Elastic Shape Analysis of Three-Dimensional Objects

Synthesis Lectures on Computer Vision

Synthesis Lectures on Computer Vision is edited by Gérard Medioni of the University of Southern California and Sven Dickinson of the University of Toronto. The series publishes 50–150 page publications on topics pertaining to computer vision and pattern recognition. The scope will largely follow the purview of premier computer science conferences, such as ICCV, CVPR, and ECCV. Potential topics include, but not are limited to:

- Applications and Case Studies for Computer Vision
- Color, Illumination, and Texture
- Computational Photography and Video
- Early and Biologically-inspired Vision
- Face and Gesture Analysis
- Illumination and Reflectance Modeling
- Image-Based Modeling
- Image and Video Retrieval
- Medical Image Analysis
- Motion and Tracking
- Object Detection, Recognition, and Categorization
- Segmentation and Grouping
- Sensors
- Shape-from-X
- Stereo and Structure from Motion
- Shape Representation and Matching

- Statistical Methods and Learning
- Performance Evaluation
- Video Analysis and Event Recognition

Elastic Shape Analysis of Three-Dimensional Objects
Ian H. Jermyn, Sebastian Kurtek, Hamid Laga, and Anuj Srivastava
2017

The Maximum Consensus Problem: Recent Algorithmic Advances
Tat-Jun Chin and David Suter
2017

Extreme Value Theory-Based Methods for Visual Recognition
Walter J. Scheirer
2017

Data Association for Multi-Object Visual Tracking
Margrit Betke and Zheng Wu
2016

Ellipse Fitting for Computer Vision: Implementation and Applications
Kenichi Kanatani, Yasuyuki Sugaya, and Yasushi Kanazawa
2016

Computational Methods for Integrating Vision and Language
Kobus Barnard
2016

Background Subtraction: Theory and Practice
Ahmed Elgammal
2014

Vision-Based Interaction
Matthew Turk and Gang Hua
2013

Camera Networks: The Acquisition and Analysis of Videos over Wide Areas
Amit K. Roy-Chowdhury and Bi Song
2012

Deformable Surface 3D Reconstruction from Monocular Images
Mathieu Salzmann and Pascal Fua
2010

Boosting-Based Face Detection and Adaptation
Cha Zhang and Zhengyou Zhang
2010

Image-Based Modeling of Plants and Trees
Sing Bing Kang and Long Quan
2009

Elastic Shape Analysis of Three-Dimensional Objects

Ian H. Jermyn, Sebastian Kurtek, Hamid Laga, and Anuj Srivastava

www.morganclaypool.com

ISBN: 9781681730271 paperback
ISBN: 9781681730288 ebook

DOI 10.2200/S00785ED1V01Y201707COV012

A Publication in the Morgan & Claypool Publishers series
SYNTHESIS LECTURES ON COMPUTER VISION

Lecture #12
Series Editors: Gérard Medioni, *University of Southern California*
Sven Dickinson, *University of Toronto*
Series ISSN
Print 2153-1056 Electronic 2153-1064

Elastic Shape Analysis of Three-Dimensional Objects

Ian H. Jermyn
Durham University

Sebastian Kurtek
Ohio State University

Hamid Laga
Murdoch University and University of South Australia

Anuj Srivastava
Florida State University

SYNTHESIS LECTURES ON COMPUTER VISION #12

MORGAN & CLAYPOOL PUBLISHERS

ABSTRACT

Statistical analysis of shapes of 3D objects is an important problem with a wide range of applications. This analysis is difficult for many reasons, including the fact that objects differ in both geometry and topology. In this manuscript, we narrow the problem by focusing on objects with fixed topology, say objects that are diffeomorphic to unit spheres, and develop tools for analyzing their geometries. The main challenges in this problem are to register points across objects and to perform analysis while being invariant to certain shape-preserving transformations.

We develop a comprehensive framework for analyzing shapes of spherical objects, i.e., objects that are embeddings of a unit sphere in $\mathbb{R}^3$, including tools for: quantifying shape differences, optimally deforming shapes into each other, summarizing shape samples, extracting principal modes of shape variability, and modeling shape variability associated with populations. An important strength of this framework is that it is *elastic*: it performs alignment, registration, and comparison in a single unified framework, while being invariant to shape-preserving transformations.

The approach is essentially Riemannian in the following sense. We specify natural mathematical representations of surfaces of interest, and impose Riemannian metrics that are invariant to the actions of the shape-preserving transformations. In particular, they are invariant to reparameterizations of surfaces. While these metrics are too complicated to allow broad usage in practical applications, we introduce a novel representation, termed square-root normal fields (SRNFs), that transform a particular invariant elastic metric into the standard L^2 metric. As a result, one can use standard techniques from functional data analysis for registering, comparing, and summarizing shapes. Specifically, this results in: pairwise registration of surfaces; computation of geodesic paths encoding optimal deformations; computation of Karcher means and covariances under the shape metric; tangent Principal Component Analysis (PCA) and extraction of dominant modes of variability; and finally, modeling of shape variability using wrapped normal densities.

These ideas are demonstrated using two case studies: the analysis of surfaces denoting human bodies in terms of shape and pose variability; and the clustering and classification of the shapes of subcortical brain structures for use in medical diagnosis.

This book develops these ideas without assuming advanced knowledge in differential geometry and statistics. We summarize some basic tools from differential geometry in the appendices, and introduce additional concepts and terminology as needed in the individual chapters.

KEYWORDS

elastic Riemannian metric, shape model, shape metric, elastic registration, shape summary, modes of shape variability

Contents

Preface . . . . **xiii**

Acknowledgments . . . . **xv**

1 Problem Introduction and Motivation **1**

1.1 Problem Area: 3D Shape Analysis 1
1.2 General Goals and Challenges 2
1.3 Past Approaches and Their Limitations 6
1.4 Our Approach: Elastic Shape Analysis 9
1.5 Organization of This Book 10
1.6 Notation 11

2 Elastic Shape Analysis: Metrics and Representations **13**

2.1 Shapes 13
2.2 Elastic Shape Analysis 14
2.2.1 Encoding of Registration 14
2.2.2 Riemannian Metric and Optimal Registration 15
2.2.3 Geometric Invariance 18
2.3 Background: Elastic Framework for Curves 19
2.3.1 Elastic Metric for Curves 20
2.3.2 Geometric Invariance 23
2.3.3 Summary of Elastic Framework for Curves 24
2.4 Elastic Framework for Surfaces 25
2.4.1 Square-Root Map 25
2.4.2 Generalizing the Elastic Metric for Curves 27
2.4.3 Elastic Metric for Surfaces 28
2.4.4 Reduced Elastic Metric: Square-Root Normal Field 31
2.4.5 Geometric Invariance 34
2.4.6 SRNF Inversion Problem 34
2.5 Summary and Next Steps 45
2.6 Bibliographic Notes 46

