Investigating Techniques for Pointwise Correspondences in Deformable Shapes

A THESIS SUBMITTED IN PARTIAL FULFILLMENT FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in Computer Science and Technology

in the Goa Business School

GOA UNIVERSITY

by

Manika BINDAL

Discipline of Computer Science and Technology,

Goa Business School, Goa University,

Taleigao Plateau, Goa

January 2025

Declaration of Authorship

I, Manika BINDAL, hereby declare that the research work entitled "Investigating Techniques for Pointwise Correspondences in Deformable Shapes" submitted to Goa University for the award of the degree of Doctor of Philosophy, is a genuine and original work conducted by me under the supervision of Prof. Venkatesh KAMAT and Prof. Jyoti D. PAWAR.

This thesis has not been submitted, either in part or full, to any other University or Institution for the award of any research degree or diploma. All sources and references used in the preparation of this thesis have been appropriately acknowledged and cited in the bibliography.

I further declare that the work presented in this thesis is the result of my independent research, and that no part of it has been plagiarized or reproduced from other sources without proper citation.

Signed:

Date:

Certificate

This is to certify that the research work entitled "Investigating Techniques for Pointwise Correspondences in Deformable Shapes" submitted by Manika BINDAL, for the degree of Doctor of Philosophy at Discipline of Computer Science and Technology, Goa University, is an original and independent work carried out under my supervision.

I further confirm that the work presented in this thesis is based on the student's own research, and all sources and references have been duly acknowledged. The student has worked conscientiously, and I have supervised the progress regularly.

I hereby recommend the thesis for consideration for the award of the degree.

Advisor: Co-Advisor:

Retd. Prof. Venkatesh KAMAT	Prof. Jyoti D. PAWAR
Goa Business School	Goa Business School
Goa University	Goa University

"There is no such thing as a new idea. It is impossible. We simply take a lot of old ideas and put them into a sort of mental kaleidoscope. We give them a turn and they make new and curious combinations. We keep on turning and making new combinations indefinitely; but they are the same old pieces of colored glass that have been in use through all the ages ..."

Mark Twain

Abstract

Humans intuitively understand shapes, as the skill is vital for interacting with the ambient environment. Machines, however find it challenging to match shapes due to factors like unique metric structures and the ability of shapes to deform in a non-rigid manner. Growing availability of high-resolution datasets, combined with technological advancements, has prompted the shape analysis community to prioritize extracting highly detailed geometric and topological information from shapes.

Establishing correspondence mapping between distinct elements of deformable shapes is a fundamental task while analysing shapes. It is essential for accurate shape matching, registration, and comparison. Aligning corresponding features enables transfer of attributes like color, texture, or deformation information more effectively, while also improving the reconstruction of objects from partial data.

This thesis aims at exploring techniques designed to capture localised details of shapes depicting sharp features, a challenge that traditional approaches often struggle to address. One proposed approach utilizes shape-aware distances to improve resilience against topological noise and enhance point-to-point accuracy across varying triangulations. Building on recent advancements in shape analysis, powerful technique - functional map framework, has been utilized for enhancing pointwise mapping. This framework leverages the spectral domain of shapes to establish correspondences between function spaces, ultimately translating to point-to-point correspondences. To preserve high-frequency features of shapes, use of Hamiltonian operator with Gaussian curvature have been proposed. This functional mapping approach, which constructs linear mappings between function spaces instead of directly matching points, relies heavily on the choice of basis functions. A comprehensive exploration of these bases highlights their impact on the precision and efficiency of correspondence techniques, offering insights into future directions for more robust shape analysis. It emphasizes the importance of selecting the right basis function tailored to the specific needs of

the application, taking into account factors such as computational efficiency, stability, and the required level of detail.

Acknowledgements

It has been a *long* roller-coaster ride of unexpected twists and turns, both personally and professionally. As they say, *it takes a village to build a dream*, I wouldn't have done it, if not for the people who were there for me along the way and continue to be. I am truly blessed.

I am eternally indebted to Kamat sir, whose belief in me has given me the strength to persevere and whose constant encouragement leave me at a loss for words to convey my gratitude. Thank you sir for letting me explore, for always asking *"Is learning still interesting?"* and for teaching me the ropes simply by being yourself – both in academia & in life. I am forever grateful.

While many seek a single guiding star, I've been graced with a constellation of mentors. I sincerely thank Jayanthan sir *(Maths dept.)* for graciously allowing me to take his course on Differential Geometry and for his consistent presence in my DRC meetings and guidance, despite not being obligated to do so. My deepest gratitude to Tatu sir *(DAIICT, Gandhinagar)* for his willingness to share knowledge as it has substantially improved my understanding of key fundamental concepts in my work. I am truly grateful for his time, patience, and the insightful discussions that have significantly shaped the direction of my thesis. Thank you sir, for your invaluable mentorship. I thank Jyoti ma'am wholeheartedly, for continuously inspiring me to embrace life's challenges with a calm and holistic perspective. I sincerely thank Ashwin Sir for his constructive feedback and steering my work in the right direction. Thank you, Karmali Sir, for reminding me to 'never fear the unknown' when I needed it the most and for asking thought-provoking questions. Thank you, Bhaskar sir, for your kind words of encouragement and sharing the idea of *deep work* with me. I am still working on it. Thank you Yma ma'am for sponsoring *Ustad Zakir Hussain Sahab's* concert and for inspiring me to be a more dedicated and conscientious teacher. Thank you Ramrao sir for your compassionate smile and for demonstrating how to accomplish more within a 24-hour period. I am grateful to Siddhartha sir and Suyash sir *(IITB)* for giving me the opportunity to spend a semester at the institution, where I gained valuable exposure to foundational concepts essential for my work. I sincerely thank Hanumant Sir for arranging hostel accommodation for me during that stay. I am grateful to Gad sir *(Electronics dept.)* for consistently facilitating our financial aid.

Behind every milestone, there is a group of people who made it possible. Thanks to Swapnil's endless nagging, I've become a reluctant scholar. His follow−up *calls* that felt more like follow−up *interrogations*. Thank you for those, it was a much needed boost to my writing phase. Also, I'll never forget the day he introduced me to the infamous Konkani word 'ghaati'. Thanks for that unforgettable linguistic adventure too. Blessed to have you Swapnil, our own Ordinance Oracle. Thank you Teja for always telling me 'do it no, yaa' and for patiently listening to me while making an effort to respond in *hindi*. My *only* lab partner, Vandana ma'am, thank you for all kinds of talks we had back then, filled with frequent canteen trips. Gaurang, my first friend in the University, thank you for also being my data-recovery wizard. Thanks, Pradnya, for your contagious positivity and constant inspiration. Heartfelt thanks to Aniket for those long conversations and managing our PhD Scheme. Thank you Nandesh for being so welcoming. Payaswini, our photographer, thank you for those car rides and your friendship. Thank you Razia for being so cheerful and encouraging, always. Thank you McNeil & Deepak for smooth operation of my in-person DRC meetings. Thank you Nutan Ma'am, Kavita, and Surya Bhaiya Ji for your constant support in handling administrative tasks, allowing me to focus on my learning. Special thanks to Nutan ma'am for introducing me to the delightful flavors of Goan cuisine. Thank you Rekha Kaki for your love & blessings.

Some foundations fortify academic pursuits, while others illuminate the path to a fulfilling life. I am deeply grateful to *mumma-papa* for blessing me with an incredible life and guiding me through every step of the journey so far... Your love is the greatest gift, and I'm eternally grateful to have you as my parents. I truly hit the jackpot twice − once with my amazing parents, and again with my wonderful in-laws. Thank you *mummyji−papaji* for your endless love, laughter and for embracing me as your own. Though our time together was short, I cherish every memory. May your soul find peace and light in heaven. Anubhav, my madcap maestro! I'm so glad I found my kind of crazy in you. Thanks for being the yin to my yang and for braving life's storms by my side. Nitara, our little fighter, you light up our world every single day. Witnessing your resilience and love, makes life's journey so much sweeter. Thank you for being you.

