as a problem of multi-block non-Euclidean data analysis. In this regard, it is promising to solve the problem by (1) mapping each non-Euclidean block to a Euclidean block via methods in shape statistics and then (2) utilizing existing Euclidean methods to jointly analyze multiple Euclidean blocks. Recently, Feng et al. (Feng et al., 2018) have proposed a statistical method, called Angle-based Joint and Individual Variation Explained (AJIVE), that extracts the joint and individual variation of multiple Euclidean blocks. This method shows notable effectiveness and robustness in HDLSS tasks. Therefore, this dissertation develops a statistical method based on AJIVE for multi-object shape analysis, combining proper Euclideanization techniques (i.e., PNS) with AJIVE.

The proposed method, called Non-Euclidean Joint and Individual Variation Explained or NEUJIVE, is designed for extracting joint and individual shape variation from MOCs where the variability across multiple objects can be largely different. In essence, NEUJIVE can be seen as an extension of AJIVE. This extension accounts for non-Euclidean properties of each object. Specifically, given landmark-based shape representations of MOCs, NEUJIVE first removes the positional and size information of each object, yielding multiple blocks of spherical data. Considering the spherical property of shape data after removal of translation and scale (Dryden and Mardia, 2016), NEUJIVE uses PNS (Jung et al., 2012) to obtain the Euclidean representations of each object. Last, NEUJIVE performs AJIVE on the multiple Euclidean blocks to obtain a) joint shape variation, which explains how the shape of nearby objects jointly varies and b) individual shape variation, which accounts for the shape variation specific to each object. By combining PNS and AJIVE, NEUJIVE can effectively extract joint and individual shape variation that are insensitive to different variability across objects.

This dissertation verifies the method NEUJIVE by using the extracted joint shape variation features to classify ASD and non-ASD samples. Each sample consists of s-rep implied boundary points of the hippocampus and of the caudate. The experiments show that the joint shape features of the hippocampus and the caudate improve the classification performance. Additionally, due to good interpretability of the results from NEUJIVE, the joint shape variation from multi-object complexes can lead to interesting findings with respect to the development of ASD.

Beyond the above joint analysis of the within-object shape in MOCs, it is also of interest to jointly analyze within- and between-object shape about an anchor object. The within-object shape characterizes features within the boundary, while the between-object shape captures relations between the anchor object and its neighbors as discussed below.

1.2.3 Explicitly Modeling Between-object Shape with Non-branching Linking Structures

Between-object shape features are important in MOCs because these features characterize how the shapes of nearby objects interrelate. It is important to have between-object shape features with good correspondence across a population for statistical analysis. The between-object shape features of interest include positional and geometric relations between multiple objects (Bossa and Olmos, 2007) – the positional relations include relative pose and orientation between any two objects, while the geometric relations include local shape captured using geometric properties such as boundary curvature.

The variation of relative pose often overwhelms the within-object shape variation in MOCs. Many previous methods (e.g., (Gorcowski et al., 2007)) proposed aligning MOCs to remove/decouple the variation of relative pose from the total variation of MOCs. However, the alignment also ignores anatomically important relations between objects. In (Damon and Gasparovic, 2017) the researchers proposed representing each object with a skeletal representation. The between-object shape can be explicitly captured by linking structures from skeletal representations in an MOC. Such

linking structures can capture both the positional and geometric relations between objects. Especially, the linking structures can also capture geometric relations between objects when the objects are abutting (Krishna et al., 2022), in which the linking structures capture the geometry of the shared boundary. However, Damon-Gasparovic’s linking structures are sensitive to relative pose between objects, leading to inconsistent generic properties of the between-object shape (see (Damon and Gasparovic, 2017) for a list of generic properties). Such inconsistent generic properties make it difficult to have correspondence for statistical analysis.

To address the above difficulties, this dissertation develops a modified version of Damon-Gasparovic’s linking structure. This modified linking structure aims to avoid inconsistent generic properties (e.g., branches) yet to retain sensitivity to positional and geometric relations between objects. The modification can be understood as a process of smoothing the boundaries of objects such that the smoothed s-reps consistently link with each other. Due to this modification, the resulting between-object shape features have consistent generic properties and good correspondence across a population. Thus, the modified linking structure is useful in statistical analysis.

In practice, it is difficult to jointly align multiple objects (Hong, 2020). To extract geometric features that are invariant to rigid transformations yet are sensitive to local deformations, this dissertation fits local affine frames via a deformation-based method. As a result, the fitted local affine frames capture rich local geometric features from the interiors of objects, which are useful in classifying the shape. Particularly, fitting local affine frames allows us to capture local geometry of an s-rep that is invariant to rigid transformations. Also, this dissertation uses local affine frames to capture between-object shape features that are invariant to rigid transformations.

The joint analysis of within- and between-object shape features about an object yields comprehensive shape analysis of the object. For example, the development of ASD alters the shape of the hippocampus as well as the relation between the hippocampus and its nearby objects (e.g., the caudate). To achieve a comprehensive shape analysis of this object complex, I jointly analyze within- and between-object shape features about the hippocampus and its relation to the nearby caudate using NEUJIVE, yielding CoWBO (Correlation of Within- and Between-Object) shape features. While the between-object shape can vary much more dramatically than the within-object shape, NEUJIVE can produce robust CoWBO shape features regardless of the different variability of the two.

1.3 Thesis and Contributions

Thesis. *Joint shape variation and between-object shape relations can be effectively extracted from multi-object complexes by deliberately combining geometric methods which capture within- and between-object shape features invariant to rigid transformations, and statistical methods which focus on shape correlation in multi-object complexes.*

Specifically, following properly relaxed Blum conditions, the fitting process results in more satisfactory s-reps with desirable geometric and topological properties for single objects. The proposed Non-EUclidean Joint and Individual

Variation Explained method (NEUJIVE) can explain multi-object shape variations with notable effectiveness and robustness. The modified linking structures provide rich and useful geometric features of multi-object complexes, including within- and between-object shape features. To verify the thesis, I conduct experiments on the hippocampus-caudate pairs in 6 months olds, classifying the ASD vs. the non-ASD group. The CoWBO shape features gives the best classification performance. I summarize the major contributions of this dissertation as follows:

Contributions in methodology:

1. I propose a refinement algorithm that takes into consideration the smoothness of the interior geometry. This algorithm improves the automatic s-reps fitting method due to (Hong, 2018) so that the fitted s-reps admit smooth interior level surfaces of radial flow.
2. I propose a method of interpolating s-reps using skeletal coordinates. This interpolation leads to smooth spokes' directions on the skeleton. I show the advantages of the proposed interpolation in modeling between-object links.
3. I propose fitting an affine frame field for s-reps. The fitted affine frame field captures local geometry of objects' interiors. Moreover, the geometric features with respect to local affine frames are invariant to rigid transformations, which is desired in the analysis of multi-object complexes.
4. I propose a modified version of Damon-Gasparovic's linking structures. My proposed method produces external linking axes of consistent topology across a population. Then I extract between-object shape features with good correspondences from the modified linking structures. Finally, I demonstrate the between-object shape features (i.e., link lengths and link vectors' directions) to be helpful in improving classification performance in two-object complexes.
5. Because ASD can change the shape of multiple brain structures, I propose utilizing correlated shape features of the hippocampus and caudate pairs in classifying the ASD vs. non-ASD group. These correlated shape features should relate to the development of ASD. The experiments in this dissertation show that the correlated shape features from the hippocampus and caudate pairs improve the classification performance.
6. To reveal the correlated shape features, I propose combining the power of a Euclideanization technique (i.e., PNS) and the power of an integrative analysis technique (i.e., AJIVE), yielding a powerful statistical tool called NEUJIVE. NEUJIVE yields robust estimation of the correlation in multi-object complexes, allowing me to extract the correlation of within- and between-object shape features from the hippocampus-caudate pairs.

Contributions in engineering:

1. I implement an automatic s-rep fitting pipeline based on Hong's dissertation. Different from Hong's definition of a discrete s-rep (d-s-rep), I re-parametrize the d-s-rep such that two ends (i.e., the head and the tail) of a single

object are better represented by taking advantage of ellipsoidal coordinates. The ellipsoidal coordinates allow for denser sampling in highly curved regions, as compared to sparser sampling in flat regions. The implementation is included in the SlicerSALT package (Vicory et al., 2018).

2. I implement NEUJIVE as a MATLAB package. This package allows for the analysis of multi-object landmarks using NEUJIVE. Also, the toy examples provided in this package demonstrates the effectiveness and robustness of NEUJIVE in analyzing blocks of different variability.
3. I implement the geometric methods in a Python package *shanapy*, including i) the implementation of the proposed non-branching linking structures given a pair of objects, ii) the improved method of s-reps fitting to single non-branching objects and iii) the method of fitting local affine frames for s-reps.
4. I provide tutorials, user manuals and development notes for the above tools. These documents help users to analyze the data in their research.

1.4 Overview of the Remaining Chapters

This dissertation is organized as follows. Chapter 2 discusses relevant methods to my research and details the ASD data that is used in this dissertation. Chapter 3 details the improved s-rep fitting methods including initialization, interpolation and refinement. The experiments demonstrate that the resulting s-reps have good geometric properties. Moreover, these s-reps can well capture the geometry of objects. Chapter 4 discusses the statistical method developed for explaining multi-object shape variations. Using the proposed method, I show that the joint shape variation features of the two-object complexes are useful in classifying the ASD group vs. the non-ASD group. Chapter 5 presents the geometric methods that capture between-object relations from multi-object complexes. The methods include modified linking flow and fitted local affine frames. I evaluate the geometric and statistical properties of the proposed methods in that chapter. Chapter 6 summarizes the contributions of this dissertation as well as the limitations of the methods. Moreover, I discuss potential directions for future research.

CHAPTER 2: BACKGROUND

In this chapter, I discuss the background related to my research. In particular, this chapter includes the following topics:

1. The MRI modalities that produce the images in which the anatomic objects studied were segmented (Section 2.1).
2. The ASD data that is related to the driving problem of this dissertation (Section 2.2).
3. Computational challenges in anatomic shape analysis (Section 2.3).
4. Basic differential geometry in representing the shape of an object (Section 2.4).
5. Theories and methods of skeletal shape models (Section 2.5).
6. Representation of the interior of an object (Section 2.6).
7. Statistical methods relevant to this dissertation in Euclidean spaces (Section 2.7).
8. Statistical methods relevant to this dissertation for non-Euclidean data (Section 2.8).
9. The classification methods and metrics used in this dissertation (Section 2.9).

Anatomic shape analysis starts from segmentation of anatomic structures in medical images. The segmentation often relies on recognizing structures in various medical imaging modalities. For example, T1- and T2-weighted MRI images are two commonly used modalities in neuroscientific research. While brain structures look different in the two modalities, it is a common practice to segment brain structures using the two modalities together, especially for underdeveloped infants brains. Section 2.1 reviews the key properties of the two modalities. In Section 2.2 I detail the subcortical shape data that is used in this dissertation.

Section 2.3 discusses the objectives of statistical analysis on the segmented anatomic shape. On the one hand, anatomic shape analysis ought to be sensitive to local morphological changes because the local morphology variation often provides important indications of pathological development. On the other hand, the analysis ought to be robust against noise and artifacts in capturing the variation of anatomic shape.

Sections 2.4 to 2.6 introduce potential methods to approach the above objectives based on the following ideas. First, Section 2.4 introduces mathematical tools in differential geometry that are useful in capturing boundary geometry of a single object. Second, Section 2.5 discusses theories and algorithms relating to skeletal representations, which allow us to capture geometric features of an object with good correspondence. Specifically, I focus on two types of skeletal representations: the Blum medial axis and s-reps. Third, Section 2.6 introduces the previous research on the

interior geometry of an object with a fitted s-rep. As shown in Section 1.1.2, objects' interiors can provide more robust geometric features than objects' boundaries.

Regarding statistical analysis of multi-object shape, it is important to incorporate the geometry of the space in which the shape data live. In this regard, Section 2.7 discusses how the relevant statistical methods are built upon the geometry of a Euclidean space. Compared to the Euclidean statistics, it is more relevant to this dissertation to analyze data in a curved manifold with Riemannian metric. Section 2.8 discusses the existing statistical methods that explicitly utilize the manifold structure of data to characterize the distribution of the data in a non-Euclidean space.

I evaluate the proposed methods in classifying the ASD vs. the non-ASD group with the shape of the hippocampus and the caudate. In Section 2.9 I describe the classification framework including the statistical objective, the relevant methods and the evaluation metrics.

2.1 Magnetic Resonance Imaging (MRI)

In neuroscientific research, structural MRI images play a critical role in analyzing morphological development of brain structures. The image contrast between brain structures and background tissues allows for the recognition and segmentation of structures. However, since the intensities of the structures are different in various modalities of MRI, it is necessary to understand the mechanism of MRI for recognizing and analyzing the structures of interest in images.

MRI takes advantage of the different magnetic resonance properties of protons in various tissues. Specifically, the scanner activates protons in the tissue in the form of water and macromolecules (e.g., proteins, fat, etc.) in the human body. Then these activated hydrogen protons relax and release signals by i) turning their magnetic directions to the direction of the static magnetic field (B_0) and ii) dephasing the spinning dipoles (Hashemi et al., 2012). The released signals are captured for imaging. Both density of these protons and tissues-specific magnetic properties are the sources of the image contrast.

T1 and T2 relaxation time are the two magnetic properties that are relevant to this dissertation. T1 refers to the time constant that is determined by the relaxation time an activated proton realigns back to B_0 , while T2 refers to the time constant that is determined by the relaxation time of dephasing the spinning protons. Accordingly, images that maximize the contrast from T1 (resp. T2) are called T1-weighted (resp. T2-weighted) images¹.

Therefore, tissues with different composition look different in T1- and T2-weighted images. For example, since water has long T1 relaxation time and short T2 relaxation time, the regions filled with water (e.g., cerebrospinal fluid) look dark in T1-weighted images but bright in T2-weighted images. Moreover, because gray matter in the adult human brain contains more cell bodies (e.g., neurons and glial cells) than white matter, the gray matter appears as darker gray in T1-weighted images than the white matter. For underdeveloped infants' brains, however, the composition of tissues

¹Technically, the maximization involves tuning the repetition time (TR) and echo time (TE). Simply put, shorter times can maximize the contrast between tissues but can also disregard a lot of signals from tissues with longer T1. A reader is referred to (Hashemi et al., 2012) for more details.

are so complicated that the image contrast is not strong enough for segmentation. Therefore, both T1- and T2-weighted images should be used in segmenting infant's brain structures, as described in the following.

2.2 Materials: the Shape of Subcortical Structures from Infant Brain Images

This dissertation conducts experiments on the image dataset that was a part of National Institutes of Health-funded study of Autism. The dataset is referred to as the Infant Brain Imaging Study (IBIS) (Wolff et al., 2012).

The participants were scanned at the ages around 6, 12 and 24 months while they were naturally sleeping. In this dissertation, I focus on 176 participants of this study who are at high familial risk (HR) of ASD because they have family members who are already diagnosed with ASD. At the age of 24 months these participants were assessed by their behaviors using the Autism Diagnostic Observation Schedule (Lord et al., 2000). Among the participants, there are 34 infants (HR-ASD) who were clinically diagnosed with ASD and 142 infants (HR-NonASD) who did not meet criteria for ASD.

The scanned images of those participants include T1-weighted (T1w) and T2-weighted (T2w) images. All the images have the same spatial resolution of $1.0 \times 1.0 \times 1.0 \text{ mm}^3$. In segmenting the brain structures, both the T1w and T2w images for 12 month scans were used to enhance the contrast because of the under-development of an infant's brain (see section 2.1). The segmented brain structures in 12 month scans were then deformed to generate surfaces for 6 month scans using within-subjects registration (Avants et al., 2008); see (Kim et al., 2016) for more details.

Among the subcortical structures, the hippocampi and the caudate nuclei in both hemispheres of the human brain were found to be associated with ASD. Hence, I focus on the left hippocampus and the left caudate nucleus in this dissertation. From the above segmentation, SPHARM-PDM surfaces have been fitted to the left hippocampus and the left caudate (Mostapha et al., 2017), producing a triangle mesh with 1002 boundary points for each surface.

2.3 Anatomic Shape Analysis

In general, statistical analysis of anatomic shape is challenged by (1) how to represent geometric properties of anatomic structures and (2) how to characterize the geometric variation of a population.

A 3D anatomic structure can be represented by various computational models, including an atlas representation (Toga et al., 2006; Qiu et al., 2010b), a mesh (Durrleman et al., 2014), a point cloud (Cates et al., 2006; Gutiérrez-Becker et al., 2021), a distance map w.r.t. the boundary (Golland et al., 2005), a skeletal-based model (Pizer et al., 2003; Pouch et al., 2020), and an implicit function (Golland et al., 2000; Nain et al., 2007; Kim et al., 2021; Reuter et al., 2006). Compared to general shape analysis, anatomic shape analysis ought to address local morphological changes of structures in the human body (Dryden and Mardia, 1998) because these often indicate systematic differences due to biological processes (e.g., diseases). For example, (Nicolson et al., 2006) suggested that autism could be associated

with subtle local reduction of hippocampal sizes. Therefore, a good anatomic shape representation should be capable of capturing such local morphological changes.

However, an overemphasis of capturing local geometric features can sacrifice the robustness of shape representations. As suggested above, segmented structures often have noise and artifacts (see e.g., fig. 1.2 left). A shape model that strictly represents boundary geometry (e.g., the Blum medial axis, as discussed in the following section) can unexpectedly overfit a noisy boundary. Such a shape model represents individual objects quite well yet loses generalizability to unseen shape.

Towards achieving a good trade-off between sensitivity and generalizability, Tu et al. (Tu et al., 2018) have compared the sensitivity and generalizability of various shape models. Among these models, skeletal representations (s-reps) show better statistical properties than landmark-based models (e.g., SPHARM-PDM (Styner et al., 2006)) because the s-reps (1) faithfully capture robust geometric features of individual anatomic structures and (2) provide good correspondences in a population.

Moreover, not every geometric feature is of interest in anatomic shape analysis. Typically, translation, orientation and size of anatomic structures are regarded as irrelevant information to the shape. The existing research removes these irrelevant information via similarity transformations on each single shape. For example, Procrustes alignment (Dryden and Mardia, 2016) aligns a population of objects by performing similarity transformation on each object’s boundary.

To be specific, given a population of objects each of which is represented by a Point Distribution Model (PDM), the alignment minimizes the sum of squared Euclidean distances between each transformed shape and the estimated mean. Let $x \in \mathbb{R}^{3 \times m}$ be a PDM that consists of m landmarks in \mathbb{R}^3 . The overall distance d of the population to the Procrustes mean μ is

$$d = \sqrt{\sum_{i=1}^n |R_i x_i \beta_i + T_i - \mu|} \quad (2.1)$$

where $i = 1, \dots, n$ indexes a sample in the population. In the notation, R_i is a rotation matrix, $\beta_i > 0$ is a scale parameter, and T_i is a translation matrix. The matrix μ is the estimated mean shape of the population. For the shape that have small variations, which is the case for every object in the dataset, the minimization of d^2 converges to a unique solution. Finally, the distance from the Procrustes mean μ to every sample in the population can be used to characterize the distance in shape space; see e.g., (Dryden et al., 2019).

Positions of the boundary points can be regarded as low-order (zeroth-order) geometry of an object. Because anatomic shape analysis can benefit from rich geometric information including curvature on the boundary, this dissertation advocates to take into account high-order geometric features. In this regard, differential geometry gives differential entities that can be used to characterize the local geometry of higher order, as discussed below.

2.4 Differential Geometry in Shape Analysis

Differential geometry treats surfaces as differentiable structures on which local geometry can be characterized by differential entities. Curvature, for example, is an important differential entity that can be determined through differential operations (i.e., second derivatives). From a geometric perspective, measures of curvature are used to describe how the normal vectors locally vary on a surface². This local variation of the normal vectors gives a notion of the local shape on a surface.

Though boundaries of objects are “almost all” curved from a global view, an open *neighborhood* gives rise to the notion of a flat local patch that allows differentiation. As an example, take a 2D surface M that resides in a 3D space, i.e., $M \subset \mathbb{R}^3$. A neighborhood of a point $p \in M$ is defined as a sub-region U on the surface M containing the point p . This sub-region U is indeed the intersection of the surface and an open ball that has a center p and radius $r > 0$, i.e.,

$$U := M \cap B_r(p) \quad (2.2)$$

where $B_r(p) = \{x \in M \mid d_M(x, p) < r\}$ is the notation of the open ball, and $d_M(\cdot, \cdot)$ represents the distance metric (e.g., Riemannian metric) posed on the surface. The ball $B_r(p)$ is open (i.e., $d_M(x, p) < r$) such that the differentiation is well-defined everywhere in the neighborhood.

The differentiation in a neighborhood U leads to various differential entities that characterize local geometry. The *shape operator* is a well-known example of such differential entities. By definition (O’Neill, 2006), the shape operator is a linear mapping from a tangent space to itself; namely³,

$$S(v) = -\nabla_v N \quad (2.3)$$

where ∇_v denotes the directional derivative w.r.t. a given velocity direction v through p . Essentially, given a velocity vector $v \in T_p U$ on the surface, the mapping $S(v)$ produces the difference of normal vectors along v , which also lives in the tangent space $T_p U$ because the derivative of a unit vector field (i.e., N) is a vector field orthogonal to that unit vector field. In case of the future confusion with the radial shape operator, this shape operator defined in eq. (2.3) is called the differential shape operator.

In fact, key entities resulting from the shape operator can be independent of the choice of basis for the tangent space. On a smooth 2D surface, a shape operator can be represented as a 2-by-2 matrix that consists of two linearly independent vectors. Moreover, because the shape operator is self-adjoint, it has real eigenvalues and orthogonal

²In this section, I focus on the theories for smooth surfaces. These theories also apply in discrete geometry, although the implementation differs between various surface representations.

³The minus sign in eq. (2.3) is due to convention; see (O’Neill, 2006).

eigenvectors (O'Neill, 2006). The two real eigenvalues are called principal curvatures, while the corresponding two eigenvectors are called principal directions. When a velocity direction happens to be in one of the principal directions at p , the normal vectors purely swing in that direction at the rate of the corresponding principal curvature.

As the two eigenvalues of a shape operator can be either positive or negative, the sign of their product (called *Gaussian curvature* (Koenderink, 1990)) indicates whether the surface is locally convex/concave (positive) or hyperbolic (negative). In addition, the sign of the average of the two eigenvalues (called *mean curvature*) is useful to distinguish local convexity from concavity. By the convention that normal vectors are always pointing outward on an object with a closed boundary, a negative mean curvature indicates local convex shape. Chapter 3 details a method that utilizes mean curvature on a surface to deform a closed surface of spherical topology to an ellipsoid.

2.5 Theories and Methods of Skeletal Representations

Regardless of the noisy boundaries of anatomic structures, the interior geometry ought to give robust and insightful understanding of the shape of an object. However, the statistical analysis of the interior geometry is challenged by defining interior correspondences across a population. To address this challenge, a skeleton-based representation described below can provide good correspondences for interior regions of objects. In this section I discuss the mathematics of two cases of skeletal representations, namely, the Blum medial axis (Blum and Nagel, 1978) and generalized skeletal representations.

Section 2.5.1 discusses important properties of the Blum medial axis. Unfortunately, the state-of-art methods of constructing the Blum medial axis of non-branching objects can have inconsistent branching structures. This inconsistency is due to the sensitivity to the nuanced geometric features (e.g., a small bump) on the boundary. Thus, a population of objects can have poor correspondence using the Blum medial axis.

Section 2.5.2 provides an analogy between a 2D Blum medial axis and a shape function. This analogy helps to understand the importance of consistent generic properties across a population of Blum medial axes. Such consistency allows us to establish good correspondence across a population of objects. A reader may skip this section if it is clear that skeletal representations with consistent topology can give good correspondences and thus are beneficial in anatomic shape analysis.

To address the inconsistent topology in Blum medial axes, general skeletal representations have been researched and applied by (Damon, 2003, 2004; Pizer et al., 2003). Thanks to the generalization, these skeletal representations with consistent branching topology give better correspondences across the population, as discussed in Section 2.5.3.

Based on the theories, many methods have been proposed to produce skeletal representations for anatomic shape analysis. Section 2.5.4 discusses the previous algorithms of fitting skeletal representations and the limitations of those algorithms.

2.5.1 Properties of the Blum Medial Axis

The Blum medial axis contains a set of points that are centers of circles (in 2D) or spheres (in 3D) that have two or more tangencies with an object's boundary. Given a continuous boundary, the Blum medial axis can be derived via a method of geometric flow, called "grassfire" flow. Specifically, all surface points simultaneously flow along the inward normal vector field until they meet at the medial place. Thus, the result medial axis can be regarded as a set of singular points of that flow. This grassfire flow gives a mapping from medial points to boundary points, i.e.,

$$\phi : B \mapsto M \quad (2.4)$$

where M denotes a medial axis and B denotes a boundary. Any boundary point $x \in B$ has a corresponding medial point $\phi(x) \in M$.

A *spoke* is a vector that points from a medial point to its boundary point. The spoke captures the relations between the boundary point x , the medial point $\phi(x)$ and other interior points between the two. Analogous to a normal vector field defined on a smooth boundary, a spoke field is defined on the Blum medial axis.

Different from a smooth boundary, a medial axis can be divided into strata (i.e., subsets) (Damon, 2003). These strata together should form the same branching topology as the boundary. A slabular boundary, for example, should be associated with a medial axis with no branches. In a non-branching medial axis, those strata can be classified into the following categories⁴.

1. A stratum with smooth points (denoted as A_1^2) that are centers of circles or spheres which have exactly two tangent points with the boundary. In discussing spokes associated with two sides of an object, I denote this smooth stratum with either M_+ or M_- , meaning that each smooth point is associated with two spokes pointing to opposite sides of the object.
2. A stratum with edge points (denoted as A_3) that are centers of spheres which are tangent with the boundary at the extreme curvature (crest) points. This stratum bounds the stratum of smooth points. In a 3D medial axis, the edge points form a closed curve, referred to as the *edge curve* in the following text.

The spoke fields defined on the above strata satisfy the following Blum conditions as a result of the grassfire flow.

1. The two spokes at a smooth point are of the equal length. This length is the radius of the circle or sphere that is tangent to the boundary.
2. The angle between the two spokes at a smooth point is bisected by the tangent vector at that point.
3. The spoke from an edge point is on the limiting tangent plane at the edge.

⁴Each category is essentially an equivalent class in terms of the generic properties discussed in Section 2.5.2

4. Every spoke is orthogonal to the boundary at the tip of the spoke. This condition is also referred to as the *partial Blum condition*.

With these properties, the Blum medial axis can well capture the individual boundary geometry. However, the correspondences across a population can be weakened due to the inconsistent branching topology across a population. For instance, a small bump on a boundary can lead to an extraneous branch of the resulting medial axis, as illustrated in fig. 2.1 left. This sensitivity of Blum medial axis jeopardizes the quality of correspondence across a population and thus harms a succeeding statistical shape analysis.



Figure 2.1: Left: a 2D object with non-branching topology except for a boundary bump induced by noise. The noise induces an extra branch on the skeleton (blue). Right: a fitted skeleton to the object that obtains the noise-free object's branching topology. This skeletal model is associated with a slightly modified (dotted) boundary.

The sensitivity of Blum medial axis is due to its strong dependence on the normal vectors of the boundary, which are sensitive to any perturbation of the boundary. To overcome this problem, (Pizer et al., 2003; Damon, 2003, 2004) studied the relations between skeletal geometry of a medial axis and differential geometry of the boundary. Particularly, (Damon, 2003, 2004) showed (1) that relaxing the aforementioned Blum conditions can still have the desired relation between the skeletal and boundary geometry (2) but that the relaxation can nicely lead to consistent skeletal representations with slightly different associated boundary, shown as the dashed curve in fig. 2.1 right. Hence, it is worthwhile to relax these conditions for better correspondences across a population. This realization gives rise to the generalized skeletal representations, as discussed in Section 2.5.3.

2.5.2 Shape Analysis as Functional Data Analysis

As skeletal representations take advantage of interior geometry of objects, this section discusses (1) the mathematical model of the shape involving the interior of objects and (2) the correspondences provided by the mathematical models across a population.

The shape of an object is beneficially understood in terms of not only the boundary but its interior. Given an object with a closed boundary and an interior point of this object, a natural mathematical model that describes the shape is the distance function from the interior point to every boundary point. The relation of these functions across objects is discussed below.

Let Ω denote an object and let $p \in \Omega$ denote an interior point of the object. The shape of Ω can be described by a distance function

$$f_p(x) = \|x - p\| \quad (2.5)$$

where x denotes a boundary point and $\|\cdot\|$ means Euclidean distance. In the literature the notation $\partial\Omega$ is used to denote the boundary of an object. Thus, the domain of $f_p(x)$ is constrained as $x \in \partial\Omega$. This shape function $f_p(x)$ is determined by the placement of p . Varying p results in different mathematical models of the shape of Ω . Fig. 2.2 illustrates the generic properties of various functions $f_p(x)$ as I vary p .

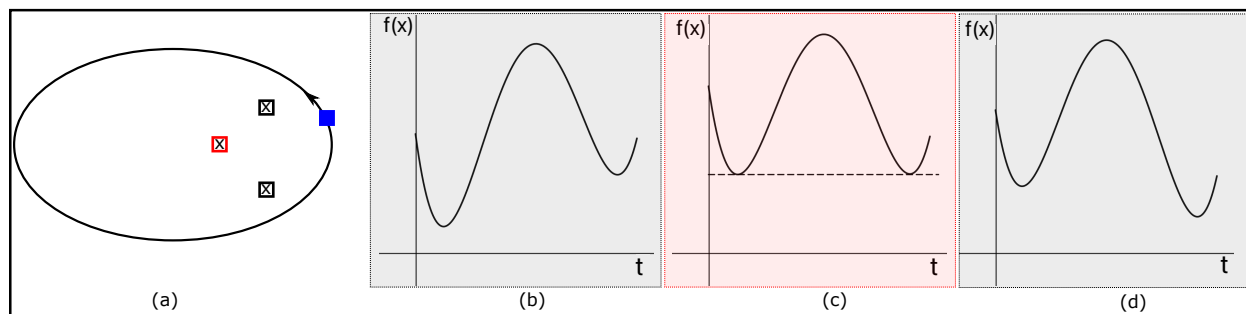


Figure 2.2: Illustration of different generic properties of distance functions $f_p(x)$ as the interior point p varies. (a) In an ellipse three interior points (three crosses in boxes) and a boundary point (blue square). As the boundary point moves along the ellipse boundary counter-clockwise, the distances between the boundary point and the three interior points are shown on the vertical axes in (b-d). (b) The distance function corresponding to the interior point in the top black box. (c) The distance function corresponding to the medial point in the red box. The dashed horizontal line indicates that the two local minima are identical for medial points. (d) The distance function corresponding to the interior point in the bottom black box. The differences of the two local minima between (b) and (d) are discussed in the text.

In placing the interior point p , the Blum medial axis (Blum and Nagel, 1978) only focuses on medial points like the cross in the red box shown in fig. 2.2 (a). Each medial point has at least two closest boundary points. Hence, each medial point is associated with a shape function of the same generic property as the function shown in fig. 2.2 (c); that is, the function has two or more equal local minima. Such medial points exist in a general shape in 2D and 3D.

The existence of the medial points can be seen from the transition of the shape functions associated with interior points. For simplicity, I discuss the existence using a 2D shape Ω (i.e., an ellipse) with a closed boundary as shown in fig. 2.2 (a). The boundary $\partial\Omega$ can be parameterized via a mapping $\alpha : \mathbb{I} \mapsto \partial\Omega$, where \mathbb{I} denotes an interval $[0, 1]$. The mapping has to satisfy $\alpha(0) = \alpha(1)$ so that the boundary is closed. Moreover, an arbitrary boundary point can be written as $x = \alpha(t)$, where $t \in [0, 1]$; see e.g., the blue filled square in fig. 2.2 (a). The shape function defined in Equation (2.5) is a function w.r.t. the continuous variable t after substituting x with $\alpha(t)$. As the interior point p moves in fig. 2.2 (a) from the top black box to the medial red box and to the bottom black box inside the object, the respective shape functions are shown in figures 2.2 (b), (c) and (d). Each shape function $f_p(x)$ is continuous and has two local minima – they are at the values of t when the connection between p and a boundary point $\alpha(t)$ is orthogonal to the

boundary at $\alpha(t)$. Let the two local minima be, respectively, $f_p(\alpha(t_1))$ and $f_p(\alpha(t_2))$, where $0 \leq t_1 < t_2 \leq 1$. The difference $f_p(\alpha(t_1)) - f_p(\alpha(t_2))$ of the two local minima changes from negative (e.g., fig. 2.2 (b)) to positive (e.g., fig. 2.2 (d)). According to the intermediate value theorem, there must exist a transition point where the two local minima are identical (e.g., fig. 2.2 (c)). Such a transition function is associated with a medial point in the Blum medial axis. The above proof of existence can also be generalized to 3D shape. I skip that generalization in this dissertation.

Due to the abovementioned consistent generic properties of the associated shape functions, medial points provide good correspondences of interior geometry across a population. When analyzing a population of objects' shape, I must align the shape functions such that important shape features are compared between equivalent parameters of shape functions. Intuitively, the peaks and valleys of a shape function are expected to align to peaks and valleys, respectively (see (Srivastava and Klassen, 2016)). In the above example, the shape function shown in fig. 2.2 (c) with valleys of equal heights are easier to align – it is easier to find the correspondence of a medial point p across a population than to find the correspondent interior points of the black boxes. Therefore, the equal local minima property is a useful generic property of a medial point that facilitates anatomic shape analysis.

The Blum medial axis of an object is a set of medial points that have two or more identical local minima or a local minimum of multiplicity 2 (see Section 2.5.1). Hence, it is of importance to have good correspondences for the Blum medial axis across a population. However, as discussed in Section 2.5.1, it is challenging to establish good correspondences on the Blum medial axis because of its inconsistent topology in a population. By contrast, a general skeletal representation is superior because it fits to a population of objects with consistent topology, as discussed below.

