

**CLASSIFICATION OF NEUROANATOMICAL STRUCTURES
BASED ON
NON-EUCLIDEAN GEOMETRIC OBJECT PROPERTIES**

Junpyo Hong

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

Chapel Hill
2018

Approved by:
Stephen M. Pizer
J.S. Marron
Martin A. Styner
Beatriz Paniagua
Heather C. Hazlett

©2018
Junpyo Hong
ALL RIGHTS RESERVED

ABSTRACT

JUNPYO HONG: Classification of Neuroanatomical Structures based on
Non-Euclidean Geometric Object Properties
(Under the direction of Stephen M. Pizer)

Studying the observed morphological differences in neuroanatomical structures between individuals with neurodevelopmental disorders and a control group of typically developing individuals has been an important objective. Researchers study the differences with two goals: to assist an accurate diagnosis of the disease and to gain insights into underlying mechanisms of the disease that cause such changes.

Shape classification is commonly utilized in such a study. An effective classification is difficult because it requires 1) a choice of an object model that can provide rich geometric object properties (GOPs) relevant for a given classification task, and 2) a choice of a statistical classification method that accounts for the non-Euclidean nature of GOPs.

I lay out my methodological contributions to address the aforementioned challenges in the context of early diagnosis and detection of Autism Spectrum Disorder (ASD) in infants based on shapes of hippocampi and caudate nuclei; morphological deviations in these structures between individuals with ASD and typically developing individuals have been reported in the literature. These contributions respectively lead to 1) an effective modeling of shapes of objects of interest and 2) an effective classification.

As the first contribution for modeling shapes, I propose a method to obtain a set of skeletal models called *s-reps* from a set of 3D objects. First, the method iteratively deforms the object surface via Mean Curvature Flow (MCF) until the deformed surface is approximately ellipsoidal. Then, an *s-rep* of the approximate ellipsoid is obtained analytically. Finally, the ellipsoid *s-rep* is deformed via a series of inverse MCF transformations. The method has two important properties: 1) it is fully automatic, and 2) it yields a set of *s-reps* with good correspondence across the set. The method is shown effective in generating a set of *s-reps* for a few neuroanatomical structures.

As the second contribution with respect to shape modeling, I introduce an extension to the current s-rep for representing an object with a narrowing sharp tail. This includes a spoke interpolation method for interpolating a discrete s-rep of an object with a narrowing sharp tail into a continuous object. This extension is necessary for representing surface geometry of objects whose boundary has a singular point. I demonstrate that this extension allows appropriate surface modeling of a narrowing sharp tail region of the caudate nucleus. In addition, I show that the extension is beneficial in classifying autistic and non-autistic infants at high risk of ASD based on shapes of caudate nuclei.

As the first contribution with respect to statistical methods, I propose a novel shape classification framework that uses the s-rep to capture rich localized geometric descriptions of the object, a statistical method called Principal Nested Spheres (PNS) analysis to handle the non-Euclidean s-rep GOPs, and a classification method called Distance Weighted Discrimination (DWD). I evaluate the effectiveness of the proposed method in classifying autistic and non-autistic infants based on either hippocampal shapes or caudate shapes in terms of the Area Under the ROC curve (AUC). In addition, I show that the proposed method is superior to commonly used shape classification methods in the literature.

As my final methodological contribution, I extend the proposed shape classification method to perform the classification task based on temporal shape differences. DWD learns a class separation direction based on the temporal shape differences that are obtained by taking differences of the temporal pair of Euclideanized s-reps. In the context of early diagnosis and detection of ASD in young infants, the proposed temporal shape difference classification produces some interesting results; for both the hippocampus and the caudate nucleus, the temporal shape difference does not seem to be as predictive as the shape alone.

To my family and friends. . .

ACKNOWLEDGEMENTS

I would like to express my sincere gratitude to everyone who has made my time at UNC memorable. I would like to thank my dissertation advisor Prof. Stephen Pizer for his mentorship, guidance, and wisdom without which this work would not have been possible. I would like to thank other members of my dissertation advisory committee: Prof. J.S. Marron, Martin Styner, Heather Hazlett, and Dr. Beatriz Paniagua for providing their expertise, discussion, and invaluable input to my dissertation research.

I thank Prof. Sungkyu Jung for providing help in using Principal Nested Spheres. I would like to thank Dr. Jared Vicory for providing his input on spoke interpolation. I would like to thank IBIS network for kindly providing their data.

I would like to thank my colleagues: Juan, Liyun, Mahmoud, Bingjie, Zhiyuan and others. Finally, I thank all of my friends and family for their continued support and love.

TABLE OF CONTENTS

LIST OF TABLES	xi
LIST OF FIGURES	xii
LIST OF ABBREVIATIONS	xiv
1 Introduction	1
1.1 Overview	1
1.2 Current Findings on ASD from Neuroimaging Studies	2
1.3 Common Object Models used in Shape Classification	3
1.4 Skeletal Representations	3
1.5 Automatic Generation of S-reps with Good Correspondence	4
1.6 Extensions of the Discrete S-rep to Model a Singular Point	5
1.7 Common Classification Methods used in Shape Classification	5
1.8 Shape Classification of a Temporal Pair of Objects	6
1.9 Thesis and Contributions	7
2 Background	9
2.1 Object Models	9
2.1.1 The Point Distribution Model	10
2.1.2 The S-Rep	10
2.2 Statistical Methods	11
2.2.1 Principal Component Analysis	12
2.2.2 Principal Nested Spheres Analysis	13
2.2.3 Commensuration of Spherical GOPs	13

2.3	Classification Methods	13
2.3.1	Support Vector Machine	14
2.3.2	Distance Weighted Discrimination.....	14
3	Materials	15
4	Automatic Generation of Case-Specific S-reps	16
4.1	Introduction.....	16
4.2	Background.....	17
4.2.1	Mean Curvature Flow	17
4.2.2	Thin Plate Splines	18
4.3	Method	18
4.3.1	Applying MCF to the Deforming Surface	19
4.3.2	Determining if the Deforming Surface is Ellipsoidal	19
4.3.3	Obtaining the S-rep of the Best-Fitting Ellipsoid	20
4.3.4	Transforming the Ellipsoid S-rep back to the Original Surface	23
4.3.5	Refining Skeletal Geometric Constraints of the Initial Fitted S-rep	24
4.4	Results	25
4.5	Discussion and Conclusion	31
5	Extensions of the S-rep to Model a Singular Point	32
5.1	Introduction.....	32
5.2	Background.....	33
5.2.1	Mathematics of S-reps	33
5.2.2	Mathematics of Spoke Interpolation	34
5.3	Method	35
5.3.1	Singular Point Primitive	36
5.3.2	Proposed Interpolation Method	36
5.3.2.1	Skeleton Interpolation	36
5.3.2.2	Spoke Interpolation	39

5.3.3	Euclideanization of the Augmented S-reps	40
5.4	Results	41
5.5	Conclusion and Discussion	43
6	Non-Euclidean Shape Classification	44
6.1	Introduction.....	44
6.2	Classification Method	45
6.2.1	Euclideanization of s-reps and basis of the transformation between s-rep space and Euclidean space	46
6.2.2	Learning the Separating Direction	46
6.2.3	Computing the Function from Projected Feature Values to the Probability of being Autistic	46
6.2.4	Classification based on Probability Produced by the Mapping Function	49
6.3	Experimental Analysis	49
6.4	Results	50
6.5	Conclusion and Discussion	51
7	Non-Euclidean Temporal Shape Difference Classification	55
7.1	Introduction.....	55
7.2	Method	56
7.2.1	Euclideanization of Temporal Differences of an S-rep	56
7.2.2	Classification	56
7.3	Results	56
7.4	Conclusion and Discussion	58
8	Conclusions & Discussion	59
8.1	Summary of Contributions	59
8.2	Discussion and Future Work	63
8.2.1	S-reps Fitting	63
8.2.1.1	Optimization	63
8.2.2	Different Geometric Flows for Automatic S-reps	64

8.2.3	Modeling Objects with a Singular Point	64
8.2.4	Non-Euclidean Shape Classification	65
8.2.5	Understanding Morphological Changes Associated with Growth	66
8.2.6	Further Clinical Studies	67
8.2.7	Conclusion	67
BIBLIOGRAPHY		68

LIST OF TABLES

6.1	Table of an AUC of the ROC of the selected classification methods and the pure random guessing in classifying the hippocampus of the autistic infants and the non-autistic infants. The result shows that as expected, the Euclideanized s-rep-based classification performs the best.	50
6.2	The parallel results to the table 6.1 in classifying the caudate nucleus. The same conclusion observed in the table 6.1 can be drawn.	50
7.1	Table of AUCs of classifying HR-ASD and HR-Neg based on 1) temporal shape difference in hippocampi at 6 months and 12 months (the top row), 2) 6 months hippocampi shapes (the 2 nd row), 3) 12 months hippocampi shapes (the 3 rd row), and 4) the pure random guessing (the bottom row).	57
7.2	Table of AUCs of classifying HR-ASD and HR-Neg based on 1) temporal shape difference in caudate nuclei at 6 months and 12 months (the top row), 2) 6 months caudate nuclei shapes (the 2 nd row), 3) 12 months caudate nuclei shapes (the 3 rd row), and 4) the pure random guessing (the bottom row).	57

LIST OF FIGURES

2.1	(a) Skeletal model of a hippocampus s-rep; (b) solid model implied by that s-rep. Yellow spheres are sample points along the skeletal surface. Solid lines extending from these sample points are spoke vectors, which are approximately normal to the boundary surface. Interpolation of a discrete s-rep into a continuous skeleton with a continuous field of spokes forms a continuous s-rep whose spokes completely fill the interior of the object they are representing.	11
4.1	(a) The solid model of the template hippocampus s-rep; (b) The solid model of the warped template hippocampus s-rep.	17
4.2	The illustration of the workflow of the method to obtain the s-rep fitted to the input boundary surface.	19
4.3	Skeletal ellipse of an ellipsoid whose principal radii are 23, 8, and 5. The red circles are the sampled internal skeletal points of the skeletal ellipse, and the black boundary curve is the end curve of the true skeletal ellipse. To support the legacy representation of the s-rep, the black ellipse is slightly eroded to the blue ellipse. The black points are the true end skeletal points of the true skeletal ellipse, and these are pulled toward to the eroded ellipse. These pulled points are blue points on the blue ellipse. In this case, the skeletal points are sampled to form 5×7 grid. After the end skeletal points are sampled along the end curve, the interior skeletal points are uniformly sampled along the blue lines until the lines intersect the medial curve of the skeletal ellipse, <i>i.e.</i> , the horizontal line at $y = 0$	21
4.4	The computed spokes of the same ellipsoid used in the figure 4.3 The colored lines are the spokes that are analytically computed using the equations. The cyan and magenta lines denote spokes pointing top and bottom side of the ellipsoid. The red lines point to crests of the ellipsoid.	22
4.5	(a) Fitted s-rep of the hippocampus using the proposed approach; (b) solid model implied by that s-rep. As can be seen in (a), the lengths of the crest spokes are longer than expected.	24
4.6	(a) Fitted s-rep of the hippocampus before the refinement on crest spokes; (b) fitted s-rep of the hippocampus after the refinement on crest spokes.	25
4.7	The histogram of the distances from each base spoke end of the fitted s-rep to the corresponding object boundary for a case in the population of 6-month olds' hippocampi.	26
4.8	The histogram of the mean spoke end distance for the population of hippocampi.	27

4.9	The histogram of the distances from each base spoke end of the fitted s-rep to the corresponding object boundary for a case in the population of 6-month olds' caudate nuclei.	28
4.10	The histogram of the mean spoke end distance for the population of caudate nuclei.	29
4.11	(a) Fitted s-rep of a neonate's left lateral ventricle using the proposed approach; (b) solid model implied by that s-rep.	29
4.12	The histogram of the distances from each base spoke end of the fitted s-rep to the corresponding object boundary for a neonate's lateral ventricle.	30
5.1	An interpolated skeleton of one of the CN s-rep with the new singular point primitive. The blue arrow seen at the lower left corner of the figure is the directional component of the singular point primitive.	39
5.2	Densely interpolated spokes of the same CN s-rep previously seen in figure 5.1 with the singular point primitive. The magenta vectors denote up spokes, the cyan vectors denote down spokes, and the red vectors denote crest spokes.	40
5.3	The histogram of the average spoke end distance of the augmented CN s-rep fitted to a population of 6-month olds' caudate nuclei.	42
6.1	Visualizations of the class likelihoods of hippocampi. The empirical histogram of the scalar projection of the non-autistic cases in the training set onto the separation direction is plotted using the blue dotted lines and with the Gaussian probability distribution as the blue solid curve. Similarly for the autistic class in green.	47
6.2	Selected frames from the sequence of the s-reps while walking along the separation direction through the pooled backward mean from the autism class to the non-autism class. Viewing the sequence as a looping movie makes the local shape changes between the two classes more noticeable.	52

LIST OF ABBREVIATIONS

ASD	Autism Spectrum Disorder
B-PDM	Boundary Point Distribution Model
CPNS	Composite Principal Nested Spheres Analysis
CN	Caudate Nucleus
PCA	Principal Component Analysis
PNS	Principal Nested Spheres
PPCA	Polysphere Principal Component Analysis
DWD	Distance Weighted Discrimination
GOPs	Geometric Object Properties
MCF	Mean Curvature Flow
MRI	Magnetic Resonance Imaging
TPS	Thin Plate Spline
TSD	Thin Shell Demons

CHAPTER 1

Introduction

1.1 Overview

The driving biomedical problem of this dissertation is an early diagnosis and detection of Autism Spectrum Disorder (ASD). This work focuses on classifying the autistic from the non-autistic among infants at high familial risk (HR) of ASD based on shapes of hippocampi or caudate nuclei. Among these infants at high risk of ASD, I refer those who are later diagnosed with ASD as *HR-ASD* and those who are not as *HR-Neg*. There are two important factors to consider for this shape classification problem:

- A choice of an object model to represent the shape of the hippocampus and the caudate nucleus.
- A choice of a method to learn a classification rule based on shape descriptions of hippocampi and caudate nuclei.

An effective shape classification is difficult because of several challenges associated with each design choice. Firstly, with numerous object models having been proposed in the literature, choosing an object model capable of providing shape descriptions of the object that may be relevant for a classification task at hand is in itself challenging. Secondly, one has to consider whether an object model can be fitted to a set of object segmentation with reasonably good correspondence to avoid any spurious classification. Finally, after settling the choice of the object model that satisfies the two aforementioned concerns, choosing a classification method capable of yielding an effective classification rule poses yet another challenge.

In order to effectively discriminate the HR-ASD from the HR-Neg infants based on morphological traits of hippocampi and caudate nuclei, the following properties are necessary for each component of the shape classification.

- An object model needs to
 - provide rich localized descriptions of those structures that are sensitive to any subtle morphological difference.
 - be fitted to an object segmentation robustly and automatically with good correspondence in the target objects across the sample population.
 - be able to adequately represent surface geometry of an object with a singular point, *e.g.*, the caudate nucleus.

- A classification method needs to
 - robustly classify cases when the input feature dimension is larger than the sample size.
 - account for the fact that GOPs abstractly live on a curved manifold.

This dissertation describes novel methodological contributions in each of these areas. In subsequent sections of this chapter, I first overview the current findings in the driving application problem. Then, I provide details on the intuition behind each of my contributions. Finally, I conclude this chapter with my thesis statement and the organization of the remaining chapters.

1.2 Current Findings on ASD from Neuroimaging Studies

ASD is a neurodevelopmental disorder characterized by a few core symptoms including repetitive behaviors and difficulties in social interaction and communication. ASD manifests very early in life. Previous behavioral studies of ASD show that clinical symptoms may appear as early as 12 months with most infants receiving a diagnosis by the age of 4. This relatively narrow window of birth to age of onset of symptoms provides opportunities to study the underlying neurodevelopmental process of infants with ASD.

Recent longitudinal infant brain imaging studies of infants have yielded many important insights about the developmental process of infants who are diagnosed with autism by the age of 2. Findings include increased cortical gray and white matter volumes, increased amygdala volumes, and atypical neurodevelopmental trajectories (Hazlett et al., 2012; Wolff et al., 2013).

However not much is known about localized morphological differences of *subcortical* structures between the individuals with and without ASD during this early developmental period. In older children, there have been a number of studies showing that subcortical structures including the caudate nuclei and the hippocampi are implicated in ASD. This has motivated me to develop classification methods based on local surface geometric properties of the hippocampus and the caudate nucleus from the perspective of shape classification.

1.3 Common Object Models used in Shape Classification

Classification methods based on a global volume of the objects (Hazlett et al., 2012) have been common due to the simplicity of computing and understanding the measure. However, such methods may not yield an effective discrimination when classifying anatomical structures between patients at the early stage of the disease and healthy control subjects because a single measure such as the global volume may not adequately explain subtle morphological differences between the two groups. In addition, those methods have weakness in providing localized insights into causes of the observed difference in volume, *e.g.*, that the volume change is driven by one or more regions of the structures. The localized analysis could yield an important biological insight into the disease.