This work would not have been possible without the generous funding provided via Visvesvaraya PhD Scheme by Ministry of Electronics and Information Technology, Government of India. I conclude my PhD journey, eager for the new experiences life holds, with a heart full, a mind enriched, and a spirit grateful.

Contents

List of Figures

List of Tables

List of Algorithms

List of Abbreviations

ADMM	**A**lternating **D**irections **M**ethod of **M**ultipliers
CMH	**C**oordinate **M**anifold **H**armonics
CMM	**C**ompressed **M**anifold **M**odes
CQHB	**C**oupled **Q**uasi **H**armonic **B**asis
GPS	**G**lobal **P**oint **S**ignature
HKS	**H**eat **K**ernel **S**ignature
HO	**H**amiltonian **O**perator
ICP	**I**terative **C**losest **P**oint algorithm
LAB	**L**andmark **A**dapted **B**asis
LBO	**L**aplace-**B**eltrami **O**perator
LMH	**L**ocalized **M**anifold **H**armonics
MADMM	**M**anifold **A**lternating **D**irections **M**ethod of **M**ultipliers
MDS	**M**ulti-**D**imensional **S**caling
POD	**P**roper **O**rthogonal **D**ecomposition
SMACOF	**S**caling by **MA**jorizing a **CO**mplicated **F**unction
WKS	**W**ave **K**ernel **S**ignature

List of Symbols

$\mathbb{R}^m$	m-dimensional Euclidean Space
$\mathcal{S}$	Shape as continuous surfaces
Ω	Parameter domain for a function
$\mathcal{X}$	Shape as discrete surface; a triangle mesh
$(\mathcal{S}, d_{\mathcal{S}})$	Shape as a metric space with distance metric $d_{\mathcal{S}}$
f_e	low-dimensional shape embedding
$\Delta_{\mathcal{S}}$	Laplace-Beltrami Operator for surface $\mathcal{S}$
$\mathcal{H}_{\mathcal{S}}$	Hamiltonian operator for surface $\mathcal{S}$
$\mathbf{D}$	diagonal matrix of vertex area weights for the mesh
$\mathbf{W}$	stiffness matrix of the cotangent mesh Laplacian
Λ	diagonal eigenvalue Matrix of mesh Laplacian
$\mathcal{F}(\mathcal{S}, \mathbb{R})$	space of real-valued functions on shape $\mathcal{S}$
η	point-to-point map between deformable shapes
η_F	functional representation of a given pointwise map η
$\mathbf{C}$	functional map expressed as a matrix in given basis
$\mathbf{\Phi}, \mathbf{\Psi}$	eigenvector or functional basis matrix
$\nabla(f)$	gradient of a function f
K	Gaussian curvature

Dedicated to my beloved parents-in-law, who now watch over me from heaven. Your desire to see me graduate has been my guiding light. This is for you both...

Chapter 1

Introduction

The fascinating world of 3D shapes has captivated scholars for centuries, with Euclid's (Euclid, 300 BC) work providing foundational insights into how simple geometric forms relate to each another. Shapes are intuitively understood by humans, as it is fundamental to our existence, however machines struggle to comprehend them. Due to the challenges in capturing, processing and storing 3D data, traditionally, computer vision community prioritized 2D data *i.e.*, images and videos. While RGB data offers efficiency and ease of processing, it lacks depth information and can only show data from a single viewpoint. 3D models, though more complex to process, provide a complete representation of an object's geometry. The journey from real-world objects to digital representations often introduces noise, missing data, and discretization errors. With advancements in the acquisition, modeling and visualization technology, object recognition shifted from image-based objects to shape-based objects having more complex geometric and topological information. 3D data has become indispensable for applications like protein structure prediction for medical advancements (Knoverek, Amarasinghe, and Bowman, 2019), architecture for optimizing design efficiency (Sawhney and Crane, 2020) and virtual & augmented reality (VR/AR) for effective immersive experience (Dammann, Steger, and Stelzer, 2022).

Shape analysis, an interesting field at the intersection of computer vision, computer graphics, and geometry processing, focuses on understanding the properties and relationships of 3D shapes. It includes retrieval of shapes

(Tangelder and Veltkamp, 2008; Gasparetto, Minello, and Torsello, 2015), reconstructing the surface of shape from data points (Berger, Levine, Nonato, Taubin, and Silva, 2013), segmenting the shape in meaningful parts (Shamir, 2008), shape compression (Maglo, Lavoué, Dupont, and Hudelot, 2015), etc. and more. Shape correspondence, in particular, refers to matching two (or more) shapes in order to produce meaningful mapping between them (Elad and Kimmel, 2003). Determining correspondences is a fundamental task and is often the first step in various applications like shape registration (Tam, Cheng, Lai, Langbein, Liu, Marshall, Martin, Sun, and Rosin, 2013), symmetry detection (Mitra, Pauly, Wand, and Ceylan, 2013), statistical shape modeling (Taylor Chris and Carole, 2008; Hasler, Stoll, Sunkel, Rosenhahn, and Seidel, 2009), shape morphing (Alexa, 2002) and transfer of various attributes like deformation (Sumner and Popović, 2004), texture (Dinh, Yezzi, and Turk, 2005), segmentation labels (Elghoul and Verroust-Blondet, 2013), and style of one group of shapes to another (Xu, Li, Zhang, Cohen-Or, Xiong, and Cheng, 2010).

The concept of shape correspondence is closely related to shape matching. According to Veltkamp et al. (Veltkamp and Hagedoorn, 2001), shape matching involves quantifying the similarity between shapes using a dissimilarity measure. For example, while a simple *match* or *not match* output suffices for distinguishing between animal and human shapes, establishing correspondences requires a detailed, explicit mapping. This highlights the distinction between matching and correspondence problems. In practice, the terms *shape matching* and *shape correspondence* are often used interchangeably, as finding correspondences is one of the methods for assessing shape similarity.

A 3D shape is the digital representation of a real-world object, encompassing both its interior and exterior. Accurately modeling interiors, requires specialized equipment such as CT scanners to obtain volumetric representation. While it's possible to infer some interior properties from external observations, this is complex and often unnecessary. In this work, the focus would be on shapes that are considered as discrete samples of smooth surfaces embedded in a three-dimensional space. Triangle meshes, which approximate

FIGURE 1.1: Shapes from Kids Dataset (Rodola, Rota Bulo, Windheuser, Vestner, and Cremers, 2014) undergo isometric deformations, primarily illustrating bending around the arms and legs

the object's boundary, are the most commonly used digital representation for 3D objects (refer Section 1.1.2).

Shape matching methods are primarily designed and categorized based on how they account for shape deformation. For example, consider a toy car and a piece of clay. The toy car represents a shape that can undergo rigid transformations, such as rotation or translation, while maintaining its geometric and structural integrity. Regardless of whether it is pushed, pulled, or moved, the shape of the car remains unchanged. In contrast, the piece of clay represents a shape that exhibit flexibility and is characterized by it's ability to undergo substantial deformation. It can be deformed in a non-rigid manner through squishing, stretching, or bending, resulting in significant changes to its shape.

Despite the development of efficient algorithms for rigid shape matching (Gelfand, Mitra, Guibas, and Pottmann, 2005a), non-rigid shape matching continues to be a key research focus due to the intricate challenges posed by elastic, topological, and structural variations in shapes (Bronstein, Bronstein, and Kimmel, 2007). Non-rigid deformations although provide greater flexibility in modeling the shapes, it also expands the space of deformations leading to the increased complexity of shape correspondence problems

(Bronstein, Bronstein, and Kimmel, 2008). This work will concentrate on specific subclasses of non-rigid deformations, such as isometric and near-isometric deformations. Isometric deformation imposes limitations on the possible ways a shape can deform in a non-rigid manner while keeping intrinsic properties intact. For instance, preserving pairwise geodesic distances *i.e.*, the distance between any two points along the surface of the shape, ensures the shape deforms without stretching. Examples of isometric deformation include bending a sheet of paper, folding a piece of cloth, or twisting a pipe, where the fundamental distances and proportions within the material remain unchanged despite alterations in shape. In Figure 1.1, the kid's shape is illustrated in various poses, demonstrating isometric deformations achieved by bending the legs and arms. These poses highlight how isometric deformations maintain relative distances, showing changes in shape while preserving a human body's proportions and geometry. To establish correspondences, the goal is to recognize that shapes can either depict the same object in various poses or distinct objects with a common underlying structure, even when viewed from different angles. In either case, there should be a natural correspondence between the two surfaces that preserves a common structure. For example, in the case of the same human in different poses, the 'natural' correspondence would match corresponding body parts, such as left hand to left hand.