2.5.3 Generalized Skeletal Representations with Good Correspondences in a Population

Generalized skeletal representations are motivated by the desire to produce topologically consistent skeletal representations for good correspondences. Taking into account the factors discussed above, fitting a skeletal representation to a boundary can lead to a generalized skeletal representation regardless of random perturbations on the boundary. To this end, it is necessary to model the mapping from the skeletal representation to the interior and to the boundary of the object. Let $\tau \in [0, 1]$ specify the radial distance (i.e., the proportion of spoke lengths) from the skeleton to the boundary. The mapping can be understood as inflating the skeleton ($\tau = 0$) to interior level surfaces ($0 < \tau < 1$) and to the boundary ($\tau = 1$); that is,

$$\psi_\tau : M \mapsto B_\tau \tag{2.6}$$

where M denotes a skeletal representation and B_τ , called an onion-skin, is a τ level surface. Fig. 1.3 shows B_τ at various τ values. The optimal skeletal representation should minimize the deviation of the *implied boundary* B_τ at $\tau = 1$ from the segmented boundary of the object.

The mapping $\psi_\tau(\cdot)$ is called the *radial flow*. A smooth radial flow produces non-self-intersecting interior and a smooth boundary of an object. There exist mathematical tools (Damon, 2004) that measure the smoothness of a radial flow, called the radial shape operator. I will explain and exploit this tool in Chapter 3.

Like the Blum medial axis, a spoke that points from a skeletal point $p \in M$ to the corresponding boundary point on $\psi_1(p) \in B_1$ in a generalized skeletal representation relates the skeletal, the interior and the boundary geometry. Such rich geometric features from the spokes have good locality (Lu et al., 2007) as well as good correspondences across cases. As such, generalized skeletal representations are favorable for statistical shape analysis.

A specific generalized skeletal representation is named as an s-rep. Since many structures in the human body have no subfigures (i.e., branches) on their boundaries, the fitting process of an s-rep starts with a non-branching skeletal structure and then deforms that skeletal structure to fit the object of interest. As a result, the fitted s-reps have consistent non-branching topology in a population. Based on such consistency, the correspondences between s-reps provide useful shape features in anatomic shape analysis.

2.5.4 Algorithms of Fitting Skeletal Representations

Siddiqi and Pizer (Siddiqi and Pizer, 2008) gave a survey on the previous methods of fitting skeletal representations. Those methods include a ridge following approach (“cores”), a boundary erosion approach, discrete morphology approaches, Voronoi approaches, continuous medial fitting approaches (“cm-reps”) and previous fitting approaches such as (Joshi et al., 2002; Pizer et al., 2003). Because those previous approaches reflected some – but not all – of the Blum medial properties, the resulting representation is called an “m-rep”. None of those approaches, with the possible exception of a recent form of cm-reps (Yushkevich et al., 2019), can achieve the goal of automatically computing skeletal representations suitable for statistical anatomic shape analysis as discussed above.

Yushkevich et al. (Yushkevich et al., 2019) modified the method due to S. Arguillère et al. (Arguillère et al., 2015) to construct a diffeomorphism between objects that mapped a finite set of points representing the medial axis and the corresponding boundary points for the first object to another finite set for the second object. Nonetheless, without taking into account C^1 or higher order information, such a diffeomorphism could fail to satisfy higher order conditions on the boundary as well as on the skeletal geometry. Specifically, it has been shown in (Damon, 2021) that the diffeomorphism that performs on the boundary and skeleton together can break the branching medial structures in both generic and non-generic cases.

To address deficits in skeletal representations, Pizer et al. (Pizer et al., 2013a) proposed a refinement stage of an m-rep within the shape space obtained from statistical analysis. The refined m-reps achieve a tight shape distribution for a population. However, that shape space is free from geometric constraints, which still cannot resolve self-overlaps. Later, the refinement algorithm studied in (Tu et al., 2016) takes into account both global consistency and individual

specificity of m-reps. However, the algorithm relies on statistics (i.e., variation modes) of a population as the refinement proceeds. It is computationally expensive and impractical for refining each single medial structure. In Chapter 3, I describe an automatic method that fits s-reps to non-branching objects, producing desirable geometric and statistical properties in anatomic shape analysis.

2.5.5 Summarize the Basics of Skeletal Representations

In this section, I have discussed two forms of skeletal representations and the statistical benefits from these skeletal representations. In addition to the rich shape features provided by skeletal representations, these shape features should also have good correspondences across a population. Specifically, I introduce a Blum case and a generalized skeletal representation, called an s-rep. These two representations both provide means of capturing interior geometry of an object. However, the Blum case can have completely different topology in a population of objects (e.g., the hippocampus). By contrast, generalized skeletal representations (e.g., s-reps) can be chosen to have consistent topological properties in a population. These skeletal representations are fitted to the objects to provide good correspondences across the population. In addition to the desired statistical properties of skeletal representations, there also exist psychological motivations to utilize skeletal representations in recognizing the shape. A reader is referred to (Siddiqi and Pizer, 2008) for more details. Last, I review the previous methods of fitting skeletal representations. These methods can produce skeletal representations with good correspondences for specific samples. However, these methods may produce inconsistent branching topology across a population, which limits the use of these methods.

On the basis of fitted s-reps, the following section shows an example in previous research which extracts interior geometric features with good correspondence.

2.6 Interior Geometry of Objects

As mentioned in Chapter 1, the interior of an object can provide more robust geometric features than its boundary. S-reps give a means of extracting interior geometric features of objects in good correspondences. This section reviews the previous research (Pizer et al., 2022) that initiates the research on the interior geometry of objects using s-reps.

2.6.1 Skeletal Coordinates of an Interior Point

Pizer et al. (Pizer et al., 2022) have defined that any interior point of an object can be recognized as a 3-tuple (θ, τ_1, τ_2) . These 3 coordinates can be straightforwardly understood in terms of relative positions to a skeleton and to its spine, which is a 1D skeletal curve of the 2D skeleton. Take an interior point $p(\theta, \tau_1, \tau_2)$ as an example. The value τ_1 captures how far the point p is from the spine on the skeleton; the value θ measures the angle circulating the edge

curve in a counter-clock direction; the value τ_2 measures the movement of p along a spoke – i.e., $\tau_2 = 0$ if p is on the skeleton, while $\tau_2 = 1$ if p is on the boundary.

2.6.2 Local Orthonormal Reference Frames

In many traditional research projects, the geometric models are established within a global coordinate system (see e.g., (Hong, 2018; Miolane et al., 2021; Vicory et al., 2018)). Such geometric features are sensitive to rigid transformations, including translation and rotation. This sensitivity can hardly be removed in multi-object shape analysis because of the difficulty of aligning multi-object complexes. In this regard, local reference frames show promise of extracting shape features that are insensitive to rigid transformations in object recognition (Guo et al., 2013), shape matching (Petrelli and Di Stefano, 2011) and registration (Yang et al., 2016; Lei et al., 2017; Malassiotis and Srinivasan, 2007).

Pizer et al. (Pizer et al., 2022) defined local orthonormal frames on s-reps for use in describing within-object geometry though I show them also to be of use in describing between-object geometry. An orthonormal local frame consists of (1) a unit vector $\widehat{\nabla\tau_1}$ (on the skeleton) that points away from the spine, (2) a unit vector $\widehat{\nabla\tau_2}$ that points away from the skeleton and (3) the cross product of the two vectors, i.e., $\widehat{\nabla\theta} = \widehat{\nabla\tau_2} \times \widehat{\nabla\tau_1}$. Geometrically, the direction $\widehat{\nabla\theta}$ (on the skeleton) accounts for circulating the skeletal edge curve in a counter-clock wise way.

Moreover, the rotation between those frames captures differential geometry within and across level surfaces of the radial flow (see eq. (2.6)). Pizer’s local frame field of a target object can be obtained by deforming a template object (e.g., an ellipsoid) along with its local orthonormal frames. Because this deformation can produce non-orthonormal frames in the target object, an orthonormalization step is needed. As a result, the orientations of these orthonormalized frames are determined by local geometry at the frames’ origins.

2.7 Statistics in Euclidean Space

This dissertation will make use of non-Euclidean statistics to analyze multi-object shape (see Section 2.8). To prepare for the discussion of non-Euclidean statistics, this section reviews how the relevant statistical methods work in the Euclidean space.

In particular, this section discusses theory, methods and applications in which the data are analyzed as being sampled from the Euclidean space. In the following sections, I first review the geometry of Euclidean space in Section 2.7.1. Then I review statistical methods that are based on the Euclidean metric and that are relevant to this dissertation. Specifically, Section 2.7.2 introduces a problem of analyzing multiple blocks of Euclidean data. Each block contains Euclidean shape descriptors of an object. The analysis focuses on the joint variation of the multiple blocks. To this end, Section 2.7.3 describes some traditional methods that seek the optimal directions in feature space maximizing the

joint variation. However, such methods may produce spurious directions in high dimension low sample size scenarios. Instead, Section 2.7.4 details a robust statistical method, called Angle-based Joint and Individual Variation Explained (AJIVE), that can extract joint and individual variation from multi-block Euclidean data.

2.7.1 Geometry of Euclidean Space and Subspaces

Distance between Euclidean data. The distance metric in the data space plays an essential role in analyzing a population. A metric can be seen as a kernel function that takes in two entities in a common space and outputs a scalar (e.g., a distance or an angle) that captures the difference between the two entities. For example, given two data points that are represented by two feature vectors in the Euclidean space \mathbb{R}^n , the distance between these two data points is measured by the Euclidean metric. Often, the Euclidean space \mathbb{R}^n is considered as a *normed vector space* with the following properties:

1. A data point in \mathbb{R}^n is an n -dimensional vector. The vector $\mathbf{0}_n$ whose entries are all zeros is a valid data point in the space.
2. Addition and scalar multiplication are valid operations that are closed in \mathbb{R}^n . In particular, these operations satisfy linear properties.
3. The similarity of any two unit vectors in \mathbb{R}^n can be measured by the Euclidean inner product of the two vectors.

Distance between Euclidean subspaces. In addition to the above relations between data points within a Euclidean space, some statistical problems require understanding the relations between two subspaces of a Euclidean space. For example, multi-block Euclidean data can be seen as samples from multiple Euclidean subspaces that can be of different ranks, units or scales. The distances among these Euclidean subspaces are related to statistical correlation of multiple blocks (Pezeshki et al., 2004).

A commonly used metric to measure the distance between Euclidean subspaces is called *principal angles*. Small principal angles correspond to strong correlation of the blocks from the subspaces. Let $\rho(\cdot, \cdot)$ denote a metric measuring the distance between two Euclidean subspaces $\mathbb{J} \subset \mathbb{R}^n$ and $\mathbb{Q} \subset \mathbb{R}^n$; let l and m denote the dimensions of \mathbb{J} and \mathbb{Q} , respectively, assuming that $l \leq m$. Let $v_i \in \mathbb{J}$ and $q_i \in \mathbb{Q}$ denote two basis vectors of dimensions n , where $i = 1, \dots, l$. The principal angles θ 's of the two subspaces are defined as acute angles $0 \leq \theta_1 \leq \theta_2 \leq \dots \leq \theta_l \leq \frac{\pi}{2}$ that are minima of

$$\theta_j = \min_{v_j, q_j} \cos^{-1} \left(\frac{\langle v_j, q_j \rangle}{\|v_j\| \cdot \|q_j\|} \right) \quad v_j \perp v_i, q_j \perp q_i, i \in [1, j-1] \quad (2.7)$$

where $\langle \cdot, \cdot \rangle$ denotes the dot product of two vectors. The pair of vectors v_j and q_j is called a pair of *principal vectors*, where $j = 1, \dots, l$. To compute principal angles, Principal Angle Analysis (PAA) utilizes the following lemma (Börck and Golub, 1973).

Lemma 2.1. *Let \mathbb{J} and \mathbb{Q} be two linear subspaces in \mathbb{R}^n , and let the rows of a matrix $M_J \in \mathbb{R}^{l \times n}$ and the rows of a matrix $M_Q \in \mathbb{R}^{m \times n}$ be orthonormal basis vectors for \mathbb{J} and \mathbb{Q} , respectively. The singular values σ_i of the matrix $M_Q M_J^T \in \mathbb{R}^{m \times l}$ relate to the solution of eq. (2.7) via*

$$\sigma_i = \cos \theta_i \quad i = 1, \dots, l \quad \text{and} \quad \sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_l \geq 0. \quad (2.8)$$

The intersection space of \mathbb{J} and \mathbb{Q} gives rise to statistical correlation of data blocks sampled from \mathbb{J} and \mathbb{Q} (Pezeshki et al., 2004). This space can be a zero space if no significant correlation found between the data blocks. In Section 2.7.4 I discuss a statistical method (AJIVE) that makes use of PAA to extract correlation between multiple Euclidean data blocks. Essentially, AJIVE estimates the basis vectors of the intersection space by focusing on the pairs of principal vectors that have small enough principal angles. The intersection space can have at most l basis vectors because the subspaces \mathbb{J} and \mathbb{Q} can have at most l pairs of principal vectors. Ideally, the intersection space is generated by the pairs of principal vectors that have principal angles 0, each of which corresponds to a singular value equal to 1 of the matrix $M_Q M_J^T$. In practice, however, the principal angles can equal 0 only with probability 0 because of noise in real data.

2.7.2 Correlation Analysis of Multi-block Euclidean Data

In this dissertation I need to analyze how the shape of multiple related objects jointly varies. This problem can be formulated as a multi-block data analysis, where each block is a data matrix and the statistical correlation is taken as the joint shape variation. In this section I assume that each block is sampled from a Euclidean space.

Because each object can be described by a different number of Euclidean features, the multi-block data can be sampled from multiple Euclidean subspaces of different ranks. In this section I use $X_1 \in \mathbb{R}^{d_1 \times n}$ and $X_2 \in \mathbb{R}^{d_2 \times n}$ to denote two blocks of observations, where d_1 and d_2 are the respective dimensions of shape features. The joint analysis of X_1 and X_2 reveals how the observations in the two blocks jointly vary. A typical joint analysis problem is influenced by many factors. First, the sample size n is often too small to reveal a generalizable pattern of the correlation. Second, the dimension $d_1 + d_2$ of the two-block features is often so large that it easily introduces notable computational error in a high dimensional space. Third, the variability of one block can be dominating and thus bias the correlation.

Section 2.7.3 describes some existing methods that deal with correlation of multi-block Euclidean data. All the methods but JIVE (short for Joint and Individual Variation Explained) require that $d_1 = d_2$. Moreover, none of these methods can effectively address the above challenges. In Section 2.7.4 I discuss a more recent method (i.e., AJIVE) that is promising in solving those issues. In AJIVE, d_1 can be different from d_2 , but the two blocks are measurements from the same n samples.

2.7.3 Traditional Methods of Correlation Analysis of Multi-block Euclidean Data

In this section I review related Euclidean data integration methods. These related methods include Consensus PCA (CPCA) (Wold et al., 2005), JIVE (Lock et al., 2013), Generalized Canonical Correlation Analysis (GCCA) (Shen et al., 2014), Partial Least Squares (PLS) (Wold, 2004), Hierarchical PCA (HPCA) (Westerhuis et al., 1998) and DISCO-SCA (Van Deun et al., 2012).

A straightforward approach to obtain the basis for the joint variation space is via a multi-block Principal Component Analysis (PCA) on the composite matrix $\{X_k\}$, which is called Consensus PCA (CPCA) (Wold et al., 2005). However, the principal directions resulting from this approach can be dominated by the block of relatively larger variability, ignoring the joint variation.

Alternatively, Lock et al. (Lock et al., 2013) proposed Joint and Individual Variation Explained (JIVE) to decompose data into three structures, namely, the joint, individual, and residual structures. With this decomposition model, the method iteratively decomposes the data via minimizing the residual components. However, because JIVE estimates the dimension of the joint variation space via a permutation test, it is likely that JIVE can overestimate the joint component by inappropriately including individual components (Feng et al., 2018).

Another direction of multi-block data analysis is based on Canonical Correlation Analysis (CCA) (Hotelling, 1936) that maximizes the correlation over all possible directions, yielding pairs of canonical loadings $\{u_k \in U_k | k = 1, \dots, K\}$, where U_k is the left singular matrix of X_k . However, the canonical loadings are not well-defined in HDLSS problems. Moreover, Shen et al. (Shen et al., 2014) pointed out that it is not necessarily beneficial to include additional datasets in classification with CCA. The researchers proposed GCCA that can jointly analyze more than 2 datasets. Moreover, they showed that GCCA can help the classification when some sufficient conditions are satisfied. That work shows an example of classification where an unsupervised learning is followed by a supervised classification. Also, in the classification of real datasets, the researchers adopted an embedding method called Multi-Dimensional Scaling (MDS). Importantly, the embedding via MDS and the dimension reduction via GCCA make no use of class labels. The classification splits training and test data from GCCA's results, which I call post-feature-selection in a classification experiment. However, the GCCA directions that maximize the correlation may not produce jointly varying patterns regarding two blocks. The inconsistent varying patterns in two blocks lead to inconsistent interpretation of various joint components. Moreover, this inconsistency complicates model selection.

Different from CCA, Partial Least Squares (PLS) (Wold, 2004) directions are the directions in feature spaces that maximize covariance instead of correlation. However, the resulting scores can be biased by structured noise, which is defined as systematic variation that is not linearly correlated between blocks (see (Jiang, 2018) Section 4.2 for more details). To avoid the bias due to structured noise in standard PLS, Hierarchical PLS (HPLS) applies PCA on the data matrix prior to fitting a PLS model. As a result, the model can have better interpretability than standard PLS (see e.g.,

(Lindström et al., 2006)). Alternatively, Trygg et al. (Trygg and Wold, 2003) proposed removing structured noise of two blocks via projections orthogonal to PLS components, yielding a method called O2-PLS. Due to the orthogonal projections, this method provides good interpretation of structured noise for each block. However, the joint variation space is different across blocks. The multiple joint variation spaces are not necessarily orthogonal to individual variation spaces, which might deteriorate the interpretability of the joint components.

Deun et al. (Van Deun et al., 2012) proposed DISCO-SCA to simultaneously decompose the joint and individual components from multi-block data. This method first extracts the residual component via Simultaneous Component Analysis (SCA). Then it rotates the SCA components towards an optimal user-defined target matrix. Compared to JIVE and AJIVE, DISCO-SCA rotates the concatenation of the joint and the individual components. This leads to orthogonal constraints between individual components across blocks. Such a restrictive constraint rarely corresponds to biological behaviors (van der Kloet et al., 2016).

2.7.4 Joint Analysis of Two-block Euclidean Data via AJIVE

Following Lock’s algorithm as reviewed above, AJIVE treats multi-block Euclidean data as a sum of joint, individual and residual components. Specifically, let $X_k \in \mathbb{R}^{d \times n}$ denote the data block of the k^{th} object, where $k = 1, 2$ in this dissertation. Let J denote the joint component shared by the two objects. Let I_k and E_k , respectively, denote the individual and the residual component specific to the k^{th} object. The feature matrix X_k can be decomposed as

$$X_k = J + I_k + E_k, \quad k = 1, 2 \quad (2.9)$$

The key improvement of AJIVE over JIVE is to provide more robustness and efficiency by estimating the joint, individual and residual spaces from observations. As noted above, the joint variation space is where statistical correlation of multiple blocks arises. To estimate the joint variation space, AJIVE uses PAA to estimate the intersection space of the multiple Euclidean data spaces, which is equivalent to the joint variation space of the multiple blocks (Pezeshki et al., 2004). Moreover, because the joint and individual components should be uncorrelated, the joint variation subspace is modeled as a space orthogonal to every individual variation space in AJIVE.

AJIVE estimates the joint and individual spaces within the score space \mathbb{R}^n . In small sample size problems where $n \ll d$, the covariance matrix of the two blocks is of low rank. This low rank problem leads to defective empirical canonical correlation in the statistical methods on the basis of \mathbb{R}^d (Pezeshki et al., 2004). Instead, because two data blocks are from the same n subjects, they can share the same score space \mathbb{R}^n , which provides the basis of seeking the joint and individual variation subspaces of the two blocks. Thus, AJIVE estimates the joint variation space within \mathbb{R}^n via PAA, as follows.

Orthonormal Scores Subspaces from Singular Value Decomposition (SVD). Because PAA requires orthonormal basis vectors of two Euclidean subspaces in \mathbb{R}^n (see Lemma 2.1), the first step of AJIVE is to obtain orthonormal score vectors from the two data blocks. This orthonormalization can be achieved via SVD. Moreover, the operations of SVD on the two blocks can avoid spurious joint variation by putting negligible signals in the residual component E_k (see eq. (2.9)). The number of negligible components is determined by choosing an *initial rank* r_k .

Specifically, let $U_k S_k V_k^T$ be the result from SVD of the data block $X_k \in \mathbb{R}^{d \times n}$, where U_k and V_k^T are orthonormal matrices. By choosing an appropriate initial rank r_k , I can simultaneously truncate (1) the columns of U_k , (2) the rows of V_k^T and (3) the rows and columns of S_k such that the negligible singular values are removed. As a result, I obtain the approximated left singular matrix $\hat{U}_k \in \mathbb{R}^{d \times r_k}$, the approximated singular value matrix $\hat{S}_k \in \mathbb{R}^{r_k \times r_k}$ and the approximated right singular matrix $\hat{V}_k^T \in \mathbb{R}^{r_k \times n}$. As a result, the low rank approximation of X_k can be written as $\hat{X}_k = \hat{U}_k \hat{S}_k \hat{V}_k^T$. Importantly, the rows in \hat{V}_k^T give orthonormal basis vectors of \hat{X}_k 's score space, which can be handled by PAA as follows.

Principal Angle Analysis on Multiple Score Spaces. PAA on multiple score spaces aims to analyze the relations between score spaces from all the blocks of interest. In particular, the intersection of multiple score spaces can be obtained from the principal vector pairs associated with small principal angles. Given the above orthonormal matrices \hat{V}_k^T , the following lemma provides a means of simultaneously evaluating the principal angles and obtaining the basis vectors of the intersection subspace (Horn and Johnson, 2012; Miao and Ben-Israel, 1992).

Lemma 2.2. *Let the columns of $V_1 \in \mathbb{R}^{n \times l}$ and $V_2 \in \mathbb{R}^{n \times m}$ be orthonormal bases for the score spaces of X_1 and X_2 , respectively. Assume $l \leq m$ and $(l+m) < n$. Let the non-zero singular values of the matrix*

$M = V_1^T V_2 = \underbrace{U_M}_{l \times l} \underbrace{\Sigma_M}_{l \times m} \underbrace{V_M^T}_{m \times m} \in \mathbb{R}^{l \times m}$ be $\sigma_1 \geq \dots \geq \sigma_l > 0$. Define a matrix

$$Q := \begin{bmatrix} V_1^T \\ V_2^T \end{bmatrix} \in \mathbb{R}^{(l+m) \times n} \quad (2.10)$$

Then the singular values of Q can be computed using the singular values of M as follows:

$$\sqrt{1 + \sigma_1}, \dots, \sqrt{1 + \sigma_l}, \underbrace{1, \dots, 1}_{m-l}, \sqrt{1 - \sigma_l}, \dots, \sqrt{1 - \sigma_1}.$$

Moreover, the right singular matrix V_M of M can be written as $V_M = \left(\overbrace{V_{M1}}^{m \times l}, \overbrace{V_{M2}}^{m \times (m-l)} \right)$. Then the left and right singular matrix U_Q and V_Q^T of Q can be obtained using U_M and V_M as follows:

$$U_Q = \begin{bmatrix} \frac{1}{\sqrt{2}}U_M & 0 & \frac{1}{\sqrt{2}}U_M \\ \frac{1}{\sqrt{2}}V_{M1} & V_{M2} & -\frac{1}{\sqrt{2}}V_{M1} \end{bmatrix} \in \mathbb{R}^{(l+m) \times (l+m)} \quad (2.11)$$

$$V_Q = \Sigma_Q^{-1} \begin{bmatrix} \frac{1}{\sqrt{2}}(V_1U_M + V_2V_{M1}) & V_2V_{M2} & \frac{1}{\sqrt{2}}(V_1U_M - V_2V_{M1}) \end{bmatrix} \in \mathbb{R}^{n \times (l+m)} \quad (2.12)$$

where $\Sigma_Q \in \mathbb{R}^{(l+m) \times (l+m)}$ is a diagonal matrix with singular values of the matrix Q . \square

Similar to eq. (2.8), large singular values of Q correspond to pairs of principal vectors that have small principal angles. Specifically, the rows in the right singular matrix V_Q^T of Q gives the basis vectors that span the intersection space. Thus, AJIVE constructs the basis of the intersection space from the first l rows of the right singular matrix V_Q^T . Essentially, each basis vector is the scaled sum of a pair of principal vectors with a small principal angle, see eq. (2.12).

The individual variation space is constructed based on the orthogonality constraint between the joint and individual variation space. The complete algorithm of AJIVE is shown in Algorithm 1.

Algorithm 1 AJIVE decomposes two block Euclidean data

Require: Blocks X_1 and X_2 ; initial ranks r_1 and r_2

- 1: Approximate X_1 and X_2 with initial ranks r_1 and r_2 . // Low rank approximation
 - 2: Compute each score subspace based on the approximated data. // Results in score subspaces of multiple blocks
 - 3: Estimate the joint variation subspace via Lemma 2.2.
 - 4: Use the orthogonality constraint to construct the individual variation space for each block.
 - 5: **return** the joint components J_1 and J_2 and block specific individual matrices I_1 and I_2 .
-

Intuitions of the Joint and Individual Components from AJIVE. Intuitively, the variation in the joint component is driven by the common function on multiple blocks. Each basis vector of the joint variation space can be seen as a basis vector of the common function. Therefore, the construction of the joint variation subspace is equivalent to discerning that common function. Moreover, the joint structures J are the projection of data in the joint variation space. This projected data reflects effects of the common function.

The individual components I_k are block-specific and are uncorrelated with the joint structure. Correspondingly, every individual variation subspace is a space orthogonal to the joint variation space. However, the individual components are not required to be orthogonal to each other, leading to better interpretability in biological applications than other methods (e.g., DISCO-SCA).

Despite the desirable statistical properties and computational efficiency of AJIVE, the linear combination in eq. (2.9) only makes sense in the Euclidean setting. That decomposition can fail in a curved non-Euclidean space, e.g., a shape space. To address this problem, Chapter 4 will generalize the method by making use of non-Euclidean statistics.

2.8 Non-Euclidean Data Analysis

In many cases data are not Euclidean; for example, shape data distribute in non-Euclidean spaces. While in the previous section the Euclidean metric of the ambient space can approximately measure the distance between two shape data, the distribution of a shape population can be more efficiently captured by using intrinsic geometry of a non-Euclidean space (see (Patrangenaru and Ellingson, 2019) for more details). To this end, Section 2.8.1 discusses the geometry of a non-Euclidean space (called a pre-shape space), in which the shape data of an object is decoupled with position and size information. Within this non-Euclidean space, previous research has provided mathematical tools for analyzing the shape distribution. For example, as discussed in Section 2.8.2, Fletcher et al. (Fletcher et al., 2004a) proposed a method called Principal Geodesic Analysis (PGA) that characterizes data distribution with geodesics on a manifold. As an alternative, Jung et al. (Jung et al., 2012) proposed a method called Principal Nested Spheres (PNS) that uses both geodesic and non-geodesic curves to capture principal components of the distribution, as detailed in Section 2.8.3. Because PNS results in better statistical characterization of shape distributions than PGA does, this dissertation uses PNS to approach multi-object shape analysis.

2.8.1 Non-Euclidean Statistics in Shape Analysis

Shape data live in a non-Euclidean space by nature because of non-linear constraints on shape data: two objects that differ up to a similarity transformation (i.e., translation, rotation and scaling) are understood as the same shape. For example, in Kendall's statistical theory (Kendall, 1984) where the shape is represented by landmarks on an object's boundary, the geometric information of the object should be invariant under similarity transformations, including translation, rotation and scaling. Based on landmark-based shape representations, this section shows the manifold structure of the shape data after the removal of position and size information.

The position variation of an object can be removed by subtracting the coordinates of its centroid from the coordinates of landmarks (Dryden and Mardia, 1998). Specifically, given a 3D object represented by m landmarks, let $X \in \mathbb{R}^{m \times 3}$ denote the shape feature matrix of the object. This object can be centered via a linear transformation $C \cdot X$, where the centering matrix $C = I_m - \frac{1}{m} \mathbf{1}_m \mathbf{1}_m^T$. In the centering matrix, $I_m \in \mathbb{R}^{m \times m}$ denotes an identity matrix and $\mathbf{1}_m \in \mathbb{R}^{m \times 3}$ is a matrix of ones.

To remove the size variation, let $S(\cdot)$ denote a function that measures the centroid-based size of an object. Specifically, let \bar{X} denote the centroid of the object. The function $S(\cdot)$ is defined as

$$S(X) = \sqrt{\sum_{i=1}^m |X_i - \bar{X}|^2} \quad (2.13)$$

where X_i denotes the coordinates of the i^{th} landmark. To take advantage of the centering matrix C , the function $S(X)$ can be written as

$$S(X) = |C \cdot X| \quad (2.14)$$

Combining the removal of position and size variation, the shape matrix X can be transformed via

$$Z_C = \frac{C \cdot X}{|C \cdot X|} \quad (2.15)$$

where $Z_C \in \mathbb{R}^{m \times 3}$ is referred to as a *pre-shape* because it is invariant to translation and scaling of the object represented by X . Accordingly, the pre-shape space is the space of all pre-shapes.

The pre-shape space is isomorphic to a unit hypersphere because any pre-shape is of unit size, i.e., $|Z_C| = 1$. Importantly, there exists no zero element in this space, making it inappropriate to center pre-shapes at zero. In some applications (e.g., (Huckemann et al., 2010)), shape features are also expected to be invariant to rotations of the object. There, a shape space is a stratified space of the pre-shape space, in which pre-shapes of different orientations are identified. Because rotations of objects are of the interest in this dissertation, I have decided to preserve orientations of objects, restricting the data space to the pre-shape space. In the next section, I describe two different methods that use proper distance metrics in the pre-shape space to analyze the shape distribution.

2.8.2 Principal Geodesic Analysis (PGA)

PGA uses geodesics to measure the distance between any two points in a pre-shape space and more generally in a Riemannian manifold. These geodesics correspond to straight lines in a (Euclidean) tangent space, where the data distribution can be analyzed via Euclidean methods (e.g., PCA). To this end, PGA (1) utilizes an Exponential operator to map data in the tangent space of a manifold onto the manifold and (2) utilizes a Logarithmic operator to map the data back in the tangent space. Specifically, an Exponential map takes in data in the tangent space $T_p M$ of a manifold M at a base point $p \in M$ and maps the data to M . Inversely, a Logarithmic map is locally defined on a Riemannian manifold M as a map from the manifold to its tangent space $T_p M$ based on a predefined point $p \in M$. These two maps relate any straight line in a tangent space with a geodesic curve on the manifold through the base point of that tangent space.

PGA estimates the Fréchet mean on a manifold by minimizing the total squared distance to every data point on the manifold. The distance is measured by geodesic distance on the manifold. In the estimation of the Fréchet mean, the algorithm repeats the process of i) mapping data to a tangent space via a Logarithmic map ii) computing the Euclidean mean value in the tangent space and iii) mapping the Euclidean mean to the manifold via an Exponential map to give a new base point of the Logarithmic map until convergence.

In PGA, the principal variation directions are straight lines in the tangent space through the Fréchet mean μ . Accordingly, the Exponential map Exp_μ of these directions are understood as geodesic curves on the manifold. However, such representations of variation can be problematic in characterizing anatomic shape distributions. It is found that the distribution of anatomic shape data follows not necessarily geodesic curves but sometimes follows non-geodesic curves (Jung et al., 2012). Using either geodesic or non-geodesic curves on the shape manifold, the following method called Principal Nested Spheres (PNS) yields a better understanding of anatomic shape statistics.

2.8.3 Principal Nested Spheres (PNS)

Compared to PGA, PNS (Jung et al., 2012) gives more statistically efficient representations of pre-shapes in the following senses. First, PNS estimates a backwards mean that is more representative of spherical population than the Fréchet mean (see Section 6.1 in (Pizer et al., 2019)). Second, PNS models the principal components in the pre-shape space via either a great circle (a geodesic curve) or a small circle (a non-geodesic curve), depending on the distribution of the data.

Specifically, PNS iteratively reduces the dimension of a unit hypersphere, yielding a hierarchy of subspheres (treated as generalized principal components) that fits the data by minimizing residuals (see Fig. 2.3). From these subspheres the pre-shapes are represented by geodesic distances (restricted to the range of $[-\pi/2, \pi/2]$). Eventually, these geodesic distances are taken as Euclidean representations of pre-shapes. I detail the process as follows.




Figure 2.3: Obtain PNS scores by fitting nested spheres to n samples, starting from the initial hypersphere \mathbb{S}^{n-1} . Blue dots illustrate pre-shapes on hyperspheres. Each subsphere A_{n-l-1} is parametrized by its center and radius (i.e., v_l and r_l) and is associated with a score vector of length n . The concatenation of the $n-1$ score vectors is taken as the Euclidean representations of the n samples.

Initial Hypersphere. Let $X \in \mathbb{R}^{d \times n}$ be the n pre-shapes of dimension d , where $d \gg n$ in this dissertation. An initial hypersphere \mathbb{S}^{n-1} can be obtained by throwing away components with zero singular values via SVD.

Fit Subspheres to Data. Instead of only fitting great circles to single object complexes (e.g., PGA), PNS fits a hierarchy of subspheres that are of decreasing dimensions from \mathbb{S}^{n-1} down to \mathbb{S}^1 , as illustrated in Fig. 2.3. At each dimension, the best fitting subsphere is obtained by minimizing the sum of squared residual geodesic distances along the surface of the sphere. This best fitting subsphere is not necessarily a great circle because sometimes a small circle can better fit the variation which is natural to s-reps (Damon and Marron, 2014; Pizer et al., 2019) and the s-rep implied boundary points. PNS computes the backwards mean on the nontrivial subsphere of the lowest dimension (i.e., \mathbb{S}^1). Then, this backwards mean is pulled back to every component in high dimensional spheres, yielding a representative mean on the nested spheres.

Compute PNS Scores from the Fitted Subspheres. In order to compute PNS scores of an object, which can be taken as Euclidean(ized) representations, PNS measures the geodesic distances from principal components that are represented by either a great or a small circle. In fact, each subsphere is the fitted principal component of its embedding sphere. As a result, all the nested subspheres of the initial hypersphere \mathbb{S}^{n-1} produce $n - 1$ score vectors, where n denotes the number of samples.