The Point Distribution Model (PDM) (Cootes et al., 1992) has been a popular object model in shape classification (Davies et al., 2003). The PDM describes a shape of the object as a tuple of enumerated points usually placed along the boundary of the object. The PDM has several desirable traits such as being easy to create from image data (Davies et al., 2003; Styner et al., 2006) and supporting a localized analysis of the shape of the objects. However, the localized analysis of shapes with PDMs is still limited due to PDMs only explicitly capturing positional geometrical properties of the objects; therefore, the PDM-based analysis is limited in accurately modeling inter-class shape variations involving a local direction and a local width.

1.4 Skeletal Representations

Skeletal representations describe the shape of an object via a collection of geometric primitives each of which has an associated skeletal locus. The s-rep is a skeletal representation whose geometric primitives are vectors called *spokes* pointing to the object boundary from each skeletal locus to

the object boundary. These spokes provide the rich localized shape properties of the object, *i.e.*, orientation, width, and position. These GOPs are capable of modeling a localized nonlinear inter-class shape difference that may be a useful insight into the classification task.

However, these GOPs, *e.g.*, a set of local directions, are non-Euclidean. Therefore, a classification method has to be chosen carefully because of non-Euclidean nature of the s-rep GOPs: without properly handling the nonlinearity of GOPs, a classification method, *e.g.*, the Support Vector Machine, designed to classify Euclidean data fails to learn a sensible classification rule (Sen, 2008). I use an alternative approach called *Euclideanization* where these non-Euclidean s-rep GOPs are properly transformed into a Euclidean feature tuple.

1.5 Automatic Generation of S-reps with Good Correspondence

For each target object in the training set, a set of s-reps are first obtained by registering an implied boundary surface of the template s-rep model to each target object boundary. The template s-rep is deformed to roughly fit each target object boundary via the Thin Plate Spline (TPS) transformation computed based on the correspondence established by Thin Shell Demons (TSD). Then, spokes of each warped template s-rep are further refined so that the s-rep tightly fits the target object boundary. This process yields a set of fitted s-reps that have a reasonable correspondence in spokes across the set.

The template s-rep is often chosen to be the s-rep fitted to the same target object in the past. However, the template s-rep model, *e.g.*, the left hippocampus s-rep of a young adult, often does not reflect shape variations observed in the given sample population of the target object, *e.g.*, a set of left hippocampi of young infants. In short, establishing a reasonable correspondence in spokes by warping a single s-rep model to each target object in the set poses another challenge when the chosen template s-rep is not representative of the sample population.

However, an analytical expression of the s-rep is known for an ellipsoid. Thus, the ellipsoid s-rep can be deformed to fit each target object if a mapping between the ellipsoid and the target object is established. I use Mean Curvature Flow (MCF) to obtain the ellipsoidal mapping of the target surface. Then, I apply a deformation that is the inverse of MCF to an ellipsoid s-rep followed by refinements

in spokes to the deformed ellipsoid s-rep to obtain a fitted s-rep. Reasonable correspondence in spokes is achieved because ellipsoids are related by scaling each axis of the ellipsoid.

1.6 Extensions of the Discrete S-rep to Model a Singular Point

My driving medical problem requires appropriately capturing geometric properties of the caudate nucleus (CN). The literature reports morphological deviations of the CN between individuals with and without ASD. Especially notable is the association of ASD patients' restricted and repetitive behaviors with enlargement of the global volume. However, detailed analysis of the local shape differences of CN between the two groups is lacking.

The input boundary surface of the CN derived from a medical image, *e.g.*, structural Magnetic Resonance Imaging, tends to have a tapering sharp tail, *i.e.*, a singular point. The singular point on the boundary of the caudate nucleus makes it difficult to represent the structure with an s-rep because it causes the skeletal surface to collapse to a point. This degeneracy of the skeletal surface causes a number of issues in modeling the surface geometry of the caudate nucleus.

To address this challenge in modeling neuroanatomical structures with a singular point, *e.g.*, the caudate nucleus, via the s-rep I introduce an additional geometric primitive dedicated to represent the singular point to account for the degeneracy. Then, I extend the current quad-based spoke interpolation method so that the interpolated boundary surface patches are C^2 continuous up to the singular point.

1.7 Common Classification Methods used in Shape Classification

Accurate classification of medically imaged objects on the basis of their GOPs is difficult. In such a setting the input feature space is a curved manifold due to the non-Euclidean nature of the GOPs. However, a classification method designed to find an optimal classification rule on the non-Euclidean feature space is lacking. Direct application of the Euclidean classification methods on GOPs yields a suboptimal classification rule; as demonstrated in (Sen, 2008) it might even be invalid. In this dissertation I apply Principal Nested Spheres Analysis (PNS) to the GOPs to transform the non-Euclidean nature into the Euclidean. PNS yields a set of the best-fitting submanifolds of the GOPs as well as their signed geodesic distances to the respective best-fitting submanifold. These

signed distances can be regarded as the Euclidean counterparts to the original GOPs; the original non-Euclidean GOPs appear now to be reasonably *Euclideanized* for the Euclidean classification methods.

The shape classification problem is often a High Dimensional Low Sample Size (HDLSS) problem in which the dimension of the input GOPs tuple is larger than the number of available data samples. In the HDLSS classification problem Marron (Marron et al., 2007) has shown that many classifiers, *e.g.*, the Support Vector Machine and the Mean Difference Classifier, lose generalizability due to a phenomenon called *data piling*. To account for this phenomenon in my shape classification tasks, I use Distance Weighted Discrimination, which is designed to learn a generalizable classification rule in HDLSS classification problems.

1.8 Shape Classification of a Temporal Pair of Objects

A number of recent psychiatric studies have shown that the brain developmental trajectories of HR-ASD infants noticeably deviate from those of the typically developing infants (Hazlett et al., 2017). However, not much is known about the correlation between ASD and the developmental trajectories of the implicated subcortical structures. This has motivated me to investigate the problem of classifying objects based on a growth information as to whether the disease is present. For the early diagnosis of ASD, I have chosen to investigate a growth information of hippocampi and caudate nuclei between 6 months and 12 months.

The main task in the shape pair classification is to accurately encode the temporal difference in GOPs of an object between the two time points into the DWD classifier's feature tuple. This raises a statistical question whether the GOPs of each of the s-rep are analyzed together or separately for PNS. In this work I chose the latter approach in which the Euclideanized input feature tuple is obtained by taking difference between the s-rep GOPs that are Euclideanized by a common polar system produced from a union of s-reps across the time points and across the classes.

I have evaluated the effectiveness of the temporal difference classification method in two scenarios: 1) ASD classification given the temporal pair of hippocampi of 6-month olds and 12-month olds, 2) ASD classification given the temporal pair of caudate nuclei of 6-month olds and 12-month olds,

1.9 Thesis and Contributions

Thesis: *Classification of medically imaged neuroanatomical structures based on their geometric object properties benefits from the following:*

1. *An object model providing local object width and orientation in addition to positional information*
2. *An object modeling procedure that efficiently yields geometric correspondence across a population*
3. *A classification method capable of learning on a curved manifold*

The major methodological contributions of this dissertation are as follows:

1. A fully automatic procedure to obtain a set of s-reps fitted to each target object boundary data with reasonably good correspondence in spokes.
2. An extension to the current s-rep modeling framework for appropriately representing an object with a singular point, *e.g.*, the caudate nucleus.
3. A non-Euclidean classification method that use PNS to obtain the Euclideanized s-rep GOPs on which DWD classifier is trained.
4. A statistical technique to compare different shape classification methods.
5. A non-Euclidean classification method for classifying a temporal difference of s-reps in which the input feature tuple to DWD is formed by taking the difference of the Euclideanized s-rep GOPs for each time point

In addition to the above methodological contributions, I have also accomplished the following engineering contributions:

1. Redesign of the computer representation of the discrete s-rep, which allows modeling objects with different topologies including the spherical (slabular), the quasi-tubular, and the toroidal
2. Integration of Pablo, the main piece of software used to fit and visualize s-reps into 3D Slicer for easier distribution of the s-rep modules to the medical image analysis community

3. Modernization of the s-rep modules
4. Implementation of the module that converts a legacy s-rep into the new s-rep
5. Implementation of the module that visualizes a legacy s-rep in 3D Slicer
6. Improvement of the numerical stability of estimating surface normals and principal curvatures from the implicit representation of the object by a signed distance image.

CHAPTER 2

Background

In this chapter I lay out background information that is common for each of the subsequent chapters. I first overview object models. Then, I overview statistical methods to understand shape features and classification methods.

2.1 Object Models

At a high level there are two categories of object models that have been proposed for statistical analysis: continuous, parameterized models modulo parameterization (Kurtek et al., 2012; Jermyn et al., 2012; Bauer et al., 2010, 2012; Durrleman et al., 2014) and discrete models. Due to the discrete models' strengths in explicitly dealing with localized features, I focus on those models. Among the discrete models are those based on deformations of an atlas (Beg et al., 2005; Miller et al., 2002; Wang et al., 2007), those based on the b-PDM (Cootes et al., 1995; Styner et al., 2006; Davies et al., 2003), and those based on skeletal models (Styner et al., 2004; Yushkevich and Zhang, 2013; Bouix et al., 2005; Schulz et al., 2016). The b-PDM-based models have been the most popular. The skeletal models were designed to add local object width features and local directional features to those provided by PDMs.

I overview the two object models that I compare in this work: the PDM and the s-rep. For each model, I provide

- a brief description of the representation
- a brief description on fitting the object model to a boundary description provided as input
- a brief description of a curved manifold on which shape features of the representation live.

2.1.1 The Point Distribution Model

The PDM is a point tuple for each object in a training set. In the boundary PDM (b-PDM) each example object in the set has a set of enumerated points along its boundary, with points with corresponding index in each object chosen so as to be in correspondence across the training set.

There are several ways to produce a set of PDMs with reasonable correspondence from a set of objects, including manual placements of landmarks, automatic placements of points based a criterion such as minimum description length (Davies et al., 2003) or spherical harmonics (Styner et al., 2004, 2006).

The PDM GOPs, *i.e.*, a tuple of point coordinates, are often interpreted in two ways: 1) as a point on a flat Euclidean space or 2) as a point on a curved manifold (Kendall, 1984). Consider the b-PDM in the training set \underline{P} with n boundary points. By scaling the entire point tuple such that the sum of squares of all the center-of-mass-relative point features has unit value, this can be thought of projecting the point tuple onto the unit hypersphere \mathbb{S}^{3n-4} . The dimensionality of $3n - 4$ comes from the fact that three degrees of freedom were used during alignment and one more degree of freedom was used to normalize a scale to unity. Therefore, as rigorously shown by Kendall in his work (Kendall, 1984), the translationally aligned b-PDM can be represented as a concatenation of this scaling factor and this normalized tuple of points: the b-PDM abstractly lives on the manifold $\mathbb{R}^+ \times \mathbb{S}^{3n-4}$.

2.1.2 The S-Rep

The discrete s-rep is a skeletal discretization of the interior of the object. It consists of a set of points sampled on the skeletal surface, which is a folded surface with the same topology, in this dissertation spherical topology, and vectors called *spokes* associated with each of the skeletal points that are approximately normal to the object boundary surface. These spokes explicitly capture local direction and local width information of the object. An example discrete s-rep of a hippocampus can be seen in figure. 2.1.

S-reps are fitted to an object via a posterior optimization that tries to match the boundary surface implied by the continuous s-rep interpolated from the discrete spokes by Vicory’s work in (Tu et al., 2016) to the input object boundary from the image data while remaining as medial as possible.




Figure 2.1: (a) Skeletal model of a hippocampus s-rep; (b) solid model implied by that s-rep. Yellow spheres are sample points along the skeletal surface. Solid lines extending from these sample points are spoke vectors, which are approximately normal to the boundary surface. Interpolation of a discrete s-rep into a continuous skeleton with a continuous field of spokes forms a continuous s-rep whose spokes completely fill the interior of the object they are representing.

The optimization is initialized by deforming a reference s-rep to an individual case via the TPS transformation that is computed based on boundary correspondence between the reference object and the target object. The boundary correspondence is established by TSD. Then, the deformed reference s-rep is refined to fit tighter to the input object boundary by lengthening or shortening spokes. This process yields a set of s-reps whose spokes are in reasonable correspondence in the training set.

S-rep GOPs, *i.e.* a tuple of positions, local width, local orientation, lie on a curved manifold. Consider a discrete s-rep \underline{s} with n spoke vectors and m skeletal points. The set of skeletal points forms a PDM that is aligned such that its center of gravity is at the origin. Additionally, this tuple of centered points is scaled by a factor making the sum of squared distances to the origin to be unity. Therefore, this PDM is described by a tuple of centered points that abstractly lives on the unit hypersphere \mathbb{S}^{3m-4} and an associated log-transformed scaling factor. The directional component of each spoke abstractly lives on the unit 2-sphere \mathbb{S}^2 , and the log-transformed associated length component of each spoke lives on the Euclidean space \mathbb{R}^1 . Thus, a single discrete s-rep abstractly lives on $\mathbb{R}^{n+1} \times \mathbb{S}^{3m-4} \times (\mathbb{S}^2)^n$.

2.2 Statistical Methods

I provide brief descriptions of statistical methods commonly used to understand the input shape data. I first overview PCA, the conventional approach. Then, I overview PNS analysis, a variant of PCA to

analyze data that live on abstract spheres. I briefly describe CPNS, a statistical analysis technique that is appropriate for analyzing the data that live on a Cartesian product of Euclidean space and hyperspheres. Finally, I finally briefly describe Polysphere PCA (PPCA), a method designed as an extension to PNS to properly understand the correlation of variables in each domain in a Cartesian product of hyperspheres, *i.e.*, polysphere.

2.2.1 Principal Component Analysis

Principal component analysis (PCA) is an important statistical method for analyzing Euclidean data. It provides a means of reducing the intrinsic dimension of data by capturing its major modes of variation. PCA has been widely used in the field of medical image analysis and computer vision because descriptions of objects of interest are often high dimensional whereas the important variations can be quite low dimensional. Those modes of variation are often quite illuminating and relevant to the task at hand.

PCA can be understood in terms of a forward or backward procedure. The forward method progressively builds up the dimension of the approximating subspace being fitted to the data, whereas the backward method progressively reduces the dimension of the subspace being fitted to the data. Both approaches yield the same result when the data lie on a Euclidean space. However, GOPs do not lie on a Euclidean space. The backward approach typically yields different results from the forward approach when applied to non-Euclidean data. As noted in (Damon and Marron, 2013), the backward approach is usually more appropriate to analyze those non-Euclidean features.

The forward PCA increases dimension by adding the component that captures the most remaining variance; at each iteration a component that best describes the data and that is orthogonal to previous components is added to form a new best fitting manifold so that the current manifold is the best fitting submanifold of the data in the original dimension. The principal component scores are found by projecting all the data onto the found submanifold.

In contrast, the backward view of PCA progressively reduces the intrinsic dimension of the manifold by removing the component of the least variance from all the data points; at the beginning of each iteration the data is projected onto the submanifold found in the previous iteration, and then the best fitting submanifold is found by minimizing the sum of squared distances of all the projected data.

2.2.2 Principal Nested Spheres Analysis

Principal Nested Spheres (PNS) analysis is a special case of backward PCA on hyperspheres. PNS progressively reduces intrinsic dimension by finding the best fitting subsphere \mathbb{S}^{k-1} that is nested in the current hypersphere \mathbb{S}^k . At each iteration, the data points are first projected onto the subsphere found in the previous iteration; then the fitting is done by minimizing the sum of squared geodesic distances of all the projected data points to the subsphere. Over the training cases PNS will yield a tuple of signed geodesic distances to the best fitting subsphere for each dimension-reduction iteration.

These signed geodesic distances are a Euclideanized form of their spherical counterparts; I call this process Euclideanization. The final result of PNS yields Euclideanized variables and a set of polar systems that provide a means of transformations between the original space and Euclideanized space and vice versa. The dimension 0 point in feature space produced at the end of this iteration is called the backwards mean. (Jung et al., 2012) provides more information on the method.

2.2.3 Commensuration of Spherical GOPs

Suppose the data of interest live on a Cartesian product of a Euclidean space and hyperspheres. Such an instance includes any model described by a combination of GOPs involving PDMs, lengths, directions, and scaling. In this case, PNS is applied independently to each GOP that lives on a hypersphere.

In the case of s-reps PNS is applied to centered scale normalized point tuple and each spoke directions separately. Each application of PNS on spherical GOPs produce their Euclidean counterparts. To make the components appropriately commensurate (Jung et al., 2010) when analyzing a set of s-reps, I multiply each Euclideanized value derived from a PDM by the geometric mean of the scale factors in the training population, and I multiply each Euclideanized value derived from a direction by the geometric mean of its associated length.

2.3 Classification Methods

At a high level there are two categories of classification methods: likelihood based methods and separating direction based methods. See (Duda et al., 2012) and (Friedman et al., 2001) for an

overview of common existing classification methods. In this dissertation I focus on the methods based on finding an optimal separating direction that points from one class to the other because the inter-class shape difference can easily be visualized by interpolating points in the feature space along the separation direction. Given a separation direction, the class label is decided based on scalar projection of data points onto the separation direction.