Shape matching techniques can be broadly categorized into geometric transformation based and descriptor based approaches. Geometric transformation-based methods focus on aligning two shapes by finding a transformation that minimizes the distance between corresponding points or features (Gelfand, Mitra, Guibas, and Pottmann, 2005b). These alignment techniques are primarily applied in 3D object reconstruction (Hu, Wen, Van Kaick, Chen, Lin, Cohen-Or, and Huang, 2018), and more recently, they have been extended to reconstruct entire landscapes using time-varying registration methods (Deng, Yao, Dyke, and Zhang, 2022). In contrast, descriptor-based methods establish correspondences using similarities between shape descriptors, without requiring shape alignment (Castellani, Cristani, Fantoni, and Murino, 2008), making them suitable for shapes with different topologies. A hybrid method combines both approaches, iteratively refining the alignment and subsequently determining the correspondences, which is useful especially for matching shapes that are undergoing non-rigid deformations. Various approaches to compute shape correspondences are surveyed in (Kaick, Zhang, Hamarneh, and Cohen-Or, 2011; Sahillioğlu, 2020); please refer Section 1.3 for a detailed discussion.

1.1 Shapes

Kendall defined shape as the geometric essence of an object, independent of its position, size, or orientation (Kendall, 1977). The non-rigid nature of many natural shapes from the intricate folds of a human face to the dynamic curves of a tree, makes their digital representation challenging. Kendall's definition is based on a rigid body assumption and is not applicable to deformable objects. An effective shape representation should be resolution-adaptive, computationally efficient, topology-preserving, scalable, robust to noise, and capable of supporting deformations. Shape representations can be broadly classified into two categories: continuous and discrete, as outlined below:

1.1.1 Continuous Shape Representations

Shapes in geometry processing are described by 2D surfaces that form the boundaries of physical 3D objects, as orientable continuous 2D manifold embedded in $\mathbb{R}^3$ (Botsch, Kobbelt, Pauly, Alliez, and Levy, 2010). Continuous shape representations model smooth surfaces using mathematical functions, capturing the true geometry of the shape without the loss of detail

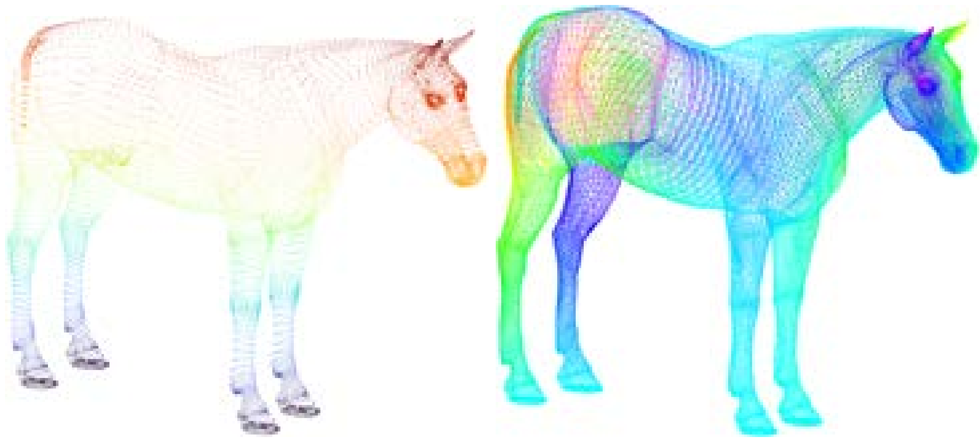

FIGURE 1.2: Geometry of a horse shape from TOSCA Dataset (Bronstein, Bronstein, and Kimmel, 2008) represented as point cloud (left) and a triangular mesh (right)

from discretization. Formally, continuous surfaces ($\mathcal{S}$) can be specified either by a parametric representation or by an implicit representation. Parametric surfaces are vector-valued function $f : \Omega \rightarrow \mathcal{S}$ that maps a parameter domain $\Omega \in \mathbb{R}^2$ to the surface $\mathcal{S} \in \mathbb{R}^3$. Spline and subdivision surfaces (Zorin, Schröder, DeRose, Kobbelt, Levin, and Sweldens, 2000) are examples of parametric representations of shape. Implicit (or volumetric) representation is the zero set of a scalar-valued function $\mathbf{S} : \mathbb{R}^3 \rightarrow \mathbb{R}$ *i.e.*, $\mathcal{S} = \{x \in \mathbb{R}^3 \mid \mathbf{S}(x) = 0\}$, where the solution to this equation represents the surface of the shape. Signed-distance fields (Jones, Baerentzen, and Sramek, 2006) and level sets (Wenger, 2013) describe implicit representations of surfaces. Continuous representations can be computationally intensive for complex shapes and may struggle to represent detailed geometries, posing challenges for various shape analysis tasks, in particular shape correspondence.

1.1.2 Discrete Shape Representations

The geometric structure of a 3D object can be described discretely using various representations such as a polygon mesh and point cloud (Botsch, Kobbelt, Pauly, Alliez, and Levy, 2010), as visualized in Figure 1.2. Raw 3D data acquired by sensors is often represented as a point cloud, a set of points in 3D space. Each point $\mathcal{V}_i$ represents a sample of the object's surface, defined by its 3D coordinates (x_i, y_i, z_i). Point clouds capture shape geometry without any connectivity information among the points (Alexa, Behr, Cohen-Or, Fleishman, Levin, and Silva, 2003). Levoy (Levoy and Whitted, 1985) introduced point clouds to enhance visual efficiency, which are now being extensively utilized to effectively represent 3D objects in learning-based shape analysis techniques (Liu, Zhang, Kadam, and Kuo, 2021; Deng, Zhang, Ding, and Zhang, 2023). While point clouds provide a flexible and efficient way to represent geometry of shape, they lack explicit surface information. A 2D polygon mesh $\mathcal{X} = (\mathcal{V}, \mathcal{E}, \mathcal{F})$ embedded in 3D is a collection of vertices $\mathcal{V}$, edges $\mathcal{E}$ and faces $\mathcal{F}$, approximating the smooth surface with a piece-wise linear mapping containing both geometric and topological information of the shape. Each vertex $v_i \in \mathcal{V}$ is a 3D coordinate (x_i, y_i, z_i), and faces are often triangles, quadrilaterals, or other simple polygons, with each edge shared by at most two faces. Polygon meshes are versatile representations capable of modeling intricate geometric forms. Triangles, the simplest convex shapes, are favoured for their hardware-compatibility and minimal vertex count required to enclose a region (see Figure 1.2). Polyhedral meshes, similar to polygon meshes, extend the concept to 3D volumes by using polyhedra instead of polygons.

The choice of shape representation significantly influences the effectiveness of shape correspondence algorithms. Different representations offer varying levels of detail and computational efficiency, and therefore necessitate tailored approaches for establishing correspondences between shapes. For instance, the alignment of point clouds is achieved via RANSAC algorithm (Fischler and Bolles, 1987), surface matching involves detecting underlying deformations (Liu, Zhang, Shamir, and Cohen-Or, 2009), matching of shape skeletons is addressed in the work by Biasotti et al. (Biasotti, Marini, Spagnuolo, and Falcidieno, 2006), etc. In this work, triangular meshes are utilized for representing shapes.