In previous studies (e.g., (Pizer and Marron, 2017)), PNS scores have shown better statistical properties than other PCA-like methods (e.g., PGA or geodesic PCA (Huckemann et al., 2010)) in representing shape data. Also, PNS is useful in analyzing data that have both non-Euclidean and Euclidean features. For example, to analyze a population of s-reps (see Section 2.5.3), Pizer et al. (Pizer and Marron, 2017) proposed extending PNS to manifolds that are products of 2-spheres and Euclidean spaces. In that work, the spherical features (e.g., spokes' directions) are separately Euclideanized with PNS. Then the researchers concatenate the Euclidean and Euclideanized features with deliberate normalization. Finally, the composite feature matrix is analyzed with PCA, yielding the principal variation components of the s-rep data.

Because PNS provides good understanding of the distribution of data on curved manifolds, PNS features are often used in statistical tasks such as classification (Hong et al., 2016) and hypothesis testing (Schulz et al., 2016). In the following section, I discuss how to use PNS features in supervised binary classification problems.

2.9 Supervised Binary Classification

To verify the association between autism and the shape features extracted from two-object complexes, I use the proposed methods (as introduced in Section 1.2) in classifying autism versus non-autism or testing the statistical significance of the group difference. The methods produce local shape features in single objects and the features

between the two objects. From these features a binary classifier can learn a separation hyperplane between autistic samples and non-autistic samples. This section focuses on the usage of the proposed features in classification.

In the following I first detail the statistical model in the classification problem (see Section 2.9.1). Then in Section 2.9.2 I review two relevant methods that approach the problem, that is, the Support Vector Machine (SVM) (Cortes and Vapnik, 1995) and Distance Weighted Discrimination (DWD) (Marron et al., 2007). The geometric interpretation of SVM can help to understand the mechanism of DWD. Because DWD is more robust than SVM in HDLSS problems, I use DWD in this dissertation. Finally, in Section 2.9.3 I describe the metrics in evaluating a classification model used in later chapters.

2.9.1 Statistical Goal

To better predict whether an object shape is autistic (positive) or non-autistic (negative), a classifier should capture the association pattern between the class labels and the shape features. In the training stage, the classifier is fed with shape features from observed objects and the associated class labels. The classifier is expected to learn the mapping from the shape features to the class labels in a way that has low misclassification.

Specifically, let $x \in \mathbb{R}^d$ be the shape features in a Euclidean space, where d is the dimension of a feature space; let $y \in \{0, 1\}^n$ be class labels of the training data. The probability that the classifier should learn is the probability that an object is autistic given its shape features, i.e., $p(y = 1 | x)$. The following section gives two methods that can estimate this probability.

2.9.2 Support Vector Machine vs. Distance Weighted Discrimination

This section discusses two similar methods that learn a linear decision boundary: the Support Vector Machine (SVM) and Distance Weighted Discrimination (DWD). Here, a decision boundary is referred to as a hyperplane that best separates two classes in a feature space. Though both SVM and DWD are based on the idea of maximizing the separability between the two classes, they differ in measuring the separability of two classes.

SVM maximizes the minimum distance from each class to the decision boundary. This minimum distance is determined by samples in each class that are closest to the decision boundary. These samples are referred to as support vectors. In order to deal with two classes with overlapped distributions, SVM provides slack variables to allow slight misclassification. This relaxation makes the algorithm more generalizable to unseen data. However, in a high-dimensional feature space, the decision boundary can be sensitive to noise, leading to overfitting to the training data. Specifically, (Marron et al., 2007) show an overfitting problem called “data-piling”. The data-piling problem refers to a situation where the observations in each class are piled at a single location on a normal vector of the decision

boundary. As a result, this learned decision boundary performs poorly on new data though it can accurately separate the training data.

To alleviate the data-piling problem, (Marron et al., 2007) proposed DWD that takes into account the distances of all the samples rather than merely the support vectors. More precisely, DWD takes into account all the distances weighted by the relative positions of samples in each class. Specifically, DWD minimizes the summation of reciprocals of distances from every sample to the decision boundary. Therefore, the objective function of DWD can be written as

$$f(\{x_i | i = 1, \dots, n\}) = \sum_{i=1}^n \left(\frac{1}{r_i} + \xi_i \right) \quad (2.16)$$

where $x_i \in \mathbb{R}^d$ is a feature vector of the i^{th} sample; ξ_i represents the slack term that allows misclassification; $r_i \in \mathbb{R}$ is the signed distance to the decision boundary, given by

$$r_i = y_i(\omega^T x_i + \beta) \quad (2.17)$$

where $y_i \in \{-1, 1\}$ represents the class label of the i^{th} sample; $\omega \in \mathbb{R}^d$ is the normal vector of the decision boundary; β gives the point of intersection of the decision boundary with the normal vector. The reciprocal of distances lends higher importance to those samples that are close to the decision boundary than those far from the decision boundary. In doing so, the learned decision boundary can be better generalized to new data, especially in an HDLSS context.

2.9.3 Evaluation Metrics

In this dissertation, the evaluation of a classifier is via testing the learned decision boundary with new data. There exist many metrics in evaluating the testing performance. A commonly used metric, the so-called Area Under the Curve (AUC), can be obtained from a Receiver Operating Characteristics (ROC) curve (see e.g., fig. 2.4).

An ROC curve depicts the relation between True Positive Rate (TPR) on the vertical axis and False Positive Rate (FPR) on the horizontal axis as the decision threshold varies. TPR characterizes the ability of a classifier to recognize true positives, while FPR characterizes the rate of a classifier raising false alarms; defined as the following:

$$TPR = \frac{TP}{TP+FN} \quad \text{and} \quad FPR = \frac{FP}{FP+TN} \quad (2.18)$$

where TP (True Positive) represents the number of predicted positives that are actually positive and FP (False Positive) represents the number of predicted positives that are actually negative. Likewise, TN (True Negative) represents the number of predicted negatives that are actually negative, and FN (False Negative) represents the number of predicted negatives that are actually positive.

Receiver Operation Characteristic curve




Figure 2.4: The blue curve is an ROC curve resulting from a classification problem. The gray curve is the ROC curve from random guess with AUC 0.5.

A larger area under the ROC curve indicates a better classifier. As the decision threshold becomes less and less strict, a good classifier can produce high TPR while the FPR remains low. On the graph, a good classifier is therefore associated with an ROC curve clinging to the top left corner, covering a large AUC. Comparing to the random guess for which the AUC is 0.5, a reasonable classifier should be associated with an AUC larger than 0.5.

CHAPTER 3: FITTING S-REPS TO NON-BRANCHING OBJECTS

3.1 Introduction

Within-object shape features are critical ingredients of multi-object complexes. For the purpose of jointly analyzing multiple objects and of extracting between-object relations, it is necessary to capture within-object shape features that have good correspondence across populations. To this end, this chapter discusses a method that automatically fits s-reps to non-branching objects such as the hippocampus and the caudate. This s-rep fitting method should produce s-reps of consistent topology so that good correspondences can be established across a population. Moreover, a desirable s-rep should give smooth geometric features within an object, as noted in Section 1.2.1.

As discussed in Section 2.5.3, a radial flow that maps a skeleton to level surfaces of the flow gives a means of capturing interior geometry of an object, which is more robust than the boundary geometry in anatomic shape analysis. To achieve a smooth radial flow, Damon has introduced general skeletal structures that relax the appropriate Blum conditions (see Section 2.5.3). This relaxation gives rise to a general case of skeletal structure, called an s-rep, which has shown promise in anatomic shape analysis. Specifically, Pizer et al. (Pizer et al., 2019) proposed using these relaxed conditions in fitting a non-branching skeletal structure (a non-branching s-rep) to a slabular object¹, as opposed to deriving the structure from the boundary. Starting from a discretely sampled version of this s-rep (“d-s-rep”), the implied region generated by the radial flow (see Sections 2.5.1 and 2.5.3) needs to have geometric features that are reasonably matched to the target object, to be non-self-overlapping and to correspond well across a population of similar objects. To achieve all three aims, this chapter describes a two-stage d-s-rep fitting framework: first, I initialize a d-s-rep that has consistent skeletal branching topology within a population; second, I refine the d-s-rep with respect to geometric match qualities.

The first stage of the fitting process is based on Hong’s dissertation (Hong, 2018), in which the topology of the population of objects is taken into account. Many human structures are diffeomorphic to ellipsoids; namely, there exists a bijection that is differentiable with differentiable inverse and that maps the boundary of such structure to that of an ellipsoid. Moreover, many of them have no branching subfigures. The ellipsoid is used as the basic slabular object because 1) its Blum medial structure is an s-rep with a flat, non-branching structure whose edge is an ellipse that directly maps to the crest of the ellipsoid; 2) the map between its boundary and its skeletal representation has an analytic form; 3) it is easy to parameterize both the boundary and the skeletal sheet in such a way that the crest curve on the boundary

¹As defined in Chapter 1, An object is referred to as a *slabular* object if its length is greater than its breadth, which is greater than its width.

of the ellipsoid divides it into two open halves, namely, a north and a south side; 4) there is a natural correspondence between eccentric ellipsoids via the ratios of their largest, middle-sized, and smallest principal radii respectively.

Based on the above observations, I initialize an s-rep via an estimated mapping between the target object and its skeletal representation given 1) a diffeomorphic map ϕ between the boundary of the target object and the boundary of an ellipsoid and 2) the s-rep of that ellipsoid, which can be analytically computed provided an ellipsoidal boundary. Specifically, I extend the inverse boundary mapping ϕ^{-1} to the skeletal sheet, producing what is called a stratified diffeomorphism ϕ' . Then, I apply ϕ' to the ellipsoid's d-s-rep to produce the initial fit.

The second stage refines the initial fit to have better geometric properties. In achieving a better fit, this refinement relaxes all the Blum conditions mentioned in section 2.5.1. Specifically, the refinement avoids self-overlapping skeletal structures and minimizes geometric differences from the target objects. The geometric properties are measured in three terms; they are the boundary positions, the tangent fields and the radial curvature (Damon, 2008) – corresponding to zero-, first- and second-order geometric properties, respectively.

The fitting method depends on the above geometric properties of a dense collection of spokes, while the initialization stage outputs a coarse representation of a d-s-rep. To fill this gap, I implement an interpolation method for d-s-reps in Section 3.4. Different from the previous interpolating methods ((Vicory, 2018; Pizer et al., 2019)), my method interpolates spokes' directions using skeletal coordinates (see Section 2.6.1), yielding smooth directional vectors on the skeleton.

The remaining of this chapter is organized as follows. Section 3.2 gives a formal definition of s-reps. Section 3.3 details the method of initializing an s-rep. Section 3.4 describes the interpolation of a d-s-rep into a continuous s-rep. Section 3.5 discusses the objectives of the refinement stage. To demonstrate the proposed framework, I fit s-reps to the hippocampus and to the caudate described in Section 2.2. Section 3.6 shows the results from both the initialization and refinement stages. Section 3.7 discusses the accomplishments of this chapter and proposes future work. The proposed framework is publicly accessible in SlicerSALT (Vicory et al., 2018). Additionally, I provide a python package called “shanapy” that includes the s-rep fitting methods.

3.2 What is an S-rep?

In this dissertation I restrict myself to s-reps of the form (M, S) , where M is formed from two smooth collocated sheets and thus has ellipsoidal topology, and where S is a field of non-crossing spoke vectors on M (see fig. 3.1 left). More precisely, (M, S) is formed from 3 components:

1. the component (M_+, S_+) , where M_+ is the south side of M and M_+ is mapped to the southern boundary of the object via S_+ ,
2. the component (M_-, S_-) , where M_- is the north side of M and M_- is mapped to its northern boundary via S_- and




Figure 3.1: Left: a 3D s-rep for a hippocampus. The grid connects skeletal points. The yellow curve is the skeletal edge. The vectors proceeding from the skeletal points are the spokes. Middle: the 3D implied boundary of this s-rep. Right: the surface mesh of this hippocampus being fitted.

3. the component (M_0, S_0) , where M_0 is an edge curve of M that divides M into two open submanifolds and M_0 is mapped to a crest curve on its boundary via S_0 .

The submanifolds M_+, M_-, M_0 compose the colocated skeletal sheets (in 3D, curves in 2D). Moreover, these maps S_+, S_- and S_0 are continuous vector-valued functions. A *spoke* is a vector $s \in S$ that has a length r and direction U , and it proceeds from a point on M_+, M_- or M_0 , respectively. The spoke s points from a skeletal point p to a point $\mathcal{B}(p)$ on the boundary.

In Blum’s view of medial models and Damon’s view of skeletal models (Damon, 2003), the two colocated sheets M_+ and M_- are thought of as a single surface M with boundary ∂M (i.e., M_0). The locations on $M \setminus M_0$ are the “smooth” points and the locations on M_0 are the edge points. In this conventional view, associated with each smooth point p there are two spokes $s_{+,p}$ and $s_{-,p}$. Notice that in a skeletal model the length of $s_{+,p}$ is not necessarily equal to that of $s_{-,p}$. In addition, the spokes swing infinitely fast as their base points approach and pass around the edge, as detailed in Section 3.5.

To establish good correspondence among a population of s-reps, I sample the skeletal sheets and the associated spokes in a consistent way as follows. I discretize M_{\pm} (i.e., the disjoint union of M_+ and M_-) into grids and M_0 into intervals, as shown in fig. 3.2. As a consequence, the vector fields S_+, S_-, S_0 are also sampled at each skeletal point. I denote the result from the discretization as a “d-s-rep” $(\widehat{M}, \widehat{S})$.

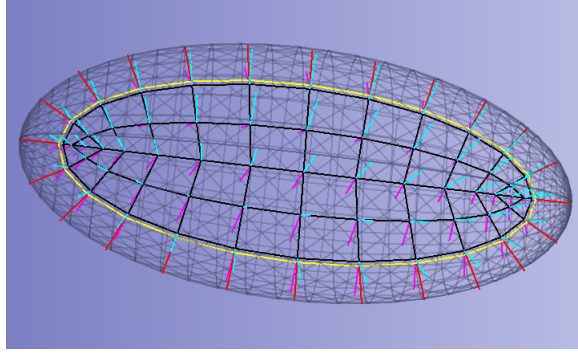


Figure 3.2: An oblique view of an ellipsoid’s skeletal representation, where the skeletal sheet is discretized as a grid. The interior black grid is the discretized collocated skeletal sheets M_{\pm} . Each grid point on M_{\pm} is associated with a spoke, pointing to either the north (magenta) or south (cyan) side of the ellipsoid boundary. The yellow curve is the edge M_0 , from which red spokes point to the crest curve of the ellipsoid.

3.3 Automatic Initialization Methodology

Since many human organs have a similar boundary and skeletal topology with a slabular object, we can treat them as different derivations from a symmetric slabular object such as an ellipsoid. Such uniform boundary and skeletal topology yields statistics-friendly (mutually corresponding) shape features. Inspired by this, I establish an initial fit to an object by deforming a skeletal representation of an ellipsoid to the skeletal representation of that object.

To automatically obtain the corresponding ellipsoid and perform the deformation, I estimate a stratified diffeomorphism on boundaries via mean curvature flow and on skeletons via the Thin Plate Splines (TPS) algorithm. Specifically, for each object I initialize a skeletal representation in three steps:

1. I establish the correspondence between the surface of the object and that of the ellipsoid (Hong, 2018).
2. I derive a continuous skeletal representation of the ellipsoid and discretize it.
3. I infer the interior mapping from the ellipsoid to the object and apply it on the discrete skeletal representation of ellipsoid (Hong, 2018).

There is an easily computed smooth normal vector field on the boundary. To establish the correspondence between the object and an ellipsoid, I perform mean curvature flow (Taubin, 1995) on the boundary at each step, given the normal vector field. After each step the normal vector field is updated on the deformed boundary. I keep track of a chain of these local diffeomorphisms at each step. Consequently, the final map from the object to the ellipsoid is the composition of these local diffeomorphisms. The above procedure is summarized in Algorithm 2.

While this mean curvature flow can ultimately converge to a sphere, on the way it becomes a near-ellipsoid, denoted as \tilde{E} . In fitting s-reps of the hippocampus and the caudate, I found that the flow using a large enough K in Algorithm 2 can produce surfaces that are close enough to ellipsoids.

Algorithm 2 Deform a closed surface to an ellipsoid (a basic implementation)

Input: A surface mesh \mathcal{B} of an object.

Result: A corresponding ellipsoid's surface.

$\mathcal{B}_1 \leftarrow \mathcal{B}$ // This surface \mathcal{B}_1 will be deformed in the process.

for $i \leftarrow 1 : K$ **do**

Compute the normal vector N at every point $p^i \in \mathcal{B}_i$.

Compute the mean curvature H at every point $p^i \in \mathcal{B}_i$.

Update the point p^i with the deformed point p^{i+1}

$$p^{i+1} = p^i + \delta H N \quad (3.1)$$

The points p^{i+1} form a new surface mesh \mathcal{B}_{i+1}

end

Next, I derive the skeletal representation for that best-fitting ellipsoid. Let r_1, r_2, r_3 be radii of the ellipsoid. Then any point (x, y, z) within or on that ellipsoid satisfies

$$\frac{x^2}{r_1^2} + \frac{y^2}{r_2^2} + \frac{z^2}{r_3^2} \leq 1 \quad (3.2)$$

where I assume $r_1 > r_2 > r_3$.

The flat skeletal (here medial) sheets M_{\pm} reside in the middle of the ellipsoid, and M_0 is an ellipse. Let a skeletal point $p \in M$ have coordinates (x_p, y_p, z_p) and let m_1, m_2 be the radii of M_0 . Then I have

$$\frac{x_p^2}{m_1^2} + \frac{y_p^2}{m_2^2} \leq 1 \quad (3.3)$$

where

$$m_1 = \frac{r_1^2 - r_3^2}{r_1} \quad \text{and} \quad m_2 = \frac{r_2^2 - r_3^2}{r_2} \quad (3.4)$$

The coordinate z_p is a constant. In this dissertation I set $z_p = 0$ for convenience.

I must now discretize the skeletal sheets of the ellipsoid. The location of the skeletal sample points was motivated by two considerations. First, the quality of the fit of the model into any object depends especially heavily on the fit at the crest and especially at its vertices (local extrema of convex curvature). For an ellipsoid these object surface loci correspond to the edge curve of the medial ellipse and to that curve's two vertices, respectively. Thus, as illustrated in fig. 3.2, on the medial ellipse I generate samples more densely at and near the vertices of the edge and less densely in the middle of the edge.

The second consideration is that, for consistency with the skeletal modeling of 3D objects, the sampling of the 2D medial ellipse should also be skeletal, or more precisely, medial. That is, the samples should be arranged along normals to the ellipse’s boundary, i.e., the edge curve M_0 . These normals are sampled at the boundary, at a position slightly eroded from the boundary, and then equally spaced between that position and the ellipse’s medial axis (including a point on the medial axis). The rate of sampling along the normals is a parameter choosable by the user.

I sample the medial points along the edge curve (the yellow curve in fig. 3.2) as follows. The points on the edge are calculated as $(m_1 \cos \theta, m_2 \sin \theta, 0)$, where θ is uniformly sampled. The result is that the arclengths between adjacent edge points in the flatter region of the ellipse are notably greater than those near the vertex.

From each of these sampled points on the edge, I compute all the interior sample points (the vertices of the black grid in fig. 3.2) as locations along the ellipse’s normals at the respective edge points. The resulting grid is made up of quadrilaterals. Those near the ends of the ellipse’s medial end-curve are quite small, and four of the quadrilaterals are degenerate, having two collinear sides.

I now must compute the spokes at each skeletal sample point. For each point $p \in M_{\pm}$, I find the nearest boundary point $\mathcal{B}(p)$ on the corresponding side of the ellipsoid analytically. The points p and $\mathcal{B}(p)$ form the endpoints of a spoke. Having both boundary points and corresponding skeletal points, I have formed the discrete skeletal representation $(\widehat{M}, \widehat{S})$. Because the boundary of the near-ellipsoid \widetilde{E} can be arbitrarily close to that of the best-fitting ellipsoid, this discrete skeletal representation can as well fit the near-ellipsoid \widetilde{E} .

Finally, I deform the discrete skeletal representation of \widetilde{E} to the target object. Having the local diffeomorphisms in the process of flowing the boundary points of the object to those of \widetilde{E} , now I apply the inverse of these diffeomorphisms in a reverse order step by step on the boundary points of \widetilde{E} . Meanwhile, I infer the deformation on M via TPS and apply it to give the base points of spokes in \widehat{S} . This yields an initial fit of an s-rep to the target object.

Example 4.1 The following example shows the process of fitting an initial s-rep to the boundary of a hippocampus.

Given the surface mesh of an object (see fig. 3.3), I build a diffeomorphism from this boundary to that of a near-ellipsoid via mean curvature flow. At each step of the flow the boundary becomes smoother. To make the steps small enough and highly likely to be correct and easy to invert, I set $K = 500$ in Algorithm 2 in this example. At the end of the flow, I can achieve a near-ellipsoid (see the middle figure in fig. 3.3) that is close enough to its best-fitting ellipsoid (shown as the yellow mesh in the right figure). The right figure in fig. 3.3 shows the discretized skeletal sheet of that best-fitting ellipsoid. The resulting s-rep also well fits to the near-ellipsoid.

Now I show the process of deforming the s-rep of the best-fitting ellipsoid back to fit the target object. To do so, I build a stratified diffeomorphism to apply to the discrete skeletal points and to the corresponding spoke ends. I estimate this diffeomorphism using TPS based on the inverse deformation from the target surface to the ellipsoidal surface. With the stratified diffeomorphism I transform both the heads and tails of spokes in the ellipsoidal s-rep, producing an initial s-rep of the target object, as shown in figs. 3.4 (a-c). This example shows that the stratified diffeomorphism results




Figure 3.3: Left: Input a 3D shape of a hippocampus represented as a surface mesh. Middle: The overlay of the target shape (red) and a near-ellipsoidal shape after mean curvature flow (black). Right: the best fitting ellipsoid (yellow transparent mesh) to the black mesh in the middle figure and its discrete s-rep (consisting of cyan and magenta spokes, a discrete skeletal mesh and an edge curve).

in the initial s-rep whose end points can well fit to the target boundary. Moreover, the non-branching topology of the skeleton is preserved. Figs. 3.4 (d-e) are the results from the interpolation and from the refinement, respectively, which are discussed in the following sections.

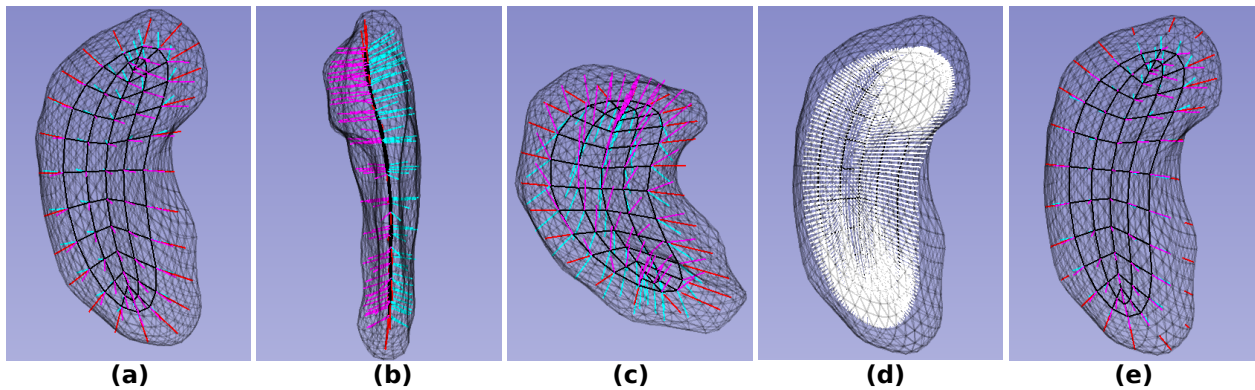


Figure 3.4: (a-c) The result of the initial fit to a hippocampus from different perspectives. The gray transparent mesh is the input target surface mesh. The black grid in the middle of the object is a discrete representation of the skeletal sheet. The cyan and magenta spokes indicate spokes pointing to different sides. The red spokes are edge spokes. (d) The white spokes are the interpolated northern spokes of the initial fit. (e) The refined s-rep.

3.4 Interpolation

The interpolation of an s-rep is a method that produces an arbitrarily dense set of spokes in an object. This method allows me to evaluate the fit of the s-rep to the object's boundary. Also, the method allows me to produce continuous interior level surfaces of radial flow (see section 1.2.1) to capture within-object shape features. This section discusses the methodology of that interpolation. Given a d-s-rep, I treat the interior elements (i.e., smooth skeletal points and the attached spokes) and others (i.e., edge points and edge spokes) differently as follows.

3.4.1 Interpolate Skeletal Sheets and Attached Spokes

Interpolating spokes at smooth skeletal points involves (1) interpolating skeletal points' coordinates, (2) interpolating spokes' directions and (3) interpolating spokes' radii.

Skeletal points. I fit a cubic Hermite patch (Han et al., 2006) to each cell on a skeleton and interpolate the cell into a continuous patch. Each cell is formed by based points of 4 primary spokes. I interpolate the skeletal points in a cell-by-cell manner.

Spokes' directions. The spokes' directions should form a smooth unit vector field on a skeleton. To achieve this smoothness, I interpolate spokes' directions using skeletal coordinates (see Section 2.6). Specifically, let $\widehat{\nabla\tau_1}$ denote the direction pointing away from the spine, and let $\widehat{\nabla\theta}$ denote the direction circulating the edge curve in a counter-clock wise way. I interpolate spokes' directions first along $\widehat{\nabla\tau_1}$ then along $\widehat{\nabla\theta}$, see fig. 3.5. Along each direction, I treat the spokes' directions as quaternions and use Spherical Linear Interpolation to interpolate the directions between primary spokes that are show as black line segments in fig. 3.5.

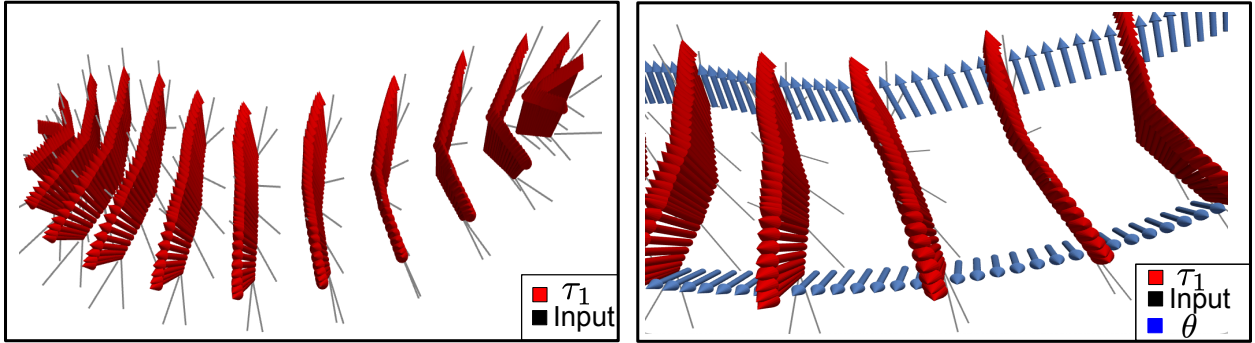


Figure 3.5: Interpolating spokes' directions using the interior coordinates defined in (Pizer et al., 2021). Left: the interpolation method first interpolates spokes' directions along $\widehat{\nabla\tau_1}$ directions. These interpolated directions are shown as red arrows. The base points of the arrows are interpolated via a Hermite patch as described in text. Right: The interpolation along $\widehat{\nabla\theta}$ direction. The interpolated spokes directions are shown as blue arrows.

Spokes' radii. I interpolate spokes' radii via a successive subdivision method. Given a cell bounded by 4 skeletal points (each associated with a spoke), I interpolate spokes' radii (Pizer et al., 2019) at centers of edges and the center of the cell via

$$r_{p_{\frac{1}{2}}} = U_{p_{\frac{1}{2}}} \cdot \left[\frac{1}{2}(S_{p_0} + S_{p_1}) \right] - \frac{d^2}{4}(S_{p_0} \cdot U_{ww,p_0} + S_{p_1} \cdot U_{ww,p_1}) \quad (3.5)$$

where subscripts $p_0, p_{\frac{1}{2}}$ and p_1 are respectively the start, middle (interpolated) and end positions. The scalar d is the distance from $p_{\frac{1}{2}}$ to p_0 . The unit vector $U_{p_{\frac{1}{2}}}$ is the interpolated spoke direction at position $p_{\frac{1}{2}}$, while $S_p = r_p \cdot U_p$. Further, U_{ww,p_0} and U_{ww,p_1} are the second derivatives of the spoke directions along the interpolating direction w at point p_0 and p_1 , respectively. The first term in eq. (3.5) can be understood as the average radius at the middle position. This

projection should be a good estimated radius of the interpolated spoke. The second term corrects this estimate based on the skeletal geometry. For the center of a sub-grid quadrilateral, I interpolate the radii first along $\widehat{\nabla\theta}$ then $\widehat{\nabla\tau_1}$ and also along $\widehat{\nabla\tau_1}$ then $\widehat{\nabla\theta}$, and I average the results.

These interpolated positions produce sub-grid cells of half size in each direction, producing a subdivision step. I repeat this process, yielding the successive subdivision for spokes' radii.

3.4.2 Interpolate Spokes in a Crest Region

After the above interpolation, the edge curve of the skeleton is densely sampled. These dense edge points are taken as base points of spokes to be interpolated in the crest region. Therefore, the interpolation in the crest region is turned into a problem of interpolating boundary points in the region of the crest curve. I first interpolate these boundary points along the crest curve and then interpolate across the crest curve on the boundary as follows.

Because I want the directions of edge spokes to align with the tangent directions of the skeleton ((Damon, 2003)), I first interpolate edge spokes that connect edge points of the skeleton and crest points on the boundary. To this end, I infer a continuous deformation from primary edge spokes, which are determined in the initialization stage, sending pairs of these edge points and their corresponding crest points to TPS. The estimated deformation via TPS can flow the interpolated edge points due to the previous interpolation (see Section 3.4.1) to the boundary, producing interpolated crest points (shown as the cyan points in the left panel of fig. 3.6).

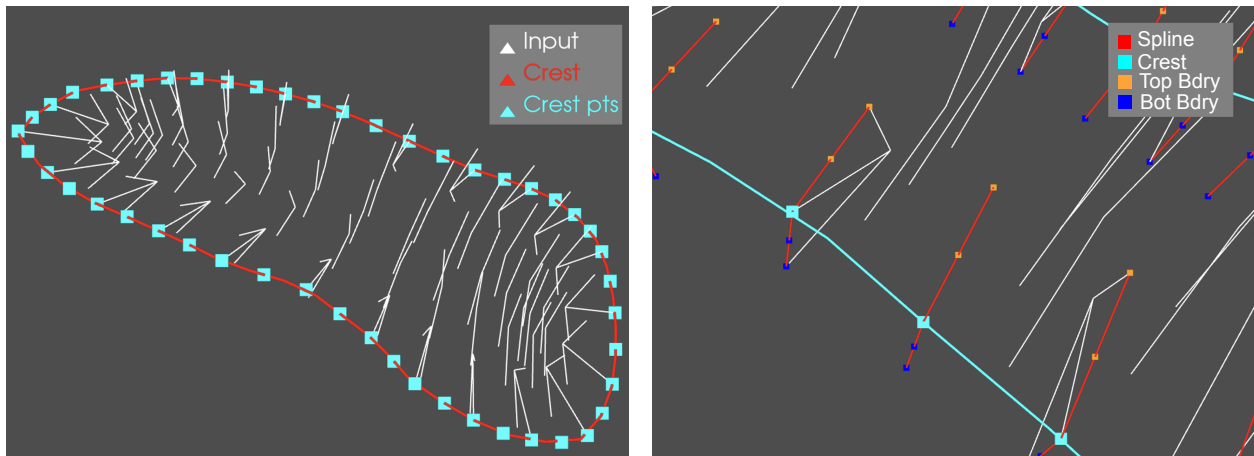


Figure 3.6: Interpolating boundary points for spokes in the crest region after interpolating spokes on the smooth skeletal sheet. The white line segments are the spokes of the input s-rep. Left: Interpolation of crest points by TPS. The ends of primary edge spokes are the input of TPS, which outputs the interpolated crest points. Right: Interpolation of boundary points around the crest. I interpolate boundary point along each spline to form spokes around the crest region. See the text for more details.

Next, in order to allow spokes to swing smoothly as approaching and passing across the edge, I fit cardinal splines (Schoenberg, 1973) on the boundary across the crest. Each spline, shown as a red curve in the right panel of fig. 3.6,

passes through (1) a top boundary point (shown as an orange point) associated with an outermost non-edge spoke (2) a crest point (shown as a cyan point) and (3) a bottom boundary point (shown as a blue point) associated with an outermost non-edge spoke. Along these splines, I interpolate boundary points on the top and bottom sides of the boundary, resulting in the orange points (on the top side) and blue points (on the bottom side). These interpolated boundary points are taken as the ends of spokes around the crest.

Were this whole interpolation to be applied continuously, it would turn a d-s-rep into a corresponding s-rep. Thus, assuming this interpolation, I can simply refer a d-s-rep as an s-rep. I will use the term “s-rep” in this way for the remainder of this dissertation.

3.5 Refinement

This section discusses the second stage of the fitting. While taking advantage of Damon’s relaxation of medial conditions, I wish to deviate only slightly from the medial conditions while maintaining the non-self-overlap and non-branching topology. I accomplish this by designing an optimization in which I penalize those deviations that can change most strongly in the initialization stage (see Section 3.5.1). In Section 3.5.2 I address the penalties involving non-edge spokes. Then in Section 3.5.3 I address the refinement on edge spokes. Finally, I discuss the implementation details in Section 3.5.4.

3.5.1 Relax Blum Conditions

Pizer et al. (Pizer et al., 2013a) have shown that the following enumerated objectives often provide a tighter fit to the geometry of the target object than an unbranching purely Blum-based fit. I will compare under two conditions the properties of the fit between the boundary \mathcal{B} of the target object and $\tilde{\mathcal{B}}$, that is created from the ends of the primary and interpolated spokes. The two conditions compared are when the full set of spokes satisfy the Blum conditions relative to $\tilde{\mathcal{B}}$ and when they are released from those conditions, according to Damon’s skeletal geometry.