3 Computing Geometrical Quantities 47

3.1 Computing in Shape Space 48
3.1.1 Optimal Registration and Alignment 49
3.1.2 Optimal Deformation 50
3.1.3 Putting it All Together 51
3.1.4 Simplifying the Computations Using SRNFs 52
3.2 Registration and Alignment Using SRNFs 53
3.2.1 Optimization over the Rotation Group 54
3.2.2 Optimization over the Reparameterization Group 54
3.3 Geodesic Computation Techniques on General Manifolds 60
3.3.1 Geodesic Computation via Path-Straightening 60
3.3.2 Geodesic Computation via Shooting 61
3.4 Elastic Geodesic Paths between Surfaces using Pullback Metrics 63
3.4.1 Path-straightening under Pullback Metrics 64
3.4.2 Shooting Geodesics under SRNF Pullback Metric 67
3.5 Elastic Geodesic Paths between Surfaces using SRNF Inversion 69
3.5.1 Geodesics using SRNF Inversion 70
3.5.2 Parallel Transport in SRNF Space 71
3.6 Elastic Geodesic Path Examples 71
3.6.1 Discretization 72
3.6.2 Path-straightening 72
3.6.3 Shooting Method 74
3.6.4 SRNF Inversion 75
3.7 Summary and Next Steps 76
3.8 Bibliographic Notes 79

4 Statistical Analysis of Shapes 81

4.1 Statistical Summaries of 3D Shapes 81
4.1.1 Pullback Metric Approach 81
4.1.2 SRNF Inversion Approach 85
4.2 Statistical Models on Shape Spaces 88
4.2.1 Tangent Space and Pullback Metric Approach 88
4.2.2 SRNF Inversion Approach 91
4.3 Clustering and Classification 91
4.4 Bibliographic Notes 95

5 Case Studies Using Human Body and Anatomical Shapes . . . 97
5.1 Clustering and Classification . . . 98
5.1.1 Attention Deficit Hyperactivity Disorder (ADHD) Classification . . . 98
5.1.2 Clustering of Identity and Pose of Human Body Shapes . . . 100
5.2 Geodesic Deformation . . . 100
5.2.1 Geodesics . . . 100
5.2.2 Deformation Transfer . . . 107
5.2.3 Reflection Symmetry Analysis and Symmetrization . . . 109
5.3 Statistical Summaries of Shapes . . . 110
5.3.1 Means and Modes of Variation . . . 110
5.3.2 Random Sampling from Shape Models . . . 111
5.4 Bibliographic Notes . . . 113

6 Landmark-driven Elastic Shape Analysis . . . 115
6.1 Problem Statement . . . 116
6.2 Landmark-guided Registration . . . 117
6.2.1 Initial Registration using Landmarks . . . 118
6.2.2 Registration using Landmark-constrained Diffeomorphisms . . . 119
6.2.3 Landmark-constrained Basis for Registration . . . 120
6.3 Elastic Geodesics under Landmark Constraints . . . 121
6.3.1 Illustration of Geodesic Paths . . . 122
6.3.2 Evaluation of Performance and Computational Cost . . . 124
6.4 Landmark-constrained 3D Shape Atlas . . . 125
6.5 Bibliographic Notes . . . 127

A Differential Geometry . . . 129
A.1 Differentiable Manifolds and Tangent Spaces . . . 129
A.2 Riemannian Manifolds, Geodesics, and the Exponential Map . . . 131
A.2.1 Geodesics . . . 133
A.2.2 Exponential Map . . . 134
A.3 Lie Group Actions and Quotient Spaces . . . 135

B Differential Geometry of Surfaces . . . 141

C Spherical Parametrization of Triangulated Meshes . . . 145
C.1 Conformal Spherical Mapping . . . 146
C.2 Coarse-to-fine Minimal Stretch Embedding . . . 148

D Landmark Detection 151

D.1 Landmark Detection using Heat Kernel Signatures 151

D.2 Landmark Correspondences 152

Bibliography 155

Authors' Biographies 167

Preface

Shape analysis of *three-dimensional objects* is a fast-growing discipline in its own right. Advances in technology for 3D scanners and 3D printers, and giant leaps in storage and transfer of large datasets, have focused attention on the need to develop tools for shape analysis of whole objects. Consequently, researchers have accelerated their efforts toward developing efficient mathematical representations and related computational tools for 3D shape analysis. This research has progressed in several directions, one of which relates to Riemannian frameworks and elastic shape analysis. This direction is the focus of this book.

One of the most difficult problems in comparing shapes, especially when seeking a comprehensive statistical solution, is the registration of points across objects. Such a registration is necessary, for instance, when measuring deformations, computing shape averages, or finding principal modes of variability in shape data. Many past and current methods presume this problem away, i.e., they assume that an optimal registration is already available, but that assumption is often not true in practice. Elastic shape analysis is a branch of shape analysis that performs both registration and shape comparison in a single, unified framework using a Riemannian approach. This book is an exposition on recent developments in mathematical representations, metrics, and computational solutions for elastic shape analysis of surfaces. To understand this approach, one needs a broad array of tools, ranging from geometry and calculus to finite-element analysis and computer programming.

This book is intended for a researcher or a graduate student in computer science, applied mathematics, statistics, biology, or a related discipline. It assumes a background in linear algebra and advance calculus, and will certainly be easier for those who have taken an introductory course in differential geometry. We have tried to keep the exposition self-contained by including some background material in appendices and throughout chapters as needed. However, this book is not intended to replace textbooks on the differential geometry of surfaces, as our focus is more on computational solutions in shape analysis.

Ian H. Jermyn, Sebastian Kurtek, Hamid Laga, and Anuj Srivastava
July 2017

Acknowledgments

Some of the research presented here was developed in collaboration with Prof. Eric Klassen of the Florida State University, Prof. Zhaohua Ding of Vanderbilt University, and Dr. Qian Xie of Amazon Inc. We are very grateful for their collaboration.

This research was supported in part by NSF grants: DMS 1621787, CCF 1617397, and DMS 1208959 to Anuj Srivastava, DMS 1613054 to Sebastian Kurtek, and Murdoch University's New Staff Startup Grant Scheme to Hamid Laga. Anuj Srivastava gratefully acknowledges support from a Durham University IAS Senior Research Fellowship (EU grant FP7-609412) and the Durham University Department of Mathematical Sciences.

The authors would like to thank: Nils Hasler for providing us with the 3D human shape models; Raif Rustamov and Maks Ovsjanikov for the discussion about functional maps and for sharing their data and results; the authors of the 2007 SHape REtrieval Contest (SHREC07), Watertight Models Track, and the authors of the TOSCA dataset for making their databases publicly available.