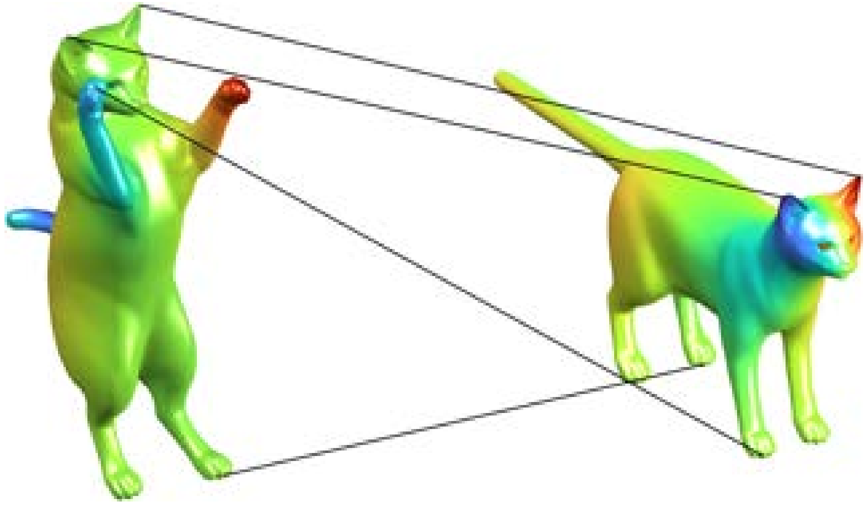

FIGURE 1.3: Meaningful point-to-point correspondences for sparse points illustrated for Cat shapes from TOSCA dataset (Bronstein, Bronstein, and Kimmel, 2008)

1.2 Shape Correspondence

Shape correspondence problem can be formally stated as: Given two (or more) shapes, say, $\mathcal{S}_1$ and $\mathcal{S}_2$, find a meaningful relation (or mapping) η among the elements (points, or components) of these shapes. In the discrete setting, determining a point-to-point map establishes correspondences between two shapes. For two meshes $\mathcal{X}_1$ and $\mathcal{X}_2$ with $|\mathcal{X}_1|$ and $|\mathcal{X}_2|$ vertices respectively, the point-to-point correspondence map is expressed as $\eta : \mathcal{X}_1 \rightarrow \mathcal{X}_2$ such that $\forall x_1 \in \mathcal{X}_1, \exists \eta(x_1) \in \mathcal{X}_2$. This map can be constrained in various ways depending on the application allowing for one-to-one, one-to-many, or many-to-many correspondences (Kaick, Zhang, Hamarneh, and Cohen-Or, 2011). Computing a meaningful mapping involves precisely matching specific points between shapes, such as ensuring that the left ear of one cat shape matches exactly with the left ear of another cat shape (see Figure 1.3). This requires understanding the shape's structure and functionality at both global and local levels. Finding a corresponding point on one shape for each point on another leads to an exponential solution space of $\mathcal{O}(n!)$ for a shape with n points. This combinatorial explosion makes the problem NP-hard, the most computationally challenging to solve.

Various aspects of shape correspondence problems have been explored,

FIGURE 1.4: Partial shape matching between a Horse shape and a Centaur shape from TOSCA dataset (Bronstein, Bronstein, and Kimmel, 2008; Bronstein, Bronstein, Bruckstein, and Kimmel, 2009), where bottom parts of both the shapes match

including full vs partial matching and sparse vs dense correspondences. In full matching, all shape elements are mapped, while in more challenging partial matching, only specific components are matched like between a horse and a centaur, though vastly different globally, share similar lower body shapes (see Figure 1.4). Sparse correspondences involve matching key components, such as in Figure 1.3, where matching the hands, legs, and head of cat shapes provide sufficient information for shape semantics (Pokrass, Bronstein, Bronstein, Sprechmann, and Sapiro, 2013). In contrast, dense correspondences are required for tasks like shape morphing (Zhang, Sheffer, Cohen-Or, Zhou, Van Kaick, and Tagliasacchi, 2008), but direct computation is infeasible, leading methods to first compute sparse correspondences and then extend them to dense ones (Lipman and Funkhouser, 2009). Below is a discussion of correspondence map computation, categorized by shape deformability:

1.2.1 Rigid Shape Matching

Determining correspondences for shapes that preserve pairwise Euclidean distances is traditionally treated as a rigid alignment (or registration) problem. Assuming rigidity, a shape can be transformed through geometric operations such as rotation, translation, reflection, or a combination of these.

(a) Deformation inducing bending with stretching, folds and creases

(b) Deformation inducing topological changes

FIGURE 1.5: Shapes of articulated wooden mannequins from SHREC 2019 Dataset (Dyke, Stride, Lai, and Rosin, 2019) subjected to various non-rigid deformations

Rigid registration seeks the best rigid transformation to align two shapes, often using the Iterative Closest Point (ICP) algorithm (Besl and McKay, 1992). This local-search method iteratively finds correspondences and minimizes distances between points, usually employing a variant of Hausdorff distance (Memoli, 2008). ICP is sensitive to initial alignment due to its greedy approach. To address occlusions and topological noise, several efficient variants of ICP (Rusinkiewicz and Levoy, 2001) have been proposed. Zhang et al. recently introduced a faster and more robust version of ICP (Zhang, Yao, and Deng, 2022), with improved convergence rates.

1.2.2 Non-rigid Shape Matching

Despite the success of rigid shape matching techniques, their utility is limited in real-world scenarios as shapes often display a diverse range of poses and deformations, causing significant variations in pairwise Euclidean distances. Shapes as depicted in Figure 1.5 can experience complex non-rigid deformations, such as stretching around the shoulders, folds, and creases on the left hand (see Figure 1.5(a)), as well as topological changes due to self-contact (illustrated in Figure 1.5(b)). These factors pose significant challenges in establishing meaningful pointwise correspondences. To enable mathematical analysis and streamline correspondence search for non-rigid deformations, traditional methods often restrict deformations to isometries or approximate isometries, as illustrated in Figure 1.1. Isometric deformations preserve geodesic distances (Kimmel and Sethian, 1998; Crane, Weischedel, and Wardetzky, 2013) between pairwise surface points, an intrinsic property exhibited by most real-world organic shapes. Elad and Kimmel (Elad and Kimmel, 2003) used pose-invariant embeddings based on intrinsic distances for correspondence. Extrinsic measures are deformation-sensitive, while intrinsic ones react to topology changes; Bronstein (Bronstein, Bronstein, and Kimmel, 2009) employed a joint similarity criterion to balance both for better correspondences. Stretching in deformations, increases the complexity of non-rigid shape matching, making tasks like semantically matching an elephant to a gorilla more challenging than matching two elephants in different poses. Kim et al. (Kim, Lipman, and Funkhouser, 2011) proposed an automatic method for smooth, low-distortion correspondence between non-isometric shapes, but it falls short for partial shape matching due to its global mapping approach without incorporating localization. For an in-depth understanding of non-rigid shape correspondences, readers are encouraged to refer the insights presented in (Bronstein, Bronstein, and Kimmel, 2008).

1.3 Shape Correspondence Approaches

Shape correspondence, a complex task, has inspired numerous techniques in shape analysis. Correspondence maps can be estimated using registration, similarity, or learning-based methods, as outlined below:

1.3.1 Registration or Alignment based approaches

Registration-based approaches align shapes by deforming one to match a target shape or by mapping both to a common domain via parametrization

(Deng, Yao, Dyke, and Zhang, 2022). Pointwise correspondence is then determined by finding points in proximity with each other, where alignment is achieved via a geometric transformation. Chang et al. (Chang and Zwicker, 2008) focus on piecewise rigid alignment by applying transformations to localized segments of the shape, rather than seeking a global transformation (refer Section 1.2.1). To register shapes undergoing non-rigid deformations with each point deformed uniquely, regularization is used to ensure that transformations are similar for points that are spatially close. This amounts to computing a deformation field for the source shape. Deng et al. (Deng, Yao, Dyke, and Zhang, 2022) examined non-rigid registration using only geometric information, while Saiti et al. (Saiti and Theoharis, 2020) addressed multi-modal shape registration with varying data structures, dimensions, densities, and noise.