1. In both conditions \mathcal{B} and $\tilde{\mathcal{B}}$ do not precisely agree, even at the primary spokes. I do want the differences between \mathcal{B} and $\tilde{\mathcal{B}}$ to be small.
2. In both conditions the tangent planes of \mathcal{B} and $\tilde{\mathcal{B}}$ at corresponding spokes do not precisely agree, even at the primary spokes. I do want the differences between the tangent planes of \mathcal{B} and those of $\tilde{\mathcal{B}}$ to be small. However, in the case of Blum conditions the spokes are orthogonal to $\tilde{\mathcal{B}}$, i.e., they satisfy the partial Blum condition; Damon’s skeletal geometry allows this condition not to hold.

3. According to (Damon, 2003), for a smooth point $p \in M_{\pm}$ in both conditions local self-overlap is prevented by the condition

$$r_p < \min\left\{\frac{1}{\kappa_{ri,p}}\right\}, \quad \text{if } \kappa_{ri,p} \in \mathbb{R}_*^+ \quad (3.6)$$

where r_p is the radius of the spoke associated with the point $\tilde{\mathcal{B}}(p)$ and $\kappa_{ri,p}$ is a positive real number denoting the i^{th} principal radial curvature at p . In my research, I wish this condition to hold. I handle this in the objective function by a ReLU function described in Section 3.5.2.3.

4. In both conditions an edge point p represents the maximum curvature point on the boundary at $\tilde{\mathcal{B}}(p)$. I wish $\frac{1}{r_p}$ to agree with the maximum curvature of the target boundary \mathcal{B} near $\tilde{\mathcal{B}}(p)$.

I also take advantage of another flexibility allowed by Damon’s skeletal geometry, namely, that the two spokes based on the skeletal point p need not be of the same length. However, because of the properties of the initialization stage, I found it unnecessary to include a term as another objective making the difference in these spoke lengths small.

I approach the objectives 1–3 by penalizing deviations from those conditions with penalties L_0, L_1 and L_2 . The penalty L_0 measures the overall squared distance from tips of spokes to the target boundary; L_1 measures the overall deviation of spokes’ directions from the perpendicularity to the boundary; L_2 measures the overall degree of violating objective 3.

3.5.2 Refinement Objectives for Interior Spokes

Taking all the above factors into account, I formulate the objective function for the refinement of interior spokes as

$$L = \alpha L_0 + \beta L_1 + \gamma L_2 \quad (3.7)$$

where weights α, β and γ are (non-negative) scalar values to weight each penalty L_0, L_1 and L_2 , respectively. Specifically, one needs to tune up α and/or β if an initial fit was far from the object’s boundary, and one needs to tune up γ if an initial fit presents self-overlaps.

In general, it is only the parameters of primary spokes \hat{S} , emanating from the grid points (i.e., \widehat{M}_{\pm}), that are optimized. The refinement, however, evaluates not only the primary but also the interpolated spokes in terms of the penalties L_0, L_1 and L_2 . By doing so, the resulting primary spokes can give good interior geometric features of an object. In the following, I will detail each penalty for a given spoke.

3.5.2.1 Penalize Boundary Distances

To compute the penalty L_0 , given the target surface mesh, I generate a 3D image of signed distances from the target object boundary. For each spoke s_p , I measure the distance d_{s_p} from its tip to the boundary according to the distance

image. Thus, the total squared boundary distance penalty is

$$L_0 = \sum_{p \in \widehat{M}_{\pm}} \sum_{p' \in \mathcal{N}(p)} d_{s_{p'}}^2 \quad (3.8)$$

where \widehat{M}_{\pm} represents the sampled skeletal points and $\mathcal{N}(p)$ denotes the interpolated skeletal points near to the skeletal point p (see Section 3.4).

3.5.2.2 Penalize Non-orthogonality to the Boundary

I penalize the non-orthogonality L_1 of spokes to a boundary. At each point of the signed distance image, I compute the direction of the gradient which is perpendicular to the distance-level-set through that point. In particular, the gradient direction at a tip of a spoke is approximately perpendicular to the target boundary. I compare this gradient direction and the spoke direction, yielding the local penalty L_1 at p , i.e., $1 - \cos \theta_{p'}$. Then I compute the overall penalty L_1 by

$$L_1 = \sum_{p \in \widehat{M}_{\pm}} \sum_{p' \in \mathcal{N}(p)} (1 - \cos \theta_{p'}) \quad (3.9)$$

where $\theta_{p'}$ is the angle between the spoke direction $U_{p'}$ and the gradient direction at the tip of the spoke.

3.5.2.3 Penalize Self-overlaps within an Object

Objective 3 aims to avoid local self-overlaps and “singular kinks” on the implied boundary. To achieve this objective, (Damon, 2003) introduces the “radial shape operator” S_{rad} at $p \in M_{\pm}$, whose eigenvalues are called the principal radial curvatures $\kappa_{ri,p}$. This operator is a modification of the differential geometric shape operator (O’Neill, 2006), using instead the unit vector field in the spoke directions. This is usually not orthogonal to the medial sheet. Consequently, S_{rad} is not a symmetric operator if the partial Blum condition does not hold.

However, S_{rad} still captures the shape of the region bounded by the implied boundary; as well the interior regions bounded by any of the level surfaces of the radial flow (formed by dilating the skeletal sheet along each spoke by τ times its spoke length, for some $0 \leq \tau \leq 1$).

Computationally, the operator S_{rad} has a matrix representation \widehat{S}_{rad} . Because the partial Blum condition does not necessarily hold, the matrix \widehat{S}_{rad} may or may not have real eigenvalues. On the one hand, for real principal radial curvatures at p from the eigen-analysis on \widehat{S}_{rad} , I can compute for each single spoke s_p of radius r_p the deviation from the condition eq. (3.6) by a ReLU function:

$$\delta_{s_p} = \max(0, r_p \kappa_{r1} - 1) \quad (3.10)$$

where $r_p \kappa_{r1}$ is the maximal eigenvalue of the matrix $r_p \widehat{S_{rad}}$ (Han, 2008). This matrix is constructed as

$$r \widehat{S_{rad}} = \left[\left(\frac{\partial S}{\partial w} - \frac{\partial r}{\partial w} U \right) Q^T (Q Q^T)^{-1} \right]^T \quad (3.11)$$

where S is the spoke vector at p ; U is a unit vector denoting the direction of the spoke; w is a velocity direction of a curve on $T_p M$ through p , and

$$Q = \frac{\partial p}{\partial w} (U^T U - I) \quad (3.12)$$

where I is the identity matrix. As described in eq. (3.13), the eigenvalues are used to form L_2 . On the other hand, if the principal radial curvatures at p are not real, I let $\delta_{s_p} = 0$.

Finally, the L_2 penalty of an s-rep is

$$L_2 = \sum_{p \in \widehat{M}_{\pm}} \sum_{p' \in \mathcal{N}(p)} \delta_{s_{p'}} \quad (3.13)$$

3.5.3 Refine Edge Spokes

The initialization may result in edge spokes that violate objective 4; that is, edge spokes may deviate from the tangent plane of the s-rep. The current solution to refine this condition is to tune spokes' lengths, followed by relocating base points of edge spokes.

To do so, I start from adjusting lengths of edge spokes such that tips are on the boundary. Then, according to the maximal curvature $\kappa_{\mathcal{B}(p)}$ at a boundary point $\mathcal{B}(p)$, I relocate the skeletal end of this edge spoke so that the spoke's length becomes

$$r_p = \frac{1}{\kappa_{\mathcal{B}(p)}} \quad (3.14)$$

In the shape with sharp ends (e.g., the caudate nuclei), the relocation can make skeletal points of edge spokes far from the outermost interior ones (see the bottom figure in fig. 3.7). This strategy can be viewed as extending M_{\pm} toward the sharp ends.




Figure 3.7: Fit of an s-rep to a caudate nucleus in the dataset described in section 2.2. Top: the initial fit. Bottom: the refined s-rep. In the refined s-rep, the skeletal points of edge spokes (red line segments) are relocated according to curvatures on the boundary.

3.5.4 Implementation

In fitting the example hippocampi shown in fig. 3.4, I interpolate the discrete s-rep to evaluate the penalties L_0 , L_1 and L_2 (corresponding to eqs. (3.8), (3.9) and (3.13), respectively) as the primary spokes being refined in the procedure. It is recommended to have 7 interpolated spokes between every two neighboring primary spokes, yielding moderately interpolated s-reps. The interpolation result is shown in fig. 3.4 (d). In the example, I set the weights α , β and γ in equation (3.7) as 0.004, 20 and 50, respectively.

The optimization method that I use is chosen with the following considerations in mind:

1. The optimization can be over as few parameters as those describing a single primary spoke and as many as those describing the full set of primary spokes. So the method must be able to handle optimization over a few or over hundreds of parameters.
2. The objective function (see eq. (3.7)) has many local optima, so a multiscale method is preferred.
3. No analytic expression for the objective function is available.
4. In each iteration of the optimization, as many as all of the primary spokes have been modified, so interpolation of many or all of the more finely sampled spokes must be recomputed. Thus, computation of the gradient of the total loss by numeric approximation at every iteration would be unacceptably inefficient. Together, considerations 3 and 4 make a method involving gradients of the objective function inappropriate.

As a result, I have chosen a derivative-free optimization algorithm named “NEW Unconstrained Optimization Algorithm” (NEWUOA) (Powell, 2004) to minimize the objective function.

3.6 Experimental Analysis

I tested the fitting method on a machine with Ubuntu 18.04.3 LTS, 15.5 GB memory, Intel® Core™i7-8750H CPU @ 2.20GHz and GeForce GTX 1060. With the above-mentioned default parameters (e.g., the sampling rate of a skeletal sheet), a direct initialization takes about 2 minutes. A refinement from a previous result and the default parameters (e.g., the subdivision level) takes about 11 minutes.

I fit s-reps to the surface meshes of the hippocampi. Fig. 3.8 shows the initial and refined implied boundaries of a typical hippocampus. The initial implied boundary, shown in fig. 3.8 (a), presents a singular kink in the middle, and it has been resolved by the refinement (see fig. 3.8 (c)). The initial implied boundary has overall small distances from the target boundary positions (see fig. 3.8 (b)), while the refinement reduces those distances around the crest region (see fig. 3.8 (d)).

To examine the quality of the initial and refined fits, I compute the distances from tips of primary and interpolated spokes to the target boundaries at the recommended subdivision level. The statistics of these distances are shown




Figure 3.8: (a): The implied boundary of the initial fit to a hippocampus. (b): The heat map of unsigned distances measured locally between the initial implied and target boundaries. (c): The implied boundary of the refined skeletal representation. (d): The local distances (in mm) from the target boundary to the implied boundary of the refined skeletal representation.

in fig. 3.9. The results show that the worst case of the initialized s-reps has an object-average distance of $0.48mm$ between the implied and the target boundaries, as compared to $0.41mm$ after refinement. Due to the refinement, the object-average distance from implied boundaries to the target boundaries is improved by 11% on average. The improvement could be more significant if case-by-case parameters were employed. With slightly tuned parameters, I show the improvement of local distances between the implied and target boundaries in an example hippocampus in figs. 3.9 (c) and (d).




Figure 3.9: (a): The object-average distances (measured in mm) between the implied and target boundaries in sorted order 177 hippocampi. (b): The kernel density estimation (KDE) of histograms of average distances between the implied and target boundaries in these 177 hippocampi before and after the refinement. (c): the box plot of local distances from the target boundary in one typical case. (d): the histogram of the local distances in the same case.

With regard to the high order boundary geometric features, singularities on implied boundaries can be resolved by the refinement (see figs. 3.8 (a) and (c)). To quantify local shape on the initial and refined implied boundaries as well as target boundaries, I measure the curvedness C (Koenderink, 1990)² on each boundary. The curvedness at a point on a surface is defined as

$$C = \sqrt{(\kappa_{max}^2 + \kappa_{min}^2)/2} \quad (3.15)$$

²The corresponding notation used in the reference is R .

where κ_{max} and κ_{min} are the maximal and minimal normal curvature, respectively. The metric C is close to 0 when the local shape is flat; C gets larger when the local shape is sharp.

I present the histograms of local curvedness measured on an initial implied boundary, its refined implied boundary and its target boundary in fig. 3.10. The smooth target boundary (see the dashed curve) resulting from SPHARM-PDM has a relatively tight distribution of curvedness. The initial fit (see the blue curve) implies many local patches with high curvedness on the boundary; the refinement smooths a lot of them (see the orange curve).

The pattern of points adopted to sample the skeletal surface was described in Section 3.3. The measurements used for figs. 3.9-3.10 were made using this adopted pattern. Previously, I had made the hippocampal skeletal fittings and the same measurements using a pattern of points where the quadrilaterals were closer to being the same size (Hong, 2018), so as compared to the adopted pattern, the central part of the basis skeletal ellipse was sampled more finely than with the final sampling pattern and it was sampled less finely near the vertices of the ellipse. The histograms corresponding to figs. 3.9-3.10 for this previous pattern showed the performance with the adopted pattern to be somewhat superior.

I have succeeded in fitting s-reps to the caudate nuclei well (see fig. 3.7 for example). I intend to apply the framework on other complex shapes. However, the conventional mean curvature flow can develop singularities in evolving local shape that have sections that are nearly rotationally symmetric and have convex mean curvature (Huisken, 1990), which happens, for example, with mandibles (see fig. 3.11). Such a problem can be solved by replacing the flow with a conformalized mean curvature flow algorithm as studied in (Kazhdan et al., 2012), as discussed in section 6.2.1.

3.7 Conclusions and Discussion

The essential motivation of fitting s-reps is to capture within-object shape features for statistical analysis. The fitting process should be automatic and independent on an explicitly pre-defined template. The implied boundary by a good s-rep should approximate the target boundary as to boundary positions, tangent fields and curvatures. To obtain such an s-rep, this chapter has described a two-stage fitting framework: the first stage constructs a stratified diffeomorphism




Figure 3.10: The histogram of the curvedness on the initial and refined implied boundaries and target boundaries of a typical example. I apply the logarithm on the count of boundary point plus 1, yielding the y-value. The curvedness resulting from equation (3.15) is the x-value.




Figure 3.11: The conventional mean curvature flow develops singularities (around the circled region) for a mandible. Left: the input surface mesh of the mandible. Right: the result from a few iterations of the flow.

between an ellipsoid and the target object, resulting in an s-rep of fixed branching topology; the second stage refines the s-rep to minimize the deviations from the Blum conditions. As shown in the results, both initial fits and refined results approximate the boundary positions quite well. The refinement improves the fits of the skeletal structures to the boundaries in terms of boundary distances and curvedness. Moreover, Section 3.6 shows some complex shapes that can fail the proposed methods (see fig. 3.11).

This work has been published in (Liu et al., 2021). I have packaged the fitting algorithm in an open-source toolbox SlicerSALT (<http://salt.slicer.org/>). SlicerSALT provides the state-of-the-art statistical shape analysis methods as extensions of a customized 3D Slicer. It includes a number of powerful tools for visualizing and processing data. Also, the consistent user interfaces in SlicerSALT make it easy to operate. Like other extensions of 3D Slicer, SlicerSALT supports modules written in C++ and Python. In addition to SlicerSALT, I have released a simplified implementation of the fitting method in Python, called shanapy.

The fitting method can also be applied on 2D skeletal modeling, i.e., to study the the shape of boundary contours in 2D images. To this end, I construct a generalized cylinder with smooth and closed boundary surface by repeating the target contour and capping it smoothly and symmetrically at both ends. Then I apply the framework discussed in this chapter to fit a 3D s-rep to this generalized cylinder. An initial fit is obtained by slicing the 3D s-rep in the middle cross-section of the cylinder.

Given fitted s-reps of the hippocampus and of the caudate, I can extract the within-object shape features from each object, forming two feature matrices for the hippocampus-caudate pairs. Specifically, in the following chapter I collect the implied boundary points from s-reps of the hippocampus and of the caudate, forming the input of the joint analysis of the two objects.

CHAPTER 4: NON-EUCLIDEAN JOINT AND INDIVIDUAL VARIATION EXPLAINED FOR MULTI-OBJECT COMPLEXES

4.1 Introduction

Shape correlation of multi-object complexes in the human body can have significant implication for understanding the development of disease. While there exist many statistical methods that can extract correlation from multi-block data, very little research can effectively extract intrinsic shape correlation. It is especially difficult to extract the correlation when the involved objects have different variability in separate non-Euclidean spaces. For instance, as discussed in Section 1.1.1, the hippocampus and the caudate have different variability due to the development of ASD. Such differences can bias the analysis of shape correlation. Because of the need of analyzing the intrinsic shape correlation of the hippocampus and the caudate subject to the development of ASD, this chapter focuses on statistical correlation of within-object shapes in multi-object complexes.

The issues that I address in this chapter are the following. First, it is difficult to obtain robust patterns of shape correlations because of the complex data space and also because of different variability across multiple objects as mentioned above. Second, because the shape correlation is often coupled with variation of relative pose and size between objects, it is difficult to extract the intrinsic shape correlation (Gorcowski et al., 2007). Third, it is often difficult to interpret the shape correlation, making the results hypothetical to researchers and clinical users.

Many existing geometric and statistical methods attempt to address the above issues. For instance, deformation-based methods (e.g., (Gori et al., 2017)) analyze the deformation of multi-object complexes that involves the deformation of ambient spaces. However, the correlation of within-object shapes can be biased by the relative pose variation between objects. In this regard, statistical methods of multi-block data analysis show promise in extracting correlation in multi-object complexes (see Section 2.7.3 for a review of the relevant statistical methods). Nonetheless, few methods can effectively extract correlation of shape data due to the different variability across various objects.

To address the issues, this chapter extends existing geometric and statistical methods for extracting intrinsic joint shape variation¹ in two-object complexes. Specifically, I use boundary points on basis of s-reps to capture intrinsic shape features. Then I extract joint shape variation from two blocks of shape features via a statistical method that I call Non-EUclidean Joint and Individual Variation Explained (NEUJIVE). The resulting jointly varying shape features are the basis of hypothesis testing and classification. Because previous research (Hong, 2018) has shown significant

¹I use the term “shape correlation” interchangeably with “joint shape variation”.

association between the development of ASD and anatomic features of either the hippocampus or the caudate, I measure the joint shape variation of hippocampus-caudate pairs using the proposed methods and find that the joint shape variation is useful in classifying the ASD vs. non-ASD group.

NEUJIVE is designed for extracting robust joint and individual variation from two-object complexes where variability between the two objects can be largely different. In essence, NEUJIVE is an extension of a recent method called Angle-based Joint and Individual Variation Explained (AJIVE) (see Section 2.7.4). Different from AJIVE, NEUJIVE accounts for non-Euclidean properties of each object. Specifically, NEUJIVE consists of two critical steps: (1) convert shape descriptors to Euclidean representations and (2) extract the joint variation from two-block Euclidean data via AJIVE. Considering the spherical property of shape data after removal of translation and scale (Dryden and Mardia, 2016), I adopt a method called Principal Nested Spheres (PNS) (Jung et al., 2012) to convert spherical data to Euclidean variables. By combining PNS and AJIVE, NEUJIVE can effectively extract joint shape variation that is insensitive to different variability across objects.

To capture intrinsic shape features of objects with good correspondence, I use skeletally implied boundary points from s-reps. These implied boundary points capture smooth geometric features of individual objects for the joint analysis. Moreover, the s-rep implied boundary points allow straightforward interpretation of the joint shape variation.

The rest of this chapter is organized as follows. In Section 4.2 I revisit the driving problem and introduce the solution that utilizes joint shape variation in classification. To produce the joint shape variation, in Sections 4.3 to 4.5 I detail the proposed methods (1) that capture intrinsic shape features via s-rep implied boundary points, (2) that Euclideanize each block of shape features and (3) that extract joint variation from two Euclidean blocks via AJIVE. In Section 4.6 I demonstrate that NEUJIVE can extract useful patterns from simulated two-block non-Euclidean data. Then I show another simulation to verify the benefit of the joint shape variation in classification. Last, the proposed methods are verified in applications (including hypothesis testing and classification) to the ASD data (see Section 2.2). Section 4.7 discusses a general form and broader impacts of the methods. This chapter concludes with remarks and discussion in Section 4.8.

4.2 Classify ASD vs. non-ASD with Subcortical Shape

I want to take advantage of classification to demonstrate the effectiveness of NEUJIVE in extracting the joint shape variation from two-object complexes. An effective joint shape variation that has strong association with the development of ASD ought to improve classification performance. To this end, this section discusses the classification framework in which I use shape correlation in two-object complexes as classifying features.

To classify the ASD and the non-ASD group, I set up a binary classification model as

$$\hat{y} = g(w^T X) \quad (4.1)$$

where $w \in \mathbb{R}^{d \times 1}$ is a learnable weight vector, d is the dimension of the feature space and T means transpose. The matrix $X \in \mathbb{R}^{d \times n}$ represents the data of n configurations in the proposed feature space. I highlight that in this dissertation the rows of X are features, while the columns are cases. The output $\hat{y} \in \{0, 1\}^n$ denotes the vector of predicted labels for the n configurations. The symbol $g(\cdot)$ denotes a mapping to the predicted binary class labels.

The contributions of this chapter concentrate on producing the feature matrix X that captures the joint shape variation of the hippocampus and the caudate. Because this joint shape variation should be strongly associated with the development of ASD (Richards et al., 2020; Qiu et al., 2010a), the use of the joint shape variation ought to produce good classification performance. I propose methods to extract the joint shape variation in an unsupervised way, making use of PNS and AJIVE, as illustrated in fig. 4.1.

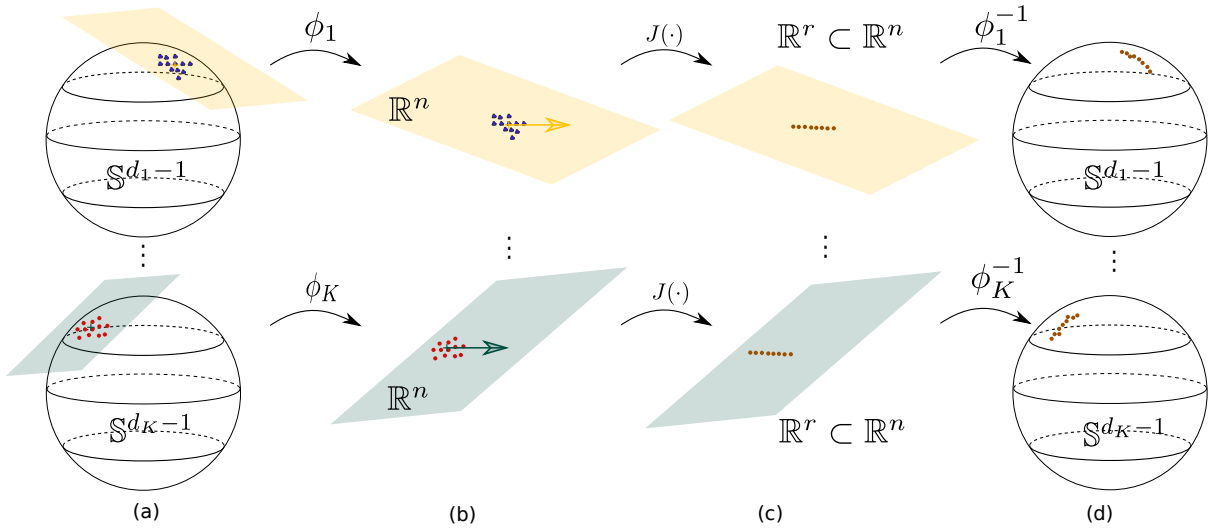


Figure 4.1: The intuitive illustration of extracting the joint components with NEUJIVE. (a) Input data live in non-Euclidean pre-shape spaces. (b) Construct the shared variation subspaces (\mathbb{R}^r represented by the yellow and green arrows) within the score spaces (\mathbb{R}^n shown as the yellow and green planes). (c) Project the data to obtain the joint components, shown as brown dots. Here, r denotes the dimension of the joint variation subspace. (d) Map the joint components (shown as the brown dots) of X_k back to the pre-shape space via ϕ_k^{-1} for each k .

As I focus on two-object complexes (i.e., the hippocampus-caudate pairs), let $X_1 \in \mathbb{R}^{d_1 \times n}$ denote the shape descriptors of the hippocampi, where d_1 denotes the dimension of the feature vector of a hippocampus and n denotes the number of complexes. Likewise, let $X_2 \in \mathbb{R}^{d_2 \times n}$ denote the shape descriptors of the caudate nuclei, where d_2 denotes the dimension of the feature space of the caudate. Both d_1 and d_2 are larger than the number of samples n .

The first step is to capture intrinsic shape features of each single object. Next, because shape features should be invariant under translation and scaling, I map the geometric features onto a unit hypersphere. Then I obtain Euclidean representations of each object via PNS. Let ϕ denote a transformation from the intrinsic shape features of an object to a block of Euclidean variables, i.e.,

$$\phi : \mathbb{R}^{d \times n} \mapsto \mathbb{R}^{(n-1) \times n} \quad (4.2)$$

where d is the dimension of the shape descriptors of the object and $d \gg n$. Because n data points lie in a hyperplane of dimension $n - 1$, higher level PNS scores are all zero and hence are not included, as detailed in Section 2.8.3.

Finally, I extract the joint shape variation from the two PNS blocks using AJIVE. I use ψ to denote the mapping from the two PNS blocks of eq. (4.2) to the low dimensional joint variation space, i.e.,

$$\psi : \mathbb{R}^{(n-1) \times n} \times \mathbb{R}^{(n-1) \times n} \mapsto \mathbb{R}^{r \times n} \quad (4.3)$$

where r denotes the dimension of the joint variation space.

Essentially, the proposed joint shape variation is obtained by composing Euclideanization and the AJIVE, i.e.,

$$X = \psi(\phi(X_1), \phi(X_2)) \quad (4.4)$$

where X_1 and X_2 denote intrinsic shape features of the two objects, respectively.

In the following, Section 4.3 discusses capturing the intrinsic shape of the hippocampus and of the caudate, resulting in feature matrices X_1 and X_2 . Section 4.4 details the transformation ϕ that projects shape descriptors onto a unit hypersphere and then Euclideanizes the spherical variables via PNS. Section 4.5 discusses the transformation ψ that extracts the joint shape variation from the PNS blocks.

4.3 S-rep Implied Boundary Points

This section discusses the shape model that I use to represent two-object complexes. Because segmented boundaries of anatomic structures are often noisy, s-reps are designed to provide robust and rich geometric features of objects, as discussed in Chapters 2 and 3. Importantly, geometric features provided by s-reps are designed to have good correspondence for joint analysis of multi-object complexes, leading to interpretable analysis results (see Chapter 2).

A fitted s-rep via the methods in Chapter 3 gives rise to a smooth mapping, so-called radial flow, from the skeleton to a set of level surfaces of radial flow (Damon, 2003). These level surfaces have their own geometry (e.g., normals) without intersections within an object, as shown in fig. 1.3. Moreover, the radial distances $\tau \in [0, 1]$ are measured as fractions of spokes' lengths from the skeleton. The boundary of an object is equivalent to a level surface at $\tau = 1$, while

the skeleton is a level surface at $\tau = 0$. In this chapter I make use of landmarks on the level surface at $\tau = 1$ for each object. These landmarks are referred to as “s-rep implied boundary points”, formed by the tips of spokes.

The s-rep implied boundary points provide good correspondence for each object across a population. The correspondence of these boundary points are determined by the correspondence of spokes. As illustrated in fig. 4.2 (c), the primary spokes in each object have good correspondences across a population. Additionally, the interpolation method of an s-rep gives an arbitrary number of spokes. As discussed in Chapter 3, this interpolation method uses the relative positions on the skeleton, which also provide correspondence across a population.

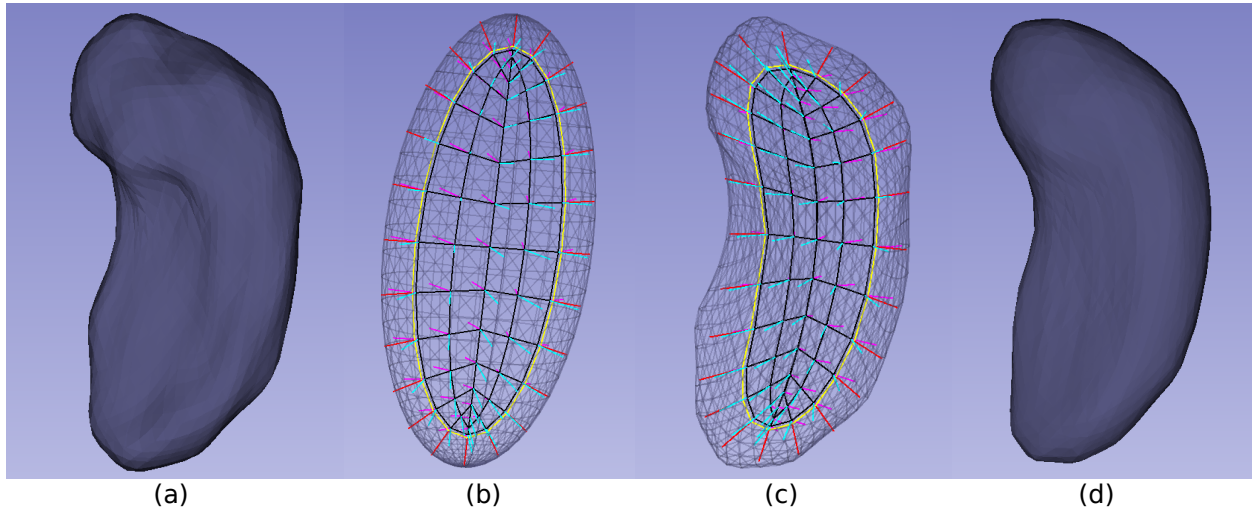


Figure 4.2: Fit an s-rep to a hippocampus and derive implied boundary from the s-rep. (a): SPHARM-PDM representation of the segmented boundary of a hippocampus. (b): A near ellipsoid (black transparent mesh) that is diffeomorphic to the hippocampal boundary. The ellipsoidal s-rep is discretized, yielding skeletal points (black grid points) and discrete spokes (colored line segments). (c): The discrete s-rep fit to the hippocampus. (d): The implied boundary derived from the hippocampus s-rep.

Fig. 4.2 shows shape representations of an example hippocampus. In the panel (a), the hippocampal shape is represented via SPHARM-PDM (Styner et al., 2006). Though SPHARM-PDM gives a straightforward representation of segmented boundary geometry, it shows weaker robustness and inferior correspondence across a population than s-rep based shape representations (see e.g., (Tu et al., 2018; Hong et al., 2016)). The panel (b) shows an ellipsoid that is diffeomorphic to the hippocampal boundary in panel (a). The ellipsoidal s-rep is consistently discretized, yielding the primary spokes shown as the line segments in red, cyan and magenta. Each skeletal point on the black grid is associated with two spokes (cyan and magenta) pointing to two sides of the ellipsoid. Each skeletal point along the edge curve (in yellow) is associated with a fold spoke (red) pointing to the crest curve of the boundary. This ellipsoidal s-rep is deformed to fit the hippocampus, as shown in the panel (c). Then the s-rep is refined to have better geometric properties – for example, the spokes will not cross within the object so that the interior geometry of the object is smooth. The

panel (d) shows the s-rep implied boundary, which is developed and used in this research. To obtain this boundary, I interpolate the refined s-rep so that I can obtain the same number of points as with the SPHARM-PDM in panel (a).

4.4 Euclideanize Skeletally Implied Boundary Points

The above s-rep implied boundary points capture intrinsic shape features coupled with pose and size variation. However, I want the joint shape analysis of the hippocampus and the caudate to be insensitive to translation and scaling of individual objects. To this end, I separately transform the shape of the hippocampus and that of the caudate to the pre-shape space using eq. (2.15), as discussed in Section 2.8.1. Fig. 4.3 shows the resulting pre-shapes of a caudate nucleus (top row) and a hippocampus (bottom row). The left column shows the s-rep implied boundary points of each object, while the right column shows the pre-shapes transformed via eq. (2.15). As shown in fig. 4.3, this transformation preserves intrinsic shape characteristics despite the translation and scaling effects. Moreover, I find in the experiments that this transformation is helpful in classifying the ASD and non-ASD group, suggesting that the position and size of the objects are non-critical factors in distinguishing ASD among infants around the same age.

As a result of the above transformation, each object is transformed to a unit hypersphere that is also referred to as a pre-shape space. The research in this chapter benefits from the transformation to pre-shapes not only because putting the data in the pre-shape space removes the translation and scaling effects of each object, but also because the spherical geometry of the pre-shape space allows for efficient representation of data. For example, the spherical coordinates of the pre-shape space are more efficient representations than the points' coordinates in the ambient space.

Compared to general spherical coordinates, I choose more statistically efficient representations – i.e., PNS features (see Section 2.8.3) – as Euclidean representations of an object. As a result of PNS, the two-object complexes are represented by two Euclidean blocks $\mathbb{R}^{(n-1) \times n} \times \mathbb{R}^{(n-1) \times n}$. In the following, I discuss extracting joint variation from these two Euclidean blocks, which forms the basis of the statistical analysis including hypothesis testing and classification in this chapter.

4.5 Extract Joint Shape Variation in Two-object Complexes

From the above methods, each PNS block $\mathbb{R}^{(n-1) \times n}$ has n columns. Moreover, the two PNS blocks should align in columns in a way that the corresponding columns across the two blocks are from the same subject. The alignment of the two blocks in the score space allows us to extract their joint variation that is useful in discriminating the ASD and non-ASD group. I explain this alignment in Section 4.5.1.

Due to the alignment of columns, I refer to the row space \mathbb{R}^n as “a score space”. Thus, I treat the two-block data as two blocks of samples from a shared score space.