Ian H. Jermyn, Sebastian Kurtek, Hamid Laga, and Anuj Srivastava
July 2017

CHAPTER 1

Problem Introduction and Motivation

1.1 PROBLEM AREA: 3D SHAPE ANALYSIS

Shape is an important physical property that characterizes the external appearance of natural and man-made objects. It plays a central role in understanding and analyzing the roles of objects in their larger environments. For example, in the case of anatomical and biochemical objects, their shapes are important predictors of their functionality within the larger complex bio-systems in which they are situated. Understanding differences between shapes and modeling the variability within and across shape classes are, thus, fundamental problems and constitute building blocks to solutions in many application areas, ranging from computer vision and computer graphics to biology and anatomy. For instance, a study of the shapes of 3D anatomical structures and their growth patterns is of particular interest in understanding physiological abnormalities that may be linked to alterations in these shapes. Similarly, in the field of computer graphics, the growing availability of 3D models via online repositories has re-focused research efforts toward data-driven techniques for generating novel 3D shapes from existing ones. Data-driven techniques are also gaining momentum in the field of 3D reconstruction and modeling from images, range scans, and noisy point clouds. This process, however, requires powerful mathematical representations and efficient computer algorithms for capturing, modeling, and exploring the shape variability in vast collections of 3D data.

Shape analysis is a relatively old topic in computer vision, with papers and methods going back several decades. However, the early years of the new millennium saw a renewed focus on the area. This focus, which was application-oriented and data-driven, brought in new directions and tools. While the new interest was fueled by many factors, the most prominent was the increasing availability of large datasets of 3D shapes, especially in the fields of computer vision and medical imaging. It was also propelled by increases in computational power and storage, a growing interest in Riemannian methods, and a favorable atmosphere for the confluence of ideas from geometry and statistics. As a result, researchers developed novel approaches, based on mathematical tools that were new to this community, and made them practical using elegant computational solutions. The goal of this textbook is to cover some recent advances in analyzing the shape of 3D objects, with an emphasis on Riemannian frameworks and statistical analysis.

Even with the restriction to 3D objects, the degree of possible shape variability is enormous. This variability can arise in several ways. Topology can vary: shapes may have several pieces, or have different numbers of handles. Even if two objects share the same topology, they may have different geometries. For example, a deformed ball is topologically identical to a unit sphere but has a different geometry. Finally, objects may have different sizes, orientations, and positions. More specifically, the objects of interest may be observed under different scales or spatial configurations. For example, brain substructures observed in magnetic resonance images (MRIs) naturally vary in this way due to a subject's position in the scanner and the size of their individual substructure. In this book, we are interested in studying geometrical variability only, and thus we restrict our attention to objects with the same topology: we will consider only objects that consist of one piece with no handles. The boundaries of these objects are genus-0 surfaces, and can be viewed as embeddings of a sphere in $\mathbb{R}^3$. Despite the restriction, this is a rich class of shapes that includes the objects found in many of the important applications mentioned above. We note also that the techniques developed here can be adapted to objects with other topologies including quadrilateral, hemispherical, and cylindrical ones.

1.2 GENERAL GOALS AND CHALLENGES

What are the main tasks or tools required for shape analysis? Depending upon the application, there are many potential goals one could formulate. However, it seems reasonable to have a core set of tasks that any comprehensive method for statistical shape analysis must be able to accomplish. The following is a list of such tasks.

1. **Shape Metric and Deformations.** Perhaps the most fundamental task in shape analysis is quantifying the difference in the shapes of two objects. To accomplish this mathematically, we define a metric on an appropriate "shape space" that gives a "distance" between any two shapes. To give meaning to this distance, it is important to be able to associate to it an "optimal deformation." By definition, a "deformation" is a path of shapes leading from one shape to the other. We can use the metric to define the length of such a path; an optimal deformation is then the deformation of minimal length. If the length of the optimal deformation is equal to the distance between the two shapes, then we can think of the distance as measuring the "size" of the deformation needed to go from one shape to another. This requirement places a restriction on the original metric; modulo some extreme cases, it becomes possible to represent it using a "Riemannian metric," about which we will have a lot more to say in the sequel. In that context, the optimal deformation is known as a "geodesic path." Figure 1.1 shows an example of a geodesic path between two surfaces, computed under a particular metric to be described later.

2. **Shape Summary.** Given a set of objects, such as the ones shown in Figure 1.2, the method should be able to provide summary statistics, and in particular the mean and the covariance, of their shapes. The covariance can further be used to find the dominant modes of variability in the given set of shapes. One needs a formal mathematical representation of shapes, equipped with a proper shape metric, and associated computational tools, to

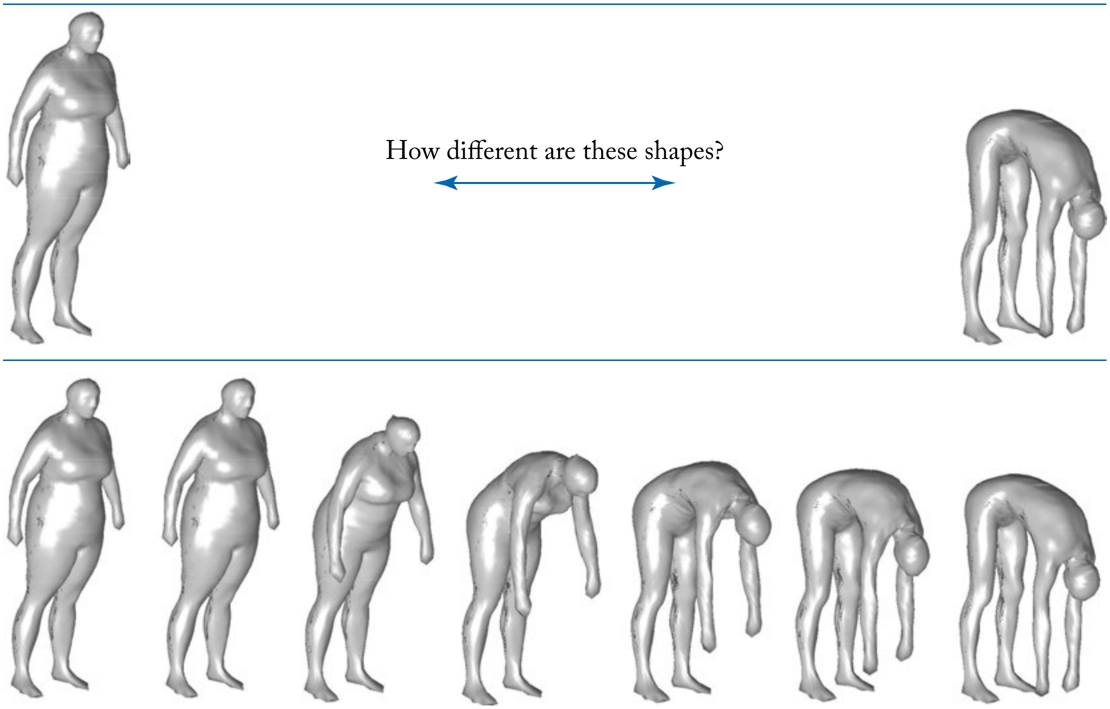

Figure 1.1: Given two surfaces, our goal is quantify the difference in their shapes and find an optimal deformation (geodesic) from one to the other.