1.3.2 Similarity based approaches

Directly computing correspondences between shapes, without first aligning or deforming them, involves leveraging geometric invariants *i.e.*, feature descriptors that remain stable under various transformations. By identifying the similarity of these descriptors, which capture intrinsic shape properties that remain unchanged under deformation, direct correspondence between shapes is achieved without explicit manipulation. These feature descriptors can be pointwise surface descriptors like Global Point Signatures (GPS) (Rustamov, 2007), Heat Kernel Signature (HKS) (Sun, Ovsjanikov, and Guibas, 2009; Bronstein and Kokkinos, 2010), Wave Kernel Signature (WKS) (Aubry, Schlickewei, and Cremers, 2011), or can be pairwise descriptors like geodesic distances (Kimmel and Sethian, 1998), biharmonic distances (Lipman, Rustamov, and Funkhouser, 2010), etc. employed for shapes undergoing isometric deformations.

Recent techniques have championed the use of learning-based descriptors including parametric spectral descriptors (Litman and Bronstein, 2013), descriptors learnt considering no known deformation model between the shapes (Corman, Ovsjanikov, and Chambolle, 2015), robust descriptors with relatively less training data learnt from raw shapes (Donati, Sharma, and Ovsjanikov, 2020). These descriptors are used to construct an objective function that maximizes the similarity between corresponding descriptors while minimizing the deformation required to align the shapes, without explicitly performing the alignment (Li and Iyengar, 2014). The objective function is optimized using established continuous optimization methods, such as linear programming as in (Berg, Berg, and Malik, 2005) or convex optimization as in (Zass and Shashua, 2008). For discrete optimization, tree-based searches as used in (Zhang, Sheffer, Cohen-Or, Zhou, Van Kaick, and Tagliasacchi, 2008) provide a viable approach. These techniques for directly computing pointwise correspondences often result in complex non-linear optimization problems that are not convex, making them difficult to solve. Adding global constraints further complicates the optimization process, introducing additional challenges.

Embeddings

A widely adopted hybrid approach involves embedding the shapes into a lower-dimensional space, where more efficient and computationally feasible mappings between them can be computed. By embedding the shapes into a shared geometric feature space, they can be aligned, enabling the computation of correspondences based on their spatial proximity. Embeddings can be computed using either isometric or spectral techniques. Isometric embedding preserves the pairwise distances on the surface of the shape during the transformation, ensuring that the original geometric relationships remain intact (Elad and Kimmel, 2003). Spectral methods, by exploiting the spectral properties of mesh operators, provide a powerful and efficient approach to capture the intrinsic geometric structure of shapes (Zhang, Van Kaick, and Dyer, 2010).

Functional Maps

The concept of functional maps, introduced by (Ovsjanikov, Ben-Chen, Solomon, Butscher, and Guibas, 2012) in the geometry processing community, was designed to match near-isometric manifolds. This approach seeks to establish a linear map between the function spaces, defined on shapes by aligning feature descriptors represented as real-valued functions. Functional maps measure the similarity of these functions and can be efficiently encoded as a low-rank matrix, depending on the chosen basis functions. Once computed, the functional map can be converted into point-to-point correspondences between shapes, providing an automatic, user-independent method that simplifies the complex problem of non-rigid shape correspondence. This efficiency has made functional maps a fundamental tool in geometry processing, serving as a cornerstone for various shape analysis tasks. Refer Chapter 2 for an in-depth discussion of functional maps.

1.3.3 Learning based approaches

Machine learning and deep neural networks are increasingly showing promise in the realm of shape correspondence (Masci, Rodolà, Boscaini, Bronstein, and Li, 2016). Xu et al. (Xu, Kim, Huang, Mitra, and Kalogerakis, 2016) introduced various data-driven approaches for determining correspondences. While supervised methods offer robust generalization capabilities, they require extensive datasets. Unsupervised methods, although class-agnostic and free from annotation requirements, often fall short of the performance achieved by their supervised counterparts (Halimi, Litany, Rodolà, Bronstein, and Kimmel, 2019). Learning-based methods can suffer from overfitting, reliance on application-specific shape priors, and the need for extensive data augmentation, making them computationally intensive and less feasible in resource-constrained environments (Donati, Sharma, and Ovsjanikov, 2020). Recently, Abdelreheem et al. (Abdelreheem, Eldesokey, Ovsjanikov, and Wonka, 2023) leveraged the capabilities of large language and vision models to address non-isometric shape matching, showcasing the advancements in this area.

1.4 Dataset

In order to assess the effectiveness of the proposed methodologies in this work, the TOSCA dataset (Bronstein, Bronstein, and Kimmel, 2008) was employed. This comprehensive collection features high-resolution, non-rigid 3D shapes across various poses, providing a robust benchmark for evaluating shape matching and correspondence tasks. The TOSCA dataset encompasses a diverse range of 80 objects, including animals (cats, dogs, wolves, horses), mythological creatures (centaurs), primates (gorillas), and human figures (male and female). Each object is represented by multiple poses, ranging from 7 to 20 poses, offering a wide variety of shape deformations and challenges for correspondence algorithms. The dataset offers precise point-to-point correspondences between shapes, serving as a benchmark for evaluating the accuracy and reliability of the proposed methods.

1.5 Contributions

In this work, the focus is on computing dense correspondence mapping between shapes undergoing non-rigid deformations, that cannot be aligned by a single rigid transformation. Shapes are represented as triangular meshes that discretize the surface of real-world objects with each surface point potentially deforming independently. The complexity of matching such shapes is addressed by exploiting surface characteristics that remain invariant under isometric deformations. The goal of this work is to investigate robust correspondence mapping techniques capable of capturing and preserving fine-grained geometric details more accurately.

1.5.1 Publications based on Thesis

During the course of this thesis, the following contributions were achieved:

Conferences

First approach is based on the functional map framework, which utilizes Hamiltonian operator modulated with Gaussian curvature that localizes and encodes sharp shape features to enhance pointwise correspondences by counteracting the Laplacian smoothing bias, while maintaining the framework's efficient linear transfer of geometry. Chapter 4 presents a detailed account of the work reported in WSCG 2023 (World Society for Computer Graphics) as follows:

> Bindal, M., Kamat, V. *Detail Preserving non-rigid shape correspondences*, 31st International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision. WSCG 2023. Czech Republic. Computer Science Research Notes, 3301:323330, 07 2023
> URL: `https://www.doi.org/10.24132/CSRN.3301.36`

Second approach improves pointwise correspondence by embedding the shape in Euclidean space via shape-aware biharmonic distances, preserving local details while being robust to noise. This enhanced metric enables construction of pose-invariant canonical forms, which, when coupled with non-rigid ICP techniques produce robust Euclidean alignment and higher correspondence accuracy over rigid approaches. Chapter 3 presents a detailed account of the work reported in CVIP 2023 as follows:

> Bindal, M., Kamat, V. (2024). Improved Metric Space for Shape Correspondence. In: Kaur, H., Jakhetiya, V., Goyal, P., Khanna, P., Raman, B., Kumar, S. (eds) Computer Vision and Image Processing. CVIP 2023. Communications in Computer and Information Science, vol 2009. Springer, Cham.
> URL: `https://doi.org/10.1007/978-3-031-58181-6_32`

Journals

In the functional map framework, basis selection is crucial in determining correspondence quality. A comprehensive comparative analysis of the diverse basis functions utilized to embed deformable shapes within the functional map framework is investigated. The bases functions are evaluated on completeness, compactness, stability, orthogonality, multiscale behavior, smoothness, robustness, and locality. The LBO eigenbasis is effective for near-isometric deformations but underrepresents high-frequency detail and lacks spatial localization. Localized bases (*e.g.*, HO, LMH) more accurately capture sharp, confined features. For non-isometric matching, CQHB and Elastic bases are preferred, whereas LAB is preferred for landmark preservation. Chapter 5 discusses the work in detail as published in the following journal and reflects a shift towards localized and extrinsic feature-aware bases functions.