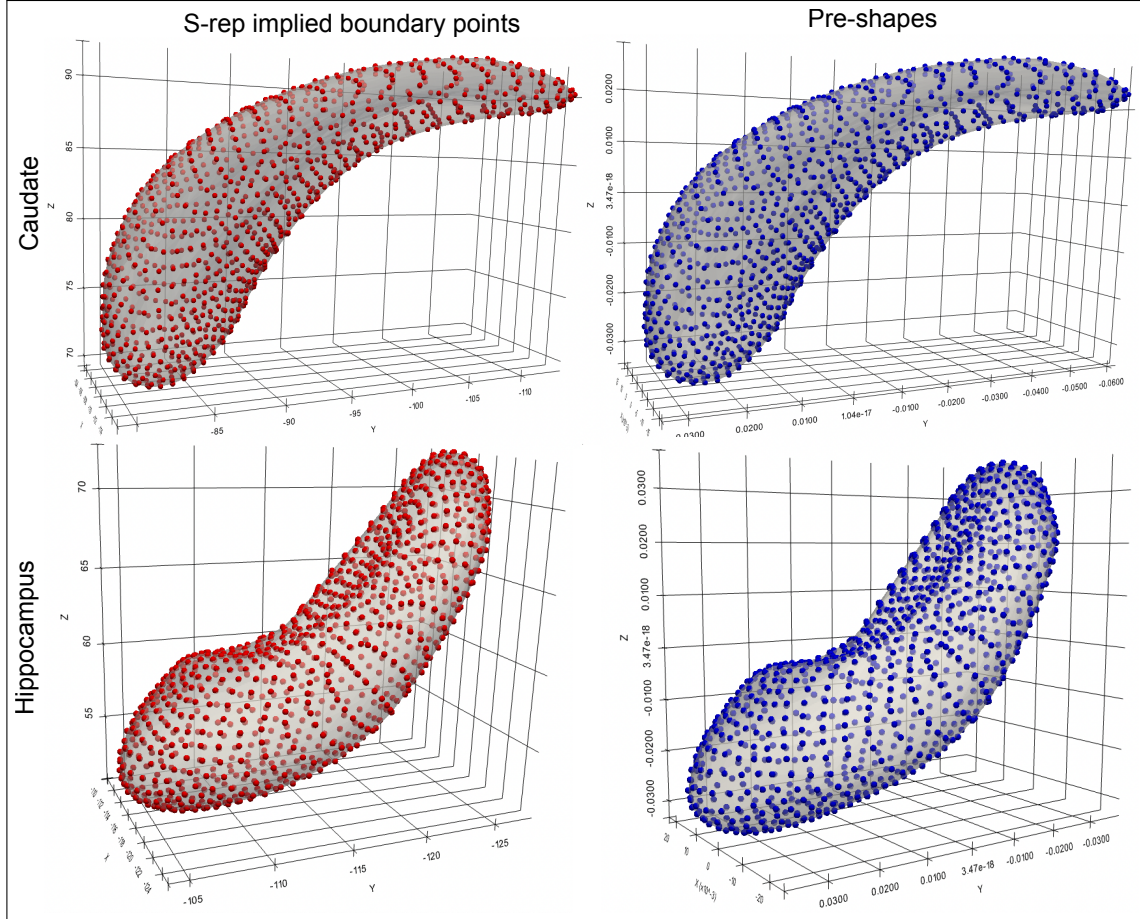


Figure 4.3: Transform s-rep implied boundary points (red points) of a caudate nucleus (top row) and those of a hippocampus (bottom row) into pre-shapes that are represented by transformed points (blue points) via eq. (2.15). Despite the translation and scaling effects, the geometry of the boundary is preserved. All these points are presented on smooth the implied boundary shown as transparent volumes underneath.

Based on the shared PNS score space, AJIVE can extract joint variation of the two blocks. Section 4.5.2 discusses the choice of ranks in NEUJIVE and the resulting joint rank from NEUJIVE.

4.5.1 Alignment in PNS Score Space

Commonality between two blocks is the basis of extracting their correlation. Though it is difficult to calibrate the two PNS feature spaces (two sets of nested subspheres), the two PNS blocks have a common score space. Moreover, as noted above, the joint shape variation of the hippocampus and the caudate should provide strong association with the development of ASD. It is natural to search the common PNS score space for the variation direction that points from the non-ASD cases to the ASD cases.

The column-wise correspondence between two PNS blocks allows us to have the shared score space \mathbb{R}^n of the two blocks. The PNS scores are geodesic distances on the respective unit hyperspheres. These scores are taken as Euclidean

representation of the shape. Two PNS score vectors (n -vectors) share a score space if the corresponding entries are about the same subjects. For example, the hippocampus score vectors and the caudate score vectors can be understood as vectors in the same score space \mathbb{R}^n if the corresponding entries are from the same n subjects. Specifically, to realize the alignment of two PNS blocks in the score space, it is necessary to consistently arrange columns in the two PNS blocks, aligning the i^{th} column in the hippocampus PNS block and the i^{th} column in the caudate PNS block both of which characterize the i^{th} subject. By doing so, the correspondence of columns is determined by the column indices.

Within the shared score space \mathbb{R}^n I search for the directions along which the hippocampus scores and the caudate scores jointly vary. These directions together give a *joint variation space*, which is of prime interest in the following discussion. I use the method AJIVE discussed in Section 2.7.4 to extract both the joint and individual variation from the two PNS blocks.

The joint component J is the projection of \hat{X}_k into the joint variation space, capturing the variation driven by common factors of the two blocks. Therefore, I expect that in the data the joint shape variation of the hippocampus and the caudate from NEUJIVE is strongly associated with the development of ASD. Moreover, because this joint component integrates information from multiple blocks, the jointly varying features are expected to be insensitive to demographic differences (Sagonas et al., 2017). Hence, the operation $\psi(\cdot)$ should enhance the discriminatory power of the two data blocks.

4.5.2 Ranks in NEUJIVE

The rank of the joint component J from NEUJIVE is the joint rank r (see eq. (4.3)). This section distinguishes the joint rank r and the initial ranks r_k . First, I discuss the implication of the joint rank r computed by NEUJIVE. Then I discuss the choosable initial rank r_k for each PNS block.

The joint rank r from NEUJIVE indicates the classes/subclasses in the data. This rank is determined as the dimension of the joint variation space. Assume that the joint variation space from NEUJIVE consists of r basis vectors in \mathbb{R}^n , where $r \leq \min(r_k) \leq n$. The number of the basis vectors has to do with (1) common characteristics of the two blocks (e.g., the severity of ASD) and (2) the degree to which the two blocks are correlated with the common characteristics. Therefore, the joint rank can be large in data of notable diversity (e.g., samples from many subtypes of the disease), while the joint rank can be small when the samples tend to be homogeneous. Especially, $r = 0$ when there is a degenerate joint variation space due to the absence of significant correlation between blocks.

As noted above, the critical choosable parameter is the initial rank in approximating each block of data. This parameter directly determines the amount of signals passed to PAA. Also, the parameter determines the rank of the residual component. Specifically, in extracting joint shape variation of the hippocampus and the caudate, an over-rank setting leads to noisy input to PAA. The extraneous noise can prevent PAA from effectively extracting joint shape

variation. Consequently, the joint shape variation may show weak association with the development of ASD. In contrast, an under-rank setting can dismiss the joint shape variation in the approximated matrix \hat{X}_k , which also prevents PAA from effectively extracting joint shape variation. In the experiments, I use cross-validation to choose an appropriate rank in the classification problem. Alternatively, Feng et al. (Feng et al., 2018) provide a means of choosing this parameter by inspecting the scree plot of each block. The optimal rank is chosen around the elbow position of the curve. This method can be useful in some applications, e.g., hypothesis testing. However, it is often difficult to manually choose the optimal rank in classification where a number of random experiments need to be done automatically.

4.6 Experiments

This section aims to verify the effectiveness and robustness of the proposed methods on both simulated and real data. In Section 4.6.1, I show that by extending AJIVE to NEUJIVE I can extract useful joint variation from two-block non-Euclidean data. Then in Section 4.6.2 I show the effectiveness of the method in extracting joint shape variation in simulated 2D objects. Different from Section 4.6.1, each block of the simulated 2D objects has two groups, formed by a common process. Last, I assess the methods in the context of the ASD dataset, as described in Section 2.2. Section 4.6.3 and Section 4.6.4 show hypothesis testing and classification of ASD and non-ASD on the subcortical structures' shape, respectively. Finally, in Section 4.6.5 I demonstrate how to interpret the joint shape variation of the structures (the hippocampus and the caudate) from NEUJIVE.

4.6.1 Joint Components of Two-block Non-Euclidean Data

This section aims to verify the effectiveness of NEUJIVE for the joint analysis of two-block non-Euclidean data in a toy example. To this end, I simulate two correlated data blocks lying on \mathbb{S}^2 , in which the joint component is created via a known function. The two blocks of simulated data can be visualized as the colored points in the left column of fig. 4.4. I simulate the data via the following steps:

1. I generate a joint component (a small circle) in the tangent space at the north pole of \mathbb{S}^2 .
2. I generate an individual component for each block in the tangent space at the north pole of \mathbb{S}^2 .
3. I map the generated data, which sum the joint component and the individual component, from the tangent space onto the sphere.
4. I rotate the above spherical data centered at the north pole to other places on the sphere.

Joint component. Because important modes of variation from s-rep based models follow small circles (see (Pizer et al., 2019) Section 6.1 for more details), I simulate a joint component along a small circle respective to each block. I

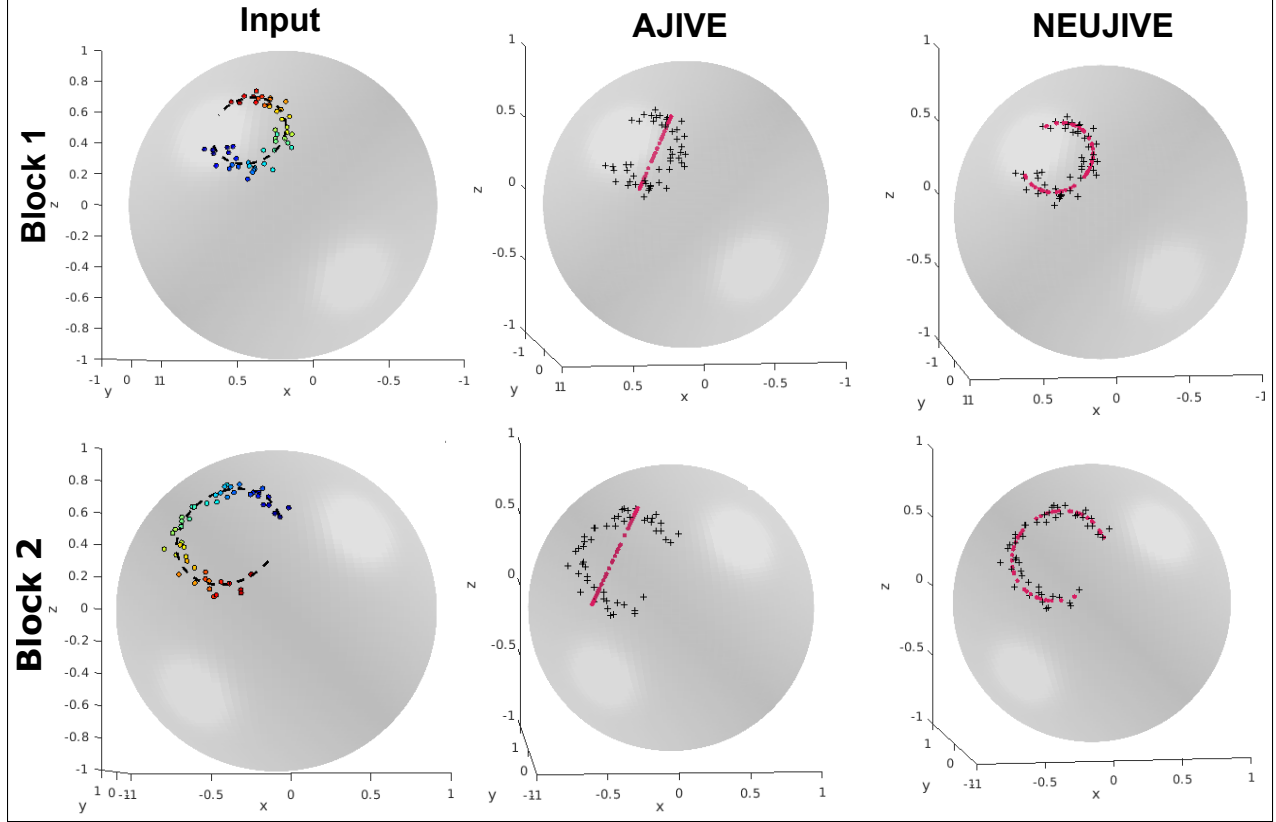


Figure 4.4: Left column shows simulated two blocks of points on \mathbb{S}^2 . The data tend to lie along a small circle with joint location on the circle indexed by the colors of points. The middle column shows the joint structure (magenta points) estimated by AJIVE in \mathbb{R}^3 . This AJIVE joint variation follows a linear pattern. The right column shows the joint structure (magenta points) estimated by NEUJIVE. This NEUJIVE joint structure captures the actual joint variation pattern in the two blocks.

sample 50 random variables θ from the uniform distribution $\theta \sim \text{Uniform}(0, \frac{3}{2}\pi)$. Then I map the θ 's onto a unit circle, yielding circular points of coordinates $(\cos \theta, \sin \theta)$. I multiply the coordinates with block-specific constants a_k to have different radii of the circles across the two blocks, where $k = 1, 2$.

Individual component. In addition to the above joint component, I generate some random variables for each block. I sample independent 2-dimensional Gaussian variables $\varepsilon_1 \stackrel{i.i.d}{\sim} N(0, 1)$ and $\varepsilon_2 \stackrel{i.i.d}{\sim} N(0, 1)$. The sum $a_k e^{i\theta} + \varepsilon_k$ of the joint and individual components forms a noisy circular distribution on a tangent plane centered at the north pole of the sphere.

Map and rotate. The simulated data in the tangent space are then mapped onto \mathbb{S}^2 via the Exponential map $\text{Exp}(\cdot)$ centered at the north pole. Finally, I rotate the data from the north pole to places that are specific to each block via rotations $g_k(\cdot)$, which result in

$$X_k(\theta) = g_k(\text{Exp}(a_k e^{i\theta} + \varepsilon_k)) \quad k = 1, 2 \quad (4.5)$$

The left column in fig. 4.4 shows the simulated two blocks X_1 and X_2 in the data space \mathbb{S}^2 . The joint component of X_1 and X_2 follows small circles on \mathbb{S}^2 , shown as the black dashed curves. The latent variables θ are indicated by colors as one moves along the small circles.

As shown in the right column of fig. 4.4, NEUJIVE can effectively extract useful joint structures (magenta points) that follow a small circle pattern. In the middle column, the joint structures (magenta points) estimated by Euclidean AJIVE are far from the underlying joint variation pattern. In particular, the joint component from AJIVE follows a linear pattern instead of a circular pattern. Thus, the result from NEUJIVE provides a mode of variation that is much more descriptive of the actual joint variation in the data. This toy example demonstrates that NEUJIVE can effectively capture the joint variation between multi-block non-Euclidean data. I owe this effectiveness to the efficient representation of PNS scores.

Moreover, the effectiveness of NEUJIVE remains when the joint component follows different patterns across blocks. In the bottom row of fig. 4.5 I show the NEUJIVE components when the simulated two blocks are of notably different variability. Different from the above example, the first block of data is generated via

$$X_1 = g_1(Exp(p(\theta) + \varepsilon_1)) \quad (4.6)$$

where $p(\theta)$ is a straight line in the tangent space at the north pole. A point on this line has coordinates $(\theta, 0.3\theta)$. The second block is still generated via eq. (4.5). Moreover, to have different noise levels across the two blocks, I set different standard deviations of ε_k across the two blocks. The results in the bottom left cell show that NEUJIVE can still give effective estimation of the joint components regardless of the different variability between the two blocks.

To understand the robustness of the joint structures estimated by NEUJIVE as I increase the noise level, I simulate two blocks of data via eq. (4.5), increasing the standard deviation of ε_k from 0.01 to 0.1. The simulated data under 3 different noise levels can be seen in fig. 4.6 left, i.e., $\sigma(\varepsilon_k) = \{0.01, 0.05, 0.1\}$. For each noise level, I compute the joint structure from NEUJIVE. I measure the Root of Mean Square Error (RMSE) between the NEUJIVE joint structures and the actual joint structures. The curves of RMSE vs. noise levels of the two blocks are shown in fig. 4.6 right. This figure shows that the RMSE is almost linearly increasing along with the increasing noise level.

In addition to the joint component, NEUJIVE also results in reasonable individual components for respective blocks, clustered around the PNS mean, shown as the blue dots in fig. 4.5. I have also tried the relevant methods that are reviewed in Section 2.7.3 including GCCA, O2-PLS (el Bouhaddani et al., 2018) and DISCO-SCA to decompose this toy dataset. However, I failed to find meaningful components on the sphere. I suspect that the failure is due to the difficulties of extracting joint variation in feature space after many non-linear transformations on θ 's in eq. (4.5).

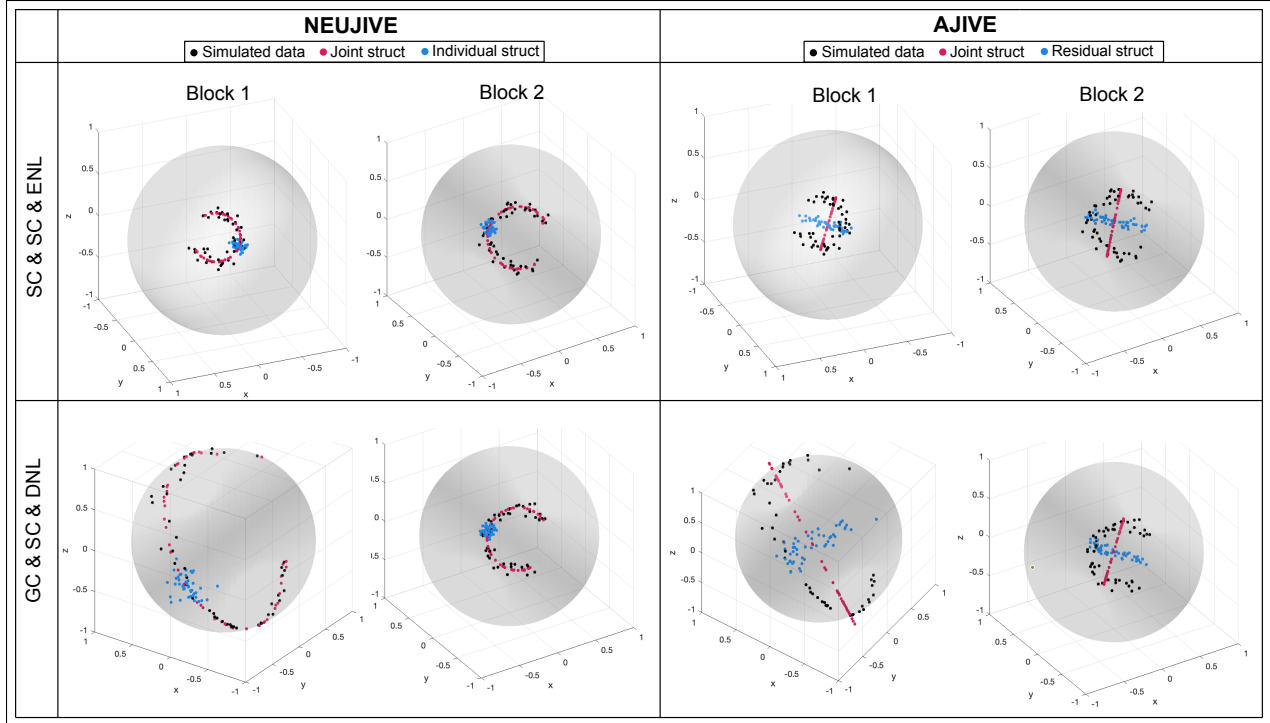


Figure 4.5: Joint and individual structures from NEUJIVE on two correlated non-Euclidean blocks. The first row shows the simulation where the joint structures are designed as circular variables along Small Circles (SC) in the two blocks (see eq. (4.5)). Moreover, the two blocks share the Equal Noise Level (ENL) in which ε_k share the same standard deviation. The second row shows a different simulation study. In this simulation, the joint component is designed to be a Great Circle (GC) in block 1. But this component follows a small circle in block 2. Moreover, the two blocks have Different Noise Levels (DNL). On each sphere, the black dots are simulated data. The magenta dots are joint structures from each algorithm. The blue dots are individual structures from NEUJIVE and residual structures from AJIVE, respectively.

4.6.2 Joint Shape Analysis in Classification

Beyond the above example, I also want to show that NEUJIVE can extract useful joint variation from two blocks each of which has two groups. To this end, I design the two groups to be formed by a common process in the two blocks. By doing so, the joint variation of the two blocks should focus on the common process and thus be useful in classifying the two groups.

In particular, I simulate two blocks of shape data formed by 2D landmarks. The simulated data is analogous to the real shape data (see Section 2.2) in the following senses. First, the simulated shape are obviously different across the two blocks, analogous to the hippocampus and the caudate shape. Second, I impose on the two blocks a common process (via Procrustes alignment²) to form two groups. This common process acts as the development of ASD, which simultaneously changes the shape of the hippocampus and the caudate. Last, it is difficult to classify the two groups within each block because the two groups have close means, which is also the case in classifying ASD vs. non-ASD

²While Procrustes alignment is normally used to preprocess shape data, this simulation instead uses Procrustes alignment to produce a geometric transformation that is sensitive to NEUJIVE and that forms the group difference.

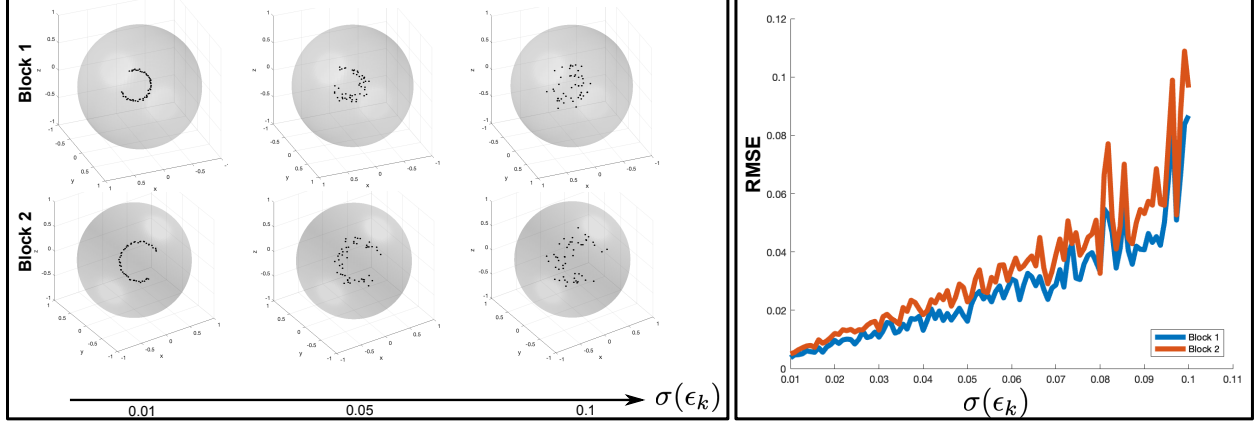


Figure 4.6: Root Mean Square Error (RMSE) between the actual joint structures and NEUJIVE estimated joint structures as increasing the noise level. Left panel: Simulated data as increasing the standard deviation $\sigma(\epsilon_k)$ of noise from 0.01 to 0.05 and to 0.1. The two circular blocks are of different radii, i.e., $a_1 = 0.25$ while $a_2 = 0.35$. Right Panel: RMSE of the NEUJIVE joint structures of the two blocks as increasing standard deviation of noise. The red and the blue curve are, respectively, the RMSE of the block 1 and the block 2.

using either the hippocampus or the caudate shape. In such scenarios, NEUJIVE can effectively extract joint shape variation that focuses on the group difference.

Fig. 4.7 shows the simulated two data blocks, in which the two groups are denoted as the red and the blue. In this simulation, I start with a dataset from (Dryden and Mardia, 2016) in which eight 2D landmarks are sampled on each of the skulls of 29 male adult gorillas, shown as the blue circles in the top left panel. Then I create another group via Procrustes alignment φ , yielding the landmarks shown as the red crosses in the top left panel.

Next, I create a secondary block of data for the purpose of joint analysis, shown in the bottom left panel of fig. 4.7. This secondary block is designed to have notably different shape from the above data block. In addition, the two groups in the secondary block are designed to be formed via the common process φ . Specifically, I first move the top-most landmarks in the top left panel to be farther away. Then I rotate all landmarks by 45 degrees. These modified landmarks are shown as the blue circles in the bottom left panel. Second, I apply a Procrustes alignment on the modified landmarks, forming the red group shown as the red crosses.

The above two blocks of landmarks with two groups form the input of NEUJIVE. Specifically, let $X_G \in \mathbb{R}^{16 \times 29}$ denote the original landmark matrix of the gorillas' skulls. The red group in the first block can be written as $\varphi(X_G)$. The concatenation of the two groups forms the first block X_1 of NEUJIVE. As I create the secondary block, I obtain the modified landmarks denoted as $\tilde{X}_G \in \mathbb{R}^{16 \times 29}$. Likewise, the red group in the second block can be written as $\varphi(\tilde{X}_G)$. The concatenation of \tilde{X}_G and $\varphi(\tilde{X}_G)$ forms the second block X_2 . In total, the inputs of NEUJIVE are organized as

$$\begin{aligned} X_1 &= [X_G \quad \varphi(X_G)] \in \mathbb{R}^{16 \times 58} \\ X_2 &= [\tilde{X}_G \quad \varphi(\tilde{X}_G)] \in \mathbb{R}^{16 \times 58} \end{aligned} \tag{4.7}$$

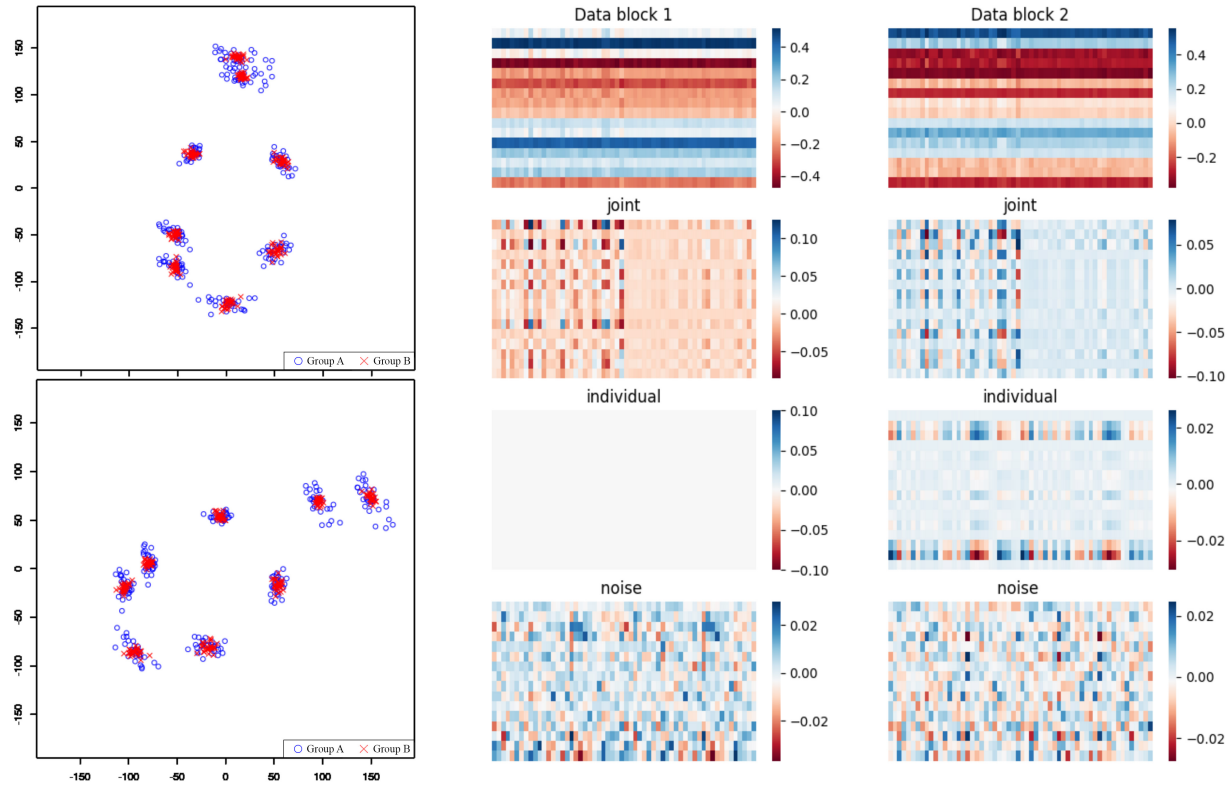


Figure 4.7: Joint analysis of two blocks of shape data formed by 2D landmarks from male gorillas' skulls. The left column shows the two blocks: block 1 (top) and block 2 (bottom). Each block has two groups (shown as the blue circles and the red crosses). I detail the generation of the two blocks in the text. The right two columns show the heat maps of the PNS matrices (the top row), the joint structures (the second row), the individual structures (the third row) and the residuals (the last row). The color represents the entry values of each matrix. See the text for detailed interpretation.

I present the resulting components from NEUJIVE in the right two columns of fig. 4.7. I obtain the pre-shapes of X_1 and X_2 by centering and normalizing each shape by the centroid size, yielding two pre-shape matrices. To each matrix, I apply PNS to obtain the PNS feature matrix shown in the top row. The colors in the heat maps represent the entry values of the PNS features of X_1 (the middle column) and X_2 (the right column). The other rows show the entry values of the matrices resulting from the joint analysis of the two PNS blocks. The resulting NEUJIVE components include (1) the joint components (shown in the second row), (2) the individual components (shown in the third row) and (3) the residual components (shown in the last row).

I can observe notable patterns from the joint components, as compared to the input matrices, the individual components and the residual components. Across all these matrices, each column is a sample, while each row contains values for a feature. I can find clearly different patterns between the two groups of samples in the joint component of every block. This visualization suggests that the joint component from NEUJIVE focuses on the common factors that form the two groups. As such, I make use of the joint component in classifying the two groups.

The individual components also are as expected in both blocks. Because of the fact (1) that the first block uses the initial rank of 2, while the second block uses the initial rank of 3 and (2) that the joint variation is rank 2, there is no individual variation specific to the first block.

To demonstrate that the joint shape variation improves the classification between the red and the blue group, I compare in the following (1) the classification using the joint component from NEUJIVE, (2) the classification using landmarks and (3) the classification using the joint component from other joint analysis methods including AJIVE.

To classify the two groups, I train a robust linear classification method called Distance Weighted Discrimination (DWD). This has been shown to have better properties than the support vector machine (see Section 2.9.2). I randomly select 80% of the data in the training while I use the remaining in the test. To avoid the bias due to splitting the data, I repeat this experiment 100 times. Table 4.1 reports the average test accuracy over all the repetitions. The classification of each block uses the following sets of features (corresponding to the rows of Table 4.1): (1) the original coordinates of landmarks (the first row), (2) the PNS scores (the second row), (3) the joint structures from AJIVE on the original coordinates (the third row), (4) the joint structures from AJIVE on the spherical coordinates (the fourth row) and (5) the joint structures from NEUJIVE (the fifth row).

From Table 4.1 I find that the joint component from NEUJIVE (the last row) improves the classification performance compared to raw landmarks (the first row) by a large margin. It is difficult for a linear classifier to learn a good decision boundary from the coordinates of landmarks. Moreover, I find that the Euclidean AJIVE can capture some group difference. Yet, due to the complexity of the data space, the group difference is relatively weak in AJIVE’s joint component, compared to those from Spherical AJIVE and NEUJIVE. Importantly, NEUJIVE’s joint component more strongly indicates the group difference than that of Spherical AJIVE. I owe this advantage to the statistically efficient data representation by PNS.

Table 4.1: Test performance using different features in classification

Features	X_1	X_2
Landmarks	0.44	0.44
PNS	0.5	0.48
Euclidean AJIVE	0.59	0.55
Spherical AJIVE	0.68	0.67
NEUJIVE	0.75	0.72

4.6.3 Hypothesis Testing on the Hippocampus-Caudate Data

This hypothesis test aims to verify whether the difference between the two groups (ASD vs. non-ASD) is statistically significant in the joint variation space. Because the sample size is small, the hypothesis test should be based on permutation testing ideas. To this end, I accomplish this task in the following steps.

I start from s-rep implied boundary points of the hippocampus-caudate pairs, each of which is associated with a class label (i.e., ASD or non-ASD). I separately convert the landmarks of the hippocampus and those of the caudate into pre-shapes as described in Section 4.4. Then I apply NEUJIVE to the pre-shapes of the hippocampus and of the caudate.

Fig. 4.8 shows for various features the Explained Variance Ratio as a function of the number of eigenmodes. I choose s-rep implied points to remove non-systematic geometric details on the boundary. As a result, the s-rep implied boundary points give smooth geometric features across a population. Also, Fig. 4.8 suggests that PNS features (green curves) are more statistically efficient representations, especially for the caudate and for the hippocampus-caudate pair, as compared to the pre-shapes (orange curves).




Figure 4.8: The cumulated Explained Variance Ratio (EVR) of the first 35 eigenmodes vs. the number of eigenmodes. The left and the middle panel show the EVR of shape features from single objects – the hippocampus and the caudate, respectively. The right panel shows the EVR of concatenated shape features of the hippocampus-caudate pairs. Best viewed in color images.

Fig. 4.8 also gives us a means of choosing initial ranks for low rank approximation as follows. Considering the Explained Variance Ratio of PNS features from s-rep implied boundary points (s-rep PNS), I choose the initial rank 50. In this setting, NEUJIVE finds a joint component of rank 2 (i.e., $r = 2$ from fig. 4.1 (c)), which suggests very few common characteristics in the data, as discussed in Section 4.5.2. It turns out that these two dimensions give statistically significant difference between the ASD and non-ASD groups.

Within the 2-dimensional joint variation space (a subspace in \mathbb{R}^n), I test the data projected on the unit vector between the means of the two groups with a permutation test, which is implemented by a method called Direction-Projection-Permutation (DiProPerm) (Wei et al., 2016). I let DiProPerm generate 1000 random permutations and then compute the univariate statistic, i.e., mean difference (MD) between the two groups for each permutation.