I concentrate on linear classification methods because 1) in High Dimensional Low Sample Size (HDLSS) setting (when the dimension of data is larger than the sample size), El Karoui demonstrated that there is little value added by Kernel tricks since the kernel method is asymptotically linear method in this setting (El Karoui et al., 2010). 2) I want to gain insights into the driving application problem, *e.g.* classifying autistic infants from non-autistic infants by directly inspecting features. I especially pay attention to the separating direction vector since large entries in the vector indicate that the corresponding feature is relevant. A good separating direction provides additional information and insight into the data by visualizing the trends between the classes by linearly interpolating and then synthesizing the data in the original feature space along the direction.

2.3.1 Support Vector Machine

The SVM (Cortes and Vapnik, 1995) is a binary classification method that yields a separation direction in the feature space by minimizing the gap between the two classes. SVM then classifies a new example by thresholding the scalar value of the projection of its feature tuple onto this direction.

2.3.2 Distance Weighted Discrimination

DWD is a classification method similar to SVM but which is more robust to noise and limited sample size. Like the SVM, DWD takes in two classes of data and yields a separating direction that can be used to classify new data points through projection and thresholding. Unlike SVM, the separating direction computed by DWD is influenced by all points in the data set. A full description of DWD can be found in (Marron et al., 2007).

CHAPTER 3

Materials

The image dataset on which my study bases was acquired from an NIH-funded study of Autism referred to as the *Infant Brain Imaging Study* (IBIS). I focus on the subset of the participants of this study who are at high familial risk (HR) of ASD because they have an older sibling already diagnosed with ASD. There are 49 infants who were clinically diagnosed with ASD at age 24 months (HR-ASD) and 149 infants who did not meet criteria for ASD (HR-Neg). MRI scans (Siemens 3T TIM Trio scanners, T1 weighted imaging at $1.0 \times 1.0 \times 1.0$ mm voxel resolution) were acquired while infants were naturally sleeping. The infants were scanned longitudinally at the age of 6 months, 12 months, and 24 months. The MRI scans were aligned to a common coordinate system. Subcortical structures were automatically segmented from the aligned MRI scans (Wang et al., 2014). Details on the original MRI dataset can be found in (Hazlett et al., 2017; Wolff et al., 2015; Lewis et al., 2014).

Bilateral hippocampi and the caudate nuclei are expected to be associated with ASD. However, I fitted the s-rep to the left hippocampus and the left caudate nucleus for method development. 15 records for the HR-ASD at 6-months and 6 records for the HR-Neg were not available to us at the time of the analysis. Thus, the dataset we use in this study consists of 34 autistic infants and 143 non-autistic infants.

In this dissertation, I study the following classification tasks between HR-ASD infants and HR-Neg infants based on 1) the left hippocampi s-reps at 6 months 2) the left caudate nuclei s-reps at 6 months 3) the temporal difference of the left hippocampi s-reps of the infants at 6 months and 12 months 4) the temporal difference of the left caudate nuclei s-reps of the infants at 6 months and 12 months.

CHAPTER 4

Automatic Generation of Case-Specific S-reps

4.1 Introduction

S-rep has been shown to be a powerful shape representation in various statistical shape analysis tasks including hypothesis testing (Schulz et al., 2016), classification (Hong et al., 2016), shape probability distribution estimation [ref], and segmentation [Jared's segmentation]. For these tasks it is important that the set of s-reps have good correspondence in spokes across the set to yield meaningful statistical insights related to the morphology of the target object in the sample population. Tu et al. (Tu et al., 2018) proposed an entropy-based method to improve the s-rep spoke correspondence. However, for a set of s-reps of moderate size (*e.g.*, 30 cases) the method cannot be applied because of its prohibitive computational cost.

As mentioned in section 2.1.2, one approach to establish the s-rep correspondence across the set is to warp a single template s-rep to fit each target object surface. However, using a single template model as a basis of the spoke correspondence tends to cause difficulties for the s-rep fitting process. This approach often yields a set of s-reps that contain geometric artifacts, *e.g.*, creases and folds, on the boundary surface implied by the fitted s-reps as can be seen in figure 4.1. These artifacts arise because of 1) the difficulty in finding the appropriate template s-rep that is representative of the population of interest and 2) the difficulty in estimating the deformation that respects the skeletal geometry of the target object.

I argue that the following properties are desirable for the optimal template s-rep:

1. It is case-specific.
2. It is obtained from an analytical expression.
3. The correspondence across the cases in the set can be achieved with a simple transformation.




Figure 4.1: (a) The solid model of the template hippocampus s-rep; (b) The solid model of the warped template hippocampus s-rep.

In this chapter I present a novel method to automatically generate a set of fitted s-reps with reasonably good correspondence across the cases. The method works as follows:

1. Deform a target object surface to an ellipsoid via Mean Curvature Flow (MCF) [ref to MCF?].
2. Obtain a skeletal model of the ellipsoid from a known analytical expression.
3. Deform the ellipsoid s-rep back to the original object via reverse MCF transformations.

The remainder of this chapter is organized as follows. First, I provide a brief description of MCF and the Thin Plate Splines (TPS). Then, I describe the method in detail. I demonstrate that the method is effective at generating s-reps for 6-month old infants' hippocampi, caudate nuclei, and neonates' lateral ventricles. Finally, I conclude this chapter with discussion.

4.2 Background

4.2.1 Mean Curvature Flow

Mean Curvature Flow (MCF) is a geometric flow method in which the input surface evolves along the surface normal proportional to the mean curvature of the surface. MCF is mathematically defined

as the following.

$$\begin{aligned}\frac{\partial}{\partial t}\mathcal{S}(x,t) &= \mathcal{H}(x,t)\mathcal{N}(x,t) \\ \mathcal{S}(x,0) &= \mathcal{S}_0(x)\end{aligned}$$

where \mathcal{S}_0 is the original surface, \mathcal{S} is the deforming surface, x is a point on \mathcal{S} , $\mathcal{H}(x,t)$ is the mean curvature of \mathcal{M} at x at time t , and $\mathcal{N}(x,t)$ is the surface normal of \mathcal{M} at x at time t . One of well-studied asymptotic behavior of MCF is that \mathcal{S} will converge to a sphere as $t \rightarrow \infty$. MCF can be formulated both continuously and discretely. In this work I use a discrete formulation of MCF: MCF on a triangle mesh. Given a triangle mesh $\mathcal{S}(V,E)$ where V and E denote vertices and edges of the mesh respectively, the mean curvature \mathcal{H} at each vertex $x \in V$ is estimated as the average of dihedral angles of the triangles that share the vertex x . More extensive introduction of MCF is available in [ref].

4.2.2 Thin Plate Splines

Thin Plate Splines (TPS) has been commonly used in non-rigid image registration given two sets of landmarks. The method generates a full 3D deformation field by interpolating the displacement of the landmarks. TPS is not guaranteed to produce a diffeomorphic transformation; it can produce folds within the deformation field. However, this can be prevented if the displacement vectors are short enough. A more comprehensive reference on using TPS for image registration is available in [Bookstein ref].

4.3 Method

Figure 4.2 shows the workflow of the method. First, I iteratively deform the input target surface via MCF. At each iteration I determine if the current deforming surface is approximately ellipsoidal. If it is, I obtain the s-rep of the best-fitting ellipsoid to it. Then, I transform the ellipsoid s-rep first to the approximate ellipsoid and then in reverse direction of the MCF so that the s-rep will roughly fit the original boundary data. I provide details of each step of the method.

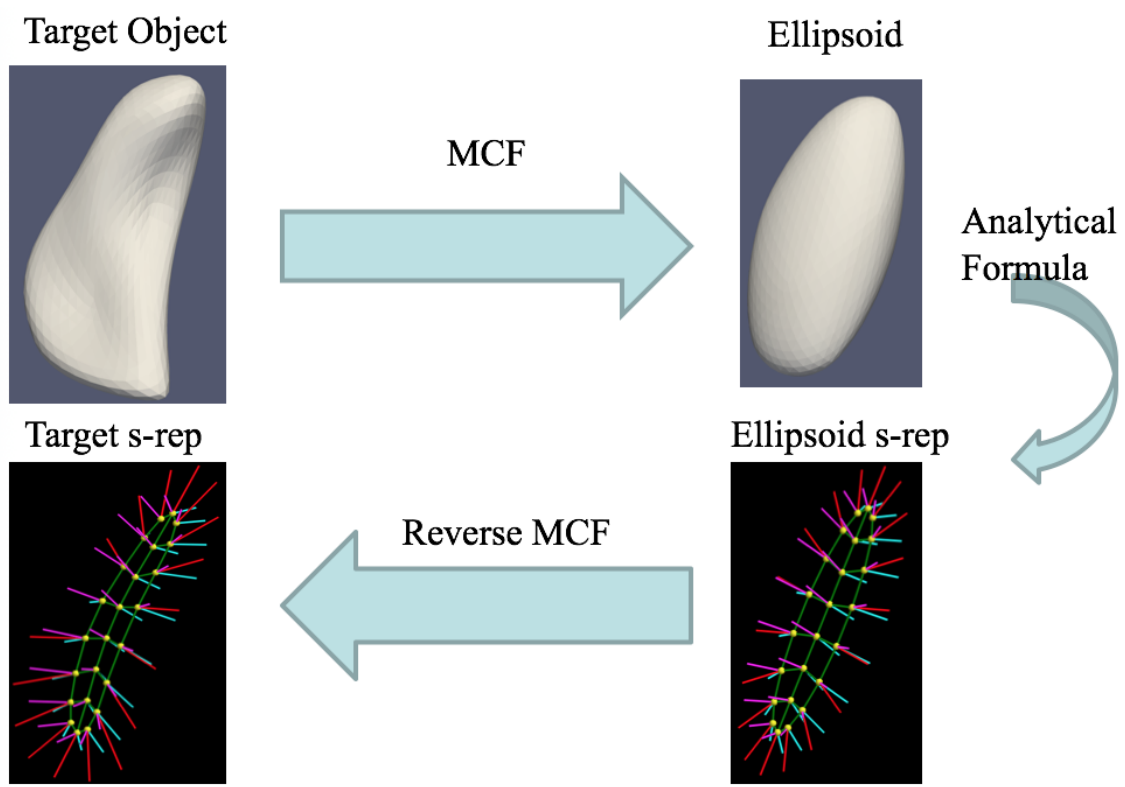


Figure 4.2: The illustration of the workflow of the method to obtain the s-rep fitted to the input boundary surface

4.3.1 Applying MCF to the Deforming Surface

I iteratively apply MCF to the triangle mesh of the object derived from the image data. In order to improve the numerical stability in computing the discrete mean curvature at each vertex, at each iteration of the process prior to applying MCF I apply Taubin smoothing [ref]; as a side effect, the smoothing yields more regularly shaped triangles.

4.3.2 Determining if the Deforming Surface is Ellipsoidal

After I apply MCF to the mesh, I examine the second moment of the mesh to see if it is ellipsoidal.

The second moments matrix of the mesh is computed as the following.

$$\begin{bmatrix} \sum_i x_i^2 & \sum_i x_i y_i & \sum_i x_i z_i \\ \sum_i y_i x_i & \sum_i y_i^2 & \sum_i y_i z_i \\ \sum_i z_i x_i & \sum_i z_i y_i & \sum_i z_i^2 \end{bmatrix}$$

where x_i, y_i, z_i are coordinates of i^{th} boundary vertex relative to the center of mass of those vertices.

The eigendecomposition of the matrix yields the three principal axes and spread of the data along each principal axis. Using the result of the eigenanalysis, I derive the analytical expression of the best-fitting ellipsoid for the mesh. Then, I compute absolute value of shortest distance between each vertex of the deforming mesh and the best fitting ellipsoid. If the 95% quantile of the distance is less than a threshold, I conclude that the current surface is ellipsoidal enough. I save all the intermediate deforming meshes S_i $i = 1, 2, \dots, n$ for later computation of the reverse MCF transformation.

4.3.3 Obtaining the S-rep of the Best-Fitting Ellipsoid

Once the surface deformed by MCF is ellipsoidal enough, the best-fitting ellipsoid's s-rep is analytically derived as the following:

1. Sampling the skeletal surface of the ellipsoid
2. Obtaining spokes of the sampled skeletal points.

Given the principal radii r_x, r_y, r_z of the ellipsoid that is aligned so that the radii are ordered as $r_x > r_y > r_z$, the skeletal sheet of the ellipsoid is an ellipse that satisfies the following equation.

$$\frac{x^2}{m_x^2} + \frac{y^2}{m_y^2} \leq 1 \quad (4.1)$$

where $m_x = \frac{r_x^2 - r_z^2}{r_x^2}$, $m_y = \frac{r_y^2 - r_z^2}{r_y^2}$.

Once the analytical expression of the skeletal ellipse is obtained, I need to produce a sampled skeleton in a way that supports classification on the set of fitted s-reps. Then, I sample n_1 points along the major axis u_1 and n_2 points along the minor axis u_2 of the ellipse to generate a grid of skeletal points on the ellipse. Both n_1 and n_2 are odd numbers. In our own research group the choice of n_1 and n_2 is made on the basis of level of detail of the surface geometry that we deem necessary for a statistical task on a given population of shape. Please refer to chapter 3 of (Vicory, 2016) to learn how the number of sampled skeletal points affects the level of detail on the surface geometry. In this dissertation I won't discuss this matter further.

Given how many points to sample on the skeleton, I want to sample the skeleton such that the quadrilaterals formed by neighboring vertices are uniformly shaped. I first uniformly sample along




Figure 4.3: Skeletal ellipse of an ellipsoid whose principal radii are 23, 8, and 5. The red circles are the sampled internal skeletal points of the skeletal ellipse, and the black boundary curve is the end curve of the true skeletal ellipse. To support the legacy representation of the s-rep, the black ellipse is slightly eroded to the blue ellipse. The black points are the true end skeletal points of the true skeletal ellipse, and these are pulled toward to the eroded ellipse. These pulled points are blue points on the blue ellipse. In this case, the skeletal points are sampled to form 5×7 grid. After the end skeletal points are sampled along the end curve, the interior skeletal points are uniformly sampled along the blue lines until the lines intersect the medial curve of the skeletal ellipse, *i.e.*, the horizontal line at $y = 0$.

the boundary of the ellipse. Then, I sample the interior skeletal points uniformly along the ellipse's normal at each boundary skeletal point other than the first and last n_1 points; the interior points are sampled until they intersect with the medial curve of the ellipse, *i.e.*, the horizontal line at $y = 0$. Figure 4.3 shows an example of a sampled skeletal ellipse.

Given a grid of sampled skeletal points on the skeletal ellipse, I need to compute the ellipsoid's spoke vectors for all the skeletal points. I compute spoke vectors by finding the ellipsoid's boundary points implied by the spokes. Let q be a boundary point implied by a spoke and p be a skeletal point to which the spoke is attached. Using the spherical parametrization of the ellipsoid, the coordinate of q is $(r_x \cos(\theta) \cos(\phi), r_y \sin(\theta) \cos(\phi), r_z \sin(\phi))$. According to the skeletal geometry, a valid spoke vector \vec{pq} satisfies the following property: q lies on a bitangent sphere centered at p maximally contained within the ellipsoid. Therefore, the unit vector $\frac{\vec{pq}}{|\vec{pq}|}$ is the surface normal of the ellipsoid at q . I need to compute θ and ϕ of q that satisfy the skeletal geometric constraint to compute the spoke

vector. Given that the skeletal point $p = (x, y, 0)$, θ and ϕ are computed as the following:

$$\tan(\theta) = \frac{ym_x}{xm_y}$$

$$\cos(\phi) = \frac{\sqrt{y^2m_x^2 + x^2m_y^2}}{m_xm_y}$$

When the value of ϕ becomes zero, namely on the boundary of the ellipse, the spoke points to the locus of minimal (negative) principal curvature called the *crest*. This spoke has two important skeletal geometric properties: 1) the spoke lies on the tangent plane of the end curve of the skeleton and 2) the length of the spoke is the radius of that principal curvature. An example of the computed spokes of the ellipsoid s-rep can be seen in the figure 4.4.




Figure 4.4: The computed spokes of the same ellipsoid used in the figure 4.3 The colored lines are the spokes that are analytically computed using the equations. The cyan and magenta lines denote spokes pointing top and bottom side of the ellipsoid. The red lines point to crests of the ellipsoid.

By mathematical definition the skeletal points on the end curve of the ellipse should have only one spoke that points to the crest of the ellipsoid. However, the true crest spokes are pulled to be attached to the nearest end skeletal points of the eroded skeletal ellipse to support the legacy computer representation of the s-rep.

For the purpose of classification I need to obtain a set of s-reps fitted to the given population of shapes with reasonable spoke correspondence. In order to ensure reasonable spoke correspondence of the fitted s-reps, I need to establish spoke correspondence among ellipsoid s-reps. Application of MCF on different cases in the population often yields ellipsoid s-reps of different principal radii. Different ellipsoids can be related with each other by simple scalings along the principal axes if the following assumptions hold:

1. Shapes in the population are rotationally aligned
2. The converged ellipsoids have distinct principal radii, which can be used to establish correspondence among the axes of the ellipsoids.