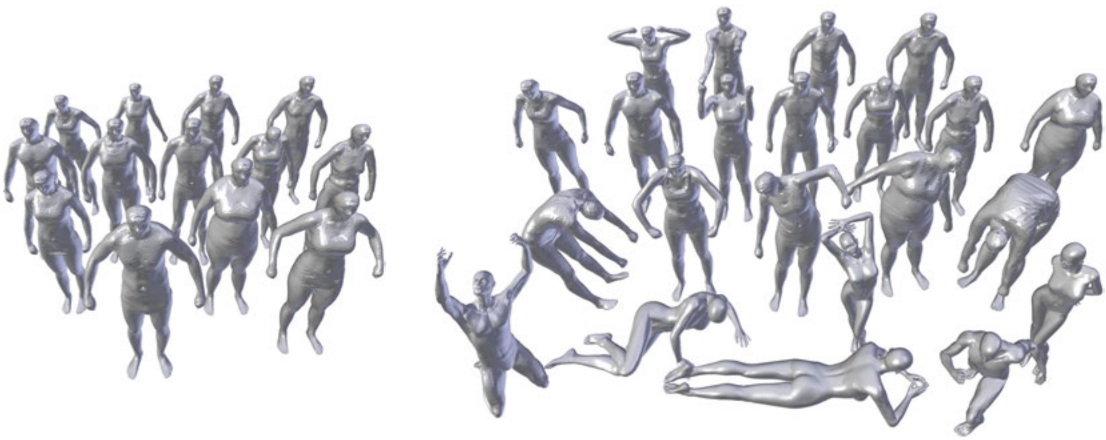

Figure 1.2: A population of human shapes in a neutral pose (left) and in various poses (right).

define, compute, and analyze these sample statistics. Figure 1.3a displays an example of summarizing a set of shapes using a metric-based approach. The summary here is in the form of a mean shape (center), and the three top principal modes of variation (shown in the three rows).

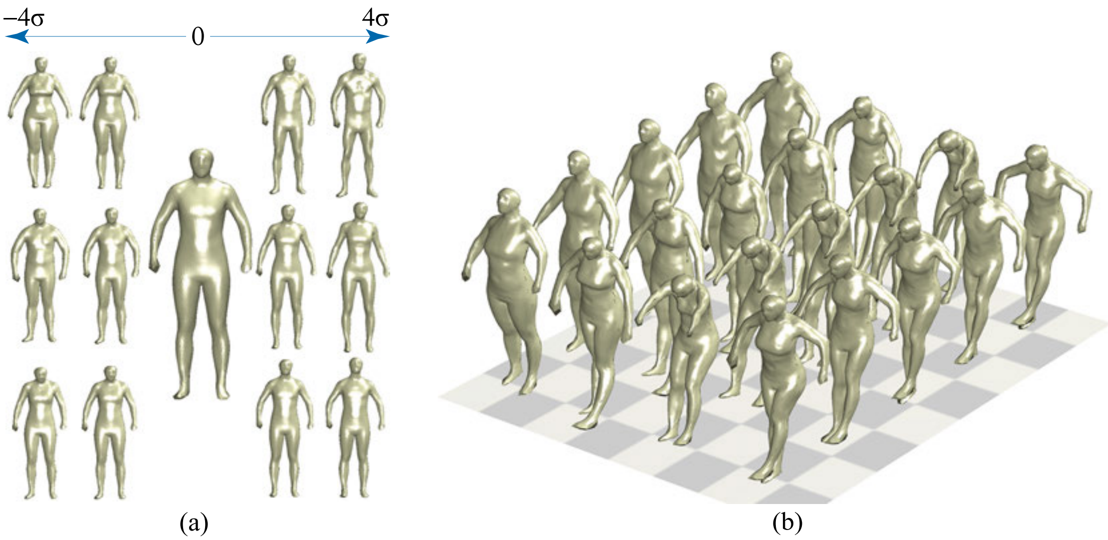

Figure 1.3: Given a set of shapes, one goal is to summarize this set using a mean shape and principal modes of variability (panel (a)). One can also learn a shape model, and then generate random samples from that model to study its effectiveness (panel (b)).

3. **Shape Modeling.** The approach should allow the definition of comprehensive, yet tractable, probability models for capturing essential variability in the shapes of a given set of objects. On the one hand, these shape models should be reasonable in complexity, leading to real-time statistical analysis of shapes; on the other, they should be powerful enough to capture the original variability and provide satisfactory class separation. One can use principal modes of variability in the data, first, to find lower-dimensional submanifolds where most of the data lies, and then to impose tractable statistical models on these submanifolds. Figure 1.3b shows examples of random shapes generated from a parametric shape model learned from training data.

4. **Shape Clustering, Classification, and Testing.** Given a shape metric, and probability models on a shape space, one can extend standard tools from pattern recognition and machine learning to the shape analysis context. For instance, one can develop techniques for shape estimation, shape-based clustering of objects, and hypothesis testing for shape classification.

5. **Real Applications.** Finally, an important goal has to be successful use of these tools on large datasets involving objects in real applications. Since a strong motivation for this work comes from the easy availability of shape databases, it is vital to demonstrate the success of any methods on these databases in order for them to be deemed useful.

Now that we have laid out a core set of tasks, the next questions are: *What makes shape analysis of 3D objects challenging? Why are the tasks of characterizing and analyzing shapes, and modeling shape variability, difficult ones*? Keep in mind that we are interested in shapes of boundaries of 3D objects, represented as 2D surfaces in $\mathbb{R}^3$. The main challenges are as follows.

1. **Registration of Surfaces.** Perhaps the most important aspect of shape comparison is registration. Registration is a matching of points across objects, i.e., it determines which point on one object matches with which point on the other (Figure 1.4). Registration is often a difficult problem to solve, especially across objects with large deformations and pose variability. A variety of techniques have been developed, depending on the context, for solving the registration problem. A major drawback of these past techniques is that registration is mostly performed independently of ensuing shape comparisons. In other words, the objective functions used for registration are different from the metrics used to compare and analyze shapes. It seems more natural to treat the two problems in a unified, comprehensive setup, using a single metric! In fact, that is one of the main accomplishments of the techniques presented in this book.