Fig. 4.9 shows the test statistics on the MDs in those permutations using two sets of features. The observed MDs of the two groups are at the positions of the vertical dashed lines. Along with the permutation statistics, I show the empirical p-values, which are the proportion of the permutations that have larger MDs than the observed MD. I also show the z-scores, which measure how many deviations that the observed MDs are above the average MD of their

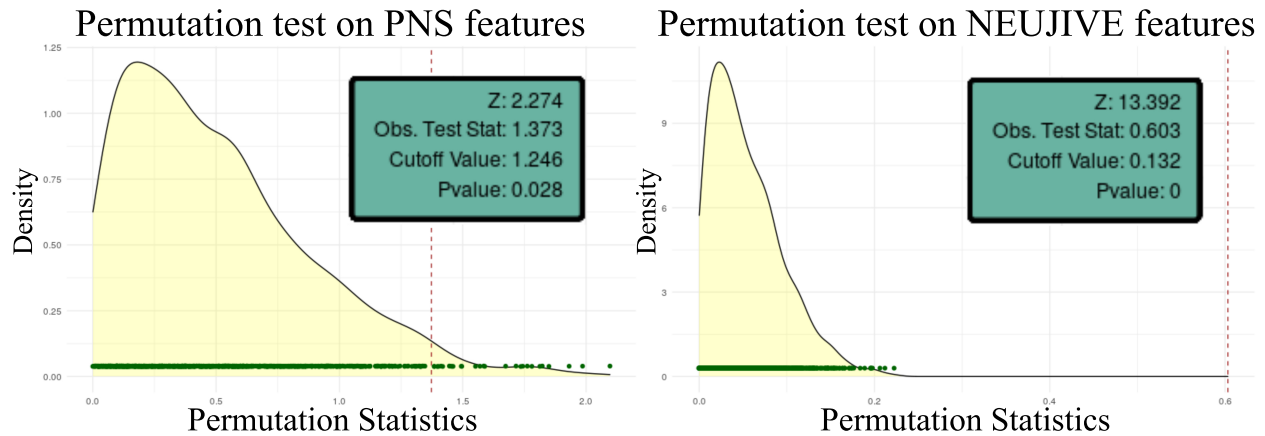


Figure 4.9: The results from the permutation test of concatenated PNS blocks (left) and the test of NEUJIVE joint components (right). The green dots on the bottom of each figure represent the test Mean Differences (MDs) from the permutations. The black curve represents the estimated density function of those MDs.

permutations. The p-values and z-scores together indicate statistical significance of group difference with respect to the test features. Smaller p-values and larger z-scores indicate more statistical significance of the group difference.

NEUJIVE joint components give statistically significant difference between the ASD and non-ASD groups, as shown in fig. 4.9 right. I found no permutations that have larger MDs than the observed MD, i.e., the empirical p-value equals 0. Moreover, the z-score (≈ 13.39) is relatively large, indicating strong statistical significance. As a comparison, I use the concatenated PNS blocks of the hippocampus and the caudate in the test, yielding less statistically significant difference, as shown in the left panel. Specifically, there is a small proportion ($\sim 2.8\%$) of permutations that have larger MDs between the ASD and non-ASD group than the observed MD.

Table 4.2 presents more comprehensive results from the hypothesis test using various sets of shape features. The first row shows the results from using the concatenation of s-rep implied boundary points. The second row shows the results from using the concatenation of Euclideanized PDMs with PNS. Comparing these two rows, I note that the Euclideanization improves the discriminatory power of the shape features. Moreover, using joint structures from AJIVE (the third row) results in a more statistically significant difference between the two groups. The fourth row shows the results of the proposed test with NEUJIVE features. It demonstrates the advantages of using the joint shape variation in the test. The last three rows of Table 4.2 show the results from the compared joint analysis methods that are reviewed in Section 2.7.3. These methods can extract the correlation of features in high dimensional feature space. Yet, the group difference from these methods are not as significant as those from NEUJIVE.

4.6.4 Classification

Because joint shape variation captures correlated shape features from multiple objects and because this correlation is closely associated with class labels, the joint components from NEUJIVE lead to statistically significant differences

Table 4.2: Hypothesis testing of ASD vs. non-ASD with different joint features

	p-values ↓ ^a	z-scores ↑ ^a
Concatenate Implied PDMs	0.3106	0.3679
Concatenate Euclideanized PDMs	0.028	2.274
AJIVE (Feng et al., 2018)	0.0136	2.7714
NEUJIVE (ours)	0	13.392
HPLS (Wold et al., 1996)	0.1135	1.206
PNS + HPLS	0	4.4263
E-GCCA (Shen et al., 2014)	0.0169	2.0031

^a Smaller p-values and larger z-scores indicate more significant differences.

in the above hypothesis test. In this section, I attempt to exploit the joint components from NEUJIVE in classifying the ASD vs. non-ASD group.

Because the joint analysis via NEUJIVE is based on the alignment of PNS score space, as discussed in Section 4.5.1, it is necessary to estimate the joint shape variation by pooling the training and test data. Oftentimes, the number of training and test samples are different. Therefore, the estimated joint shape variation space in the training domain is not readily transferable to the test domain. To address this problem, I pool the training data with the test data for the joint analysis via NEUJIVE. The resulting joint shape variation is the basis of training and testing a classifier. Because the process of producing joint shape variation does not make use of class labels, it is natural to compare with other similar classification methods such as the method using GCCA (Shen et al., 2014).

I obtain the joint shape features following steps similar to Section 4.6.3. I again use DWD as the classifier in this section. I compute the Area Under the ROC (Receiver Operating Characteristics) Curve (ROC-AUC) as the evaluation metric. Fig. 4.10 (a) shows the ROC-AUCs as a function of initial ranks. From Fig. 4.10 (a) I can see that the optimal performance results from the initial rank of about 50. Fig. 4.10 (a) also shows that both the under-ranks ($r_k < 47$) and over-ranks ($r_k > 54$) can lead to poor performance in classification.

To select a good initial rank for the test data, I use 5-fold cross-validation to select the optimal initial rank from 20 sets of NEUJIVE joint features. Specifically, I apply NEUJIVE to the pooled data with 20 different initial ranks. By observing the Explained Variance Ratio via s-rep PNS features in Fig. 4.8, I set the same initial rank for the hippocampus and the caudate PNS block. The 20 parameters result in 20 sets of NEUJIVE joint features. I apply the same partition of samples³ to every set of NEUJIVE joint features into the training, validation and test data. To avoid bias due to splitting data, I generate 1000 random partitions of samples.

I compared test AUCs from NEUJIVE with that from other multi-block analysis methods, including Hierarchical PCA (HPCA) (Westerhuis et al., 1998), Consensus PCA (CPCA) (Wold et al., 2005) and Generalized CCA (GCCA)

³I create a random partition of the samples into 10 roughly equal-sized subsets of the ASD group and likewise with the non-ASD group. I set aside one of the subsets from each group for testing and use the remaining subsets for training and validation.

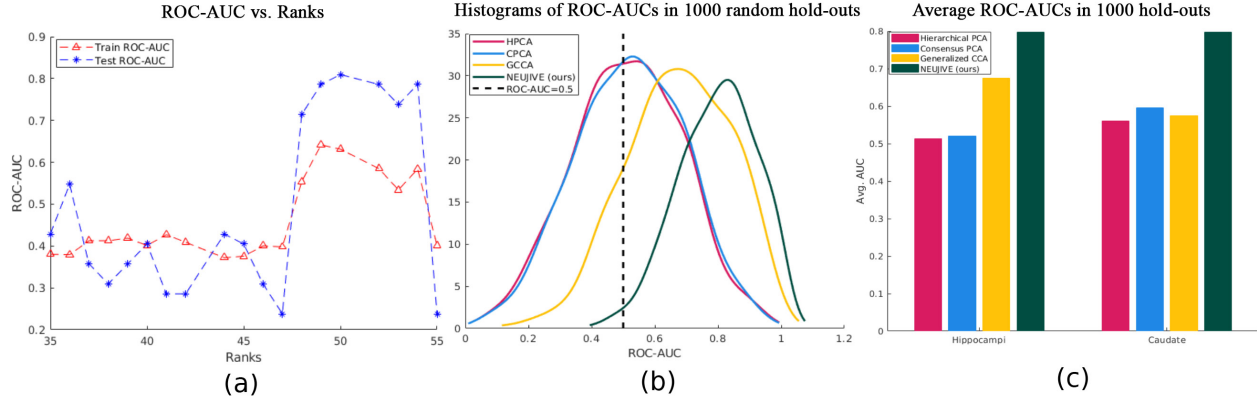


Figure 4.10: Results from the classification using joint structures. (a) The AUCs in classifying ASD vs. non-ASD with NEUJIVE joint features vs. the chosen initial ranks. (b) The histograms of test AUCs using joint structures of the hippocampus from different methods. (c) The average test AUCs from various methods using the joint structures of hippocampi (left part) and caudate nuclei (right part).

(Shen et al., 2014). I show the histogram of the test ROC-AUCs in Fig. 4.10 (b). NEUJIVE (shown as the green curve) performs better than the compared methods.

Fig. 4.10 (c) shows the average test ROC-AUCs (from the 1000 partitions) from NEUJIVE (the green bars) and the compared methods. In the classification using the joint structures of either the hippocampus or the caudate, NEUJIVE significantly improves the average test ROC-AUC. This improvement benefits from both the Euclideanization via PNS and the focus on joint structures provided by AJIVE. Also, I find that the joint structures from NEUJIVE are almost consistent across blocks, shown as the green bars. As a comparison, the other methods yield notably different joint structures for the hippocampus block vs. for the caudate block.

4.6.5 Shape Differences Between Groups in the Joint Variation Subspace

In medical applications, it is important to interpret the results from anatomic shape analysis. The results from NEUJIVE can be pulled back to the pre-shape space for understanding the joint shape variation patterns. In this section I demonstrate the interpretation of the joint shape variation of the hippocampus and the caudate from NEUJIVE.

Specifically, I show how to interpret the difference between the ASD and the non-ASD group in NEUJIVE’s joint variation space. Given the NEUJIVE joint structure $J \in \mathbb{R}^{r \times n}$, I first separately average the columns of the ASD group and those of the non-ASD group, yielding a matrix $Y \in \mathbb{R}^{r \times 2}$. Then I map the matrix Y to the PNS hippocampus block and to the PNS caudate block, respectively, resulting in the hippocampus PNS features of two groups and the caudate PNS features of two groups. The two PNS blocks are both of dimension $(n - 1) \times 2$. I apply the inverse PNS to obtain pre-shapes of the hippocampus and those of the caudate. Finally, I compute Euclidean distance for each object between the pre-shapes of the two groups.

The distances between the corresponding landmarks are shown in fig. 4.11 as heat maps overlaid on a non-ASD hippocampus-caudate pair. Such heat maps suggest that some regions (the red) in the hippocampus and the caudate are more jointly associated with ASD than other regions (the blue). This visualization provides a means of understanding joint shape variation of multi-object complexes in the development of ASD.

4.7 Broader Impact

Though this research focuses on multi-object shape analysis, there are many possibilities to generalize the proposed methods to other applications. This section discusses a more general form of the methods. The general form leads to broader impact beyond multi-object shape analysis. Specifically, I discuss possible generalization (1) of data representations (2) of manifolds that the representations live on and (3) of Euclideanization methods. Then, I discuss potential applications of the individual variation component from NEUJIVE.

First, it is possible to have different representations of data in some applications. I only use s-rep implied boundary points in this research because these landmarks have good correspondence and also because it is rather simple to map the data to a unit hypersphere. However, other shape features from s-reps (e.g., spokes' directions) also provide useful information of an object. Such shape features often require a different mapping to manifolds and to the Euclidean representations (e.g., via Composite PNS (Pizer et al., 2013b)). Moreover, although this dissertation focuses on applications in two-object complexes, the methods are generalizable to multi-object complexes each of which contains more than 2 objects. Because PAA can accept multiple input Euclidean subspaces in \mathbb{R}^n and output their joint variation space (also in \mathbb{R}^n), there are no theoretical barriers preventing the analysis of more than 2 objects. However, care is needed in the selection of objects. All the selected objects should be (strongly) associated with at least a common factor (e.g., disease of few subtypes). Otherwise, it is difficult to find the joint variation with PAA.

Second, the proposed joint analysis method is applicable for the data not only on the cross-product of unit hyperspheres, but also on the cross-product of general Riemannian manifolds (e.g., polysphere (Eltzner et al., 2015)) or stratified sets. The geometry of the manifolds determines the Euclideanization method, as discussed below.




Figure 4.11: The point-wise differences (in millimeters) between a reconstructed ASD complex and a reconstructed non-ASD complex. The reconstructed ASD and non-ASD complexes have more difference in the red regions than in the blue regions.

Third, there exist various Euclideanization methods that map data on a manifold (or a stratified set) to a Euclidean space. The choice of the method should depend on the data properties. For instance, PGA is used in capturing principal components of data in a Riemannian manifold, as discussed in Section 2.8.2. The Euclidean representation from PGA is expressed in terms of an orthonormal basis in the tangent space centered at the Fréchet mean. Regardless of the different choices of the basis, AJIVE can achieve consistent results from multiple PGA blocks because of invariance properties of PGA score spaces under rotations of the basis (i.e., orthogonal transformations multiplied on the left of $\phi(X_1)$ or $\phi(X_2)$ in eq. (4.4)).

Last, although I focus on the joint shape variation in this research, the individual shape variation can also be useful in some applications. In contrastive analysis (see e.g., (Abid et al., 2018)), multiple datasets are used to extract the distinctive variation. In this regard, the individual shape variation from NEUJIVE is useful to extract robust and meaningful patterns of interest. I leave this as the future research.

4.8 Conclusions and Discussion

This chapter aimed to provide a statistical method for extracting joint and individual shape variation from multi-object complexes. Despite the recent advances in the shape statistics field and in the multi-block data analysis field, there previously existed no methods that effectively extract joint and individual shape variation. This chapter bridges the two fields, yielding a new method that can effectively extract joint shape variation from multi-object complexes.

In this chapter I focused on the joint shape variation in two-object complexes. Because ASD simultaneously changes the shape of multiple brain structures including the hippocampus and the caudate, it is natural to assume that the shape of these related structures are jointly varying in the development of ASD. Under this assumption, an effective joint shape variation should be associated with the development of ASD. Therefore, I verify the effectiveness of NEUJIVE in the hypothesis test and classification between the ASD and non-ASD group. The interpretation of the joint shape variation may lead to interesting findings regarding the development of ASD.

Not only in the experiments with real shape data, I also verified the method in toy examples, showing notable advantages over the baseline method AJIVE. In analyzing the example with multi-block homogeneous data, NEUJIVE can effectively recognize the joint variation component (see section 4.6.1), regardless of the complexity of the feature spaces. In the toy example with multi-block heterogeneous data (of two groups), NEUJIVE focuses on the group difference in the joint structures and thus results in higher classification accuracy. These toy examples explain the advantages of NEUJIVE on multi-block non-Euclidean data analysis.

These experimental advantages result from the choice of the methods in this chapter. First, to capture intrinsic shape features of the hippocampus and of the caudate, I developed s-rep implied boundary points, which yield smooth geometric features with good correspondences across a population. Second, to obtain shape features that are insensitive

to the pose and size variation, I converted s-rep implied boundary points into respective pre-shape spaces. Third, I adopted PNS to convert pre-shapes into Euclidean variables. The resulting PNS scores are insensitive to orthogonal transformations of features, making the Euclidean representations robust against rotations of objects. Moreover, the score spaces \mathbb{R}^n of two PNS blocks share basis vectors with aligned columns, assuming columns are observations. This alignment allows us to extract joint shape variation of the two objects. Fourth, I used AJIVE to estimate the joint variation of the two PNS feature matrices, yielding low-dimensional jointly varying features.

In addition to the potential applications discussed in Section 4.7, I also want to extend this research by explicitly accounting for the relations between objects. Among selected objects that are likely associated with a disease, it is possible that within-object shape features of an anchor object correlate with its relations with other objects during the development of the disease. In such a scenario, the joint shape variation of within- and between-object shape features can provide interesting information about multi-object complexes, as discussed in the following chapter.

CHAPTER 5: CORRELATION OF WITHIN- AND BETWEEN-OBJECT FEATURES

5.1 Introduction

In the development of any disease, it is reasonable to assume that multiple structures affected by the disease are correlated. This correlation may occur in within-object shape features of each structure and its relations with nearby objects. For example, the development of ASD causes shape changes of both the hippocampus and the caudate (see Chapter 1). Within-object shape of the hippocampus and its relations with the caudate as well as within-object shape of the caudate may correlate with the development of ASD. As the previous chapter has shown that there is a correlation of within-object shape between these structures, this chapter focuses on the correlation of within- and between-object shape features in two-object complexes. To decouple within-object shape features with between-object shape features in multi-object complexes, this chapter proposes a multi-object shape model that allows us to analyze positional and geometric relations between these interrelated brain structures (Bossa and Olmos, 2007) – the positional relations include relative pose and orientation between the two objects, while the geometric relations include local shape captured using geometric properties such as boundary curvature.

Specifically, this chapter aims for the following challenges in extracting the positional and geometric relations between 3D anatomic objects. First, instead of correlations between global morphological measures (e.g., volume-to-volume correlations), the brain structures interrelate in a subtle way and in local regions (Hazlett et al., 2017). This requires not only sensitive shape representations that can well capture local shape features but also good correspondences across samples. Second, the anatomically important interrelations between objects are often non-linear. The average of two-object complexes, for example, is by no means the concatenation of two averages of single objects. Third, such interrelations between objects can hardly be extracted via the analysis of the concatenated single objects because the concatenation can ignore certain interrelations between objects. Fourth, it is difficult to separate geometric from positional relations in multi-object complexes due to the difficulties of alignment (see e.g., (Gollmer et al., 2012)).

To address the above challenges, this chapter proposes a geometric model that explicitly captures within- and between-object shape of two-object complexes. The features of the model are examined in a statistical analysis to find significant correlation. These features are the basis of classification and hypothesis testing.

The geometric model is based on s-reps of single objects (see Chapter 3), which have been designed to provide good correspondences across populations. In addition to within-object shape features provided by each s-rep, I also

focus on how two s-reps link to each other. The link features are designed to capture both positional and geometric relations between the two objects, as described below.

To obtain desirable link features, I propose a modified version of the linking flow given by Damon and Gasparovic in (Damon and Gasparovic, 2017). The key idea is to create a smooth mapping from skeletons to an external linking surface that is formed to bisect the two boundaries. The proposed linking flow takes in two s-reps and outputs a link vector field defined on each s-rep. From a link vector field, I sample discrete link vectors that have good correspondences across a population. These link vectors yield between-object link features, including (1) the linking distances (lengths of link vectors) between a skeleton and the external linking surface and (2) the directions of the link vectors. Link lengths can capture the positional relation between the two objects, while the link directions can capture the relations of local geometry. This chapter explores and exploits the above link features in statistical analysis. Different from the Damon-Gasparovic linking flow, the proposed method yields a smooth non-branching external linking surface between the objects (the hippocampus and the caudate) in the ASD dataset that is described in Section 2.2. Such consistent topology allows us to have link vectors in good correspondences across a population, which is specially advantageous in 3D anatomic shape analysis where the between-object shape often varies dramatically.

Moreover, using *local affine frames* I capture geometric features that are invariant to rigid transformations yet are sensitive to local deformations. Fitting local affine frames allows us to extract differential geometry of an s-rep. I map s-rep features to these local affine frames to obtain shape features that are invariant to rigid transformations of the multi-object complex. Not only for within-object shape features, the affine frames can be also used in capturing between-object shape features that are invariant to rigid transformations.

Finally, I jointly analyze within- and between-object shape features using NEUJIVE discussed in Chapter 4, yielding Correlation of Within- and Between-Object (CoWBO) shape features. While between-object shape can vary much more dramatically than within-object shape, this method can reveal the correlation regardless of the different variability of the two.

The proposed methods show statistical advantages in classifying the autism vs. non-autism group with the shape of subcortical objects including the hippocampus and the caudate. I found that the CoWBO features made from (1) s-rep features of the hippocampus and (2) link features between the hippocampus and the caudate yield the best classification performance among various choices of shape features. In addition to the application, the contributions of this chapter in methodology can be summarized as follows.

1. I propose explicitly capturing positional and geometric relations between objects with modified linking structures.
2. I propose using affine frame fields on s-reps. These affine frame fields are fitted via a deformation-based method to better characterize the geometry of objects. I map within- and between-object shape features to the fitted local affine frames such that the shape features are invariant to rigid transformations.

3. I propose joint analysis of within- and between-object shape features using NEUJIVE. I find that the resulting correlated shape features give good performance in classifying ASD vs. non-ASD.

The remainder of this chapter is organized as follows. Section 5.2 discusses the joint analysis of within- and between-object shape features, which yields the CoWBO features. Then in Section 5.3 I detail the between-object shape features based on linking structures. In Section 5.4 I detail within- and between-object shape features using local affine frames. In Section 5.5 I evaluate the proposed methods. Section 5.6 summarizes the contributions and discusses the limitations of this approach.

5.2 CoWBO Features: Correlation of Within- and Between-Object Features

As noted above, the joint analysis of multi-object complexes gives comprehensive understanding of the population. This chapter focuses on the joint analysis of within- and between-object geometry via NEUJIVE (see Chapter 4). In using NEUJIVE, I formulate within- and between-object shape features as two input blocks of data.

Specifically, in a population of n samples, assume each sample has two non-penetrating objects. I use $X \in \mathbb{R}^{d_1 \times n}$ to denote the feature matrix from n s-reps of an object; this feature space can be embedded in a d_1 -dimensional Euclidean space. Likewise, I use $Y \in \mathbb{R}^{d_2 \times n}$ to denote the between-object shape feature matrix from the n samples; this feature space can be embedded in a d_2 -dimensional Euclidean space. Also, I use $J(\cdot, \cdot)$ to denote the operation of extracting the correlated features by NEUJIVE. Altogether, the CoWBO features are written as

$$J(X, Y) \in \mathbb{R}^{r \times n}, \quad (5.1)$$

where r is the dimension of the joint variation subspace, in which the correlated features are extracted. The dimension r is determined by NEUJIVE depending on the degree to which X and Y are correlated. I only use within-object shape features from the target object because the shape features of the other object induce distracting information due to the fact that those features are not discriminative and thus harm the classification performance.

CoWBO features represent how the within- and between-object shape jointly vary. These correlated shape features are often in a rather low dimensional feature space, as compared to the original feature space of either X or Y . Moreover, between-object shape often varies more dramatically than within-object shape. For example, the between-object shape represented by Damon-Gasparovic's linking structures can develop branches due to small positional changes of individual objects. Despite that different variability, CoWBO features can give robust estimation of the correlation.

As I will show in Section 5.5, I classify two groups of samples using CoWBO features. Section 5.3 and Section 5.4 detail the methods for generating block Y and X in eq. (5.1).

5.3 Modified Linking Flow Captures Between-object Shape Features

At first sight of multi-object complexes humans' eyes tend to pay more attention to relations between objects rather than to the details on each boundary. In a two-object complex, the two objects are roughly perceived as cylinder-like shape at first sight. Inspired by this psychological study, I propose in this section a modified version of Damon-Gasparovic's linking flow, which focuses more on between-object relations than on boundary details of individual objects.

In fig. 5.1 I simulate a two-object complex in 2D. As shown on the top, Damon-Gasparovic's method (Damon, 2019) can develop self-links that are formed by two spokes from the same object, as highlighted in the figure. Such self-links are found in notable dent regions. Because these self-links are barely relevant to between-object relations, and also because they are inconsistent across a population, I modify such links to be less sensitive to boundary details (e.g., dents), as illustrated on the bottom of fig. 5.1.

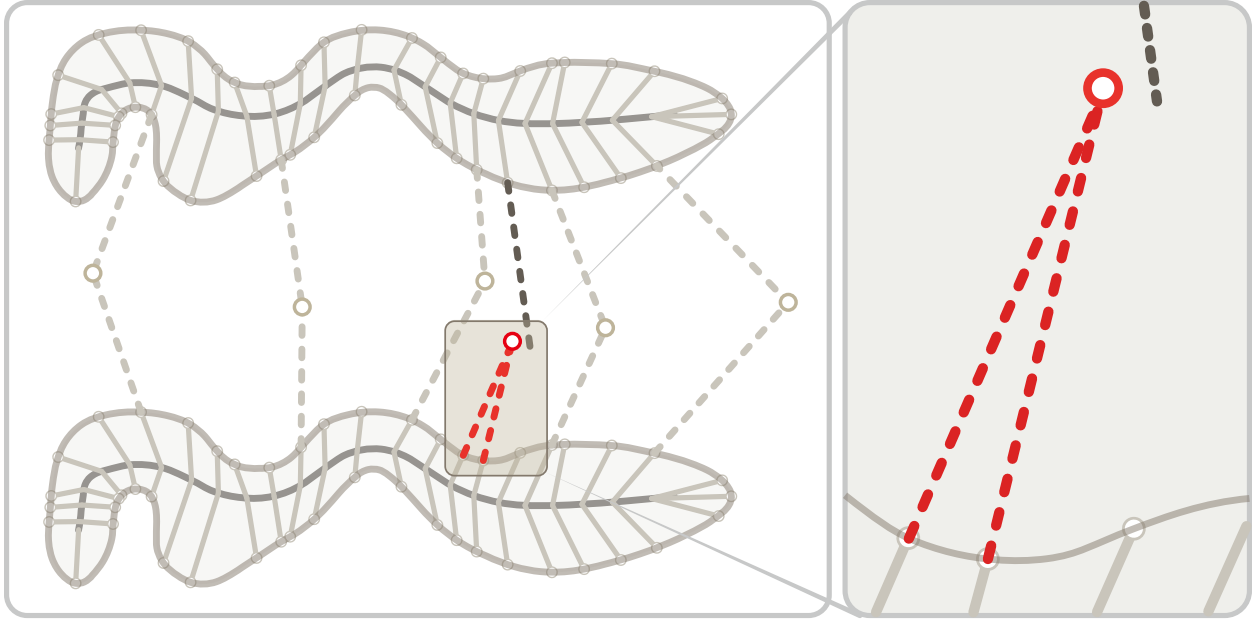
In the following, Section 5.3.1 reviews Damon-Gasparovic's linking structures on which the proposed linking structures are based. In Section 5.3.2 I first give a formal definition of the modified linking flow. Then I describe the method and algorithm of the modified linking flow. Finally, I discuss the properties of the linking flow. Section 5.3.3 discusses between-object shape features that are extracted from the modified linking flow.

5.3.1 Damon-Gasparovic's Linking Structures


This chapter aims to capture positional and geometric relations that are anatomically important. For example, a pair of abutting objects (Krishna et al., 2022) has close relations in terms of local geometry on their boundaries. Additionally, the variation of each single object should be constrained by the non-penetration condition. Instead of analyzing separate objects, it is useful to explicitly analyze within- and between-object shape.

However, there exist very few mathematical models that define the geometric relations between objects. Many existing research projects address this problem through statistical analysis of a population of data. Without explicitly capturing between-object shape, those statistical relations often ignore important and consistent variations of between-object relations due to sampling bias.

To solve the problem, Damon and Gasparovic (see (Damon and Gasparovic, 2017)) have proposed linking structures that can be built upon general skeletal structures and hence specifically upon s-reps. In essence, the researchers regard between-object relations as shape features of the between-object space (i.e., the bounded complementary space of multiple objects). Accordingly, the linking structures are equivalent to skeletal representations of the between-object shape. The computation of an s-rep linking structure involves multiple neighboring objects, each of which has a fitted s-rep. First, a target object of interest is selected. Then, each spoke in the target object is extended to potentially intersect other extended candidate spokes. Each intersected pair (including the target spoke and a candidate spoke)



Damon-Gasparovic's Linking



Proposed Linking

Figure 5.1: Simulated links between two 2D objects. Compared to Damon-Gasparovic's linking structure (top), the modified linking structure (bottom) changes the link directions to avoid self-linking, as highlighted in the yellow boxes.

is associated with an extension. Finally, the linking spoke is selected among all the candidate spokes for which the extension is minimum.

While Damon-Gasparovic’s linking structures (see the top of fig. 5.1) between s-reps provide a means of relating the geometry of adjacent objects, it is challenging to use these linking structures in statistical analysis of 3D multi-object complexes. Due to relaxation of the Blum conditions, spokes in s-reps may have no candidate links at all, yielding inconsistent empty links across a population of s-reps. In addition, the *self-linking spokes* (see the top of fig. 5.1) that are two linked spokes from the same object introduce inconsistent branching topology of between-object shape, which can harm the quality of correspondences across a population. Such self-linking spokes often are associated with dents on a boundary. Therefore, I regard them as unnecessary links in analyzing between-object relations. In the following, I develop a modified version of Damon-Gasparovic’s linking structure for statistical analysis of multi-object shape.

5.3.2 Modified Linking Flow

Modified linking flow aims to relate the geometry from multiple objects, focusing on how the objects link with each other. The flow leads to an external linking surface that is analogous to a skeletal representation of between-object space.

Definition. The modified linking flow ψ is defined as a smooth bijective mapping from a skeleton \mathcal{M} to the external linking surface \mathcal{L} between two objects, i.e.,

$$\psi : \mathcal{M} \mapsto \mathcal{L}. \quad (5.2)$$

In the modified linking flow, the resulting external linking surface is a smooth non-branching surface. Given a skeletal point $p \in \mathcal{M}$, I call the image $\psi(p) \in \mathcal{L}$ of p a link point. Accordingly, I call the vector pointing from p to $\psi(p)$ a link vector, whose magnitude and direction I call a link length and a link direction, respectively. The following methods can construct a mapping ψ such that (1) the mapping ψ is well-defined everywhere on an s-rep and (2) the external linking surface has no subfigures.

Approach. The method for constructing a desirable mapping ψ for every configuration has 3 steps:

1. Find *regularly linked spokes* using the method given by Damon (Damon, 2019), as described below;
2. Infer the mapping ψ from the regularly linked spokes;
3. Apply the mapping ψ to the skeleton.

A pair of regularly linked spokes is formed by two spokes from the s-reps of two objects; the two spokes link at a place between the two objects. Specifically, in a complex of two objects Ω_1 and Ω_2 , let \mathbf{s}_1 and \mathbf{s}_2 denote two spokes that are in the objects Ω_1 and Ω_2 , respectively. The directions of the two spokes are \mathbf{u}_1 and \mathbf{u}_2 ; the end points of the two spokes are \mathbf{y}_1 and \mathbf{y}_2 . The two spokes are candidate linked spokes *only if* there exists a non-negative scalar t such that

the extensions of the two spokes lie outside the objects and intersect at their endpoints, i.e.,

$$\mathbf{y}_1 + t\mathbf{u}_1 = \mathbf{y}_2 + t\mathbf{u}_2, \quad s.t. \quad t \geq 0, \quad (5.3)$$

where t denotes the length of the extension of the two spokes.

That intersection point, the link point $\mathbf{y}_1 + t\mathbf{u}_1$ (or equivalently, $\mathbf{y}_2 + t\mathbf{u}_2$) is on the external linking surface. Importantly, there can exist zero or many candidate spokes in Ω_1 and Ω_2 that link to \mathbf{s}_1 , satisfying eq. (5.3) with different extensions t . A linking point is the one where the length “ t ” for the pair of specific candidate linked spokes is minimal among all candidate linking points involving either of the given pair of spokes.

Next, I infer the mapping ψ using the regularly linked spokes found by the above method. Let S_1 denote a set of regularly linked spokes in Ω_1 . The skeletal points of S_1 give a subset $\widehat{\mathcal{M}}$ of the skeleton \mathcal{M} . In addition, the link points of S_1 give a subset $\widehat{\mathcal{L}}$ of the external linking surface \mathcal{L} . To derive a smooth mapping $\psi: \widehat{\mathcal{M}} \mapsto \widehat{\mathcal{L}}$, I use the Thin Plate Spline (TPS) algorithm with landmarks on $\widehat{\mathcal{M}}$ and $\widehat{\mathcal{L}}$, thereby avoiding self-links in Damon-Gasparovic’s method.

Last, I apply ψ to the complement of $\widehat{\mathcal{M}}$ to form a smooth external linking surface \mathcal{L} . By doing so, I have a link point $\psi(p)$ for every skeletal point $p \in \mathcal{M}$. Algorithm 3 gives the detailed algorithm.

Algorithm 3 Modified linking flow

Require: A pair of s-reps fitted to two nearby objects.

- 1: Select one object as the target object.
 - 2: Compute Damon-Gasparovic’s linking structure P according to (Damon, 2019) for the target object.
 - 3: Create an empty set $S_1 = \emptyset$ of pairs of points.
 - 4: **for** each link l in P **do**
 - 5: **if** l regularly links to another object **then**
 Add the two ends of l to S_1 .
 - 6: **end if**
 - 7: **end for**
 - 8: Apply TPS to infer deformation ψ using pairs of point in S_1 .
 - 9: Apply ψ to all skeletal points on \mathcal{M} of the target object.
 - 10:
 - 11: **return** the result of $\psi(\mathcal{M})$.
-

The above modification makes use of regularly linked spokes to modify both the self-linking and unlinked spokes such that two-object complexes consistently yield non-branching smooth external linking surfaces. This consistency allows me to extract between-object shape features with good correspondences. Also, because the modified linking flow is derived from regularly linked spokes of two objects, the mapping ψ can capture positional and geometric relations between objects. In order to represent these relations for statistical analysis, I discretize the link vector field and extract corresponding link features from the discrete link vectors, as described below.

5.3.3 Link Vectors and Between-object Shape Features

Like spokes in an s-rep, link vectors can be understood as a discretized vector field defined on a smooth manifold (i.e., the skeleton \mathcal{M}). To facilitate statistical analysis, these link vectors should be sampled to have good correspondences across a population.

To this end, I sample link vectors at each skeletal point in a discrete s-rep. These link vectors have good correspondence because the base manifold is sampled with good correspondence. As discussed in Chapter 3, each s-rep is deformed from an s-rep of an ellipsoid. The skeletal points are sampled with respect to the relative positions within the ellipsoid. Such consistent discretization of skeletons gives good correspondences across samples.