Given that the above assumptions hold, to provide correspondence among different ellipsoids I compute one single ellipsoid whose principal radii are the respective geometric means of principal radii of all the ellipsoids. After I obtain the mean ellipsoid s-rep, I stretch its skeletal axes of the s-rep by the respective ratios of principal radii of each ellipsoid and principal radii of the mean ellipsoid to obtain a case-specific ellipsoid s-rep.

4.3.4 Transforming the Ellipsoid S-rep back to the Original Surface

Given a case-specific ellipsoid s-rep, I need to deform the ellipsoid s-rep to match the original object boundary at the spoke ends. To accomplish this, I deform the skeletal points and the spoke ends of the best-fitting ellipsoid s-rep first to the approximate ellipsoid and then in reverse direction of the MCF. I compute the deformation as a series of the TPS transformations. I treat the intermediate surfaces' triangle mesh vertices saved throughout the MCF as the landmarks to compute each successive TPS transformation: $T^i : S^i \rightarrow S^{i-1}$ where T^i denotes i -th transformation, S^i and S^{i-1} denote i -th and $i - 1$ -th surfaces from MCF. Figure 4.5 is an example of the s-rep fitted to one of hippocampi in the dataset.

Because the reverse MCF deformation is computed purely based on boundary points, some skeletal geometric properties are not properly propagated throughout the deformation. One such property is the length of the crest spokes. For example, in figure 4.5 many of the crest spokes are too long. This is problematic for the two reasons:




Figure 4.5: (a) Fitted s-rep of the hippocampus using the proposed approach; (b) solid model implied by that s-rep. As can be seen in (a), the lengths of the crest spokes are longer than expected.

1. it narrows the quadrilaterals formed by neighboring skeletal points, which can cause numerical instability for spoke interpolation. It can even cause the skeletal surface to almost collapse to a curve, which can give wrong interpretation about the topology and skeletal geometry of the object.
2. it yields a too small curvature at the end of the crest spoke.

In the subsequent section I describe a procedure to obtain the final s-rep that has more reasonable crest spoke length.

4.3.5 Refining Skeletal Geometric Constraints of the Initial Fitted S-rep

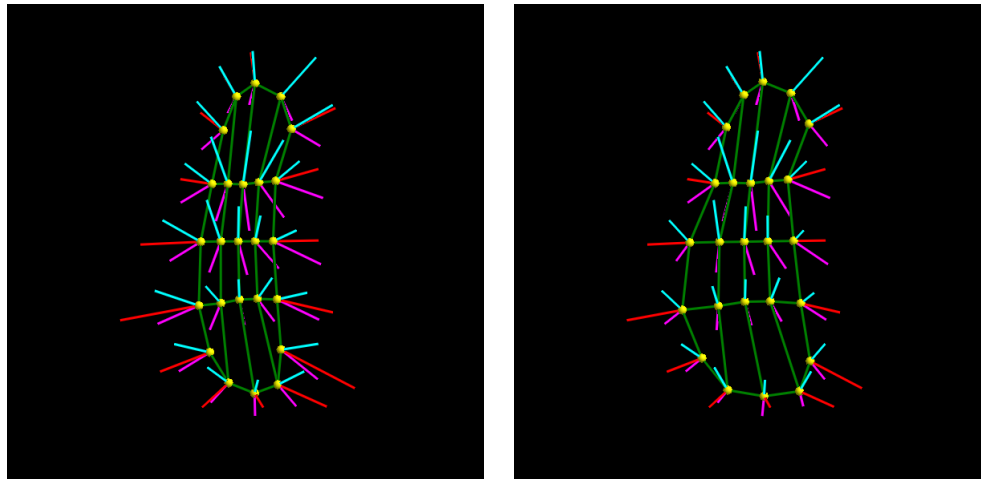
To mitigate the problem of the initial fitted s-rep having unexpectedly longer crest spokes, I place an additional constraint when computing the final TPS so that the computed deformation respects the curvature. To achieve this, for the final TPS of the reverse MCF deformation I add an additional pair of landmarks at each skeletal point on the end curve of the skeleton. The pair's target landmarks are the crest implied loci of the end skeletal points.

The target landmarks need to be the radius of the curvature away from the crest of the object along its normal. In order to estimate the radius of the curvature, I use the implied boundary of the initially fitted s-rep. I compute the radius of the curvature of a curve that spans across the spokes of

the end skeletal points; I fit a cubic B-spline to boundary points obtained from spokes interpolated from the top spoke to the crest spoke and from the crest spoke to the bottom spoke.

However, I have found the estimated radius of curvature to be an underestimate of the desired length. In contrast, the current crest spoke length based on boundary points alone is an overestimate of the desired length. As a result, I take the geometric mean of the estimated radius of curvature and the current crest spoke length as a reasonable estimate for the desired crest spoke length.

Finally, I recompute the final TPS transformation with the additional pair of landmarks so that the final warped s-rep has more appropriate crest spoke length. Below is a figure that shows before and after the refinement on the crest spokes of the same hippocampus s-rep seen in figure 4.6.



(a)

(b)

Figure 4.6: (a) Fitted s-rep of the hippocampus before the refinement on crest spokes; (b) fitted s-rep of the hippocampus after the refinement on crest spokes.

As a side effect of refining the crest spokes, the quadrilaterals formed by neighboring skeletal points are enlarged. This improves the numerical stability in obtaining the smooth implied boundary of continuously interpolated spokes.

4.4 Results

In this section I present several results of applying the proposed method to automatically obtain an s-rep of an object. I investigated the method's effectiveness in producing a set of s-reps for 6-month olds' hippocampi and caudate nuclei; for both the hippocampus and the caudate nucleus, there are

34 autistic cases and 143 non-autistic cases. In addition, to assess the method's applicability to thin objects, I applied the method to obtain s-reps for 3 neonates' lateral ventricles.

In order to get a sense on the accuracy of the proposed method, I examine how well the fitted s-rep matches the original object boundary. I examine the distribution of the distance from each spoke end of the fitted s-rep to the original object boundary; the distance is measured in voxel units where the resolution of the image is supersampled to $0.05 \times 0.05 \times 0.05$ mm. Figure 4.7 shows the histogram of the distances of a case in the population of 6-month olds' hippocampi.

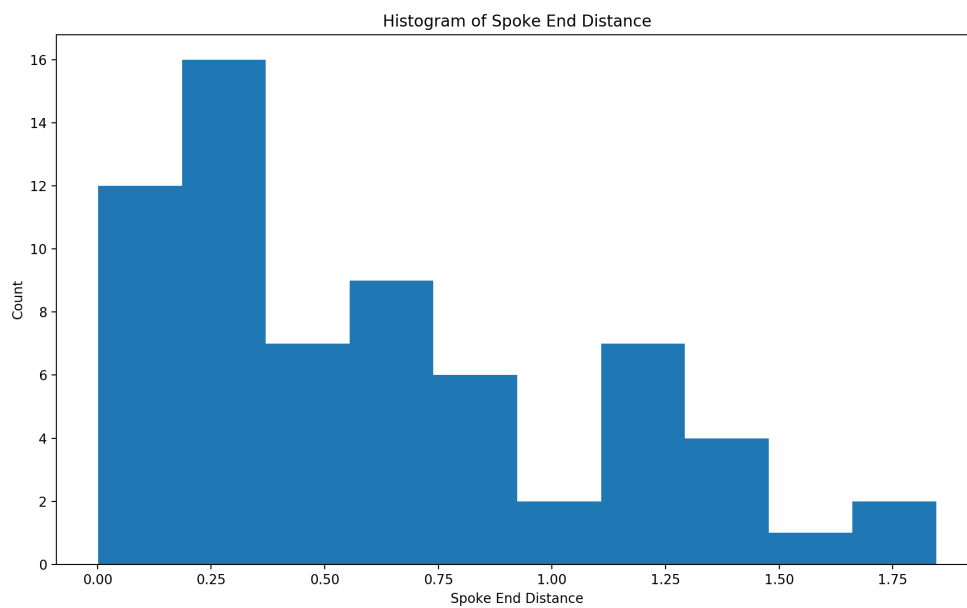


Figure 4.7: The histogram of the distances from each base spoke end of the fitted s-rep to the corresponding object boundary for a case in the population of 6-month olds' hippocampi.

As can be seen from figure 4.7, the majority of the distances is less than a voxel: most of the spokes match the original boundary within 0.05 mm. The mean and the maximum distances are 0.61 voxels and 1.85 voxels respectively, *i.e.*, 0.0305 mm and 0.0925 mm respectively.

In order to assess the quality of s-reps fitted to the population, I examine the distribution of the mean spoke end distances of each hippocampus s-rep. Figure 4.8 shows the histogram of the mean spoke end distances.

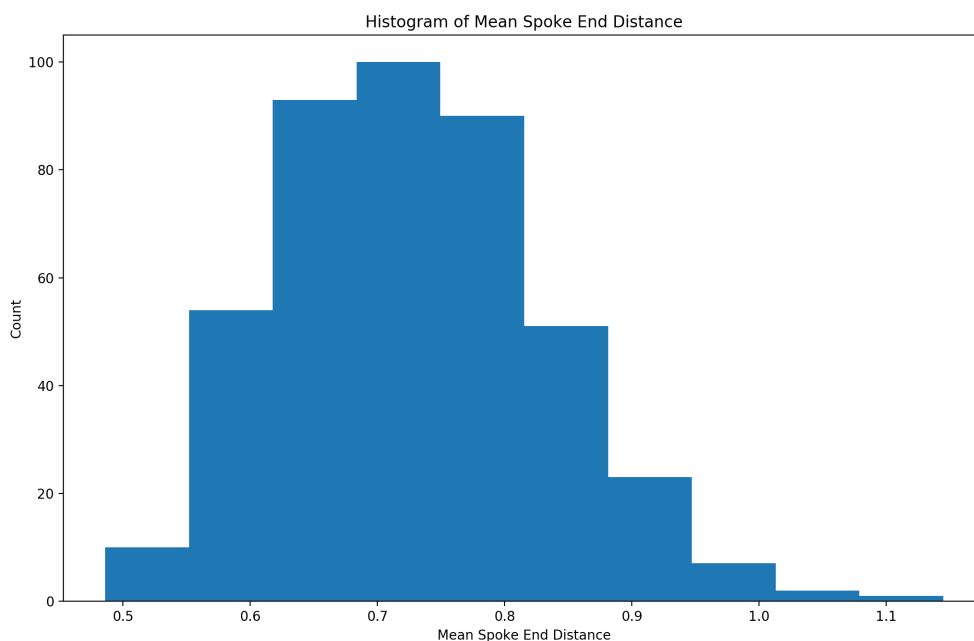


Figure 4.8: The histogram of the mean spoke end distance for the population of hippocampi.

From figure 4.8 the mean spoke end distance is less than a voxel; therefore, the fitted s-reps on average achieve sub-voxel accuracy in matching the original object boundary. The mean and the maximum of the mean spoke end distance is 0.73 voxels and 1.14 voxels respectively.

The distribution of spoke end distances for an s-rep of the caudate nucleus of an 6 months olds infant can be seen in figure 4.9 with the mean and the maximum being 0.66 voxels and 1.72 voxels respectively.




Figure 4.9: The histogram of the distances from each base spoke end of the fitted s-rep to the corresponding object boundary for a case in the population of 6-month olds' caudate nuclei.

Figure 4.10 shows the distribution of the mean spoke end distances of the fitted caudate nucleus s-reps. The mean and the maximum of the mean spoke end distance are 0.65 voxels and 0.70 voxels respectively.

In order to evaluate the method's general applicability in producing an s-rep of an object other than the hippocampus and the caudate nucleus, I applied the method to obtain the s-rep of the lateral ventricle of a neonate. Thin structures like the lateral ventricle have been challenging to adequately model with s-reps. Also, the lateral ventricle is a more strongly curving structure than both the hippocampus and the caudate nucleus. Figure 4.11 shows the s-rep fitted to one of the three lateral ventricles I fitted; in each one the fitted s-rep is capable of producing a smooth boundary, which indicates that the s-rep adequately represents the lateral ventricle's surface geometry.




Figure 4.10: The histogram of the mean spoke end distance for the population of caudate nuclei.




Figure 4.11: (a) Fitted s-rep of a neonate's left lateral ventricle using the proposed approach; (b) solid model implied by that s-rep.

Figure 4.12 shows the distribution of the chosen lateral ventricle s-rep's spoke end distances with the mean and the maximum being 0.54 and 1.40 respectively.




Figure 4.12: The histogram of the distances from each base spoke end of the fitted s-rep to the corresponding object boundary for a neonate's lateral ventricle.

In this section I empirically show that the proposed method is effective in producing a set of fitted s-reps for the hippocampus and the caudate nucleus. In addition, I demonstrate that the method's applicability is not limited to these two structures; it can be used to produce a reasonably well fitted lateral ventricle s-rep. However, the method failed to produce an s-rep for the rather complicated shaped mandible. Throughout MCF, the vertices of the deforming meshes get clustered causing numerical instability in computing the reverse MCF deformation.

4.5 Discussion and Conclusion

In this chapter I have presented a novel method to automatically generate a set of fitted s-reps with reasonably good spoke correspondence by first deriving a case-specific ellipsoid s-rep via MCF and then applying the deformation to warp the ellipsoid s-rep to match the original object boundary. It has produced a set of the reasonably well fitted s-reps for hippocampi and caudate nuclei in 6 months olds. In addition, the method produces a reasonably well fitted s-rep for the lateral ventricle that the previous s-rep fitting approach often fails to produce a satisfactorily fitted s-rep.

As it currently stands, the method fails to produce a fitted s-rep for the mandible. Mean curvature flow fails in two aspects. First, the vertices on the condyles get clustered around the poles throughout the surface evolution, which causes the reverse MCF deformation to fail. Secondly, the sharp bend at the chin fails to stably unbend failing to converge to an ellipsoid. I speculate that employing different geometric flows, *e.g.*, Wilmore Flow or Ricci Flow, together with remeshing can address these problems.

The fitted s-reps can be further improved in three aspects. First, the length of the spokes can be adjusted to tighten the s-rep's fit to the original object boundary. Secondly, the spoke directions can be adjusted to be more approximately orthogonal to the original object boundary. Lastly, different means of mapping the ellipsoid to the original object can be explored; especially, learning-based surface registration methods to learn reverse MCF mappings may improve performance.

The method currently does not handle the inherent singularity of MCF. A surface evolution via the mean curvature is known to have a singularity when the shape of the surface resembles that of a dumbbell. The method would fail in two aspects: 1) obtaining an ellipsoid s-rep of the best-fitting ellipsoid that resembles the deforming surface close enough and 2) computing the reverse MCF deformation so that the deformation field does not have any folds.

CHAPTER 5

Extensions of the S-rep to Model a Singular Point

5.1 Introduction

In the literature morphological deviation in the caudate nucleus (CN) between individuals with and without ASD has been reported (Wolff et al., 2013; Hazlett et al., 2012). Especially notable is the association of a core symptom of ASD, *i.e.*, restricted and repetitive behavior, with enlargement of the structure's global volume. However, detailed analysis on local morphological differences of the CN between the two groups is lacking.

The CN is a subcortical structure with a long tail that forms the dorsal striatum together with the putamen. However, the segmentation of the CN tends to be cut off at the tail due to a low contrast between the CN's tail and the surrounding tissue.

As it currently stands, an s-rep has difficulties in modeling an object with a singular point, *e.g.*, the CN. The singular point on the boundary of the object causes the skeletal surface to collapse to a point; this degeneracy of the skeletal surface causes a number of issues in modeling the surface geometry of the object of interest. In addition, the s-rep cannot properly model the singular point because of the s-rep's hard constraint that no spoke cross can cross each other; this hard constraint is too hard to maintain at the singular point.

To address this challenge in modeling neuroanatomical structures with a singular point via the s-rep, I introduce an additional geometric primitive dedicated to represent the singular point to account for the degeneracy of the skeletal surface. Then, I extend the current quad-based spoke interpolation method so that the interpolated boundary surface patches are C^2 continuous up to the singular point.

This chapter is organized as follows. First, I overview the underlying mathematics of the s-rep and the current quad-based spoke interpolation method. Then, I describe the method in detail as follows:

1. The geometric primitive to represent the singular point
2. The extension of the current spoke interpolation method to obtain the continuous s-rep
3. The method to properly Euclideanize a population of the s-reps augmented with the primitive

I demonstrate that the new primitive is beneficial to modeling the 6-month olds' caudate nuclei with s-reps.

5.2 Background

In this section I briefly overview background necessary to understand the proposed extension to the s-rep to model the singular point and its interpolation into a continuous s-rep. First, I overview the mathematical definition of a continuous s-rep. Then, I overview a geometric interpolation method to obtain an interpolated s-rep from a discrete legacy s-rep.

5.2.1 Mathematics of S-reps

A 3D, continuous s-rep is defined by two parametric functions:

- $p(u_1, u_2)$: a 2D, non-branching skeletal surface
- $S(u_1, u_2)$: a field of non-crossing spokes emanating from $p(u_1, u_2)$; each spoke approximately orthogonally intersects the object boundary

$S(u_1, u_2)$ can be viewed as a product of two additional parametric functions: $U(u_1, u_2)$, a unit normalized vector field of $S(u_1, u_2)$ and $r(u_1, u_2)$, a scalar field of $|S(u_1, u_2)|$, the distance from $p(u_1, u_2)$ to the object boundary along the $U(u_1, u_2)$. Accordingly, a point on the implied boundary of the s-rep can be expressed as $B(u_1, u_2) = p(u_1, u_2) + r(u_1, u_2)U(u_1, u_2)$.