> Bindal, M., Kamat, V. (2024). Diverse bases for functional spaces for non-rigid shape correspondence. Ingénierie des Systèmes d'Information, Vol. 29, No. 4, pp. 1405-1422.
> URL: `https://doi.org/10.18280/isi.290415`

1.6 Organization

The thesis is organized into the following chapters, as outlined below. Relevant terminologies are introduced and defined as they are encountered throughout the text.

In Chapter 1 the problem of shape correspondence along with its crucial applications has been introduced. Beginning with different shape representations *i.e.*, continuous and discrete, an account of various approaches used to address the correspondence problem is provided. A list of research articles contributed during the PhD period is also included.

In Chapter 2, the embedding of high-dimensional complex geometric shapes in low-dimensional spaces is discussed in detail. Various embedding techniques are explored, demonstrating their effectiveness in managing non-rigid deformations and clarifying the distinction between *intrinsic* and *extrinsic* shape geometries.

In Chapter 3 shape is considered as a metric space and describes how it can be improved upon to capture local details that refines the overall pointwise correspondences.

In Chapter 4 the introduction of localization within the functional map pipeline is discussed, aiming to improve the accuracy of point-to-point correspondences for deformable shapes.

In Chapter 5 a comparative analysis of the bases functions utilized to embed the shapes via functional maps has been discussed.

In Chapter 6 the thesis is concluded by highlighting the contributions and providing future directions.

In this thesis, similarity-based approaches are emphasized to establish pointwise correspondences for shapes undergoing non-rigid deformations, majorly isometric deformation. Chapter 2 offers an in-depth examination of embedding-based techniques, presenting them as a pivotal method for attaining precise correspondences.

Chapter 2

Shape Embedding

2.1 Introduction

Shapes, represented as surfaces, are fundamentally *embedded surfaces* existing within the ambient Euclidean space $\mathbb{R}^3$. Embedded surfaces played a pivotal role in the development of non-Euclidean geometries in the 19th century. Recognizing the foundational value of these surfaces, the eminent German mathematician, Riemann, introduced the concept of intrinsic properties. These are inherent characteristics of a surface that can be determined solely by measurements made on the surface itself, without any consideration of the surrounding space. Conversely, extrinsic properties depend on the surface's relationship to it's ambient space. These properties are influenced by the surface's position and orientation within that space. To illustrate, consider Figure 2.1, which demonstrates the measurement of the distance between a cat's left and right front paws. If the distance is measured along the cat's body, following its curves, that is the geodesic distance (depicted with red curved line) – it stays on the surface and doesn't change if the cat sits, stands or even moves. But if the distance is measured in a straight line, as if poking a ruler through the air between its paws, it is the Euclidean distance (depicted with pink line). The straight line changes length as the cat moves its paws, because it is not stuck to the cat's surface. Geodesic distance represents an intrinsic measure, whereas Euclidean distance reflects an extrinsic one. Therefore, intrinsic properties of a deformable shape are unique to its internal structure and remain constant despite any changes in its position or

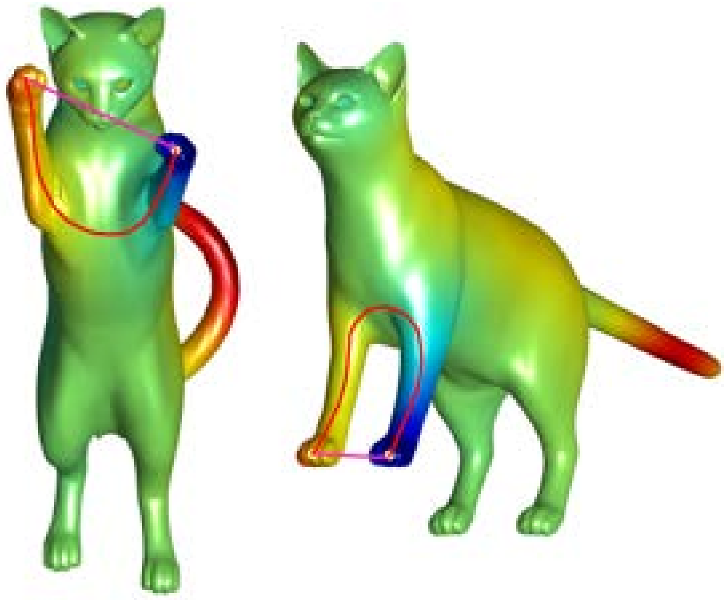

FIGURE 2.1: Illustration of cats in various poses showing intrinsic geodesic distance (red curve) and extrinsic Euclidean distance (pink line) between two corresponding points on the front paws

orientation. Whereas extrinsic properties are subject to modification as the surface's position or orientation in the space changes.

Determining correspondences for shapes undergoing non-rigid deformations, focuses on identifying stable intrinsic invariants that remain consistent, regardless of how the shape changes during the deformation process. High-dimensional shape representations often include redundant information, which increases computational complexity and slows processing times. By transforming shapes into alternative spaces, the matching process can be simplified, especially if the matching task can be more effectively modeled or solved within the constraints of the new domain. This led to the idea of embedding shapes in lower-dimensional spaces, potentially turning non-rigid shape matching challenge into known rigid matching task. For example, consider two deformable 3D shapes like human poses captured in different postures. By embedding them into a low-dimensional space that captures key features like joint angles, the matching process becomes simpler, as it focuses on fewer, more meaningful details. Therefore, *Shape embedding* involves mapping complex shapes with high-dimensional structure into lower-dimensional spaces while retaining their crucial features, usually intrinsic.

This process streamlines shape representation, making it more efficient for computing correspondences while preserving essential geometric and topological properties. This approach works effectively well for non-rigid shape correspondence, as it accommodates deformations like stretching and bending (refer Figure 1.5). The need to embed shapes in an alternate domain arises not only from computational complexity but also because two non-rigidly (or isometrically) deformed shapes belong to separate metric spaces. A metric space defines distances between points on a shape, with the metric structure of a shape X given by the pair (X, d_X), where d_X measures distances between points. Each non-rigid shape has it's own metric structure. For instance, consider two identical balls of modeling clay, depicting perfect spheres. Overlay a uniform grid of points onto each sphere and then deform one sphere into an irregular oval shape. This non-rigid transformation distorts the grid, altering the distances between the points so that they no longer correspond to their original values, highlighting that the spherical and oval shapes now exist in separate metric spaces. This difference in metric spaces makes direct distance comparisons between points on the original and deformed spheres invalid. To establish correspondences, non-rigid shapes must be mapped to a common metric space. In contrast, rigid shapes naturally belong to a shared Euclidean space, usually $\mathbb{R}^3$. Mapping non-rigid shapes into *functional* or *metric spaces* facilitates a more structured and efficient approach to matching, alignment, and comparison by transforming complex geometric data into a domain where relationships between points, features, or regions can be analyzed more effectively. Refer Chapter 3 for a detailed discussion on representing shapes as metric space.

The key is to obtain an intrinsic representation of the shape that remains invariant under isometric deformations, allowing it to be embedded into a new space where its intrinsic geometry is reflected in its extrinsic form. This reduces the computational complexity by representing transformations in the common domain with fewer parameters. In the machine learning terminology, this non-linear technique of reducing the dimensionality of the data, is referred to as *manifold learning* as well as *dimensionality reduction*. This chapter explores various approaches to achieving shape embedding.

2.2 Related Work

Metric embeddings offer a valuable technique for shape correspondence by embedding shapes into a metric space, simplifying complex problems by leveraging the shape's inherent geometry, especially when preserving its structure is crucial. Schwartz et al. (Schwartz, Shaw, and Wolfson, 1989) introduced flat embedding using MDS *Multidimensional Scaling* to flatten cortical surfaces, enabling uniform parameterization and easier surface matching. Zigelman et. al. (Zigelman, Kimmel, and Kiryati, 2002) extended it for texture mapping. Elad and Kimmel (Elad and Kimmel, 2003) proposed to compute a pose-invariant shape form, a low-dimensional representation in shared metric space, also known as *embedding space*, describing the original shape effectively. Since this shared space is Euclidean, it is challenging to perfectly represent the shape's true geometry due to the difficulty of mapping a curved surface onto a flat one. Memoli et. al. (Mémoli and Sapiro, 2005) proposed to use isometric-invariant Gromov-Hausdorff distance that can theoretically avoid representation errors. Bronstein et al. (Bronstein, Bronstein, and Kimmel, 2006) proposed a generalized framework utilizing Gromov-Hausdorff distance to compute partial embedding distances for surface matching. Memoli (Mémoli, 2007) further improved this by introducing probability measures, enabling more natural shape discretizations.