From the discrete link vectors, I focus on (1) magnitudes of the vectors that indicate link lengths between a skeleton and the external linking surface and (2) directions of link vectors that may or may not align with spokes' directions. Those link directions that align with spokes' directions correspond to the regular links, which can be regarded as geometric relation of two objects. The link directions that differ from spokes' directions are produced by the deformation ψ as described in Algorithm 3.

In order to obtain statistical correlation with within-object shape, it is necessary to Euclideanize between-object shape features in NEUJIVE. The Euclideanization of between-object shape features should consider both Euclidean variables (i.e., log link lengths) and non-Euclidean directional data (i.e., link directions). For the link lengths, I convert these non-negative scalar values to Euclidean variables distributed in \mathbb{R} with the log map (Tu et al., 2018). For the directional data that live on a unit 2-sphere, I use Principal Nested Spheres (Jung et al., 2012) to convert the data into Euclidean coordinates.

A subtlety of these features is their sensitivity to positions and orientations of objects. To overcome this problem, I capture between-object shape features in local affine frames to achieve invariance properties, as described below.

5.4 Local Affine Frames on Skeletal Representations

It is often challenging to align objects for analysis of geometric features, especially when there are multiple objects in every sample. For a population of multi-object complexes, aligning individual objects can sacrifice between-object relations, while aligning the combination of multiple objects can bias the analysis of individual objects.

To address the above dilemma, this section describes geometric features that are with respect to local frames. Specifically, Section 5.4.1 develops local affine frames to capture rich local geometric features of objects. These features, including within- and between-object shape features (see Section 5.4.2), are insensitive to rigid transformations. Importantly, because the local affine frames are fitted via a deformation-based method, they also capture local geometry, which is useful in distinguishing the shape.

5.4.1 Local Affine Frames

In this dissertation, a local affine frame refers to a general frame centered at a point on a 2D surface. Such a frame consists of 3 basis vectors that are not necessarily orthogonal to each other. Moreover, the 3 frame vectors are not necessarily of unit lengths.

Specifically, each skeletal point is associated with an affine frame that allows us to map features at the skeletal point. Let $\mathbf{x} \in \mathbb{R}^3$ be a unit directional vector (e.g., a spoke's direction) associated with a skeletal point p ; assume \mathbf{x} is represented in global coordinates. An affine frame at p can be represented as a matrix $A \in \mathbb{R}^{3 \times 3}$. In a special case where this frame is an orthonormal frame, the matrix A is an orthogonal matrix of determinant 1. Generally, the directional feature \mathbf{x} can be mapped into the local affine frame via

$$\mathbf{x}' = A\mathbf{x}, \tag{5.4}$$

where $\mathbf{x}' \in \mathbb{R}^3$ is a feature vector that is insensitive to rigid transformations.

To better capture geometry of objects, I take into account the following objectives in constructing local affine frames on s-reps:

1. the origins of these local frames should have good correspondences across samples;
2. the basis vectors of local frames within a sample should reflect within-object geometry;
3. the basis vectors of local frames across samples should capture geometric difference.

As skeletal points in s-reps provide good correspondence across a population (see Section 5.3.3), the first objective can be satisfied by fitting a local frame at each skeletal point in an s-rep. The second and third objectives are achieved by fitting local frames via a deformation-based method as follows. I choose an ellipsoid as a template shape. Then I construct orthonormal local frames (see Section 2.6) on the ellipsoidal s-rep. I construct a diffeomorphism (here via TPS) to deform this template s-rep to the s-rep of each target object, mapping skeletal points on the ellipsoidal s-rep to corresponding skeletal points on the target s-rep. Such a deformation also stretches and rotates the orthonormal frame vectors of every local frame, resulting in basis vectors of a local affine frame.

In a single 3D object like a hippocampus (see fig. 5.2), the transformation between the affine frames is achieved by a linear operator that combines the rotation and the stretching of basis vectors. Section 5.5.3 demonstrates that the local geometry captured by local affine frames is useful in classifying the ASD and non-ASD group.

An affine frame field is promising in multi-object shape analysis not only because it captures geometric features within objects invariant to rigid transformations, but also because it enriches between-object shape features, as detailed in the following.

5.4.2 Within- and Between-object Shape Features in Local Affine Frames

Due to the above methods, there are 3 geometric entities associated with each skeletal point in an s-rep: a spoke, a link vector and an affine frame. To obtain geometric features invariant to rigid transformations, I map vectors at each skeletal point into the respective local affine frame using the linear transformation eq. (5.4).

Within-object shape features from spokes include the skeletal points, the spokes' lengths and directions. To make the coordinates of skeletal points independent of the coordinate system, I convert the global coordinates of skeletal points into the coordinates relative to the skeletal center. Moreover, I map the unit direction vectors to the associated local frames. In order to retain the unit length of the projected directional vector, I normalize the three coordinates of \mathbf{x}' . Likewise, I map the link features from the link vectors to the local affine frames.

The resulting within- and between-object shape features are sent to X and Y , respectively, in eq. (5.1) for the joint analysis. In NEUJIVE, the Euclideanization of within-object shape features is the same with the Euclideanization of between-object shape features, as discussed in Section 5.3.3.

5.5 Experiments

This section demonstrates the proposed methods using the database described in Section 2.2. In Section 5.5.1, I show the between-object link features in two-object complexes, including the hippocampus and the caudate. I collect such link features and within-object s-rep features, pooling the two groups. These pooled features form two data blocks, as discussed in Section 5.2. Then I apply NEUJIVE to the two blocks. Section 5.5.2 shows (1) the generalizability of a linear classifier learned on the proposed features and (2) the statistical significance of the group difference found in proposed feature space. To demonstrate the benefit from local affine frames, Section 5.5.3 describes the features from local affine frames and shows the classification results using these features. Finally, Section 5.5.4 lays out an extensive comparison of classification performance using various features of multi-object complexes.




Figure 5.2: Fitted local affine frames on the skeleton of a hippocampus. The skeleton is shown as the orange surface. At each skeletal point, there is an affine frame consisting of a $\nabla\tau_1$, $\nabla\theta$ and normal vector (see the zoomed picture above). The vectors of $\nabla\tau_1$ and $\nabla\theta$ are not necessarily orthonormal vectors.

5.5.1 Visualization of Between-object Shape Features from Modified Linking Structures

As described in Chapter 3, I fit an s-rep to every triangular mesh. In the resulting s-rep, I have 72 smooth skeletal points on each side and 24 skeletal edge points. A smooth skeletal point is associated with a spoke pointing toward one side of an object, while an edge point is associated with an edge spoke pointing away along the tangent direction of the skeleton.

Between-object shape features from modified linking structures are link features defined on object skeletons. Each skeletal point is associated with a spoke length, a spoke direction, a link length and a link direction. To discount the object widths in the links, I show the extensions of spokes (i.e., the values of t 's in eq. (5.3)), subtracting the link lengths by the spokes' lengths. Though the extensions are distances from the boundary to the external linking surface, I can map such scalar features onto the corresponding skeletal points, as shown in fig. 5.3.

In fig. 5.3 the top row shows the extensions as heat maps on skeletons of the hippocampus (bottom) and the caudate (top) in 5 autistic cases. These heat maps share similar patterns: (1) because the caudate bends away from the hippocampus, the links in the head and the tail have smaller extensions than the links in the body and (2) because the caudate and the hippocampus are not parallel to each other and because one side of the caudate is farther from the hippocampus than the other side, the extensions gradually vary from the near side to the far side of the caudate skeleton. The bottom row shows the link directions in the same 5 autistic cases. While these directions can be slightly different from spoke directions, the results suggest that the link vectors smoothly swing along the skeletons.

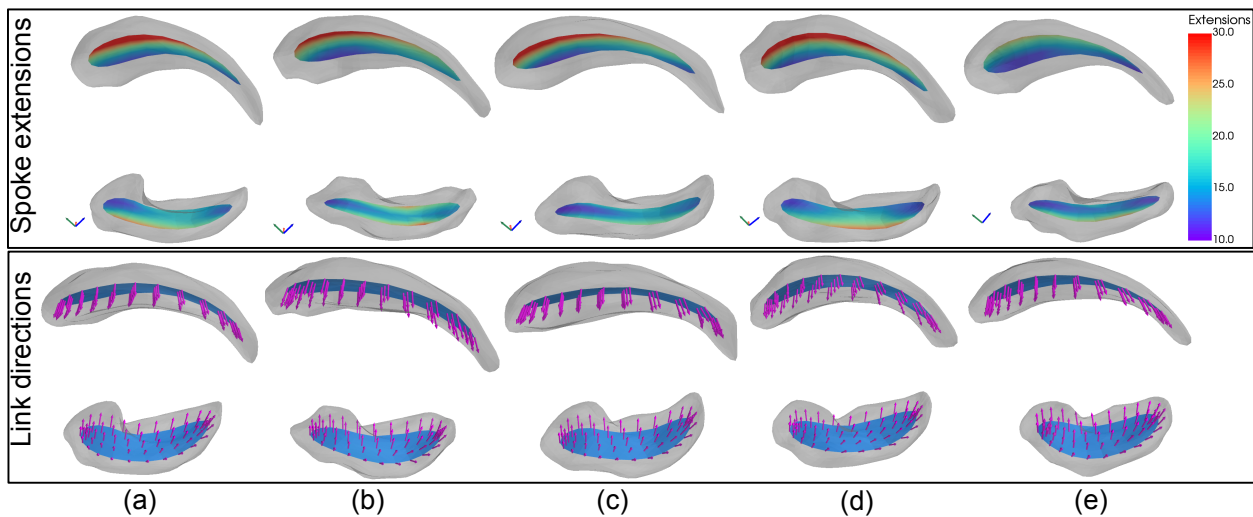


Figure 5.3: Linking features of 5 caudate-hippocampus pairs: (a)-(e). The transparent surfaces are the caudate (above) and the hippocampus (below). Top row: spoke extensions (in mm) from a boundary to the external linking surface. The heat maps on the skeletons show the values of spoke extensions. Bottom row: link directions of link vectors at discrete skeletal points, shown as the arrows in magenta. The skeletons are shown as blue surfaces.

5.5.2 Classify Autism vs. Non-autism with Two-object Complexes

In this section, I show the results from classifying ASD vs. non-ASD using the shape of the hippocampus and the caudate. Specifically, I intend to investigate the generalizability of a linear classifier trained in the proposed feature space. To this end, I train the classifier with varying training data size and evaluate the corresponding test performance. Moreover, I investigate the statistical significance of the group difference using the proposed features. I perform hypothesis testing with a large number (1000) of permutations to avoid bias of sampling.

The proposed methods take as input the SPHARM-PDM models of objects. I fit s-reps to these surface meshes, extracting the within-object shape features $\mathbf{x}_i \in \mathbb{R}^{d_1 \times 1}$ for the i^{th} sample (see Section 5.4.2). Then I develop the modified linking structure of each sample, extracting the between-object shape features $\mathbf{y}_i \in \mathbb{R}^{d_2 \times 1}$ (see Sections 5.3.3 and 5.4.2). In developing the modified linking structure, I interpolate the s-reps using the method in Chapter 3 to have a moderate number (about 10) of regular links in every sample. This operation is similar to smoothing the external linking surface because an even finer interpolation tends to “overfit” the shape of individual objects. Finally, I apply NEUJIVE to extract the correlated shape features of $X = \{\mathbf{x}_i\}$ and $Y = \{\mathbf{y}_i\}$ (see Section 5.2), where $i = 1, \dots, 176$. I set the initial ranks for both X and Y as 10 in NEUJIVE, yielding the joint shape feature matrix of dimension 4×176 (i.e., $r = 4$ in eq. (5.1)). Following the post-feature-selection idea (see Chapter 4), I perform classification and hypothesis testing using that feature matrix as follows.

Classification. I train and test a linear classifier called Distance Weighted Discrimination (DWD) (Marron et al., 2007) within the proposed feature space. I randomly select k samples from the autistic samples and k samples from the non-autistic samples to form the training dataset. The remaining are used as the test data. To avoid the bias introduced by the partitioning of training and test data, I perform 1000 random partitions and classifications for each k . From these 1000 experiments, I obtain an average test Area Under Curve (AUC) (see Section 2.9.3).

To evaluate the generalizability of the classifier. I vary k from $k = 6$ (20% of the total autism samples) to $k = 27$ (80% of the total autism samples). The resulting average test AUCs of the various k 's are shown in the rightmost panel of fig. 5.4. In this figure the outperforming (top) curve results from the CoWBO features. The other two features are obtained similarly to the above process but using different blocks in NEUJIVE. Specifically, for NEUJIVE-Sreps I use the correlated features of (1) the s-reps features of the hippocampus and (2) the s-reps features of the caudate. For NEUJIVE-PDMs I use the correlated features of (1) the SPHARM-PDMs of the hippocampus and (2) the SPHARM-PDMs of the caudate.

As I vary the training size, I observe (1) that the proposed features can learn useful patterns of the ASD vs. the non-ASD even in a small data size, e.g., when the training sample size is 20% and (2) that the increase of the test performance is proportional to the increase of training data size, which is superior to the classification using NEUJIVE-Sreps and NEUJIVE-PDMs features that are solely from within-object shape features.




Figure 5.4: Left: Hypothesis testing using the correlated shape features of SPHARM-PDMs of the hippocampus and caudate. Middle: Hypothesis testing using the correlated shape features of the hippocampal s-rep features and the link features. Right: Average test AUCs over varying training data size. The error bar at each point shows the standard error of the mean.

Hypothesis testing. To verify if the above group difference statistically significant, I test the (null) hypothesis that the two groups have the same mean value in the proposed feature space. I use a method called Direction-Projection-Permutation (DiProPerm) (Wei et al., 2016) to perform the hypothesis test for a large number of permutations.

I let DiProPerm generate 1000 permutations of the class labels. The resulting test statistics (mean differences) are shown in the left two panels of fig. 5.4 as the green dots. Specifically, the middle panel results from the hypothesis test using the CoWBO features. The observed statistic is shown as the vertical dashed line, which has p-value < 0.001 and z-score 10.074. The p-value and z-score together suggest that the group difference observed is statistically significant in the feature space formed by the proposed geometric and statistical methods.

As a comparison, the left panel results from another feature space that contains the correlated SPHARM-PDM features of the hippocampus and the caudate. Both the p-value (0.018) and z-score (2.425) suggest the group difference there is less statistically significant.

5.5.3 Local Affine Frames Capture Differentiating Local Geometry

This section demonstrates the benefit from local affine frames in the classification problem. Because ASD affects multiple local regions of a brain structure, geometric features should have different importance in classifying ASD and non-ASD across various local regions. Typically, a classifier assigns different weights to the features to best classify the two groups in a dataset. In my research, the projection to local affine frames (see eq. (5.4)) is analogous to weighting local geometric features for a better classification because the fitted local frames can capture important local geometry between groups.

Specifically, I fit every local affine frame in an object by deforming a local orthonormal frame in the template ellipsoid¹. To demonstrate the local geometry captured by the frames in the object, I focus on the deformation captured

¹I choose the Procrustes Mean ellipsoid as a template shape for each object. I use the General Procrustes Analysis method (Dryden and Mardia, 2016) to obtain this Procrustes Mean of the ellipsoids resulting from Mean Curvature Flow on each object.

along the $\widehat{\nabla\tau_1}$ and $\widehat{\nabla\theta}$ directions on the skeleton (see Section 2.6). As a result of the deformation, the unit vectors $\widehat{\nabla\tau_1}$ and $\widehat{\nabla\theta}$ in the template ellipsoid become general vectors $\nabla\tau_1$ and $\nabla\theta$ in the object. The vectors $\nabla\tau_1$ and $\nabla\theta$ together capture the skeleton deformation. I compute the cross product $\mathbf{n} = \nabla\tau_1 \times \nabla\theta$ at every skeletal point. The length of \mathbf{n} gives the area of the parallelogram defined by $\nabla\tau_1$ and $\nabla\theta$.

Then I extract the features from the cross-products \mathbf{n} to classify ASD and non-ASD. The results are shown in Table 5.1. Table 5.1 reports the average test AUCs over 1000 random holdouts, each of which randomly takes 80% of data as training data and the remaining as test data. In the table header, $\|\mathbf{n}\|$ denotes using the lengths as classifying features. The symbol “ \oplus ” denotes the concatenation of features. The last column shows the average test AUCs from the classification using the concatenation of the lengths of \mathbf{n} ’s and the coordinates of the skeletal points as features.

Object	$\ \mathbf{n}\ $	\oplus Skeletal point
Hippocampus	0.589	0.614
Caudate	0.518	0.52

Table 5.1: Average test AUCs from using affine frames.

These results suggest that the cross-products of basis vectors $\nabla\tau_1$ and $\nabla\theta$ already have some discriminatory power (AUC > 0.5). Moreover, this discriminatory power is slightly boosted by adding the skeletal points at which those local affine frames are centered.

5.5.4 Comparing with Various Feature Spaces

In order to investigate the difference between using within-object and between-object shape features in classification, I conducted extensive experiments using various sets of features, as listed in Table 5.2.

Table 5.2: Average test AUCs with features from two-object complexes

Features	Global C.S.*		L.A.F.*		
	Mean	SEM*	Mean	SEM*	
Hippocampus	① SPHARM-PDM	0.591	0.097	-	-
	② Spokes	0.596	0.103	0.624	0.01
	③ PDMs of skeletons	0.598	0.102	0.627	0.096
Caudate	④ SPHARM-PDM	0.54	0.088	-	-
	⑤ Spokes	0.514	0.088	0.552	0.01
	⑥ PDMs of skeletons	0.502	0.089	0.521	0.102
Two objects	① & ④ PDMs	0.588	0.089	-	-
	② & ⑤ Spokes	0.594	0.094	0.591	0.008
	③ & ⑥ PDMs of skeletons	0.593	0.095	0.586	0.01
NEUJIVE	① & ④ PDMs’ joint features	0.563	0.009	-	-
	② & ⑤ Spokes joint features	0.552	0.01	0.631	0.007
Proposed	⑦ Between-object linking features	0.679	0.009	0.68	0.009
	② & ⑦ CoWBO features	0.668	0.009	0.698	0.008

*C.S. stands for Coordinate System

*L.A.F. stands for Local Affine Frames

*SEM stands for Standard Error of the Mean

In the column of Global Coordinate System (Global C.S.) in Table 5.2, the points and vectors are expressed with global coordinates after Procrustes alignment, whereas the Local Affine Frames (L.A.F.) columns use the invariant features described in Section 5.4.

The rows in the table compare average test AUCs over 1000 random holdouts involving various sets of features. Fig. 5.5 illustrates some of these features. In particular, I categorize these features into 5 groups shown in the leftmost column. The first two groups (“Hippocampus” and “Caudate”, respectively) use various choices of “within-object shape features of the hippocampus and the caudate, respectively. Each choice is indexed with a circled number. Here, the PDMs of skeletons consist of 3 points sampled at every local affine frame position; they are the origin and two other points that are ε -distance away from the origin along the $\nabla\theta$ and $\nabla\tau_1$ axes.

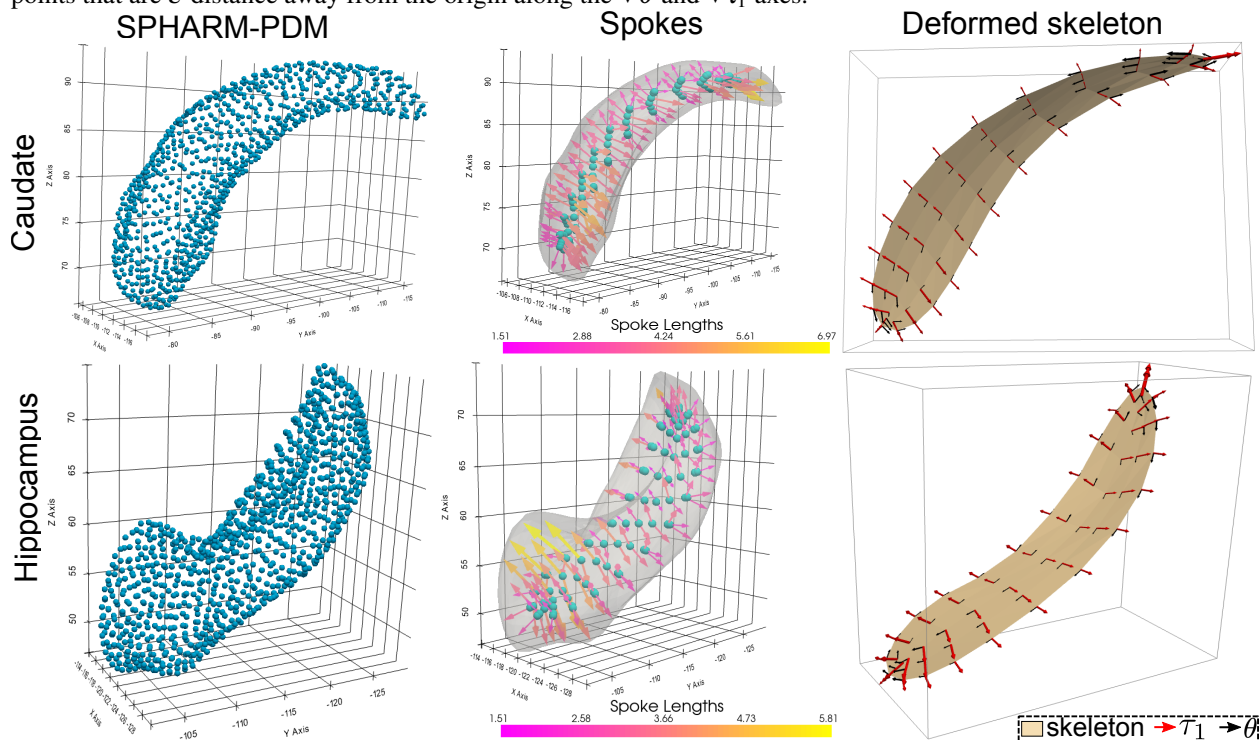


Figure 5.5: Various shape features of a two-object complex consisting of a caudate (top row) and a hippocampus (bottom row). The left column shows boundary landmarks from SPHARM-PDM (Styner et al., 2006). The middle column shows spoke features from (Liu et al., 2021) – including coordinates of skeletal points, unit directions of spokes and lengths of spokes. The right column shows deformation field from an ellipsoidal skeleton to the skeleton of a target object (the caudate or the hippocampus). The arrows τ_1 and θ indicate the deformation in two directions.

The third group uses shape features from two-object complexes. Specifically, the row “① & ④ PDMs” concatenates coordinates of SPHARM-PDMs of the hippocampus and the caudate; the row “② & ⑤ Spokes” concatenates spokes’ features of the two objects; the row “③ & ⑥ PDMs of skeletons” concatenates the skeletal PDMs.

The fourth group uses the resulting joint components from NEUJIVE. I tuned the initial ranks of NEUJIVE for each row in this group, as described in Section 4.6.4. For example, the PDMs’ joint features (“① & ④ PDMs’ joint

features”) refer to the NEUJIVE joint components taking the SPHARM-PDMs of the hippocampus and the caudate as two input blocks.

My major proposal is shown in the last group. The row “⑦ Between-object linking features” contains the hippocampal skeletal points and the link features sampled at these skeletal points. The last row results from the CoWBO features, as discussed in Section 5.2. From this table, I have some key observations as follows.

Observations from Table 5.2.

1. The hippocampal shape features give notably stronger discriminatory power than caudate features, as seen by comparing the corresponding rows in the “Hippocampus” group with the “Caudate” group. Due to this observation, I choose to neglect within-object shape features of the caudate in forming CoWBO features.
2. Using hippocampal skeletal features is slightly better than using hippocampal boundary features. However, this is not the case for classification using caudate shape.
3. Using concatenated shape features from two objects is not very much different from using hippocampal shape features alone (comparing the corresponding rows in the “Hippocampus” group with the “Two objects” group). This is because the increased classification power that might come from adding weakly classifying caudate features can not compensate the loss due to increasing feature dimensions.
4. The between-object shape features (i.e., ⑦) contain critical discriminatory information. Particularly, I used the modified linking structures of the hippocampus that capture the geometric relations with its neighboring caudate, yielding notably better classification performance than those simple concatenated shape features.
5. The global coordinates give inferior classification results compared to coordinates with respect to local affine frames, at least when within-object shape features are included. Because the template ellipsoids are deformed to each target object separately, the between-object relations are ignored in fitting local affine frames. The lack of those relations can explain the limited advantages of local affine frames in classification with multi-object shape features (see the “L.A.F.” column in the “Two objects” group).
6. **The CoWBO features give notably the best classification performance.** Compared with the result (0.624) from the spokes features of the hippocampus and the result (0.68) from the between-object link features, the average test AUC (0.698) from the CoWBO features increases by about 12% and 3%, respectively.

5.6 Conclusions and Discussion

In this chapter, I proposed a combination of geometric and statistical models to extract correlated shape features from two-object complexes. The proposed methods have been implemented in <https://github.com/ZhiyLiu/shanapy>.

5.6.1 Geometric Models

The geometric models including modified linking structures and local affine frames are both beneficial in classifying ASD and non-ASD. Especially, the classification using link features between the hippocampus and caudate outperforms the classification using single-object shape features by a large margin. This superior performance due to link features suggests that the geometric relations between the hippocampus and caudate are important in the development of ASD.

Moreover, in the classification using single-object shape features, I found that the fitted local affine frames notably boost the performance. Such a boost is due to two reasons: first, the angles and lengths of basis vectors characterize local geometry; second, the projection into local affine frames yields shape features that are irrespective of poses and positions of objects.

5.6.2 Statistical Analysis

The joint analysis of within-object s-rep features and between-object link features can further improve the classification performance. Considering the different variability of within- and between-object shape features, I used NEUJIVE to obtain their correlated shape features. To produce robust and generalizable correlations, NEUJIVE makes use of the geometry of the complex data spaces in both within- and between-object features. The resulting shape features show superior statistical power in both hypothesis testing and classification.

A limitation of NEUJIVE is that the test data and training data have to be pooled to construct the joint variation subspace. This requirement slows the inference process. However, this problem is negligible in the field of medical image analysis where data sizes tend to be small.

In addition to the effectiveness of the proposed methods in shape analysis of two-object complexes, the following chapter discusses more potential applications in future research.

CHAPTER 6: CONCLUSIONS & DISCUSSION

In this chapter I summarize the contributions and thesis of this dissertation in Section 6.1. In Section 6.2 I discuss possible extensions and limitations of the projects completed in this dissertation.

6.1 Contributions and Thesis

This section summarizes the major contributions in Chapters 3 to 5. Then I restate the thesis that was laid out in Chapter 1. In principle, this dissertation aims to solve problems in analyzing multi-object complexes for medical image analysis. To this end, I made theoretical and technical improvements, as detailed in previous chapters. These improvements include developing both geometric and statistical methods for multi-object shape analysis. The target problems and the key contributions I have made are listed below.

Problems and Solutions:

1. *Boundaries are often not reliable in anatomic shape representation. Therefore, the interior geometry of objects should be taken into consider when represent anatomic shape with s-reps.*

I discussed the s-rep fitting work in Chapter 3. In that work, a non-branching object is deformed to an ellipsoid through a series of local diffeomorphisms. Then the s-rep of that ellipsoid can be deformed and fitted to the object via the reverse deformation, forming the initial fit of the target object. Because such initialization solely relies on the boundary of the object, it can end with non-smooth geometry as a result of the reverse deformation. To address this problem, I proposed the refinement stage that makes use of the interior geometry of objects (see Section 3.5), considering how spokes smoothly swing within an object. As a consequence, the fitted s-reps capture smooth geometry within and on the boundary. Such smooth geometry can be seen in the level surfaces of radial flow, as shown in fig. 6.1.




Figure 6.1: Top: a caudate with partially corrupted boundary (highlighted by the red circle) from SPHARM-PDM (Styner et al., 2006). Bottom: The level surfaces of the radial flow of an s-rep of a caudate. These level surfaces describe the interior geometry of the object.

2. *Aligning objects can remove distracting information in shape analysis such as translation and rotation. However, it is difficult to align multi-object complexes. To address this problem and to comprehensively represent local*

geometry, I proposed local affine frames to capture geometric features that are insensitive to positions and orientations of objects but sensitive to local stretching and bending.

In Chapter 5 I extended the idea of local orthonormal frames of s-reps due to Pizer (Pizer et al., 2021) by relaxing the orthonormal conditions. The resulting non-orthonormal local frames are referred to as *local affine frames*. I fit these local frames to each object by deforming orthonormal frames in a template ellipsoid via a deformation that deforms ellipsoidal skeleton to the object's skeleton. In each object the within- and between-object shape features can be transformed to these local affine frames. The use of local affine frames removes the sensitivity to positions and orientations of objects. Moreover, these local affine frames are useful in classifying objects' shape because they also capture local geometry.

- 3. The representation of geometric relations between objects should provide good correspondences across a population for statistical analysis. To this end, I proposed and realized a modified version of Damon's linking structure to explicitly capture geometric relations between objects.*

I discussed the modified linking structures in Chapter 5. Following the philosophy of NEUJIVE as mentioned in item 5, the linking behaviors of adjacent objects can be understood as searching the intersection space of multiple shape functions (s-reps) that describe those objects. This realization inspired me to capture geometric relations by linking structures. To facilitate statistical analysis, I modified Damon-Gasparovic's linking flow (Damon and Gasparovic, 2017) to avoid many-to-one mapping from the skeleton to the external linking surface. By doing so, the resulting link vector field is well-defined everywhere on the skeleton. Moreover, the external linking surfaces are endowed with consistent non-branching topology. Such consistency allows me to sample link features with good correspondences across a population. Though this dissertation demonstrates the modified linking structures in a typically positioned two-object complex, i.e., the two objects are separated from each other, our group attempts to apply this method in other types of two-object complexes (e.g., abutting objects). I discuss this research in Section 6.2.

- 4. The correlation between structural morphology in the human brain and the development of ASD has been unclear. In this regard, I investigated the correlation in a classification task, in which I extracted shape features from two-object complexes and trained a classifier using those shape features.*

In Chapters 4 and 5 I conducted experiments on classifying the autism vs. non-autism group with the proposed shape features. Because autism can simultaneously cause shape changes of multiple brain structures I hypothesized that the integrated features of the hippocampus and the caudate would yield better classification results, compared to single-object shape features. In the experiments, I found that both the statistical and geometric correlation features (so-called CoWBO, short for Correlation of Within- and Between-Object features) are useful in classifying the ASD vs. non-ASD group. This finding suggests that between-object relations are important in

the development of ASD. Moreover, I test the statistical significance of the group differences. The results suggest that the group differences found in the correlated features are statistically significant.

5. *Shape correlation is a critical component in understanding multi-object complexes. From the statistical point of view, the joint analysis of multi-object complexes can yield robust shape correlation of a population. To this end, I developed a method called Non-EUclidean Joint and Individual Variation Explained (NEUJIVE) to extract shape correlation from multi-object complexes.*

I detail the theory and applications of NEUJIVE in Chapter 4. The underlying philosophy of NEUJIVE and the statistical methods on which it depends, AJIVE and JIVE, is that the shared subspace of multiple variation spaces should be the joint variation space in which the correlated features live. The key contribution of NEUJIVE was to extend the Euclidean method AJIVE to a method for shape data, which live in a non-Euclidean space. I noticed that the resulting statistical correlation depends on the data. The correlation can be biased due to the sampling of data and thus might deviate from anatomically important between-object relations. To address this problem, I developed explicit geometric models for multi-object shape representations, as discussed in items 4 and 5.

6. *There exist many possible combinations of features from two-object complexes. This dissertation conducted extensive experiments to select the features that give good classification performance.*

Both single-object shape features and composite shape features have discriminatory power in the IBIS dataset. Especially, the hippocampal shape features give notably stronger discriminatory power than caudate features. Moreover, the hippocampal s-rep features show slightly better performance than the boundary features. Importantly, the correlation of within- and between-object shape features (CoWBO) yield the best classification performance. The reader is referred to Chapter 5 for a comprehensive discussion.

7. *Previously, there existed no public tool for anatomic shape analysis using s-reps. In my research, I develop skeletal representation module in the toolkit SlicerSALT, which is frequently used in the clinical research community. Recently, I have released a Python package for SHape ANALysis using s-reps, called SHANAPY. The public package can be found in my GitHub repository.*

The implementation in python has been reported as a useful tool for computational geometry and topology (Miolane et al., 2021). I also provide documents so that the package can better serve the research community.

Based on all the above development and experiments, I reiterate my thesis as follows.

Thesis. *Joint shape variation and between-object shape relations can be effectively extracted from multi-object complexes by deliberately combining geometric methods which capture within- and between-object shape features invariant to rigid transformations, and statistical methods which focus on shape correlation in multi-object complexes.*

Specifically, following properly relaxed Blum conditions, the fitting process results in more satisfactory s-reps with desirable geometric and topological properties for single objects. The proposed Non-EUclidean Joint and Individual Variation Explained method (NEUJIVE) can explain multi-object shape variations with notable effectiveness and robustness. The modified linking structures provide rich and useful geometric features of multi-object complexes, including within- and between-object shape features. To investigate the thesis, I conducted experiments on the hippocampus-caudate pairs in 6 months olds, classifying the ASD vs. the non-ASD group. The joint analysis of the within- and between-object shape features gave the best classification performance among various choices of shape features in the experiments.

6.2 Limitations and Future Work

This section discusses limitations and several possible extensions of this dissertation. Some of these ideas have been started in our group and our collaborators' groups.

Specifically, Section 6.2.1 discusses limitations of the current s-rep fitting framework mentioned in fitting more objects of complex shape (e.g., mandibles). Accordingly, I discuss extending the fitting framework by making use of a more stable curvature flow called conformalized mean curvature flow.

In Section 6.2.2 I discuss the benefit from using the interior geometry of objects. To this end, I show the advantage of the improved interpolation method in Chapter 3 that utilizes the skeletal coordinates. Due to this improved interpolation method, a discrete s-rep can produce smooth interior geometry with good correspondences, yielding more discriminative between-object features than the previous interpolation method (Vicory, 2018).

Section 6.2.3 discusses potential research ideas that were alluded to in Chapters 4 and 5. These ideas focus on capturing more comprehensive geometric and statistical relations.

Last, Section 6.2.4 discusses other possible applications of my methods, varying from morphometric analysis of Alzheimer's Disease development to deep learning with CoWBO features and to robotics.