This continuous s-rep is discretized to support various statistical tasks, *e.g.*, classification. The legacy discrete s-rep used in this work has a grid of $m \times n$ discretely sampled skeletal surface with

either two (on the interior) or three (along the fold curve of the skeleton) spokes emanating from the skeleton.

However, it is necessary to ensure that s-reps properly model the surface geometry of target objects prior to the classification. To assess how well the fitted s-rep matches the original object boundary, I need to be able to quantify the difference between the implied boundary of the s-rep and the object boundary. The implied boundary is obtained by continuously interpolating the base spokes of the s-rep.

5.2.2 Mathematics of Spoke Interpolation

Interpolating a discrete s-rep into a continuous s-rep is divided this into 1) a skeletal surface interpolation operation and 2) a spoke interpolation on the interpolated skeletal surface. The skeletal surface interpolation method uses standard polynomial-based methods.

The spoke interpolation method depends on first interpolating the spoke directions U and then interpolating the spoke lengths r . Our experience shows that this strategy is preferable because we have found a good way to interpolate the directions and have found that the overall result is then very robust against approximation errors in the spoke length interpolation.

Let the skeletal surface be parameterized by (u_1, u_2) , where both parameters are integers at the corners of a quadrilateral in the grid on which the discrete s-rep is specified. Thus, the discrete s-rep gives both r and U at these quadrilateral corners. Consider interpolation of the spoke directions U at any point $p(u_1, u_2)$ within any grid quadrilateral on the skeletal surface. Our plan for interpolation of r is based on a 2nd-order Taylor series, for which we need not only the spoke directions U but also their first and second order derivative values U_{u_i} and $U_{u_i u_i}$ for $i = 1, 2$ at arbitrary points in the quadrilateral. Spoke directions live on the unit 2-sphere \mathbb{S}^2 . Thus, the sort of finite difference calculations that must be used in order to compute U at our discrete skeletal points should be done on the sphere. These calculations are done by representing the discrete spokes U as unit quaternions and thus its derivatives with respect to u_i as derivatives on the sphere. Using these derivatives, Vicory applies the *squad* method (Shoemake, 1987) of interpolating quaternions to estimate the spoke direction U at an arbitrary point interior to a quadrangle of discrete points by fitting Bezier curves to the quaternions on the surface of the sphere. This approximation allows the computation of

not only the U values but also their directional derivatives of both first and second order in either u_1 or u_2 .

Given the ability to evaluate U and its derivatives in a quadrilateral, we need to interpolate the r values in a way consistent with skeletal geometry. Spokes can be written $S = rU$. The derivatives of the spoke at a skeletal location p with respect to a step in direction v in either of the two orthogonal directions u_1 or u_2 must follow $S_v = rU_v + r_vU$ from which it follows that $r_v = S_v \cdot U$. Also, $S_{vv} \cdot U = (S_v)_v \cdot U = r_{vv} + rU$. From this a Taylor series in the length d of a small step in direction v from a skeletal position p together with three forward distance derivative approximations yields the following expression.

$$r(p + dv) = \frac{1}{2} (S(p) + S(p + 2dv)) \quad (5.1)$$

Because the same mathematics works using a Taylor series in the backwards direction about $p + 2dv$, for symmetry and to reduce approximation error the results of the two versions should be averaged, yielding the final formula as $r(p + dv) = U(p + dv) \cdot \left(\frac{1}{2} (S(p) + S(p + 2dv)) \right) - \frac{d^2}{4} (S(p) \cdot U_{vv}(p) + S(p + 2dv) \cdot U_{vv}(p + 2dv))$

This formula allows computing the spoke half-way between two horizontally adjacent quadrilateral corners and using successive subdivision with it by halving as many times as necessary to get the desired small spacing in u_1 . Applying the same method separably (on the results) in the u_2 direction, yields a spoke at any successively subdivided point within the quadrilateral. Finally, since the method gives different results when you apply it first in u_2 and then in u_1 , we compute using both orders and average the results.

At a skeletal fold, the skeletal surfaces lack of smoothness prevents the direct application of the aforementioned method. We solve this problem by dilating the fold curve into a tube of very narrow radius, treating the spoke at the curve as its value at the place it intersects the tube, and then using the method for smooth surfaces to compute the continuation of the spoke to the object boundary.

5.3 Method

As previously mentioned, the current s-rep modeling framework is limited in geometrically modeling an object with a singular point because of the degenerate skeleton and the hard constraint of non-crossing spokes. I address this challenge as follows:

1. I introduce a new geometric primitive to represent the singular point.
2. I modify the current interpolation method to produce interpolated s-reps with the new primitive.
3. I introduce a statistical method to analyze a set of s-reps that contain the new primitive.

5.3.1 Singular Point Primitive

The radius of the curvature at the singular point of the CN is zero, which yields a skeletal surface that not only collapses to a point but also intersects the CN's boundary. Based on this observation, I explicitly represent the singular point as a point on the boundary with a direction.

5.3.2 Proposed Interpolation Method

Given the new primitive, I need to be able to produce a continuously interpolated s-rep for solid modeling of an object with a singular point, *e.g.*, the CN. Vicory's quad-based interpolation method (Vicory, 2016) is not designed to handle the degeneracy of the skeleton that I aim to model via the new primitive. I adopt a different interpolation strategy while following the general workflow of the current interpolation method. First, I describe how to interpolate a skeleton emanating from the base skeleton's end curve and converging to the singular point in its designated direction; the resulting skeleton is smooth up to the singular point. Then, I describe how I interpolate the spokes given the interpolated skeletal positions. Finally, I address how I interpolate the spokes that go around the crest in this context.

5.3.2.1 Skeleton Interpolation

In this section I describe how to construct a smooth, continuous skeleton up to the singular point. I construct the skeleton from a smoothly varying series of non-crossing curves. As shown in figure 5.1, these curves emanate from the respective positions interpolated along the base skeleton's end curve and converge to the singular point in its designated direction. I construct each individual curve as follows:

1. Interpolating the curve's starting point along the base skeleton's end curve
2. Determining an emanating direction of the curve at the starting point

3. Limiting the speed along that direction to enforce the non-crossing property of the curve

A starting point of each curve is first interpolated along the base skeleton's end curve. The point is interpolated from a cubic Hermite spline fitted to two adjacent skeletal points between which the point lies in s-rep's parametric space. Given the same parametrization of the skeleton in (u_1, u_2) , let the interpolated starting point along the end curve be $p(u_1^*, u_2^0)$, and let the two adjacent skeletal end points along the end curve be $p(u_1^i, u_2^0)$ and $p(u_1^{i+1}, u_2^0)$. Hermite interpolation requires 4 control values: two positional values of each point, $p(u_1^i, u_2^0)$ and $p(u_1^{i+1}, u_2^0)$, and two vectors $p_{u_1}(u_1^i, u_2^0)$ and $p_{u_1}(u_1^{i+1}, u_2^0)$; the partial derivatives are computed via finite differences. The vector H_c containing these control values is

$$H_c = (p(u_1^i, u_2^0), p(u_1^{i+1}, u_2^0), p_{u_1}(u_1^i, u_2^0), p_{u_1}(u_1^{i+1}, u_2^0))$$

Let $H(x) = (H_1(x), H_2(x), H_3(x), H_4(x))$ where the H_i s are the cubic Hermite spline basis functions:

$$\begin{aligned} H_1(x) &= 2x^3 - 3x^2 + 1 \\ H_2(x) &= -2x^3 + 3x^2 \\ H_3(x) &= x^3 - 2x^2 + x \\ H_4(x) &= x^3 - x^2 \end{aligned} \tag{5.2}$$

Then, $p(u_1^*, u_2^0)$ can be computed by the following equation:

$$p(u_1^*, u_2^0) = H(u_1^* - u_1^i) \cdot H_c^\top$$

The same spline-based interpolation strategy is used to obtain the curve that originates from $p(u_1^*, u_2^0)$ and converges to the singular point in the designated direction. As previously noted, the interpolation requires two positional values as well as the two vectors at $p(u_1^*, u_2^0)$ and the singular point; I need to specify the vector at $p(u_1^*, u_2^0)$ so that the fitted curve at $p(u_1^*, u_2^0)$ does not cross other curves.

The idea is to specify the emanating vectors at the discrete skeletal end points and to do so in a way that allows the emanating vectors at intermediate points to be interpolated via quaternion

splines. Given the two adjacent end skeletal points $p(u_1^i, u_2^0)$ and $p(u_1^{i+1}, u_2^0)$ between which the starting point lies, the direction of the vector is obtained from a quaternion-based interpolation of $p_{u_2}(u_1^i, u_2^0)$ and $p_{u_2}(u_1^{i+1}, u_2^0)$. These partial derivatives are computed via finite differences.

Given the direction of the vector at $p(u_1^*, u_2^0)$, I need to decide the speed along the emanating vector so that the produced curve does not cross other curves. Two curves at two distinct interpolated starting points can cross each other when the emanating vectors are facing each other; however, the crossing can be prevented if the speeds along the respective vectors are constrained to be less than the half of the arc length between the two starting points. A similar intuition is often used in image registration when obtaining a dense displacement field from a sparse set of displacement vectors. Based on this observation, I constrain the magnitude of emanating vectors for the $p(u_1^*, u_2^0)$ s as follows:

1. Finding the minimal arc length of the two adjacent base skeletal end points; let this value be L_{min}
2. Finding the maximal partial derivative of the skeleton at the discrete skeletal end points with respect to u_2 ; let this value be S_{max}
3. Computing a scale factor γ as a ratio of $\frac{L_{min}}{2}$ to S_{max}
4. Adjusting the speed along the emanating directions at the discrete skeletal end points by γ
5. Interpolating the speed along the intermediate emanating vectors

Figure 5.1 shows a skeleton of one of the CN s-rep with the new singular point primitive.




Figure 5.1: An interpolated skeleton of one of the CN s-rep with the new singular point primitive. The blue arrow seen at the lower left corner of the figure is the directional component of the singular point primitive.

5.3.2.2 Spoke Interpolation

In this section I describe how to obtain a continuous s-rep from a discrete s-rep given the singular point primitive. I construct a dense vector field of spokes on the skeleton by interpolating spokes along each skeletal curves. As shown in figure 5.2 spokes on each curve smoothly vary from the respective spokes interpolated along the base skeleton's end curve to the specified direction at the singular point. The spoke vectors along each skeletal curve is obtained as follows:

1. Interpolating the initial spoke along the base skeleton's end curve
2. Interpolating the desired spoke from the initial spoke to the singular point.

An initial spoke is first interpolated along the base skeleton's end curve. Let the initial spoke be $S(u_1^*, u_2^0)$ whose tail position is the interpolated end skeletal point $p(u_1^*, u_2^0)$. The spoke direction

$U(u_1^*, u_2^0)$ is estimated from a quaternion-based interpolation. The spoke length $r(u_1^*, u_2^0)$ is then interpolated using the aforementioned equation in section 5.2.2.

Given the end initial spoke $S(u_1^*, u_2^0)$, I need to interpolate the desired spoke $S(u_1^*, u_2^*)$ along the skeletal curve from $p(u_1^*, u_2^0)$ to the singular point. The interpolation of $S(u_1^*, u_2^*)$ follows the same workflow: 1) the interpolation of the spoke direction $U(u_1^*, u_2^*)$ that is obtained from a quaternion-based interpolation between the initial spoke direction $U(u_1^*, u_2^0)$ and the singular point direction and 2) the interpolation of the spoke length $r(u_1^*, u_2^*)$ that is obtained using the same expression presented in section 5.2.2. Figure 5.2 shows the densely interpolated spokes of one of the CN s-rep.




Figure 5.2: Densely interpolated spokes of the same CN s-rep previously seen in figure 5.1 with the singular point primitive. The magenta vectors denote up spokes, the cyan vectors denote down spokes, and the red vectors denote crest spokes.

5.3.3 Euclideanization of the Augmented S-reps

To classify HR-ASD infants and HR-Neg infants based on the caudate shape, I need to Euclideanize GOPs provided by CN s-reps augmented with the singular point primitive. In this section I describe

how to Euclideanize the augmented s-rep GOPs: the base s-rep GOPs and the singular point GOPs, *i.e.*, its boundary point and direction. First, I describe how to apply PNS to Euclideanize a set of the augmented s-reps. Then, I discuss how to make the Euclideanized GOPs of singular point primitive commensurate with the Euclideanized GOPs of the base s-rep for classification.

The Euclideanization of the augmented s-reps is very similar to that of the regular s-reps. First, the primitive's boundary point is concatenated with the base skeletal points, which yields a bigger point tuple to which PNS is applied. Then, the primitive's direction is handled by PNS just like spoke directions.

Of importance for any statistical analyses is to make each feature commensurate. In this context I need to make the Euclideanized primitive's direction commensurate with other Euclideanized GOPs. To be consistent with the commensuration of other directional primitives of the s-rep, I associate a width feature as a multiplicative scale factor with the Euclideanized direction of the singular point.

Consider a section of the continuously interpolated augmented s-rep from the base end skeletal curve to the singular point; the majority of the volume is concentrated at the base. Based on this observation, I associate the singular point primitive's direction with the geometric mean of the lengths of all base spokes along the base skeleton's end curve. Then, similar to how the Euclideanized GOPs of spoke directions are commensurated, I commensurate the Euclideanized GOP of the singular point primitive' direction by multiplying PNS-generated features by the geometric mean of end spokes' length.

5.4 Results

In this section I present couple of results of applying the proposed method to modeling the CN. The method currently requires a user to manually provide the boundary point and the direction for the singular point primitive. In this experiment, I manually specify the singular point. I specify the direction that 1) lies on the tangent plane of the base end skeleton and 2) points toward the chosen singular boundary point.

I investigate how well the interpolated CN s-reps with the singular point match the original CN boundary. In order to get a sense on the accuracy of the proposed method, I examine how well the interpolated spokes of the augmented CN s-reps match the original object boundary. I examine the

distribution of the average distances between each interpolated spoke end and the original boundary. The average distance is measured in voxel units where the resolution of the image is upsampled to $0.05 \times 0.05 \times 0.05\text{mm}$.

Figure 5.3 shows the histogram of the average spoke end distances of the augmented CN s-reps fitted to a population of 6-month olds' caudate nucleus; there are 32 caudates of HR-ASD infants and 137 caudates of HR-Neg infants.




Figure 5.3: The histogram of the average spoke end distance of the augmented CN s-rep fitted to a population of 6-month olds' caudate nuclei.

As can be seen from figure 5.3, the majority of the distances is less than 1: most of the spokes match the original boundary within 0.05 mm. The mean and the maximum distances are 0.40 voxels and 1.25 voxels respectively.

In addition, the newly added singular point primitive is beneficial to the discrimination of HR-ASD infants from HR-Neg infants; the classification based on the CN s-reps with the primitive yields 3% better performance than the one without the primitive. In subsequent chapter I shall describe the classification method in detail.

5.5 Conclusion and Discussion

In this chapter I have presented a novel method to represent an object with a singular point with an s-rep. I first introduce additional geometric primitives, *i.e.*, a boundary point and a direction, to explicitly represent the singular point. Then, I describe a method to obtain a solid model of the object by continuously interpolating the spokes of the s-rep augmented with the singular point primitive. I have demonstrated the method is useful in modeling the CN of 6-month old infants; I have shown that the added primitive is beneficial to discriminating the autistic infants and the non-autistic infants. I have shown some benefit of using the CN s-reps augmented with the singular point primitive in the context of autism classification.

The method currently requires a user to specify the boundary point and the direction. An optimization-based strategy can be used to refine both the singular point primitive and the base discrete s-rep to tightly fit the original object boundary.

As it currently stands, the method represents the cuspidal singular point. However, there are other work (Cohen et al., 2014) that model different types of the singular point based on parametrized boundary mesh representation. The underlying mathematics of their method can be adapted for s-rep-based modeling.

The proposed method can be adapted to interpolate s-rep spokes whose skeleton is a triangle mesh. To interpolate an unknown spoke on a triangle, the similar interpolation strategy can be used. Intermediate spokes can first be interpolated along each edge of the triangle. Then, the final spoke can be interpolated as a average of the spokes interpolated from each intermediate spokes.

CHAPTER 6

Non-Euclidean Shape Classification

6.1 Introduction

Binary classification of objects of interest based on medical imaging has been a common objective (e.g., (Kurtek et al., 2011; Gorczowski et al., 2010; Zhao et al., 2014)). Researchers often wish to determine if a subject has a disease or not based on geometric features of an anatomical structure from a medical image. Beyond simply providing a rule for classification is the desire to gain deeper scientific insights into phenomena underlying the disease.