The search for an appropriate shared space motivated the idea of not restricting to Euclidean domains as the shared space, rather exploring non-Euclidean spaces such as a spherical domain (Cox and Cox, 1991), or can be one of the shapes itself (Bronstein, Bronstein, and Kimmel, 2008). Lipman et al. (Lipman and Funkhouser, 2009) proposed a six-parameter Möbius transformation to conformally embed shapes into disks, preserving angles but not necessarily lengths or areas. Building on this, Kim et al. (Kim, Lipman, and Funkhouser, 2011) enhanced the approach by constructing globally consistent matchings derived from local conformal maps, ensuring improved coherence across the embedded shapes. However, these methods can introduce potential embedding errors and uncertainties, such as unbounded conformal factors and sign ambiguities. To address the correspondence ambiguities in symmetric shapes and the challenges posed by increasingly high-resolution meshes, these approaches leverage higher-order structures, hierarchical solvers (Raviv, Dubrovina, and Kimmel, 2012), or reduce computational complexity by focusing on matching shape segments rather than computing dense point-to-point correspondences. By exploiting the eigenspace of a surface-defined linear operator to embed shapes has inspired numerous advancements in shape analysis. This approach has refined pointwise correspondences and driven innovations in diverse applications such as shape compression, parameterization, clustering, segmentation, and remeshing (Zhang, Van Kaick, and Dyer, 2010). Shtern et al. (Shtern and Kimmel, 2015) introduced spectral embedding techniques leveraging the gradient fields of the Laplacian operator's eigenspace, significantly enhancing pointwise correspondences. By defining a space of functions for each shape, functional maps introduced by Ovsjanikov et al. (Ovsjanikov, Ben-Chen, Solomon, Butscher, and Guibas, 2012) aim to determine correspondences by aligning these functional spaces. This paradigm shifts focus from pointwise matching to comparing entire function spaces, enhancing flexibility and robustness in shape analysis.

2.3 Explicit Embedding

In explicit embedding, shapes are directly embedded into a target space, usually a lower-dimensional Euclidean space, where point-to-point correspondences between the shapes are established based on proximity. For instance, a sphere can be explicitly embedded onto a plane using spherical projection, mapping each point on the sphere's surface to a corresponding point on the plane. This involves constructing a map that explicitly places the shape in a specific coordinate system. While this approach emphasizes geometric precision and direct point matching, it can be less flexible when dealing with shape deformations. In this type of embedding, identifying pointwise correspondences is straightforward, as they can be directly inferred from the spatial relationships within the embedded domain. The following sections present an overview of different explicit embedding techniques.

2.3.1 Isometric Embedding

Isometric embedding involves embedding the shape in an alternate geometric space, such that the pairwise distances between all points on the original shape are preserved, ensuring its intrinsic geometry remains unchanged. Within the *m-dimensional* Euclidean domain $\mathbb{R}^m$, the embedded shapes are equivalent upto the isometries, *i.e.*, the shapes in the embedded domain are related by rigid motions, including rotations, translations, or reflections only. For instance, a flat sheet of paper represents a 2D surface with intrinsic geometry defined by the distances between points. Bending it into a cylinder, cone, or another curved shape without stretching or tearing is an example of isometric embedding. The distances between points remain unchanged, preserving the sheet's intrinsic geometry even as its form changes in 3D space. This illustrates how isometric embedding maintains intrinsic geometry while altering its embedding in a higher-dimensional space. Formally, to embed a shape $(\mathcal{S}, d_{\mathcal{S}})$ into an *m-dimensional* Euclidean domain $\mathbb{R}^m$ involves determining a mapping f_e:

$$f_e : (\mathcal{S}, d_{\mathcal{S}}) \rightarrow (\mathbb{R}^m, d_{\mathbb{R}^m}) \tag{2.1}$$

such that the distances between any two points p and q on the shape are preserved in the embedding *i.e.*,

$$d_{\mathcal{S}}(p, q) = d_{\mathbb{R}^m}(f_e(p), f_e(q)) \quad \forall p, q \in \mathcal{S}$$

Here, function f_e represents the desired isometric embedding, and $f_e(\mathcal{S})$ is the invariant representation of shape in $\mathbb{R}^m$, also known as *canonical form* of the shape; $d_{\mathcal{S}}$ denotes the intrinsic distance between points p and q on the

FIGURE 2.2: Isometric embedding obtained through the multidimensional scaling method, showing the original shape on the left and its computed canonical form on the right

original shape, while $d_{\mathbb{R}^m}$ represents the Euclidean distance between their corresponding images $f_e(p)$ and $f_e(q)$ in the embedded space (Bronstein, Bronstein, and Kimmel, 2008). For the shape $\mathcal{S}$, in order to determine an isometric embedding, first step is to establish its metric structure $d_{\mathcal{S}}$, which incorporates defining a discrete linear operator for the shape. This operator represented as a matrix is built from the shape's intrinsic geometry and topological features. The matrix can be created as a pairwise affinity matrix, reflecting the strength of connection between any two points on the shape. Alternatively, it can identify the neighbouring points for each element or represent any distance measure across the shape's surface. Next task is to compute a canonical form for the shape such that isometry is maintained, satisfying the constraint in the Equation (2.1). Identifying such embedding by ensuring that distances are accurately preserved is essential for various other shape analysis tasks too like deformation analysis, attribute transfer, etc. However, every cartographer faces the impossible task of flattening the Earth without distorting its intrinsic distances, a dilemma that parallels the challenges faced by the shape analysis research community as well, as they navigate the complexities of deformable shapes. Linial's work (Linial, 2002) on metric spaces, shows that even a simple shape like a sphere, sampled at four points, cannot be embedded in finite-dimensional Euclidean space without distorting distances. For more complex shapes, achieving an exact isometric embedding is even less feasible. This also aligns with *Gauss Theorema Egregium* (Do Carmo, 2016), which states that a surface with positive Gaussian curvature, like a sphere, cannot be isometrically embedded into a plane, which has zero Gaussian curvature. Therefore, more complex shapes cannot be perfectly embedded isometrically while preserving exact distances between points. The goal then shifts to minimizing distortions in distances as much as possible to maintain the shape's intrinsic geometry. This is achieved via powerful multidimensional scaling (MDS) techniques, as detailed in (Borg and Groenen, 2005). These methods aim to find an embedding f_e into $\mathbb{R}^m$ that minimizes distortion, under the condition:

$$f_e = \underset{f_e:\mathcal{S}\to\mathcal{R}^m}{\operatorname{argmin}} \sum_{\mathrm{p,q}\in\mathcal{S}} \mid \mathrm{d}_{\mathbb{R}^\mathrm{m}}\left(\mathrm{f_e}\left(\mathrm{p}\right), \mathrm{f_e}\left(\mathrm{q}\right)\right) - \mathrm{d}_{\mathcal{S}}\left(\mathrm{p}, \mathrm{q}\right) \mid^2$$

which describes a distance-preserving mapping, aiming to optimize for a function f_e in an iterative manner that maps points from $\mathcal{S}$ into $\mathbb{R}^m$ such that the distances between points in the target embedded space $\mathbb{R}^m$ closely resemble their original distances in $\mathcal{S}$. MDS uses a distance matrix of all pairwise points on shape to embed them in the m-dimensional Euclidean space, preserving their original distance relationships by minimizing discrepancies through optimization. Figure 2.2 illustrates an isometric embedding of a human shape, where the canonical form is obtained using the MDS algorithm, representing a bending-invariant signature of human shape. For details on the various versions of multidimensional scaling techniques proposed such as classic MDS, least-squares MDS, etc. please refer to (Bronstein, Bronstein, and Kimmel, 2008).