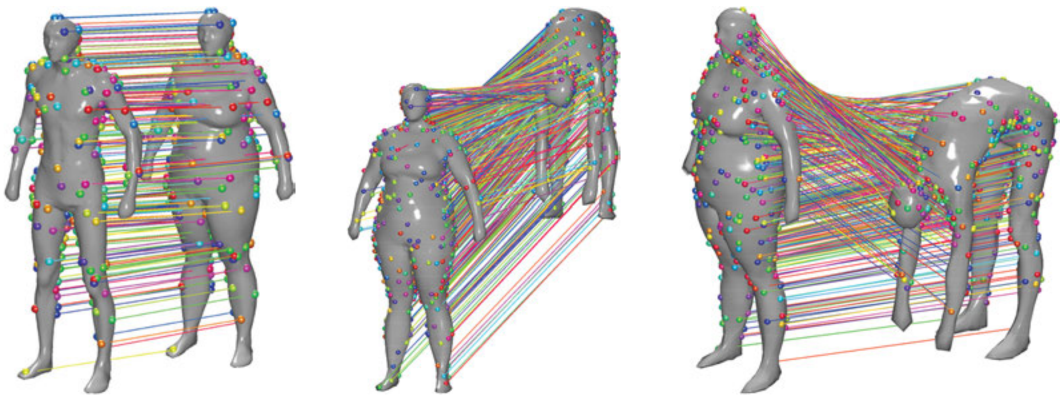

Figure 1.4: Registration of points across two surfaces that bend and stretch.

2. **Infinite-Dimensionality of Shape Spaces.** The objects of interest in this book, genus-0 surfaces, are elements of an infinite-dimensional space, and this requires the use of a Hilbert space structure to analyze that function space. Using this structure, one can represent any surface using a basis expansion and truncate the expansion to reach a finite-dimensional, albeit approximate, representation of a surface. These finite-dimensional representations can then be used for statistical analysis. The important questions are: What

should be the representation space of surfaces, and what Hilbert space structure should be used for shape analysis?

3. **Nonlinearity of Representations.** Shape representations frequently lead to nonlinear spaces, for a variety of reasons. For instance, shape is a property that is invariant to certain transformations such as rigid motion and global scaling. In the case of parametrized surfaces, it is also invariant to any reparametrization of a surface. This implies that shape spaces are quotient spaces of the original representation manifolds, in order to remove these nuisance transformations. The quotient spaces are not vector spaces and, thus, do not allow standard multivariate statistics to be applied directly. One has to utilize the geometry of these spaces, in conjunction with the chosen metric, to compute distances, take averages, or to perform Principal Component Analysis (PCA).

1.3 PAST APPROACHES AND THEIR LIMITATIONS

Perhaps the earliest known efforts in formalizing shape analysis came from D'Arcy Thompson who tried to relate the shapes of functionally similar objects. He explored the possibility of making shapes visually similar by applying transformations and making them closer than they originally appeared. His treatment of shapes appears in the form of a 1917 book titled *On Growth and Form*. Figure 1.5 shows examples, albeit using 2D objects, of using non-rigid transformations for matching two seemingly different but functionally similar objects. The left example considers Albrecht Dürer's face transforms, which were among Thompson's inspiration for studying shapes. Here, the transformation is applied to the coordinate system in which the object is represented and not to the object itself; the appearance of the object changes accordingly.

A large number of approaches in shape analysis abstract the important geometric properties of 3D objects into a set of numerical descriptors; examples are shape distributions [Osada et al., 2002], harmonic coefficients [Kazhdan et al., 2003], Zernike moments [Novotni and Klein, 2003], wavelet descriptors [Laga et al., 2006, 2007], and covariance descriptors [Tabia and Laga, 2015, Tabia et al., 2014]. Shapes are then compared using distance measures defined on the descriptor space [Laga and Nakajima, 2008]. The main shortcoming is that these representations are not invertible. That is, given an arbitrary point in the feature space, it is not possible to determine a shape or a set of shapes that correspond to that feature value. As a consequence, it is difficult to perform statistical analysis in the feature space and then map it back to the object space for inference.

In recent decades, perhaps the most prominent idea that laid foundations for statistical shape analysis in general is the work of Kendall [1977]. Instead of using descriptors, Kendall's approach represents objects with a finite set of n ordered points, called *landmarks*, sampled from the object's boundary and put in correspondence across objects. He then goes one step further and defines *shape as the property of an object that remains after variations due to translation, scale*

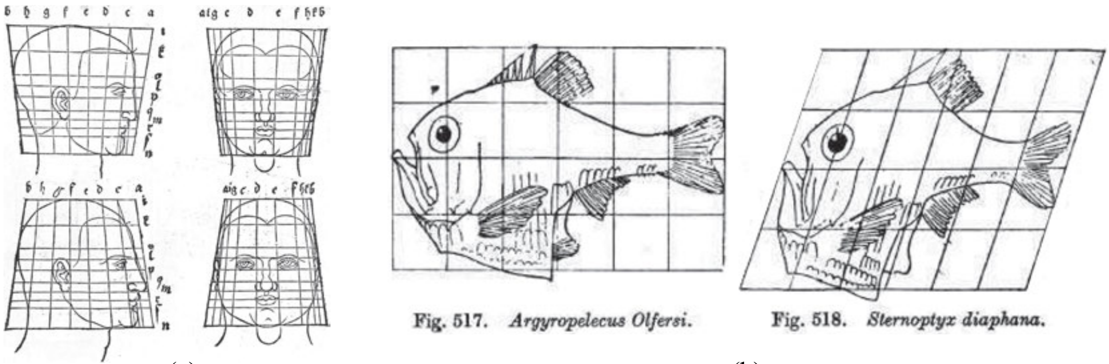

(a) (b)

Figure 1.5: Examples from D'Arcy Thompson's work on measuring differences in shapes of related objects by means of simple mathematical transformations. The example in (a) is Albrecht Dürer's face transforms, which were among D'Arcy Thompson's inspiration for studying variation in shapes using non-rigid transformations. In panel (b) he compares the shape of an Argyropelecus olfersi with that of a Sternoptyx diaphana. The data are courtesy of Wikipedia Commons.

and rotation are factored out. This is known as Kendall's "shape space theory," introduced in 1977 [Kendall, 1977] and advanced by many others since [Dryden and Mardia, 1998, Kendall et al., 1999, Le and Kendall, 1993, Small, 1996], leading to multiple developments and applications, including the Active Shape Models of Cootes et al. [1995], in the case of 2D shapes, and to the 3D morphable models of Blanz and Vetter [1999] for 3D shapes.

There has also been remarkable success in the use of medial axis representations [Siddiqi et al., 2008], especially in medical image analysis [Bouix et al., 2005, Gorczowski et al., 2010]. These representations encode the pose-invariant topological properties of shapes. The main limitation here is that these techniques have been developed for pre-registered objects only. They do not solve for registration of points across objects using medial representations; instead, they use some off-the-shelf method for registering objects before analyzing their shapes.