These geometric features are often provided by object models and should be analyzed by statistical methods suitable for shapes. One of the most popular forms of shape representation is the *Point Distribution Model (PDM)* (e.g., (Cootes et al., 1995; Styner et al., 2006; Davies et al., 2003)). A boundary PDM is a tuple of boundary points on an object, with points corresponding across the training cases. Frequently, studies using PDMs capture shape variations through the statistical method of *Principal Component Analysis (PCA)* (Cootes et al., 1992, 1995), and classification is done using *Linear Discriminant Analysis (LDA)* or the *Support Vector Machine (SVM)* (Davies et al., 2003).

In this chapter I investigate the possible improvements in classification that can arise from two modifications in the above methods. The first is to statistically analyze the object representation data in the realization that, per (Kendall, 1984), even PDMs can be understood as lying on a curved manifold. I apply PNS-based Euclideanization for this purpose. The second modification I consider is to use rich s-rep GOPs.

The method of classification I propose is DWD on GOPs that are Euclideanized using PNS. I demonstrate that, both with PDMs and with s-reps, this statistical method produces more effective

classification than those making less use of the geometry of the manifold in which the representation lies.

I apply the proposed method to the problem of classifying 3D hippocampi and caudate nuclei as autistic or non-autistic based on their GOPs. I have evaluated my method on a dataset that consists of 143 non-autistic cases and 34 autistic cases. In this application I measure performance by calculating the area under the ROC curve (AUC). As mathematically expected, the results show that

- the proposed method based on s-reps is superior to the classification based on s-reps without Euclideanization.
- Euclideanization is also beneficial to the PDM-based classification; however, the observed improvement is subtle.
- Euclideanized shape-based classification is better than the conventional global volume-based classification; the observed improvement is more notable in the hippocampal shape-based classification than in the caudate shape-based classification.

This chapter is organized as follows. In section 6.2 I first describe the proposed classification method. In section 6.3 I provide details on the experimental setup and the classification evaluation method. Then, I present the result of the proposed classification method in section 6.4. Finally, I conclude this chapter with a discussion in section 6.5.

6.2 Classification Method

The novelty of the proposed classification method comes from the fact that I recognize that s-rep GOPs abstractly live on a curved manifold and that I appropriately take that into account during classification. The proposed classification method works as follows.

1. Apply PNS to Euclideanize GOPs that live on a sphere and commensurate those features to millimeters.
2. Learn the separation direction from these features concatenated with the originally Euclidean features in the training data using DWD.

3. Compute the function that maps values projected onto the separation direction to the probability of belonging to the autistic group based on Bayes' Theorem (figure. 6.1).
4. Classify each case in the test set based on the probabilities computed using the function from the previous step.

In the following subsections, I provide a detailed description of each step.

6.2.1 Euclideanization of s-reps and basis of the transformation between s-rep space and Euclidean space

As previously noted, a discrete s-rep has some spherical GOPs, i.e., each spoke's direction and the PDM formed by its skeletal sample points. I apply PNS separately to each spherical GOP, producing corresponding Euclideanized variables.

I consider both great subspheres and small subspheres at each iteration of PNS to Euclideanize spherical GOPs of the representation. Chi-squared based hypothesis testing was performed to decide which subspheres to use at each iteration of PNS. Along with the Euclideanized variables, PNS yields a polar system to be used as the basis of a transformation from the original s-rep space to the corresponding Euclidean space, and vice versa.

I concatenate the already Euclidean and Euclideanized variables and scale each so that they are commensurate. These variables form the feature space on which classifiers are trained and tested. I denote these concatenated variables as the composite data matrix.

6.2.2 Learning the Separating Direction

The composite matrix computed via PNS is the input to DWD. That method learns a feature space separating direction between the two classes, i.e., the autistic and the non-autistic group, via the training set of discrete s-reps Euclideanized as described in the previous section.

6.2.3 Computing the Function from Projected Feature Values to the Probability of being Autistic

Given a separation direction and a case with an unknown class label, my objective is to compute that case's posterior probability of belonging to the autistic group. Using Bayes' Theorem, I can express

this probability in terms of a prior and a likelihood of each class. I derive likelihood probabilities, i.e., the probability distributions of each class, given the s-rep features, by forming a pair of histograms each describing statistics of the respective class as seen in figure. 6.1.




Figure 6.1: Visualizations of the class likelihoods of hippocampi. The empirical histogram of the scalar projection of the non-autistic cases in the training set onto the separation direction is plotted using the blue dotted lines and with the Gaussian probability distribution as the blue solid curve. Similarly for the autistic class in green.

Using the trained polar system, I first transform the s-rep of interest into a point in feature space. Let d_X be the scalar value resulting from projecting that data point X onto the separation direction; let $\{d_{\text{autism}}\}$ be projection values of positive training examples, and let $\{d_{\text{non-autism}}\}$ be projection values of negative training examples. I form a pair of empirical histograms of d_{autism} and $d_{\text{non-autism}}$ as shown in figure 6.1. By treating d_{autism} and $d_{\text{non-autism}}$ as random variables, I derive a probability distribution for each class from the respective histograms.

Given the two histograms, I need to derive the class likelihood probabilities to compute the posterior probability. I fit Gaussians to the two empirical distributions as the respective class likelihood probability distributions. To simplify the calculation of the posterior, I examined if the two fitted Gaussians have a common standard deviation; the F-test failed to reject the null hypothesis that the two distributions are Gaussian with a common standard deviation. Given this outcome, the class likelihood distributions are the two Gaussians whose mean is the sample mean of the

respective histograms, and whose common standard deviation is obtained from the unbiased least square estimate of the pooled variance of both histograms.

With these two distributions, $p(d_X|\text{autism})$ and $p(d_X|\text{non-autism})$, I can infer a class label of an unknown case if the projection value of that case d_X is given. It can be formulated by using Bayes' Theorem as follows.

By Bayes' theorem,

$$p(\text{autism}|d_X) = \frac{p(\text{autism})p(d_X|\text{autism})}{p(\text{autism})p(d_X|\text{autism}) + (1 - p(\text{autism}))p(d_X|\text{non-autism})} \quad (6.1)$$

This can be reduced to

$$p(\text{autism}|d_X) = \frac{p(\text{autism})}{p(\text{autism})(1 - R(d_X)) + R(d_X)} \quad (6.2)$$

where

$$R(d_X) = \exp \left[-\frac{1}{2} \left\{ \left(\frac{d_X - \mu_{\text{non-autism}}}{\sigma} \right)^2 - \left(\frac{d_X - \mu_{\text{autism}}}{\sigma} \right)^2 \right\} \right] \quad (6.3)$$

where the pooled standard deviation estimate is computed as

$$\sigma^2 = \frac{(n_{\text{autism}} - 1)\sigma_{\text{autism}}^2 + (n_{\text{non-autism}} - 1)\sigma_{\text{non-autism}}^2}{(n_{\text{autism}} - 1) + (n_{\text{non-autism}} - 1)} \quad (6.4)$$

where n_{autism} denotes the number of observations for the autistic observations, σ_{autism} denotes the standard deviation of the scalar projections onto the direction for the autistic observations, and similarly for the non-autistics with $n_{\text{non-autism}}$ and $\sigma_{\text{non-autism}}$.

In summary, I compute the posterior $p(\text{autism}|d_X)$ given a scalar projection value d_X along the separation direction and a prior $p(\text{autism})$. The $p(\text{non-autism}|d_X)$ is the complement of $p(\text{autism}|d_X)$. Not only does this probability communicate intuitively to a user how certain a classification of a new case is but also basing it on parameterized probability distributions allows stable predictions in the tails of the distribution.

6.2.4 Classification based on Probability Produced by the Mapping Function

I decide the class label of an unknown case given a projected value d_X and the prior $p(\text{autism})$ by comparing $p(\text{autism}|d_X)$ and $p(\text{non-autism}|d_X)$.

6.3 Experimental Analysis

To evaluate the proposed method, I use repeated random hold-out so that I do not introduce bias in the testing procedure. I first randomly partition the positive example set into 10 roughly equal size subsets and likewise with the negative example sets. I set aside one of the subsets from each class for testing and used the remaining subsets to collect statistics necessary for the classification method.

I have conducted 2,000 rounds of random hold-out with the prior $p(\text{autism})$ being set to 1/4; in the literature, one out of four infants at high familial risk for ASD is later diagnosed with ASD. Given the prior and a trained separation direction, each hold-out round yields the posterior probability $p(\text{autism}|d_X)$ for each test sample; I collect the $p(\text{autism}|d_X)$ s for the positive test samples separately from the negative test samples. I obtain two distributions of $p(\text{autism}|d_X)$ for the positive (negative respectively) test samples across the hold-out rounds. Then, I derive a ROC curve from these two $p(\text{autism}|d_X)$ distributions. I report the areas under these ROC curves (AUC) in section 6.4.

S-rep based method compared to boundary PDM-based methods

The boundary PDM is a common approach to represent a shape via a collection of points along the object's boundary. I wish to compare the qualities of classification when hippocampal shapes are represented by s-reps vs. boundary PDMs to see if the rich geometric information provided by s-reps increases discriminative power over classification based on boundary information. And similarly for caudate nuclei.

In order to make a fair comparison between boundary PDMs and s-reps, I need boundary PDMs that can be compared directly to s-reps. Recall that s-reps are a collection of spoke vectors pointing from skeletal sample points to the object's surface and that s-reps are fitted such that the spoke vectors

are in approximate correspondence across all cases in the training population; I form boundary PDMs from these spoke endpoints. I will denote these boundary PDMs as srep-PDMs.

I classify srep-PDMs in two different ways. First, we applied our DWD-based method directly to the point coordinate features. Second, in order to understand advantages of the Euclideanization on that type of the shape data, we applied PNS to the point tuples to yield Euclideanized features as well as a commensurated scale, and then we applied our DWD-based method to these features. The same validation strategy used with s-reps was applied to each of these methods. For each method, I report the final AUC in table 6.1 and in table 6.2.

6.4 Results

In this section I present the empirical results of applying the proposed method to the problem of classifying autistic and non-autistic 6-month-old infants at high risk of ASD.

Methods	AUC
s-reps + PNS + DWD	0.6400
s-reps + DWD	0.6123
boundary srep-PDMs + PNS + DWD	0.6062
boundary srep-PDMs + DWD	0.6050
global volume + DWD	0.5560
random guessing	0.5000

Table 6.1: Table of an AUC of the ROC of the selected classification methods and the pure random guessing in classifying the hippocampus of the autistic infants and the non-autistic infants. The result shows that as expected, the Euclideanized s-rep-based classification performs the best.

Methods	AUC
s-reps + PNS + DWD	0.5708
s-reps + DWD	0.5419
boundary srep-PDMs + PNS + DWD	0.5400
boundary srep-PDMs + DWD	0.5365
global volume + DWD	0.5372
random guessing	0.5000

Table 6.2: The parallel results to the table 6.1 in classifying the caudate nucleus. The same conclusion observed in the table 6.1 can be drawn.

Table 6.1 reports the performance of all the aforementioned methods in terms of the AUC in classifying the hippocampus of the autistic infants and the non-autistic infants. For s-reps, classifica-

tion using Euclideanization is superior to that without Euclideanization. For PDMs, classification using Euclideanization yields a higher AUC, but the difference is subtle. With Euclideanization both forms of model yield similar if not better classification than the common approach in the literature, a volume-based classification of autistic hippocampi. S-rep-based classification with Euclideanization is superior to all the other methods of classifying autistic infants from non-autistic infants based on hippocampal shape.

Table 6.2 reports the AUC of the same methods in classifying the caudate nucleus. The same conclusion as the hippocampal shape-based classification can be drawn.

6.5 Conclusion and Discussion

In this chapter I have presented a novel classification method that recognizes that rich geometric information is provided by s-reps and that relevant shape information does not live in Euclidean space. I have shown benefit to the classification performance when all of the GOPs of either s-reps or boundary PDMs derived from s-reps are Euclideanized via PNS analysis.

In the context of the autism classification based on 1) the hippocampal shape and 2) the caudate shape, I have shown that both the s-rep-based classification and the PDM-based classification provide an advantage over a volume-based classification; therefore, I claim that shape information adds additional discriminative power. I have also shown improvement when using s-reps over b-PDMs when both GOPs are appropriately Euclideanized; I also show that local object directions and local object width add discriminative power. I conclude that 1) shape descriptions add additional discriminative power over global volume, and 2) local object directions and local object width that s-reps provide add additional discriminative power over boundary position.

The proposed method yields a separating direction through the pooled backward mean in the feature space of the Euclideanized s-reps. Each point on this vector can be used to generate an s-rep using the polar system. Viewing the sequence of the s-reps as an animation yields understanding of the shape changes between the two classes. Figure 6.2 shows selected frames from the sequence. Our group's paper on hypothesis testing on shapes using PNS-Euclideanization (Schulz et al., 2016) analyzes the discriminability between these two classes of hippocampi locality by locality and GOP by GOP.

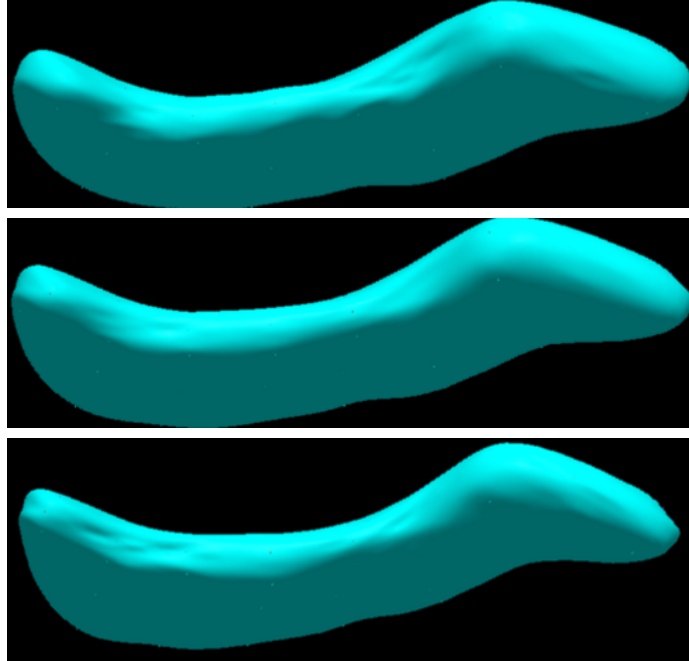


Figure 6.2: Selected frames from the sequence of the s-reps while walking along the separation direction through the pooled backward mean from the autism class to the non-autism class. Viewing the sequence as a looping movie makes the local shape changes between the two classes more noticeable.

There are still some further questions to be investigated.

- As previously noted, deviations in the global volume of neuroanatomical structures, *e.g.*, the caudate nucleus, have been observed in individuals with ASD [refs]. Measures such as the intracranial volume (ICV) and the body length (Hazlett et al., 2017) have been used to normalize volumes in these structures to remove factors not related to ASD, *e.g.*, the nutrition and the genetic factor. As it currently stands, such volume normalizations have not been performed for analyses described in this work; it would be interesting to see if the same conclusions can be drawn after normalizing volumes of hippocampi and caudate nuclei.
- As for the caudate shape-based classification reported in the table 6.2, the boundary point of the singular point primitive described in the section 5 has not been included for the boundary PDM-based analysis. However, I expect that adding the boundary position of the singular point to the current point tuple would not change the result significantly. In my separate experiment not included in this work, I measured the caudate classification performance when using boundary PDMs derived from spherical harmonics (SPHARM-PDMs), which includes the

singular point's boundary position. The performance of SPHARM-PDM-based classification without PNS-based Euclideanization is 0.5400, which does not differ significantly from that of s-rep-PDM-based classification reported in the table 6.2.

- To see if our results extend to other anatomic objects and diseases, we would like to apply the method to different application problems, e.g., classification of Alzheimer patients based on shapes of the neuroanatomical structures.
- In Euclideanizing a spoke direction using PNS, we apply PNS separately because we are making the naive assumption that each direction is independent. However, because an object surface is continuous and smooth, each direction is highly correlated to its neighbors. We would like to produce a Euclideanization method that reflects this correlation. Also, others are suggesting methods for statistical analysis directly on the curved shape-feature space manifold (Eltzner et al., 2015; Arnaudon et al., 2017), and it would be interesting to evaluate classification methods using these ideas.
- As previously mentioned in section 2 the method we used to achieve spoke correspondence in s-reps across the training set could be improved. In separate work, reported in (Tu et al., 2018), we created a method to improve the correspondence by spoke shifting on each training case, so as to minimize an entropy measure. This entropy measure reflects both shape probability distribution tightness and uniformity of coverage of the spokes in each training case. The shape probability distribution used is derived from the same PNS approach used in this paper. The correspondence was shown to be improved in a set of lateral ventricles and in a subset of the hippocampi used in this paper. It would be interesting to see whether the classification of hippocampi and caudate nuclei could be improved using these correspondence improved s-reps. Finally, (Tu et al., 2018) also showed improved PDM correspondence when using the spoke tips as the PDM as compared to a PDM derived from spherical harmonics and then improved in correspondence by the entropy-based method of (Cates et al., 2006) . This further justifies our decision to use the s-rep derived PDM instead of SPHARM-PDM in the classification study reported in this paper.