6.2.1 Fitting S-reps to Objects of More Complex Shape

In Chapter 3 I have developed the methods and software for fitting non-branching s-reps to single objects. While the methods can produce good fit to many structures in the human brain, there is still a need to fit s-reps to objects of more complex shape for statistical analysis. However, it is difficult to apply the methods in Chapter 3 to objects of complex shape such as the mandible, the cortex, etc.. For example, it is difficult to fit s-reps to the mandible because the mean curvature flow can develop singularities as illustrated in fig. 3.11. A promising solution to this problem is to replace the mean curvature flow with conformalized mean curvature flow (Kazhdan et al., 2012). Conformalized mean curvature flow is more robust in computing discrete curvature on a surface mesh. That is, this algorithm can

avoid numerical instabilities that are encountered in the previous mean curvature flow. Fig. 6.2 illustrates the process of deforming a mandible to a near-ellipsoid via conformalized mean curvature flow.




Figure 6.2: Deforming a mandible shape using conformalized mean curvature flow.

Despite the promise of conformalized mean curvature flow, there still exist many problems in producing an initial fit to a mandible. First and foremost, it is difficult to deform an s-rep back to fit the mandible via the stratified diffeomorphism on the boundary and the skeleton because the boundary deforms much faster than the skeleton. To address this problem, my colleague N. Tapp-Hughes proposed a special manipulation on the spine so that the skeleton is deformed closely along with the boundary. This idea is still being worked on but has led to some promising results. Second, the surface points on the near-ellipsoid (see the rightmost figure) cluster near its two ends, which correspond to the sharp regions around the two condyles. Such an uneven distribution of boundary points makes it difficult to fit a standard ellipsoid. To solve this problem, I tried to re-triangulate the surface of the near-ellipsoid (Valette and Chassery, 2004), allowing me to fit an s-rep to the near-ellipsoid.

Moreover, conformalized mean curvature flow has shown promise in deforming complex shape such as the colon into ellipsoids in Wang's work. In the future, it is of interest to explore the fitting to other complex shapes such as the brain cortex.

6.2.2 Benefit from Interior Geometric and Appearance Features of Objects

As I have noted in previous chapters, the interior geometric features are useful in anatomic shape analysis. Additionally, medical images' appearance features (e.g., intensities) at interior positions of objects can also be informative. The fitted s-reps provide geometric features and coordinates in correspondence for extracting these appearance and geometric features from an anatomic object. In this section, I discuss potential use of s-reps in extracting interior features as follows.

An s-rep fitted to the object can produce level surfaces of radial flow (see fig. 6.1). The interpolation method in Chapter 3 allows extracting geometric features on the level surface in good correspondence. Moreover, the interpolation using skeletal coordinates produces a smoother spokes' directions, as compared to the previous interpolation method (Vicory, 2018).

Also, the construction of linking structures benefits from the interpolation using skeletal coordinates. Because it is necessary to interpolate s-reps when finding regularly linked spokes (as detailed in Chapter 5), the modified linking

structures benefit from this improvement of the interpolation. Therefore, the classification performance using modified link features is boosted, as shown in Table 6.1. The details of both “Between-object linking” features and “CoWBO” features in the table can be found in Chapter 5.

Features	Vicory’s method (Vicory, 2018)	New interpolation
Between-object linking	0.633/0.008	0.679/0.009
CoWBO	0.652/0.009	0.668/0.009

Table 6.1: Test AUCs/SEMs using different interpolation methods.

Another example that takes advantage of interior of objects is the local affine frames proposed in Chapter 5. While I noticed that a local affine frame transforms to another frame via a linear transformation, it is of interest to explore such transformations within an object. In the future, I expect more research on capturing the interior features of objects by s-reps.

6.2.3 Geometric and Statistical Correlation for Multi-object Shape Analysis

Chapters 4 and 5 describe geometric and statistical methods that capture correlation of multi-object complexes. I realize that these methods can be extended in many directions, as described below.

First, the experiments in both Chapters 4 and 5 focus on joint components from NEUJIVE. I made this decision based on the intuition that the brain disorder ASD can simultaneously cause shape changes of multiple brain structures. However, it is still useful to investigate the individual variation of each brain structure in other applications. I leave this investigation in future research.

Second, NEUJIVE can analyze not only multi-object shape data, but also multi-block non-Euclidean data. Though I tested the effectiveness of NEUJIVE with multi-object complexes, it should be also useful in analyzing directional data (e.g., wind directions across the globe). Moreover, because the basic principle of designing NEUJIVE is to extract statistical correlation in an HDLSS context, these methods ought to shed light on correlation in other scenarios, for example, in mini-batches in training deep neural networks.

Third, the modified linking structure developed in Chapter 5 can be generalized to complexes that contain more than two objects. There, the linking flow of an object should be determined by the relation of this object and all its nearby objects. Therefore, the modified linking structure of that object should capture its relations with other nearby objects.

Fourth, the modified linking structure can also extract geometric relations between abutting objects. In the regions where two objects touch with each other, the spoke extensions are constantly 0, while link directions can provide the relations of local geometry along the shared boundary.

Fifth, the modified linking structure can also be generalized to analyze geometric relations between different parts of an object. Take a U-shaped object as an example, the object can be divided into two bar-shaped parts. While the within-object shape features can only extract the geometric features within each part, the links between the two parts of

the object can reveal additional information than the analysis of within-object shape. For example, it can be useful to obtain the variation of relative positions between the two parts.

6.2.4 Other Applications

This dissertation conducted the majority of experiments on the data from the IBIS network, described in Section 2.2. I expect more research projects and applications using the methods in this dissertation.

Recently, our collaborators in Harbin Institute of Technology have shown a research project on the hippocampal atrophy in the development of Alzheimer’s Disease (AD) (Gao et al., 2022). The researchers proposed local morphometric features on the basis of the s-reps fitted to the hippocampus. They use these features to evaluate the development of AD from progressive Mild Cognitive Impairment (pMCI) and stable Mild Cognitive Impairment (sMCI). The researchers conducted experiments on the public Alzheimer’s Disease Neuroimaging Initiative (ADNI) database. The results are consistent with histological findings about the hippocampal atrophy during the progression of AD. Such analysis of structural atrophy benefits from s-reps because s-reps can provide rich morphological features, such as objects’ widths, in good correspondence.

The correlated shape features extracted from multi-object complexes using the proposed methods can also be applied in deep geometric learning. Despite the promising results from recent research on geometric learning of single objects via deep neural networks, there is still a need to better understand geometric relations between objects. Especially in anatomic shape understanding, the underlying mapping from segmented shape data to the knowledge space (e.g., class labels) is so complex that a deep neural network needs a large amount of training data to learn that mapping. In this regard, my dissertation gives an explicit modeling of multi-object shape and their statistical correlation. These features can be used in training deep neural networks to alleviate the demands of training data.

In addition to the above statistical analysis, the modified linking structures can also be used in motion planning. In robotics, for example, the paths of robots should be deliberately planned so that the movement of robots will not collide with the environment or critical objects (see e.g., (He et al., 2020)). To this end, it is promising to model the objects with s-reps and the between-object space with the modified linking structures. By doing so, the path planning for robots can be efficient and robust.

BIBLIOGRAPHY

- Abid, A., Zhang, M. J., Bagaria, V. K., and Zou, J. (2018). Exploring patterns enriched in a dataset with contrastive principal component analysis. *Nature communications*, 9(1):1–7.
- Akhoundi-Asl, A. and Soltanian-Zadeh, H. (2007). Nonparametric entropy-based coupled multi-shape medical image segmentation. In *2007 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 1200–1203. IEEE.
- Arguillère, S., Trélat, E., Trouvé, A., and Younes, L. (2015). Shape deformation analysis from the optimal control viewpoint. *Journal de Mathématiques Pures et Appliquées*, 104(1):139 – 178.
- Avants, B. B., Epstein, C. L., Grossman, M., and Gee, J. C. (2008). Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical image analysis*, 12(1):26–41.
- Baird, G., Simonoff, E., Pickles, A., Chandler, S., Loucas, T., Meldrum, D., and Charman, T. (2006). Prevalence of disorders of the autism spectrum in a population cohort of children in south thames: the special needs and autism project (snap). *The lancet*, 368(9531):210–215.
- Blum, H. and Nagel, R. N. (1978). Shape description using weighted symmetric axis features. *Pattern recognition*, 10(3):167–180.
- Börck, a. and Golub, G. H. (1973). Numerical methods for computing angles between linear subspaces. *Mathematics of computation*, 27(123):579–594.
- Bossa, M. N. and Olmos, S. (2007). Multi-object statistical pose+ shape models. pages 1204–1207.
- Cates, J., Fletcher, P. T., Styner, M., Shenton, M., and Whitaker, R. (2007). Shape modeling and analysis with entropy-based particle systems. In Karssemeijer, N. and Lelieveldt, B., editors, *Information Processing in Medical Imaging*, pages 333–345, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Cates, J., Meyer, M., Fletcher, T., and Whitaker, R. (2006). Entropy-based particle systems for shape correspondence. In *1st MICCAI workshop on mathematical foundations of computational anatomy: geometrical, statistical and registration methods for modeling biological shape variability*, pages 90–99.
- CDC (2016). Prevalence and characteristics of autism spectrum disorder among children aged 8 years—autism and developmental disabilities monitoring network, 11 sites, united states, 2012. *Morbidity and mortality weekly report. Surveillance summaries (Washington, DC: 2002)*, 65(3):1–23.
- Corkin, S. (2002). What’s new with the amnesic patient hm? *Nature reviews neuroscience*, 3(2):153–160.
- Cortes, C. and Vapnik, V. (1995). Support-vector networks. *Machine learning*, 20(3):273–297.
- Damon, J. (2003). Smoothness and geometry of boundaries associated to skeletal structures I: Sufficient conditions for smoothness. In *Annales de l’institut Fourier*, volume 53, pages 1941–1985.
- Damon, J. (2004). Smoothness and geometry of boundaries associated to skeletal structures, II: Geometry in the blum case. *Compositio Mathematica*, 140(6):1657–1674.
- Damon, J. (2008). *Medial Representations: Mathematics, Algorithms and Applications*, chapter Geometry and medial structure, pages 69–123. Springer.
- Damon, J. (2019). Extending smooth and discrete medial/ skeletal structures to linking structures. *Preliminary preprint*.
- Damon, J. (2021). Rigidity properties of the blum medial axis. *Journal of Mathematical Imaging and Vision*, 63(1):120–129.
- Damon, J. and Gasparovic, E. (2017). Modeling multi-object configurations via medial/skeletal linking structures. *International Journal of Computer Vision*, 124(3):255–272.

- Damon, J. and Marron, J. (2014). Backwards principal component analysis and principal nested relations. *Journal of Mathematical Imaging and Vision*, 50(1):107–114.
- Deserno, M., Borsboom, D., Begeer, S., and Geurts, H. (2018). Relating ASD symptoms to well-being: Moving across different construct levels. *Psychological Medicine*, 48(7):1179–1189.
- Dryden, I. L., Kim, K.-R., Laughton, C. A., and Le, H. (2019). Principal nested shape space analysis of molecular dynamics data. *The Annals of Applied Statistics*, 13(4):2213–2234.
- Dryden, I. L. and Mardia, K. V. (1998). *Statistical Shape Analysis*. Wiley, Chichester.
- Dryden, I. L. and Mardia, K. V. (2016). *Statistical shape analysis: with applications in R*, volume 995. John Wiley & Sons.
- Durrleman, S., Prastawa, M., Charon, N., Korenberg, J. R., Joshi, S., Gerig, G., and Trouvé, A. (2014). Morphometry of anatomical shape complexes with dense deformations and sparse parameters. *NeuroImage*, 101:35–49.
- Eilam-Stock, T., Wu, T., Spagna, A., Egan, L. J., and Fan, J. (2016). Neuroanatomical alterations in high-functioning adults with autism spectrum disorder. *Frontiers in neuroscience*, 10:237.
- el Bouhaddani, Said, Uh, Won, H., Jongbloed, Geurt, Hayward, Caroline, Klari’c, Lucija, Kielbasa, M., S., Houwing-Duistermaat, and Jeanine (2018). Integrating omics datasets with the omicspls package. *BMC Bioinformatics*, 19(1):1–9.
- Eltzner, B., Jung, S., and Huckemann, S. (2015). Dimension reduction on polyspheres with application to skeletal representations. pages 22–29.
- Feng, Q., Jiang, M., Hannig, J., and Marron, J. S. (2018). Angle-based joint and individual variation explained. *Journal of Multivariate Analysis*, 166:241–265.
- Ferguson, J. N., Young, L. J., Hearn, E. F., Matzuk, M. M., Insel, T. R., and Winslow, J. T. (2000). Social amnesia in mice lacking the oxytocin gene. *Nature genetics*, 25(3):284–288.
- Fletcher, P. T., Lu, C., Pizer, S. M., and Joshi, S. (2004a). Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE transactions on medical imaging*, 23(8):995–1005.
- Fletcher, P. T., Pizer, S. M., and Joshi, S. (2004b). *Statistical variability in nonlinear spaces: Application to shape analysis and DT-MRI*. Citeseer.
- Gao, N., Liu, Z., Deng, Y., and Ma, T. (2022). Intensive atrophy of the anterior hippocampal subiculum and superior hippocampal tail are potential indicators for pmci-to-ad conversion. *In submission*.
- Gasparovic, E. (2012). *The Blum medial linking structure for multi-region analysis*. PhD thesis, The University of North Carolina at Chapel Hill.
- Golland, P., Grimson, W. E. L., Shenton, M. E., and Kikinis, R. (2000). Small sample size learning for shape analysis of anatomical structures. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 72–82. Springer.
- Golland, P., Grimson, W. E. L., Shenton, M. E., and Kikinis, R. (2005). Detection and analysis of statistical differences in anatomical shape. *Medical image analysis*, 9(1):69–86.
- Gollmer, S. T., Simon, M., Bischof, A., Barkhausen, J., and Buzug, T. M. (2012). Multi-object active shape model construction for abdomen segmentation: preliminary results. In *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 3990–3993. IEEE.
- Gorcowski, K., Styner, M., Jeong, J., Marron, J. S., Piven, J., Hazlett, H. C., Pizer, S. M., and Gerig, G. (2007). Statistical shape analysis of multi-object complexes. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8.

- Gori, P., Colliot, O., Marrakchi-Kacem, L., Worbe, Y., Poupon, C., Hartmann, A., Ayache, N., and Durrleman, S. (2017). A bayesian framework for joint morphometry of surface and curve meshes in multi-object complexes. *Medical image analysis*, 35:458–474.
- Groen, W., Teluij, M., Buitelaar, J., and Tendolkar, I. (2010). Amygdala and hippocampus enlargement during adolescence in autism. *Journal of the American Academy of Child & Adolescent Psychiatry*, 49(6):552–560.
- Guo, Y., Sohel, F., Bennamoun, M., Lu, M., and Wan, J. (2013). Rotational projection statistics for 3d local surface description and object recognition. *International journal of computer vision*, 105(1):63–86.
- Gutiérrez-Becker, B., Sarasua, I., and Wachinger, C. (2021). Discriminative and generative models for anatomical shape analysis on point clouds with deep neural networks. *Medical Image Analysis*, 67:101852.
- Han, Q. (2008). *Proper Shape Representation of Single Figure and Multi-Figure Anatomical Objects*. PhD thesis. Computer Science dissertation, Univ. of North Carolina at Chapel Hill.
- Han, Q., Pizer, S. M., and Damon, J. N. (2006). Interpolation in discrete single figure medial objects. In *2006 Conference on Computer Vision and Pattern Recognition Workshop (CVPRW'06)*, pages 85–85.
- Hashemi, R. H., Bradley, W. G., and Lisanti, C. J. (2012). *MRI: the basics: The Basics*. Lippincott Williams & Wilkins.
- Hazlett, H. C., Gu, H., Munsell, B. C., Kim, S. H., Styner, M., Wolff, J. J., Elison, J. T., Swanson, M. R., Zhu, H., Botteron, K. N., et al. (2017). Early brain development in infants at high risk for autism spectrum disorder. *Nature*, 542(7641):348–351.
- He, L., Pan, Z., Jia, B., and Manocha, D. (2020). Efficient multi-agent motion planning in continuous workspaces using medial-axis-based swap graphs. *arXiv preprint arXiv:2002.11892*.
- Hitti, F. L. and Siegelbaum, S. A. (2014). The hippocampal ca2 region is essential for social memory. *Nature*, 508(7494):88–92.
- Hollander, E., Anagnostou, E., Chaplin, W., Esposito, K., Haznedar, M. M., Licalzi, E., Wasserman, S., Soorya, L., and Buchsbaum, M. (2005). Striatal volume on magnetic resonance imaging and repetitive behaviors in autism. *Biological psychiatry*, 58(3):226–232.
- Hong, J. (2018). *Classification of Neuroanatomical Structures Based on Non-Euclidean Geometric Object Properties*. PhD thesis. Computer Science dissertation, Univ. of North Carolina at Chapel Hill.
- Hong, J., Vicory, J., Schulz, J., Styner, M., Marron, J., and Pizer, S. (2016). Non-euclidean classification of medically imaged objects via s-reps. *Medical Image Analysis*, 31:37–45.
- Hong, S. (2020). *Regression Analysis of Continuous Medial Representations on a Riemannian Manifold for the Analysis of Anatomical Shape Change*. PhD thesis.
- Horn, R. A. and Johnson, C. R. (2012). *Matrix analysis*. Cambridge university press.
- Hotelling, H. (1936). Relations between two sets of variates. *Biometrika*, 28(3-4):321–377.
- Huckemann, S., Hotz, T., and Munk, A. (2010). Intrinsic shape analysis: Geodesic pca for riemannian manifolds modulo isometric lie group actions. *Statistica Sinica*, pages 1–58.
- Huisken, G. (1990). Asymptotic behavior for singularities of the mean curvature flow. *J. Differential Geom.*, 31(1):285–299.
- Jeong, J.-Y., Pizer, S. M., and Ray, S. (2006). Statistics on anatomic objects reflecting inter-object relations. In *1st MICCAI Workshop on Mathematical Foundations of Computational Anatomy: Geometrical, Statistical and Registration Methods for Modeling Biological Shape Variability*, pages 136–145.
- Jiang, M. (2018). *Statistical Learning of Integrative Analysis*. PhD thesis, The University of North Carolina at Chapel Hill.

- JL, L., N., L., and JA, O. (2012). Functional neuroanatomy of the basal ganglia. *Cold Spring Harb Perspect Med*.
- Joshi, S., Pizer, S., Fletcher, P. T., Yushkevich, P., Thall, A., and Marron, J. (2002). Multiscale deformable model segmentation and statistical shape analysis using medial descriptions. *IEEE transactions on medical imaging*, 21(5):538–550.
- Jung, S., Dryden, I. L., and Marron, J. S. (2012). Analysis of principal nested spheres. *Biometrika*, 99(3):551–568.
- Katuwal, G. J., Cahill, N. D., Baum, S. A., and Michael, A. M. (2015). The predictive power of structural MRI in autism diagnosis. In *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 4270–4273.
- Kazhdan, M., Solomon, J., and Ben-Chen, M. (2012). Can mean-curvature flow be modified to be non-singular? In *Computer Graphics Forum*, volume 31, pages 1745–1754. Wiley Online Library.
- Kendall, D. G. (1984). Shape manifolds, procrustean metrics, and complex projective spaces. *Bulletin of the London Mathematical Society*.
- Kendall, D. G. (1989). A survey of the statistical theory of shape. *Statistical Science*, pages 87–99.
- Kim, K.-R., Dryden, I. L., Le, H., and Severn, K. E. (2021). Smoothing splines on riemannian manifolds, with applications to 3d shape space. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 83(1):108–132.
- Kim, S. H., Lyu, I., Fonov, V. S., Vachet, C., Hazlett, H. C., Smith, R. G., Piven, J., Dager, S. R., Mckinstry, R. C., Pruett Jr, J. R., et al. (2016). Development of cortical shape in the human brain from 6 to 24 months of age via a novel measure of shape complexity. *NeuroImage*, 135:163–176.
- Koenderink, J. J. (1990). *Solid Shape*, pages 319–323. MIT Press, Cambridge, MA, USA.
- Krishna, A. et al. (2022). Incorporating the geometric relationship of adjacent objects in multi-object shape analysis. *In submission*.
- Le, H., Kendall, D. G., et al. (1993). The riemannian structure of euclidean shape spaces: a novel environment for statistics. *Annals of statistics*, 21(3):1225–1271.
- Lei, H., Jiang, G., and Quan, L. (2017). Fast descriptors and correspondence propagation for robust global point cloud registration. *IEEE Transactions on Image Processing*, 26(8):3614–3623.
- Lindström, A., Pettersson, F., Almqvist, F., Berglund, A., Kihlberg, J., and Linusson, A. (2006). Hierarchical pls modeling for predicting the binding of a comprehensive set of structurally diverse protein- ligand complexes. *Journal of chemical information and modeling*, 46(3):1154–1167.
- Liu, Z., Hong, J., Vicory, J., Damon, J. N., and Pizer, S. M. (2021). Fitting unbranching skeletal structures to objects. *Medical Image Analysis*, 70:102020.
- Lock, E. F., Hoadley, K. A., Marron, J. S., and Nobel, A. B. (2013). Joint and individual variation explained (JIVE) for integrated analysis of multiple data types. *The annals of applied statistics*, 7(1):523.
- Lord, C., Risi, S., Lambrecht, L., Cook, E. H., Leventhal, B. L., DiLavore, P. C., Pickles, A., and Rutter, M. (2000). The autism diagnostic observation schedule—generic: A standard measure of social and communication deficits associated with the spectrum of autism. *Journal of autism and developmental disorders*, 30(3):205–223.
- Lu, C., Pizer, S. M., Joshi, S., and Jeong, J.-Y. (2007). Statistical multi-object shape models. *International Journal of Computer Vision*, 75(3):387–404.
- Malassiotis, S. and Srinivasan, M. G. (2007). Snapshots: A novel local surface descriptor and matching algorithm for robust 3d surface alignment. *IEEE Transactions on pattern analysis and machine intelligence*, 29(7):1285–1290.
- Marron, J. S., Todd, M. J., and Ahn, J. (2007). Distance weighted discrimination. *Journal of the American Statistical Association*, 102(480):1267–1271.

- Miao, J. and Ben-Israel, A. (1992). On principal angles between subspace in R^n . *Linear Algebra and Its Applications*, 171:81–98.
- Miolane, N., Caorsi, M., Lupo, U., Guerard, M., Guigui, N., Mathe, J., Cabanes, Y., Reise, W., Davies, T., Leitão, A., et al. (2021). Iclr 2021 challenge for computational geometry & topology: Design and results. *arXiv preprint arXiv:2108.09810*.
- Mostapha, M., Vicory, J., Styner, M., , and Pizer, S. (2017). A segmentation editing framework based on shape change statistics. In *SPIE Medical Imaging*, volume 10133, page 101331E.
- Murphy, C. M., Deeley, Q., Daly, E., Ecker, C., O’Brien, F., Hallahan, B., Loth, E., Toal, F., Reed, S., Hales, S., et al. (2012). Anatomy and aging of the amygdala and hippocampus in autism spectrum disorder: an in vivo magnetic resonance imaging study of asperger syndrome. *Autism Research*, 5(1):3–12.
- Nain, D., Styner, M., Niethammer, M., Levitt, J. J., Shenton, M. E., Gerig, G., Bobick, A., and Tannenbaum, A. (2007). Statistical shape analysis of brain structures using spherical wavelets. In *2007 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 209–212. IEEE.
- Nicolson, R., DeVito, T. J., Vidal, C. N., Sui, Y., Hayashi, K. M., Drost, D. J., Williamson, P. C., Rajakumar, N., Toga, A. W., and Thompson, P. M. (2006). Detection and mapping of hippocampal abnormalities in autism. *Psychiatry Research: Neuroimaging*, 148(1):11–21.
- O’Neill, B. (2006). *Elementary differential geometry*. Elsevier.
- Patrangenaru, V. and Ellingson, L. (2019). *Nonparametric statistics on manifolds and their applications to object data analysis*. CRC Press.
- Petrelli, A. and Di Stefano, L. (2011). On the repeatability of the local reference frame for partial shape matching. In *2011 International Conference on Computer Vision*, pages 2244–2251. IEEE.
- Pezeshki, A., Scharf, L. L., Azimi-Sadjadi, M. R., and Lundberg, M. (2004). Empirical canonical correlation analysis in subspaces. In *Conference Record of the Thirty-Eighth Asilomar Conference on Signals, Systems and Computers, 2004.*, volume 1, pages 994–997. IEEE.
- Pizer, S. and Marron, J. (2017). *Object Statistics on Curved Manifolds*, pages 137–164.
- Pizer, S. M., Fletcher, P. T., Joshi, S., Thall, A., Chen, J. Z., Fridman, Y., Fritsch, D. S., Gash, A. G., Glotzer, J. M., Jiroutek, M. R., et al. (2003). Deformable m-reps for 3d medical image segmentation. *International journal of computer vision*, 55(2):85–106.
- Pizer, S. M., Hong, J., Vicory, J., Liu, Z., Marron, J. S., et al. (2019). Object shape representation via skeletal models (s-reps) and statistical analysis. *Riemannian Geometric Statistics in Medical Image Analysis*, pages 233–272.
- Pizer, S. M., Jung, S., Goswami, D., Vicory, J., Zhao, X., Chaudhuri, R., Damon, J. N., Huckemann, S., and Marron, J. (2013a). *Innovations for Shape Analysis: Models and Algorithms*, chapter Nested Sphere Statistics of Skeletal Models, pages 93–113.
- Pizer, S. M., Jung, S., Goswami, D., Vicory, J., Zhao, X., Chaudhuri, R., Damon, J. N., Huckemann, S., and Marron, J. (2013b). Nested sphere statistics of skeletal models. In *Innovations for Shape Analysis*, pages 93–115. Springer.
- Pizer, S. M., Krishna, A., Liu, Z., and Taheri, M. (2021). Fitted frames to object interiors. *Journal of Mathematical Imaging and Vision*, Under review.
- Pizer, S. M., Krishna, A., Liu, Z., and Taheri, M. (2022). Fitted frames to object interiors. *Submitted to Journal of Mathematical Imaging and Vision*.
- Pouch, A. M., Yushkevich, P. A., Aly, A. H., Woltersom, A. H., Okon, E., Aly, A. H., Yushkevich, N., Parameshwaran, S., Wang, J., Oguz, B., et al. (2020). Automated meshing of anatomical shapes for deformable medial modeling: Application to the placenta in 3d ultrasound. In *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pages 1842–1846. IEEE.

- Powell, M. J. D. (2004). The newuoa software for unconstrained optimization without derivatives.
- Qiu, A., Adler, M., Crocetti, D., Miller, M. I., and Mostofsky, S. H. (2010a). Basal ganglia shapes predict social, communication, and motor dysfunctions in boys with autism spectrum disorder. *Journal of the American Academy of Child & Adolescent Psychiatry*, 49(6):539–551.
- Qiu, A., Brown, T., Fischl, B., Ma, J., and Miller, M. I. (2010b). Atlas generation for subcortical and ventricular structures with its applications in shape analysis. *IEEE Transactions on Image Processing*, 19(6):1539–1547.
- Reuter, M., Wolter, F.-E., and Peinecke, N. (2006). Laplace–beltrami spectra as ‘shape-dna’ of surfaces and solids. *Computer-Aided Design*, 38(4):342–366.
- Richards, R., Greimel, E., Kliemann, D., Koerte, I. K., Schulte-Körne, G., Reuter, M., and Wachinger, C. (2020). Increased hippocampal shape asymmetry and volumetric ventricular asymmetry in autism spectrum disorder. *NeuroImage: Clinical*, 26:102207.
- Sagonas, C., Panagakis, Y., Leidinger, A., and Zafeiriou, S. (2017). Robust joint and individual variance explained. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5267–5276.
- Schoenberg, I. J. (1973). *Cardinal spline interpolation*. SIAM.
- Schulz, J., Pizer, S., Marron, J., and Godtliebsen, F. (2016). Nonlinear hypothesis testing of geometric object properties of shapes applied to hippocampi. *Journal of Mathematical Imaging and Vision*, 54:15–34.
- Schwarz, T., Heimann, T., Lossnitzer, D., Mohrhardt, C., Steen, H., Rietdorf, U., Wolf, I., and Meinzer, H.-P. (2010). Multiobject segmentation using coupled shape space models. In *Medical Imaging 2010: Image Processing*, volume 7623, page 76233V. International Society for Optics and Photonics.
- Sears, L. L., Vest, C., Mohamed, S., Bailey, J., Ranson, B. J., and Piven, J. (1999). An mri study of the basal ganglia in autism. *Progress in neuro-psychopharmacology & biological psychiatry*.
- Shen, C., Sun, M., Tang, M., and Priebe, C. E. (2014). Generalized canonical correlation analysis for classification. *Journal of Multivariate Analysis*, 130:310 – 322.
- Siddiqi, K. and Pizer, S. (2008). *Medial Representations: Mathematics, Algorithms and Applications*. Springer Publishing Company, Incorporated, 1st edition.
- Srivastava, A. and Klassen, E. P. (2016). *Motivation for Function and Shape Analysis*, pages 1–19. Springer New York, New York, NY.
- Styner, M., Oguz, I., Xu, S., Brechbühler, C., Pantazis, D., Levitt, J., Shenton, M., and Gerig, G. (2006). Statistical shape analysis of brain structures using SPHARM-PDM. *The insight journal*, 1071:242–250.
- Taubin, G. (1995). Curve and surface smoothing without shrinkage. In *Proceedings of IEEE International Conference on Computer Vision*, pages 852–857.
- Toga, A. W., Thompson, P. M., Mori, S., Amunts, K., and Zilles, K. (2006). Towards multimodal atlases of the human brain. *Nature Reviews Neuroscience*, 7(12):952–966.
- Trygg, J. and Wold, S. (2003). O2-PLS, a two-block (X-Y) latent variable regression (LVR) method with an integral OSC filter. *Journal of Chemometrics*, 17:53 – 64.
- Tu, L., Styner, M., Vicory, J., et al. (2018). Skeletal shape correspondence through entropy. *IEEE Transactions on Medical Imaging*, 37(1):1–11.
- Tu, L., Vicory, J., Elhabian, S., Paniagua, B., Prieto, J. C., Damon, J. N., Whitaker, R., Styner, M., , and Pizer, S. M. (2016). Entropy-based correspondence improvement of interpolated skeletal models. *Computer Vision and Image Understanding*.
- Valette, S. and Chassery, J.-M. (2004). Approximated centroidal voronoi diagrams for uniform polygonal mesh coarsening. In *Computer Graphics Forum*, volume 23, pages 381–389. Wiley Online Library.

- van der Kloet, F. M., Sebastián-León, P., Conesa, A., Smilde, A. K., and Westerhuis, J. A. (2016). Separating common from distinctive variation. *BMC bioinformatics*, 17(5):271–286.
- Van Deun, K., Van Mechelen, I., Thorrez, L., Schouteden, M., De Moor, B., Van Der Werf, M. J., De Lathauwer, L., Smilde, A. K., and Kiers, H. A. (2012). Disco-sca and properly applied gsvd as swinging methods to find common and distinctive processes. *PLoS one*, 7(5):e37840.
- Van Kaick, O., Zhang, H., Hamarneh, G., and Cohen-Or, D. (2011). A survey on shape correspondence. In *Computer graphics forum*, volume 30, pages 1681–1707. Wiley Online Library.
- Van Rooij, D., Anagnostou, E., C. A., Auzias, G., Behrmann, M., GF, B., Calderoni, S., Daly, E., C. D., A. D., Dinstein, I., FLS, D., Durston, S., Ecker, C., Fair, D., Fedor, J., Fitzgerald, J., Freitag, C., Gallagher, L., and Buitelaar, J. (2017). Cortical and subcortical brain morphometry differences between patients with autism spectrum disorder and healthy individuals across the lifespan: Results from the enigma asd working group. *American Journal of Psychiatry*, 175.
- Vicory, J. (2018). Shape deformation statistics and regional texture-based appearance models for segmentation. *Dissertation in Department of Computer Science, University of North Carolina at Chapel Hill*.
- Vicory, J., Pascal, L., Hernandez, P., Fishbaugh, J., Prieto, J., Mostapha, M., Huang, C., Shah, H., Hong, J., Liu, Z., et al. (2018). SlicerSALT: Shape analysis toolbox. In *International Workshop on Shape in Medical Imaging*, pages 65–72. Springer.
- Wei, S., Lee, C., Wichers, L., and Marron, J. (2016). Direction-projection-permutation for high-dimensional hypothesis tests. *Journal of Computational and Graphical Statistics*, 25(2):549–569.
- Westerhuis, J. A., Kourti, T., and MacGregor, J. F. (1998). Analysis of multiblock and hierarchical PCA and PLS models. *Journal of Chemometrics*, 12(5):301–321.
- Wold, H. (2004). *Partial Least Squares*.
- Wold, S., Geladi, P., Esbensen, K., and Öhman, J. (2005). Multi-way principal components and PLS analysis. *Journal of Chemometrics*, 1:41 – 56.
- Wold, S., Kettaneh, N., and Tjessem, K. (1996). Hierarchical multiblock PLS and PC models for easier model interpretation and as an alternative to variable selection. *Journal of Chemometrics*, 10.
- Wolff, J., Cody, H., Lightbody, A., Reiss, A., and Piven, J. (2013). Repetitive and self-injurious behaviors: Associations with caudate volume in autism and Fragile X syndrome. *Journal of neurodevelopmental disorders*, 5:12.
- Wolff, J. J., Gu, H., Gerig, G., Elison, J. T., Styner, M., Gouttard, S., Botteron, K. N., Dager, S. R., Dawson, G., Estes, A. M., et al. (2012). Differences in white matter fiber tract development present from 6 to 24 months in infants with autism. *American journal of Psychiatry*, 169(6):589–600.
- Yang, J., Cao, Z., and Zhang, Q. (2016). A fast and robust local descriptor for 3d point cloud registration. *Information Sciences*, 346:163–179.
- Yushkevich, P. A., Aly, A. H., Wang, J., Xie, L., Gorman, R. C., Younes, L., and Pouch, A. M. (2019). Diffeomorphic medial modeling. In *Information Processing in Medical Imaging - 26th International Conference, IPMI 2019, Hong Kong, China, June 2-7, 2019, Proceedings*, pages 208–220.