There has been a tremendous amount of work on treating each object as a manifold and using its geometry to derive certain local features. These features are then matched across objects for registration and comparisons [Bronstein et al., 2009, 2010, Lipman et al., 2010]. These representations encode, simultaneously, the geometry and the topology of shapes. Their main advantage is that they are invariant under certain types of deformations (isometric or affine), and thus are suitable for finding correspondences between those kinds of shapes, e.g., shapes that only differ by bending. However, these methods do not provide a deformation (or geodesic) or any statistical modeling of variability across objects.

As we will see in subsequent chapters, treating the boundaries of objects as continuous surfaces, rather than discretizing them into point sets at the outset, provides more comprehensive solutions. Some approaches such as SPHARM or SPHARM-PDM [Brechbühler et al., 1995, Styner et al., 2006] tackle this problem by choosing special parameterizations, usually analogous to arc-length parameterization in the case of curves. Others restrict the type of reparameterizations that a surface can undergo to those that are area-preserving, such as Möbius transforms [Lipman and Funkhouser, 2009]. These represent major restrictions, and do not result in good registrations.

Once registrations have been computed, the next step is to find the set of transformations that align the surfaces. When dealing with rigid shapes, the transformations of interest are translation, scaling, and rotation. The problem is then referred to as rigid alignment. When using the L^2 metric for measuring closeness (and thus distances), and if the registration between f_1 and f_2 is known, such a rigid transformation can be found analytically using Singular Value Decomposition (SVD), leading to Procrustes analysis [Gower and Dijksterhuis, 2004], which forms the basis of the 2D Active Shape Models (ASM) of Cootes et al. [1995] and the 3D morphable models of Blanz and Vetter [1999].

In practice, however, such registrations are unknown, and better results are obtained if one simultaneously solves for the best registration and for the optimal rigid alignment. A popular solution to this problem is the Iterative Closest Point (ICP) algorithm [Besl and McKay, 1992, Chen and Medioni, 1992]. The algorithm iterates between two steps: (1) a matching step where registration is estimated by nearest-neighbor search using either point-to-point [Besl and McKay, 1992] or point-to-plane [Chen and Medioni, 1992] distances; and (2) an optimal alignment step using the estimated registration. Since its introduction, many variants of the ICP algorithm have been proposed. They aimed at improving various aspects of the original method, such as the speed and quality of convergence, by, for example, combining multiple distance measures with some local descriptors [Rusinkiewicz and Levoy, 2001]. Nevertheless, the ICP algorithm and its variants provide good results only when the poses of the shapes being aligned are initially close to each other.

When the poses are very different, the problem becomes very challenging since the search space for optimal alignment and registration is very large. In theory, however, three pairs of corresponding points are sufficient to define a rigid transformation that best aligns two shapes. This fact has been used in Chen et al. [1999], Papazov and Burschka [2011], and Rodolà et al. [2013] to derive RANSAC-based solutions. Their complexity, however, is high, of order $O(n^3)$ where n is the number of points being aligned. Aiger et al. [2008] introduced the four point congruent sets (4PCS) algorithm, which reduced the complexity of RANSAC algorithms to $O(n^2)$. It has been later optimized to achieve a complexity of order $O(n)$ [Mellado et al., 2014].

Many objects in nature, however, undergo non-rigid deformations composed of bending and stretching. In that case, the registration problem becomes even more challenging. Examples of methods that tried to solve the non-rigid registration problem include non-rigid variants of the ICP algorithm [Amberg et al., 2007] and deformation-driven techniques [Alhashim et al., 2015, Chang and Zwicker, 2008, Zhang et al., 2008].

1.4 OUR APPROACH: ELASTIC SHAPE ANALYSIS

As advocated earlier, a comprehensive approach to shape analysis involves performing registration and deformation in a single **unified** fashion. **Elastic shape analysis (ESA)** is a Riemannian approach that accomplishes exactly that for objects such as curves and surfaces. In ESA, one identifies an appropriate representation space for parameterized versions of the objects, and endows it with a Riemannian metric. The metric allows the definition of a geodesic, i.e., a shortest path, between two objects in the representation space. This path specifies the optimal deformation under the metric.

Registration is then handled by specifying an action of the reparameterization group on the representation space. The key idea is that the parameterizations of two surfaces define their registration; one changes the registration by reparameterizing the surfaces. This means that if we find the shortest path, not between the original two parametrized objects, but between all possible reparameterizations of those objects, then we simultaneously identify both the optimal registration of the objects, and the optimal deformation between them (as described by the shortest path). We also render the resulting registration and deformation independent of the original parameterizations. The result is therefore a registration and deformation of the corresponding *geometric* (i.e., unparameterized) objects.

The optimization over all possible reparameterizations can alternatively be viewed as finding a geodesic in the quotient of the original space by the group of reparameterizations (i.e., a geodesic between the orbits of the original objects under the reparameterization group), using a quotient Riemannian metric derived from that on the original space. As a matter of fact, one can also choose to introduce further invariances to, for example, translations and rotations, thereby incorporating alignment into the deformation as well. The space generated by taking these quotients is often known as a "shape space."

Central, then, to the implementation of the ESA program is the definition of a Riemannian metric on the space of parameterized objects that is preserved by the action of the relevant transformations, typically reparameterizations, translations, rotations, and perhaps scale. Key to the program's success in practical terms is the definition of a Riemannian metric that "measures" the types of shape changes that are important, and that renders computationally feasible the corresponding geodesic calculations.

These goals have been achieved in the elastic shape analysis of **curves** by using a particular member of the family of **elastic metrics**, in conjunction with a representation called the **square-root velocity function** (SRVF). The form of the elastic metric is extremely simple when expressed in terms of the SRVF: it becomes the L^2 metric [Srivastava et al., 2011]. This greatly facilitates computations, enabling sophisticated statistical analyses that require many geodesic calculations. Consequently, this method has been used in many practical problems [Laga et al., 2012, 2014]. Critical to its utility is the fact that the mapping from the space of curves to the SRVF space is a bijection (up to a translation). Solutions found in the SRVF space using the L^2 metric can be uniquely mapped back to the original curve space, using an analytical expression. This is significantly more efficient than performing analysis in the curve space itself.

In the next chapter, we will review the elastic shape analysis of curves, and then use it as a starting point for the development of the elastic shape analysis of surfaces.

1.5 ORGANIZATION OF THIS BOOK

The rest of this book is organized as follows. We start Chapter 2 with a brief summary of the elastic shape analysis of curves in Euclidean spaces, and then introduce the elastic Riemannian metric for analyzing shapes of surfaces. This leads to the concept of square-root normal fields (SRNFs), an efficient framework for performing shape analysis using the L^2 norm. This chapter also describes the difficult problem of converting SRNF representations back to parameterized surfaces, and an approximate and computationally efficient solution to this problem.