- Other work is in progress comparing different statistical methods against DWD. It would be interesting to see how DWD for our purpose compares to other statistical methods such as Random Forests, and Deep Learning.
- It would be also interesting to measure the relative power of classification via other shape representations that have been used in the anatomic shape analysis literature, including but not limited to parameterized surface representations used in (Kurtek et al., 2012; Jermyn et al., 2012; Bauer et al., 2010, 2012; Durrleman et al., 2014), deformation fields used in (Lancaster et al., 2003; Villalon-Reina et al., 2012), the spherical harmonic coefficients used in (Gerig et al., 2001), spherical wavelet coefficients used in (Nain et al., 2007), and atlas deformation representations such as LDDMM momentum (Beg et al., 2005; Miller et al., 2002; Wang et al., 2007).

CHAPTER 7

Non-Euclidean Temporal Shape Difference Classification

7.1 Introduction

So far in this dissertation I have discussed how to classify a set of subcortical structures in the context of the early diagnosis of infants at high familial risk to develop ASD. In the previous chapter, I have described the classification method that uses rich s-rep GOPs whose non-Euclidean nature is handled by PNS. I have shown that the method achieves 64% and 51% in classifying HR-ASD infants and HR-Neg infants at 6 months of age based on either the hippocampus s-reps or the caudate nucleus s-reps.

However, an important question is whether the classification of the autistic infants and the non-autistic infants can be improved if multiple s-rep GOPs are used to learn the classification rule. Specifically, how much does the ASD classification benefit by incorporating temporal changes in the shape as compared to the shape alone?

As previously shown, the s-rep based classification benefits from PNS-based Euclideanization. Therefore, I need to properly handle the non-Euclidean nature of the GOPs of the s-rep pair. Among many possible Euclideanization methods I consider how to apply PNS to obtain temporal s-rep differences of the hippocampus and the caudate nucleus.

The remainder of this chapter is organized as follows: First, I describe how to apply PNS to Euclideanize two distinct s-reps, *i.e.* the hippocampus and the caudate nucleus. In addition, I describe how to obtain Euclideanized temporal differences of an object s-reps. I conclude this chapter with the results followed by the discussion.

7.2 Method

In this section I describe the non-Euclidean temporal shape difference classification method. As previously shown, the s-rep based classification benefits from Euclideanization due to highly non-Euclidean nature of s-rep GOPs. Therefore, I need to Euclideanize the GOPs of the s-rep pairs prior to learning the classification rule. First, I describe how to use PNS to obtain temporal differences of s-reps at two time points. Then, I describe the classification scheme given the Euclideanized GOPs.

7.2.1 Euclideanization of Temporal Differences of an S-rep

As mentioned in chapter 2, ASD is a lifelong neurodevelopmental disorder. In the literature brain developmental trajectories of individuals with ASD are different from those of typically developing individuals. Would the developmental trajectories of the implicated brain structures such as the hippocampus and the caudate nucleus have predictive power?

Given a set of object pairs at two time points, I use PNS-based Euclideanization to obtain temporal shape differences of an object between the two time points. I adopt the same Euclideanization strategy used in [Jared's dissertation], which works as follows:

1. I apply PNS-based Euclideanization on the union of all training shapes across the classes and across the time points to obtain a common polar system.
2. Given the common polar system, I derive the Euclideanized temporal shape differences by subtracting the starting shapes from the ending shapes of their Euclideanized values.

7.2.2 Classification

Given the set of temporal differences of s-rep GOPs, I use DWD to learn to classify HR-ASD infants from HR-Neg infants. The proposed classification method is evaluated as described in section 6.3 in classifying HR-ASD infants and HR-Neg infants based on hippocampi or caudate nuclei.

7.3 Results

In this section I present results of classifying HR-ASD infants and HR-Neg infants based on temporal shape changes from 6 months to 12 months in hippocampi and caudate nuclei. I compare

discriminative power of the temporal shape difference against 6-month shape only and 12-month shape only.

Methods	AUC
temporal difference s-reps + PNS + DWD	0.5427
6 month s-reps + PNS + DWD	0.6400
12 month s-reps + PNS + DWD	0.5837
random guessing	0.5000

Table 7.1: Table of AUCs of classifying HR-ASD and HR-Neg based on 1) temporal shape difference in hippocampi at 6 months and 12 months (the top row), 2) 6 months hippocampi shapes (the 2nd row), 3) 12 months hippocampi shapes (the 3rd row), and 4) the pure random guessing (the bottom row).

Methods	AUC
temporal difference s-reps + PNS + DWD	0.5041
6 month s-reps + PNS + DWD	0.5087
12 month s-reps + PNS + DWD	0.5605
random guessing	0.5000

Table 7.2: Table of AUCs of classifying HR-ASD and HR-Neg based on 1) temporal shape difference in caudate nuclei at 6 months and 12 months (the top row), 2) 6 months caudate nuclei shapes (the 2nd row), 3) 12 months caudate nuclei shapes (the 3rd row), and 4) the pure random guessing (the bottom row).

Table 7.1 reports the performance of classifying HR-ASD infants and HR-Neg infants based on 1) temporal shape differences of hippocampi between 6 months and 12 months, 2) 6 months hippocampi shapes, and 3) 12 months hippocampi shapes. As for the ASD classification based on the hippocampal shapes, the hippocampal shapes at 6 months seems to be the most discriminative predictor.

Table 7.2 reports the AUC of the same methods in the same classification task using caudate nuclei shapes. When classifying HR-ASD infants and HR-Neg infants based on caudate shapes, the caudate shapes at 12 months seems to be the most discriminative factor.

For both of the ASD classification tasks, the temporal shape difference seems to be the least discriminative feature.

7.4 Conclusion and Discussion

In this chapter I have adapted the idea of shape difference statistics from (Vicory, 2016) to classify HR-ASD infants and HR-Neg infants based on temporal shape changes. I believe that this idea can be further extended to perform classification based on developmental trajectories of shapes involving more than two time points.

I have investigated whether a temporal shape difference between 6 months and 12 months is an important predictor in classifying HR-ASD infants and HR-Neg infants. The empirical results suggest that the temporal shape difference is not as discriminative as compared to 6 months shapes or 12 months shapes. However, I speculate that the information captured by the temporal shape differences is complementary to that captured by 6 months shapes or 12 months shapes alone; I expect the prediction to improve if the temporal difference and either of 6 months or 12 months are used together.

However, as noted in (Vicory, 2016), when the shapes undergo a relatively large deformation, a pooling-based PNS Euclideanization strategy would be less effective, since the learned basis of Euclideanization would not be particularly representative of shapes at either time points. Instead, one could employ an Euclideanization strategy in which PNS is applied to Euclideanize shapes for each class at each time point separately and then these Euclideanized shape features are transported to a common place on the shape manifold.

CHAPTER 8

Conclusions & Discussion

In this chapter I provide a summary of my contributions in section 8.1. Then, I conclude this chapter with a further discussion of my contributions and potential future work in section 8.2.

8.1 Summary of Contributions

This dissertation proposes a shape classification method that uses rich Geometric Object Properties (GOPs) of s-reps, Principal Nested Spheres (PNS) based Euclideanization, and Distance Weighted Discrimination (DWD) to classify autistic and non-autistic infants at high familial risk of Autism Spectrum Disorder (ASD) based on shapes of hippocampi or caudate nuclei. In my driving biomedical problem I demonstrate the following:

1. The classification benefits from rich localized GOPs provided by s-reps. As shown in chapter 6, the AUC of the s-rep-based classification is 5% higher than that of PDM-based classification in the context of the autism classification based on the hippocampal shape.
2. The classification benefits from PNS-based Euclideanization for both s-rep-based classification and PDM-based classification. Therefore, recognizing shapes as a point on a curved manifold is important. As shown in table 6.1 and table 6.2, the AUC of the Euclideanized s-rep-based classification is 4.5% and 2.3% higher than the classification method treating s-rep GOPs as Euclidean features. The benefit is also shown for the PDM-based classification; the PNS-based Euclideanization yields 0.2% and 0.7% higher AUC.
3. The classification benefits from shape information; both Euclideanized PDM-based classification and Euclideanized s-rep-based classification perform better than volume-based classification. As shown in chapter 6, the Area Under the Receiver Operator Characteristics

(ROC) curve (AUC) of the Euclideanized s-rep-based classification is 15% higher than that of the volume based classification in the hippocampal shape-based classification. The AUC of the Euclideanized Point Distribution Model (PDM) based classification is 0.5% higher than that of the volume based classification in the caudate shape-based classification.

The major methodological contributions of this dissertation are as follows:

1. Development of a fully automatic procedure to obtain a set of s-reps fitted to each target object boundary data with reasonably good correspondence in spokes. As detailed in chapter 4, I described a general framework that utilizes Mean Curvature Flow (MCF) to establish a mapping from the target object boundary to an ellipsoid. Given a set of ellipsoids for each target boundary, I derive a mean ellipsoid. An s-rep for the mean ellipsoid is obtained from the known analytical expression. Then, the mean ellipsoid s-rep is mapped to each individual ellipsoid and each deformed mean ellipsoid s-rep is mapped back to the original target object boundary. I argue that the resulting set of fitted s-reps have reasonable spoke correspondence because each s-rep is mapped from a common reference s-rep. The method is shown to be effective in generating a set of fitted s-reps for hippocampi, caudate nuclei, and lateral ventricle. However, the method failed to produce a fitted s-rep of a mandible, which has a rather complicated-shape. I discuss below improvements to the flow that may ameliorate this problem.
2. Development of an extension to the current s-rep modeling framework for appropriately representing an object with a singular point, *e.g.*, the caudate nucleus. As pointed out in chapter 5, the current s-rep modeling is limited in representing objects with a cuspidal singular point. To address this limitation, I used an additional point on the boundary and a direction to represent the singular point. Then, I described a method to obtain a continuously interpolated s-rep augmented with the singular point primitive. As shown in chapter 5, the primitive is beneficial in modeling the sharp tail of the caudate nucleus. In addition, the caudate classification task also benefits from the additional primitive; the AUC of the classification is 3% higher.
3. Development of a non-Euclidean classification method that uses PNS to obtain the Euclideanized s-rep GOPs on which the DWD classifier is trained. In chapter 6 I showed

that shape information adds discriminative power by showing that both PDM-based classification and s-rep-based classification are superior to the volume-based classification; the AUC of the s-rep-based classification is 15% higher than that of the volume-based classification, and the AUC of the PDM-based classification is 0.5% higher than that of the volume-based classification. I demonstrated that both PDM-based classification and s-rep-based classification benefit from PNS-based Euclideanization; the benefit of the PNS-based Euclideanization was most notable for the s-rep-based classification. The observed benefits to the AUC of the s-rep-based classification when Euclideanized by PNS are 4.5% and 2.3% for the hippocampal shape classification and the caudate shape classification respectively. Last but not least, I further showed that the classification benefits from rich localized Euclideanized s-rep GOPs, *i.e.*, local object widths and local object directions; as shown in table 6.1, the AUC of the s-rep-based classification is the highest.

4. A novel technique to compare different shape classification methods. In section 6 I employed a repeated hold-out strategy to measure the performance of the shape classification methods, *i.e.*, the s-rep-based classification and the PDM-based classification. For each fold I first use the training samples to obtain a polar system from PNS, which yields Euclideanized GOPs from which DWD learns a class separation direction. From the test sample in each fold I use the training histograms to derive posterior probabilities of test samples being autistic separately for the positive test samples and for the negative test samples. Then, I compute the final AUC by sweeping the collected posterior probabilities. The method allowed a more credible comparison of shape classification methods than comparing classification methods on the basis of AUC estimated from a single hold-out round.
5. Development of a non-Euclidean classification method for classifying a temporal pair of s-reps in which the input feature tuple to DWD is formed by taking the difference of the Euclideanized s-rep GOPs for each time point. In chapter 7 I studied whether the temporal shape difference is more discriminative than the shape alone in the context of autism classification. I have adopted the idea of shape difference statistics proposed in [Jared's dissertation] to obtain temporal changes in s-reps; s-reps are first Euclideanized by pooling across the classes and the time points, and then the temporal shape differences are obtained by taking differences

of the Euclideanized s-reps at two-time points. I examined discriminative powers of the temporal changes of the hippocampal shape and of the caudate shape between 12 months and 6 months. This produced interesting results that could not have been produced without our comparison technique. The temporal shape differences of the hippocampus and the caudate nucleus do not seem to be as predictive as compared to 6-month shape alone or 12-month shape alone. However, I speculate that information captured by the temporal shape difference is complementary to information captured by 6-month shape alone or 12-month shape alone; therefore, I expect the overall prediction can be improved if the temporal shape difference is used together with either 6-month shape or 12-month shape.

The above contributions support the following thesis:

Thesis: *Classification of medically imaged neuroanatomical structures based on their geometric object properties benefits from the following:*

1. *An object model providing local object width and orientation in addition to positional information*
2. *An object modeling procedure that efficiently yields geometric correspondence across a population*
3. *A classification method capable of learning a classification rule on a curved manifold*

In addition to the above methodological contributions, I have also accomplished the following engineering contributions:

1. Integration of Pablo, the main piece of software used to fit and visualize s-reps into 3D Slicer for easier distribution of the s-rep modules to the medical image analysis community. During my involvement in NIH funded project called *slicerSALT* (Slicer Shape AnaLysis Toolbox), I achieved following software engineering accomplishments:
 - Modernization of the s-rep modules.
 - Redesign of the computer representation of the discrete s-rep. I redesigned the computer representation of a discrete s-rep as networks of spokes using the VTK file format. The redesigned s-reps allow modeling objects with different topologies including the spherical

(slabular), the quasi-tubular, and the toroidal. In addition, the redesigned s-reps are more consistent with the underlying skeletal mathematics at the skeletal fold.

- Implementation of a module that converts a legacy s-rep into the redesigned s-rep.
 - Implementation of a module that visualizes a legacy s-rep in 3D Slicer. My colleague, Zhiyuan Liu has recently extended the visualization module to visualize the redesigned representation of s-reps.
 - The code will be publicly released by the time of the 2018 MICCAI shape workshop and will be available to download at <https://github.com/Kitware/SlicerSALT>
2. Improvement of the numerical stability of estimating surface normals and principal curvatures from the input representation of the object. In Pablo objects are represented as the implicit representation of the object by a signed distance image. Pablo used the finite difference method to estimate isovalue surface normals and principal curvatures. I notably improved the numerical stability in estimating surface normals and principal curvatures by fitting a local B-spline patch from which those geometric quantities are computed analytically.

8.2 Discussion and Future Work

8.2.1 S-reps Fitting

8.2.1.1 Optimization

A set of the fitted s-reps obtained from the method described in chapter 4 was improved via spoke length refinement to tighten their fits to the original object boundary. The spoke length refinement is formulated as a posterior optimization problem. The optimization is carried out via conjugate gradient descent where the gradient is numerically estimated. Recently a gradient-free optimization method called *NEWUOA* (Powell, 2006) is being employed to replace conjugate gradient descent. Some preliminary results of using *NEWUOA* to refine lengths of the spokes of the fitted s-reps are encouraging. My colleague, Zhiyuan Liu, is currently implementing *NEWUOA* to optimize skeletal point locations, spoke directions, and spoke lengths all together.

8.2.2 Different Geometric Flows for Automatic S-reps

As noted in section 4, my MCF-based method fails to produce the fitted s-rep for a rather complicated shaped mandible. MCF fails in two aspects: 1) throughout the surface evolution vertices on the condyles end up being clustered on the ellipsoid, and 2) the sharp bend at the jaw does not unbend fast enough before introducing singularities on other places on the mesh. These problems could potentially be solved by employing different types of flow, *e.g.*, Willmore Flow or Ricci Flow, combined with remeshing to prevent vertices from clustering. Interesting alternative curvature measures to use in evolving the surface are C and S described in (Koenderink, 1990).

Another aspect of the method that can be improved is the use of different means of computing the reverse MCF deformation that maps from an ellipsoid to the original object. The TPS transformation is not guaranteed to yield a diffeomorphic transformation, which is why the method currently works with a small time step. Evolving the surface by a small time step increases the number of the pairwise TPS transformations to be computed. Instead, different diffeomorphic surface registration methods can be explored to allow faster convergence to an ellipsoid. Recently a learning-based method has been reported (Yang et al., 2017) to produce a transformation that is comparable to the one produced by traditional LDDMM with a fraction of the computational cost. Using such a learning-based method could improve the computational performance.

8.2.3 Modeling Objects with a Singular Point

In chapter 5 I introduced an additional geometric primitive to model an object that has a sharp narrowing tail such as the caudate nucleus with an s-rep. The primitive can be used to model other sharp points such as the tip of the horn of the lateral ventricle or a sharp place on condyles of a mandible. When modeling these sharp places, the emanating directions need to be interpolated along a different side of the skeleton.

The method currently requires a user to specify the boundary point and the direction. An optimization-based strategy can be used to refine both the singular point and the s-rep fit.

As it currently stands, the primitive is designed to model a cuspidal singular point; the object narrows very fast toward the cuspidal singular point. However, there exist different types of singular points, *e.g.*, a corner. There are other factors to consider when modeling an object with a corner. First,

the rate at which an object gets narrowed is not as fast as compared to the one with a cuspidal singular point. Therefore, one may need to specify the desired rate of narrowing at each base skeletal points when modeling this type of singular point. In addition, different emanating curves can converge to the corner in different directions. Therefore, one may also need to specify a set of desired converging directions at the corner.