Chapter 3 takes this fundamental framework and develops efficient algorithmic solutions for certain basic tasks, such as performing alignment and registration, and computing geodesics. Chapter 4 takes these basic tools and develops comprehensive solutions for statistical analysis, i.e., for capturing variability of shapes within and across shape classes, finding dominant modes of variability, and modeling this variability using parametric families on submanifolds of shape spaces.

Chapter 5 illustrates these ideas using various case studies, for example the statistical shape analysis of human bodies to model their variability, and the statistical shape analysis and modeling of anatomical parts for use in medical diagnoses. Examples are taken from well-known and well-used datasets: the TOSCA dataset, the SHREC07 watertight database [Giorgi et al., 2007], and the human shape database from Hasler et al. [2009a], and thus demonstrate the real applicability of elastic shape analysis.

Chapter 6 discusses and demonstrates some extensions of these ideas to shape analysis of surfaces with landmarks.

Finally, since the approach advocated here for shape analysis of surfaces is primarily a geometric one, the reader will need a certain amount of background in geometry to follow the details. In the appendices, we provide some background material on fundamental subtopics: differential and Riemannian geometry (Appendix A); in particular of surfaces in $\mathbb{R}^3$ (Appendix B); spherical parameterization of triangulated genus-0 surfaces (Appendix C); and landmark detection and matching on surfaces (Appendix D).

1.6 NOTATION

In order to help the reader, in this section we lay out the notation used in this manuscript. There is notation for general mathematical operations, and also specialized notation for representing objects, their spaces, and mappings between such spaces.

GENERAL MATHEMATICS

- Partial differentiation of a map f by a variable u will be denoted in various ways: $\frac{\partial f}{\partial u}$; $\partial_u f$; f_u.
- Composition of maps will be denoted by a circle $\circ$: if $f : X \to Y$ and $g : Y \to Z$, then $g \circ f : X \to Z$.
- We will often denote tangent vectors to a manifold at the point x by the notation δx, indicating an infinitesimal displacement of the point.

GEOMETRIC SPACES, OBJECTS, AND MAPS

Table 1.1: Spaces

Symbols	Explanation
D	Domain space of parameterized surfaces (usually $\mathbb{S}^2$) or curves (usually $\mathbb{S}^1$).
F	Set of parameterized surfaces (curves): smooth embeddings of D in $\mathbb{R}^3$ ($\mathbb{R}^2$).
C	Pre-shape space of surfaces (curves).
$C_f = C$, S_f	Pre-shape and shape spaces in f representation.
$\mathcal{Q}$	Space of SRNFs.
C_q, S_q	Pre-shape and shape space in SRNF representation.
Γ	Diffeomorphism (reparameterization) group of D.
$\mathcal{G}$	Space of Riemannian metrics on D.

Table 1.2: Objects

Symbols	Explanation
δX	Tangent vector to space $\mathcal{X}$ containing objects of type X, at X.
M, $T_p(M)$	A general manifold and its tangent space at $p \in M$.
$f \in \mathcal{F}$, $S \in \mathcal{S}$	Parameterized and unparameterized surfaces (curves).
$s = (u, v)$, ds	Coordinate system on D; corresponding Lebesgue measure.
n, $\hat{n}$	Un-normalized and unit normal vector fields to a surface (curve).
$\gamma \in \Gamma$	Diffeomorphism (reparameterization) of D.
g	Riemannian metric induced on D by parameterized surface or curve (first fundamental form).
$\lvert g \rvert = \det(g)$, $\hat{g} = g/\lvert g \rvert$	Determinant of induced metric and "normalized" induced metric.
$r = \sqrt{\det(g)}$	Density with respect to ds of measure induced on D by parameterized surface or curve.
$J[\gamma]$	Jacobian of γ.
$\{b_i\}_{i \in I}$	Basis of $T_{id}(\Gamma)$.
$[f]$, $[q]$	Orbits of f and q under rotation and reparameterization.
α_f, α_q	Geodesic paths in $\mathcal{F}$ and $\mathcal{Q}$, respectively.

Table 1.3: Maps and operators

Symbols	Explanation
$\lvert \cdot \rvert$, $\lVert \cdot \rVert$	Norm of a vector, L^2 norm of a function.
$\langle \cdot, \cdot \rangle$, $\langle\langle \cdot, \cdot \rangle\rangle_f$	Euclidean metric in $\mathbb{R}^3$, Elastic Riemannian metric on $\mathcal{F}$.
$\langle \cdot, \cdot \rangle_2$	L^2 inner product between functions.
(f, γ), (q, γ)	Actions of Γ on $\mathcal{F}$ and $\mathcal{Q}$, respectively.
$Q : \mathcal{F} \to \mathcal{Q}$	Forward mapping to SRNFs.
SRM	Square-root mapping.
E_{inv}, E_{reg}, E_{path}	SRNF inversion energy, registration energy, path-straightening energy.
$d_{\mathcal{F}}$, $d_{\mathcal{Q}}$, $d_{\mathcal{S}_f}$, etc.	Distance function on the space $\mathcal{F}$ of surfaces, on the space $\mathcal{Q}$ of SRNFs, on the shape space $\mathcal{S}_f$, etc.

CHAPTER 2

Elastic Shape Analysis: Metrics and Representations

In Chapter 1, we took a look at the long history of shape analysis, and the varied approaches that have been taken to it. We argued that elastic shape analysis (ESA), the approach taken in this book, has significant theoretical, and as a result, practical advantages over these older approaches. We explained how, by using parameterized objects, and an appropriate measure of distance between them, ESA produces a unified approach to the problems of shape registration and comparison, in contrast to other approaches. In this chapter, we will describe ESA, and the objects and structures involved, in mathematical detail.

The mathematics involved is differential and Riemannian geometry. For those not familiar with this material, Appendix A summarizes the key ideas and definitions; terms defined in the appendix appear in boldface the first time they appear in this or later chapters. However, we will not use these ideas in a very rigorous way; geometric intuition will be more important.[1] Let us begin!

2.1 SHAPES

It goes without saying, perhaps, that the objects of interest in shape analysis are "shapes." This word is used in many ways, and we too will be guilty of overloading its meaning. For our purposes, a shape will be an m-dimensional sub-**manifold**, possibly with boundary, of $\mathbb{R}^n$. This book is concerned with the case in which $m = 2$ and $n = 3$, and, in addition, in which the $m = 2$-dimensional submanifold, the "surface," is:

- closed, i.e., is the boundary of a $(m + 1) = 3$-dimensional submanifold; and
- simply connected, i.e., it has only one piece and no handles; in other words, it is connected and has genus 0.

We are thus concerned with 3D shapes that are "blobs." As an introduction, we will also discuss in some detail the case with $m = 1$ and $n = 2$ and subject to the same constraints, i.e., the case of closed planar curves. We will also briefly describe the general situation to emphasize that the approach is not limited to these cases.