8.2.4 Non-Euclidean Shape Classification

In chapter 6 I have shown that the autism classification based on hippocampal shape or on caudate shape benefits from rich localized s-rep GOPs when the non-Euclidean nature of s-rep GOPs is handled by PNS. However, there are three aspects that can be further investigated: 1) object models, 2) Euclideanization methods, and 3) classification methods.

In this work I have compared the s-rep-based classification against the PDM-based classification. However, an interesting comparison on the same data would be the s-rep-based classification against other shape classification methods that use object models, *e.g.* a tuple of normals, spherical harmonics coefficients, spherical wavelets, and atlas deformation representation. Preliminary work has been conducted in comparing the s-rep-based classification against the classification method that uses a tuple of normals representation called Square Root Normal Fields (*SRNF*) (Jermyn et al., 2012) in a different application problem. The s-rep-based classification seems to be more discriminative than the *SRNF*-based classification in that particular application problem. However, it still remains to be seen if this is the case in the autism classification problem. Another interesting comparison would be comparing our discrete s-reps against the continuous medial representation called *c-mreps* (Yushkevich et al., 2006).

In this dissertation I only considered PNS as a means of handling the non-Euclidean nature of s-rep GOPS. However, different means of Euclideanization can be further explored. Benjamin Eltzner (Eltzner et al., 2015) has proposed Polysphere PCA (PPCA) in order to account for the correlation directly among the Cartesian product of spheres, *i.e.*, polyspheres. There is an on-going work to compare the s-rep-based classification Euclideanized by PNS against the one Euclideanized by PPCA. An interesting perspective of analyzing non-Euclidean GOPs in all dimensions at once has been proposed by Xavier Pennec (Pennec, 2015); a Euclideanization method based on this perspective would be an interesting future project.

Instead of learning a classification rule based on Euclideanized GOPs, an alternative approach is to directly work on the shape manifold. (Arnaudon et al., 2017) has proposed a method of estimating the shape probability distribution directly on the shape manifold obtained via a stochastic process. Then, two shape probability distributions can be estimated for each class. Given the two shape probability distributions, a shape can be classified on the basis of class likelihoods. In addition, there has been an early work on developing DWD that directly works on the original manifold by Suman Sen (Sen, 2008).

In this dissertation I have used DWD to learn a classification rule in the Euclideanized feature space. However, different classification methods, *e.g.*, Random Forest and Deep Learning, can be explored. Especially, deep learning based methods have been quite successful in various medical imaging tasks, *e.g.*, image segmentation, classification, and registration [refs]. Deep learning has been reported to learn and extract lower dimensional, task-relevant features from original high dimensional features. (Hazlett et al., 2017) used deep learning to learn a lower dimensional latent space of a high dimensional cortical shape feature vector. Recently, there has been work to apply deep learning for non-Euclidean data (Masci et al., 2015; Boscaini et al., 2016; Bronstein et al., 2017). For Alzheimer classification, (Mostapha et al., 2018) have applied a geodesic Convolutional Neural Network (geodesic CNN) to learn discriminative features from cortical brain surfaces with cortical shape features defined at each vertex. An interesting future direction is to investigate whether a deep learner can learn and extract discriminative Euclidean features from highly non-Euclidean data such as s-rep GOPs.

8.2.5 Understanding Morphological Changes Associated with Growth

In this dissertation I have adapted the idea of shape difference statistics from (Vicory, 2016) for classifying autistic and non-autistic infants based on temporal shape changes. I speculate that the idea can be extended to perform a classification based on developmental trajectories of shapes involving more than two time points.

However, as noted in (Vicory, 2016), when the shapes undergo relatively large deformations between the two time points, pooling-based PNS Euclideanization would not be particularly representative of respective classes and therefore would not be an effective means of Euclideanizing the s-rep GOPs. Instead, a Euclideanization method that computes respective polar systems for each

class at each time point and then transports the polar systems to a common place, *e.g.*, a mean of all shapes, on the manifold can be used. This approach can be also viewed as a type of domain adaptation problem.

Instead of Euclideanizing temporal shape pairs, longitudinal shape analysis frameworks (Nithammer et al., 2011; Fletcher, 2013) that recognize the non-Euclidean nature of shapes can be utilized to understand morphological differences in developmental trajectories between the two groups of subjects.

8.2.6 Further Clinical Studies

There are several interesting future directions with respect to the driving application problem behind this dissertation. Firstly, as previously noted in section 6.5, the effects of the global volume of hippocampi and caudate nuclei in the autism classification could have been controlled using measures such as the intracranial volume (ICV) and the body length to remove factors not related to ASD, *e.g.*, the nutrition. It would be interesting to see if the same conclusion can be drawn after controlling the effect of the global volumes in these structures. Secondly, in the literature, the gender of ASD subjects is predominantly male; therefore, it would be interesting to investigate the gender effect in classifying HR-ASD infants and HR-Neg infants. Lastly, it would be an interesting future project to use s-rep GOPs to predict clinical measures such as the Autism Diagnostic Observation Schedule (ADOS).

8.2.7 Conclusion

While this chapter discusses limitations and potential future directions for proposed methodologies, the proposed shape classification method that uses DWD to learn a class separation based on rich s-rep GOPs that are Euclideanized by PNS seems sound and effective in discriminating autistic and non-autistic infants based on shapes of the hippocampus or the caudate nucleus. It would be interesting to see if the conclusion that I have drawn in this dissertation holds for different shapes and different shape classification tasks.

BIBLIOGRAPHY

- Arnaudon, A., Holm, D. D., Pai, A., and Sommer, S. (2017). A stochastic large deformation model for computational anatomy. In *International Conference on Information Processing in Medical Imaging*, pages 571–582. Springer.
- Bauer, M., Harms, P., and Michor, P. W. (2010). Sobolev metrics on shape space of surfaces. *arXiv preprint arXiv:1009.3616*.
- Bauer, M., Harms, P., and Michor, P. W. (2012). Almost local metrics on shape space of hypersurfaces in n-space. *SIAM Journal on Imaging Sciences*, 5(1):244–310.
- Beg, M. F., Miller, M. I., Trouvé, A., and Younes, L. (2005). Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision*, 61(2):139–157.
- Boscaini, D., Masci, J., Rodolà, E., Bronstein, M. M., and Cremers, D. (2016). Anisotropic diffusion descriptors. In *Computer Graphics Forum*, volume 35, pages 431–441. Wiley Online Library.
- Bouix, S., Pruessner, J. C., Louis Collins, D., and Siddiqi, K. (2005). Hippocampal shape analysis using medial surfaces. *Neuroimage*, 25(4):1077–1089.
- Bronstein, M. M., Bruna, J., LeCun, Y., Szlam, A., and Vandergheynst, P. (2017). Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42.
- 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.
- Cohen, E., Haimes, R., and Riesenfeld, R. (2014). A curvature smooth lofting scheme for singular point treatments. In *International Conference on Curves and Surfaces*, pages 129–150. Springer.
- Cootes, T. F., Taylor, C. J., Cooper, D. H., and Graham, J. (1992). Training models of shape from sets of examples. In *BMVC92*, pages 9–18. Springer.
- Cootes, T. F., Taylor, C. J., Cooper, D. H., and Graham, J. (1995). Active shape models-their training and application. *Computer vision and image understanding*, 61(1):38–59.
- Cortes, C. and Vapnik, V. (1995). Support-vector networks. *Machine Learning*, 20(3):273–297.
- Damon, J. and Marron, J. S. (2013). Backwards Principal Component Analysis and Principal Nested Relations. *Journal of Mathematical Imaging and Vision*.
- Davies, R. H., Twining, C. J., Allen, P. D., Cootes, T. F., and Taylor, C. J. (2003). Shape discrimination in the hippocampus using an MDL model. *Information processing in medical imaging : proceedings of the ... conference*, 18:38–50.
- Duda, R. O., Hart, P. E., and Stork, D. G. (2012). *Pattern classification*. 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.

- El Karoui, N. et al. (2010). The spectrum of kernel random matrices. *The Annals of Statistics*, 38(1):1–50.
- Eltzner, B., Jung, S., and Huckemann, S. (2015). Dimension reduction on polyspheres with application to skeletal representations. In *International Conference on Networked Geometric Science of Information*, pages 22–29. Springer.
- Fletcher, P. T. (2013). Geodesic regression and the theory of least squares on riemannian manifolds. *International journal of computer vision*, 105(2):171–185.
- Friedman, J., Hastie, T., and Tibshirani, R. (2001). *The elements of statistical learning*, volume 1. Springer series in statistics New York, NY, USA:.
- Gerig, G., Styner, M., Jones, D., Weinberger, D., and Lieberman, J. (2001). Shape analysis of brain ventricles using spharm. In *Mathematical Methods in Biomedical Image Analysis, 2001. MMBIA 2001. IEEE Workshop on*, pages 171–178. IEEE.
- Gorcowski, K., Styner, M., Jeong, J. Y., Marron, J. S., Piven, J., Hazlett, H. C., Pizer, S. M., and Gerig, G. (2010). Multi-object analysis of volume, pose, and shape using statistical discrimination. *IEEE transactions on pattern analysis and machine intelligence*, 32(4):652–61.
- 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.
- Hazlett, H. C., Poe, M. D., Lightbody, A. A., Styner, M., MacFall, J. R., Reiss, A. L., and Piven, J. (2012). Trajectories of early brain volume development in fragile x syndrome and autism. *Journal of the American Academy of Child & Adolescent Psychiatry*, 51(9):921–933.
- Hong, J., Vicory, J., Schulz, J., Styner, M., Marron, J. S., and Pizer, S. M. (2016). Non-euclidean classification of medically imaged objects via s-reps. *Medical image analysis*, 31:37–45.
- Jermyn, I. H., Kurtek, S., Klassen, E., and Srivastava, A. (2012). Elastic shape matching of parameterized surfaces using square root normal fields. In *European conference on computer vision*, pages 804–817. Springer.
- Jung, S., Dryden, I. L., and Marron, J. S. (2012). Analysis of principal nested spheres. *Biometrika*, 99(3):551–568.
- Jung, S., Marron, J. S., and Pizer, S. M. (2010). A backward generalization of pca for image analysis. In *International Symposium on the Occasion of the 25th Anniversary of McGill University Centre for Intelligent Machines*, pages 111–124.
- Kendall, D. G. (1984). Shape manifolds, procrustean metrics, and complex projective spaces. *Bulletin of the London Mathematical Society*, 16(2):81–121.
- Koenderink, J. J. (1990). *Solid shape*, volume 2. Cambridge Univ Press.
- Kurtek, S., Klassen, E., Ding, Z., Avison, M. J., and Srivastava, A. (2011). Parameterization-invariant shape statistics and probabilistic classification of anatomical surfaces. In *Biennial International Conference on Information Processing in Medical Imaging*, pages 147–158. Springer.

- Kurtek, S., Klassen, E., Gore, J. C., Ding, Z., and Srivastava, A. (2012). Elastic geodesic paths in shape space of parameterized surfaces. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 34(9):1717–1730.
- Lancaster, J. L., Kochunov, P. V., Thompson, P. M., Toga, A. W., and Fox, P. T. (2003). Asymmetry of the brain surface from deformation field analysis. *Human brain mapping*, 19(2):79–89.
- Lewis, J. D., Evans, A., Pruett, J., Botteron, K., Zwaigenbaum, L., Estes, A., Gerig, G., Collins, L., Kostopoulos, P., McKinstry, R., et al. (2014). Network inefficiencies in autism spectrum disorder at 24 months. *Translational psychiatry*, 4(5):e388.
- Marron, J. S., Todd, M. J., and Ahn, J. (2007). Distance-Weighted Discrimination. *Journal of the American Statistical Association*, 102(480):1267–1271.
- Masci, J., Boscaini, D., Bronstein, M., and Vandergheynst, P. (2015). Geodesic convolutional neural networks on riemannian manifolds. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 37–45.
- Miller, M. I., Trouvé, A., and Younes, L. (2002). On the metrics and euler-lagrange equations of computational anatomy. *Annual review of biomedical engineering*, 4(1):375–405.
- Mostapha, M., Kim, S., Wu, G., Zsembik, L., Pizer, S., and Styner, M. (2018). Non-euclidean, convolutional learning on cortical brain surfaces. In *Biomedical Imaging (ISBI 2018), 2018 IEEE 15th International Symposium on*, pages 527–530. IEEE.
- Nain, D., Haker, S., Bobick, A., and Tannenbaum, A. (2007). Multiscale 3-d shape representation and segmentation using spherical wavelets. *Medical Imaging, IEEE Transactions on*, 26(4):598–618.
- Niethammer, M., Huang, Y., and Vialard, F.-X. (2011). Geodesic regression for image time-series. In *International conference on medical image computing and computer-assisted intervention*, pages 655–662. Springer.
- Pennec, X. (2015). Barycentric subspaces and affine spans in manifolds. In *International Conference on Networked Geometric Science of Information*, pages 12–21. Springer.
- Powell, M. J. (2006). The newuoa software for unconstrained optimization without derivatives. In *Large-scale nonlinear optimization*, pages 255–297. Springer.
- Schulz, J., Pizer, S. M., Marron, J. S., and Godtliobsen, F. (2016). Non-linear hypothesis testing of geometric object properties of shapes applied to hippocampi. *Journal of Mathematical Imaging and Vision*, 54(1):15–34.
- Sen, S. K. (2008). *Classification on Manifolds*. PhD thesis, University of North Carolina at Chapel Hill.
- Shoemake, K. (1987). Quaternion calculus and fast animation, computer animation: 3-d motion specification and control. Siggraph.
- Styner, M., Lieberman, J. A., Pantazis, D., and Gerig, G. (2004). Boundary and medial shape analysis of the hippocampus in schizophrenia. *Medical image analysis*, 8(3):197–203.
- Styner, M., Oguz, I., Xu, S., Brechbühler, C., Pantazis, D., Levitt, J. J., Shenton, M. E., and Gerig, G. (2006). Framework for the statistical shape analysis of brain structures using spharm-pdm. *The insight journal*, (1071):242.

- Tu, L., Styner, M., Vicory, J., Elhabian, S., Wang, R., Hong, J., Paniagua, B., Prieto, J. C., Yang, D., Whitaker, R., et al. (2018). Skeletal shape correspondence through entropy. *IEEE transactions on medical imaging*.
- 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*, 151:72–79.
- Vicory, J. (2016). *Shape deformation statistics and regional texture-based appearance models for segmentation*. PhD thesis, The University of North Carolina at Chapel Hill.
- Villalon-Reina, J., Prasad, G., Joshi, S., Jalbrzikowski, M., Toga, A., Bearden, C., and Tompson, P. (2012). Statistical analysis of maximum density path deformation fields in white matter tracts. In *Medical Image Computing and Computer-Assisted Intervention Workshop on Novel Imaging Biomarkers for Alzheimer’s Disease and Related Disorders*, pages 198–209.
- Wang, J., Vachet, C., Rumble, A., Gouttard, S., Ouziel, C., Perrot, E., Du, G., Huang, X., Gerig, G., and Styner, M. (2014). Multi-atlas segmentation of subcortical brain structures via the autoseg software pipeline. *Frontiers in neuroinformatics*, 8.
- Wang, L., Beg, F., Ratnanather, T., Ceritoglu, C., Younes, L., Morris, J. C., Csernansky, J. G., and Miller, M. I. (2007). Large deformation diffeomorphism and momentum based hippocampal shape discrimination in dementia of the alzheimer type. *Medical Imaging, IEEE Transactions on*, 26(4):462–470.
- Wolff, J. J., Gerig, G., Lewis, J. D., Soda, T., Styner, M. A., Vachet, C., Botteron, K. N., Elison, J. T., Dager, S. R., Estes, A. M., et al. (2015). Altered corpus callosum morphology associated with autism over the first 2 years of life. *Brain*, 138(7):2046–2058.
- Wolff, J. J., Hazlett, H. C., Lightbody, A. A., Reiss, A. L., and Piven, J. (2013). Repetitive and self-injurious behaviors: associations with caudate volume in autism and fragile x syndrome. *Journal of neurodevelopmental disorders*, 5(1):12.
- Yang, X., Kwitt, R., Styner, M., and Niethammer, M. (2017). Quicksilver: Fast predictive image registration—a deep learning approach. *NeuroImage*, 158:378–396.
- Yushkevich, P. A., Zhang, H., and Gee, J. C. (2006). Continuous medial representation for anatomical structures. *IEEE transactions on medical imaging*, 25(12):1547–1564.
- Yushkevich, P. A. and Zhang, H. G. (2013). Deformable modeling using a 3d boundary representation with quadratic constraints on the branching structure of the blum skeleton. In *Proceedings of the 23rd International Conference on Information Processing in Medical Imaging, IPMI’13*, pages 280–291, Berlin, Heidelberg. Springer-Verlag.
- Zhao, Q., Okada, K., Rosenbaum, K., Kehoe, L., Zand, D. J., Sze, R., Summar, M., and Linguraru, M. G. (2014). Digital facial dysmorphology for genetic screening: Hierarchical constrained local model using {ICA}. *Medical Image Analysis*, 18(5):699 – 710.