[1] In particular, we will not concern ourselves with characterizing precisely the mathematical spaces that are needed to render everything rigorously well defined. It is worth remembering that it is only recently that the correct spaces were identified in the case of curves [Lahiri et al., 2015], but this did not impede progress in applications.

The shapes we have just described are unparameterized, i.e., they are purely geometric, whereas it was stated in the previous chapter that the starting point of elastic shape analysis is a space of *parameterized* objects of interest. Why introduce a parameterization, when there is none to start with? We will answer this question indirectly, by describing the framework of elastic shape analysis and seeing how parameterization helps.

2.2 ELASTIC SHAPE ANALYSIS

In elastic shape analysis, the genus-0 surfaces we have just mentioned will be represented by **embeddings**[2] of the 2-sphere $\mathbb{S}^2$ in $\mathbb{R}^3$, i.e., by parameterized surfaces. The unparameterized surface is then the image of $\mathbb{S}^2$ under such an embedding: if the embedding is $f : \mathbb{S}^2 \to \mathbb{R}^3$, then the surface is $f(\mathbb{S}^2) \subset \mathbb{R}^3$. (We will use the term "surface" to refer to both the embedding and its image, except where this would cause confusion.)

Clearly, there is an infinite set of embeddings, all of which share the same surface as an image. This multiplicity of representations may be dealt with by realizing that all these embeddings are related by the **group action** of the **Lie group** of **reparameterizations**, Γ, the **diffeomorphisms** of $\mathbb{S}^2$. The group Γ acts on $\mathcal{F}$ by composition: $\gamma : \mathbb{S}^2 \to \mathbb{S}^2$ acts on $f : \mathbb{S}^2 \to \mathbb{R}^n$ to give $(f, \gamma) \equiv f \circ \gamma : \mathbb{S}^2 \to \mathbb{R}^n$.[3] Notice that f and $f \circ \gamma$ represent the same unparameterized surface in $\mathbb{R}^3$ because $f \circ \gamma(\mathbb{S}^2) = f(\gamma(\mathbb{S}^2)) = f(\mathbb{S}^2)$. We can thus recover the space of genus-0 surfaces (unparameterized objects) from the space of embeddings (parameterized objects), as the **quotient space** of the latter by the action of Γ, the points of which are the **orbits** of the group action.

All of this remains true, *mutatis mutandis*, no matter what the domain D of the embeddings: it applies to surfaces ($D = \mathbb{S}^2$), but equally to closed planar curves ($D = \mathbb{S}^1$), or in general to $(m + 1)$-dimensional shapes in $\mathbb{R}^n$. We thus do not need to be specific at this point, and we will use the generic D for the domain of the embeddings from now on, unless its identity is important.

The quotient construction shows that we can use parameterized surfaces instead of unparameterized surfaces, but at the cost of an infinitely redundant representation. Why would we do this?

2.2.1 ENCODING OF REGISTRATION

One of the core tasks of shape analysis, as described in Chapter 1, is shape registration: creating a one-to-one correspondence between the points of two shapes. Parameterized objects are very natural in this context, because two parameterized surfaces f_1 and f_2 automatically provide a correspondence between the corresponding unparameterized surfaces: a point $x \in f_1(D)$ in one

[2]In practice, it will be hard to insist on embeddings as opposed to **immersions** (embeddings are bijective immersions); in fact, judging by the curve case, it is not even clear that we require the maps to be immersions; nevertheless we will persist in talking of embeddings rather than these more general possibilities.

[3]We denote the generic action of a group element g on an object f by (f, g) and function composition by $f \circ g$.

surface is registered to a point $y \in f_2(D)$ in another surface only if they are both images of the same point in the domain D, i.e., if and only if there exists $p \in D$ such that $x = f_1(p)$ and $y = f_2(p)$. We can thus encode different registrations of two shapes by picking different parameterized representatives for each of them. This turns out to be a very convenient way to encode registrations, and is one of the main reasons that parameterized objects are useful. The encoding of registrations in this way is illustrated in Figure 2.1.

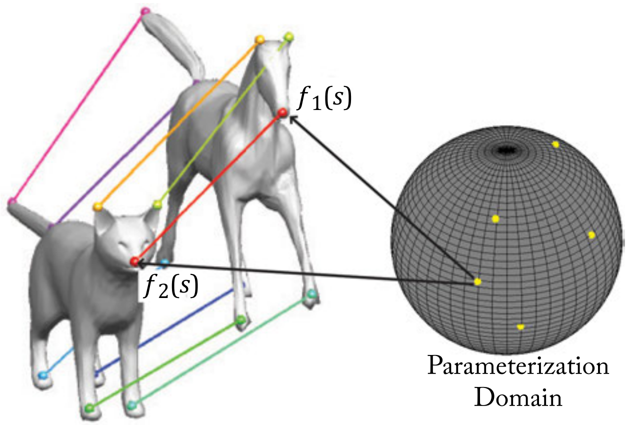

Figure 2.1: Registrations are encoded by parameterized objects via their domain. The fact that the points $f_1(s)$ and $f_2(s)$ are both images of the point $s \in \mathbb{S}^2$ means that they are registered.

At this point, though, let us note something important about this representation of registrations. Take a pair of registered points x and y as before, which by definition can be written $x = f_1(p)$ and $y = f_2(p)$ for some $p \in D$, and consider the two embeddings $f_1 \circ \gamma$ and $f_2 \circ \gamma$ produced by the action of a reparameterization. We note that $x = f_1(p) = f_1(\gamma \circ \gamma^{-1}(p)) = f_1 \circ \gamma(q)$ and $y = f_2(p) = f_2(\gamma \circ \gamma^{-1}(p)) = f_2 \circ \gamma(q)$, where $q = \gamma^{-1}(p)$. Thus, we see that if there is a $p \in D$ such that $x = f_1(p)$ and $y = f_2(p)$, then there is also a $q \in D$ such that $x = f_1 \circ \gamma(q)$ and $y = f_2 \circ \gamma(q)$. In other words, the pair f_1 and f_2 describe the same registration as the pair $f_1 \circ \gamma$ and $f_2 \circ \gamma$, in the sense that if two points are registered under one description, then they are registered under the other. Registration is thus invariant to the simultaneous action of Γ on the two embeddings involved. This point will be important as we move to the second key ingredient of ESA: the definition of a Riemannian metric on the space of parameterized objects.

2.2.2 RIEMANNIAN METRIC AND OPTIMAL REGISTRATION

As described in Chapter 1, ESA is based upon a **Riemannian metric** on the space of parameterized objects. A Riemannian metric measures the distance between infinitesimally close objects; alternatively, it measures the "size" of an infinitesimal change in (or "deformation of") an object.