2.3.2 Spectral Embedding

While MDS-based techniques offer approximate isometric embedding, their iterative nature makes them computationally expensive. Spectral embedding refers to the process of determining a bending-invariant representation of a shape by utilizing the eigenspace of a linear operator. Isospectral properties of eigenvectors, a fundamental concept in linear algebra, have been instrumental in advancing the theoretical foundations of shape analysis. Spectral correspondence techniques, which align eigenmodes based on their eigenvalue magnitudes, have been widely employed in this domain. For instance, Belkin et al. (Belkin and Niyogi, 2003) introduced Laplacian Eigenmaps to facilitate shape correspondence, leveraging the spectral decomposition of the Laplacian matrix. Similarly, Jain et al. (Jain and Zhang, 2006) utilized Principal Component Analysis (PCA) on normalized geodesic affinity matrices to compute eigenmodes, further enhancing spectral-based shape analysis. While spectral correspondence techniques offer robust tools for shape analysis, practical challenges arise in handling high-resolution shapes. Storing pairwise relationship matrices incurs substantial memory overhead, and canonical forms based on geodesic distances often exhibit sensitivity to topological changes. To address these limitations, Karni et al. (Karni and Gotsman, 2000) pioneered the use of eigenvectors derived from the topological Laplacian for mesh compression, similar to JPEG compression for images. This innovation highlighted the potential of spectral graph theory (Chung, 1997), which focuses on sparse representations that exploit local neighbourhood information instead of dense pairwise relationships, enabling more efficient eigenvector computation. Building on these foundations, Rustamov et al. (Rustamov, 2007) extended the eigendecomposition of the Laplace-Beltrami Operator, as introduced by Levy (Levy, 2006) and Reuter et al. (Reuter, Wolter, and Peinecke, 2006), to define GPS (*Global Point Signature*) embeddings. These embeddings effectively captured intrinsic geometric properties while demonstrating robustness against topological

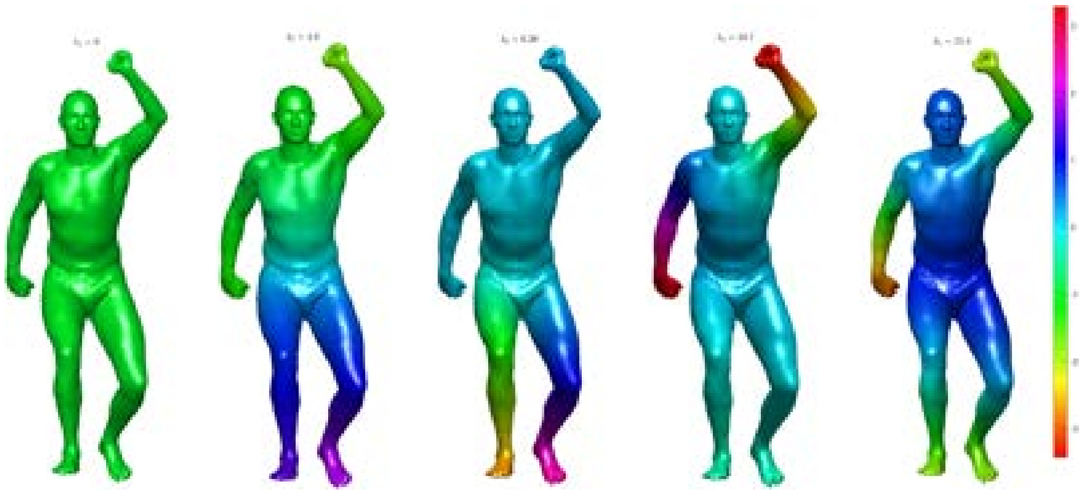

FIGURE 2.3: Laplace-Beltrami eigenfunctions in the ascending order of respective eigenvalues

noise. Spectral methods for mesh processing operate by constructing matrices that represent the mesh's structure and analyzing their eigenvalues and eigenvectors to derive meaningful solutions. For an in-depth discussion on spectral analysis via Laplace-Beltrami eigenfunctions and their applications in shape analysis, please refer (Lévy and Zhang, 2010).

Laplace-Beltrami Operator

In Euclidean space, the Laplacian of a function is the sum of its second-order partial derivatives with respect to each Cartesian coordinate. Laplace-Beltrami Operator, the *Swiss Army Knife of Geometry Processing* (Solomon, Crane, and Vouga, 2014), is the generalization of the Laplacian to Riemannian manifolds. It accounts for curvature as the divergence of the gradient on the manifold. For a shape $\mathcal{S}$ and a scalar field $f : \mathcal{S} \rightarrow \mathbb{R}$, the self-adjoint Laplace-Beltrami operator (LBO) $\Delta_{\mathcal{S}}$ captures intrinsic geometric properties of shape and is defined as the negative divergence of the gradient of scalar field as $\Delta_{\mathcal{S}} f = -div \nabla f$. This operator admits an eigendecomposition, yielding non-negative eigenvalues λ and corresponding orthonormal eigenfunctions ϕ,

$$\Delta_{\mathcal{S}} \phi = \lambda \phi$$

The eigenfunctions of the Laplace-Beltrami operator, as visualized on a human shape in the increasing order of eigenvalues in Figure 2.3 and commonly referred to as *manifold harmonics* (Vallet and Lévy, 2008), capture the geometric and topological properties of shapes. These eigenfunctions are in principle similar to Fourier bases on Euclidean domains. The Laplace-Beltrami operator, $\Delta_{\mathcal{S}}$, is independent of any global coordinate system and the underlying shape representation, whether it be meshes, point clouds, or volumes.

For a triangle mesh $\mathcal{X}$, the Laplace-Beltrami operator is effectively discretized using the cotangent scheme, yielding the representation $\mathbf{L} = \mathbf{D}^{-1}\mathbf{W}$. Here $\mathbf{W}$ captures the local geometry based on edge angles, while $\mathbf{D}$ is a lumped mass matrix consolidating the vertex area information (Meyer, Desbrun, Schröder, and Barr, 2003). The eigendecomposition of the discrete Laplace-Beltrami operator leads to the generalized eigenvalue problem $\mathbf{W}\Phi = \mathbf{D}\Phi\Lambda$ which can be numerically formulated as the following minimization problem:

$$\underset{\Phi}{\arg\min} \quad \mathrm{tr}\left(\Phi^T \mathbf{W} \Phi\right) \quad \text{s.t.} \quad \Phi^T \mathbf{D} \Phi = \mathbf{I} \tag{2.2}$$

where $\Phi \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{X}|}$ is the eigenvector matrix, defined as $\Phi = (\phi_1, \phi_2, ..., \phi_{|\mathcal{X}|})$ where each column ϕ_i is the eigenfunction arranged by increasing eigenvalues. Equation (2.2) is also equivalent to the generalized eigenvalue problem $\mathbf{W}\Phi = \mathbf{D}\Phi\Lambda$ where $\Lambda = (\lambda_1, \lambda_2, ..., \lambda_{|\mathcal{X}|})$ is the diagonal eigenvalue matrix. Despite being constructed in a local manner, eigenfunctions of this geometric operator capture global information of the shape and is considered as the workhorse of geometry processing especially for isometrically deformed shapes. Diverse discretizations have been proposed for various shape representations, include one for manifold meshes (Wardetzky, Mathur, Kaelberer, and Grinspun, 2007), one that is free from underlying mesh topology (Petronetto, Paiva, Helou, Stewart, and Nonato, 2013) and other for non-manifold meshes (Sharp and Crane, 2020). Recently, Bunge et al. (Bunge and Botsch, 2023) conducted a comprehensive study on the discrete Laplace-Beltrami operator for general polygon meshes while Belkin et

FIGURE 2.4: Visualization of real-valued functions on the TOSCA wolf shape (Bronstein, Bronstein, and Kimmel, 2008), showing geodesic distances from two different source points (left, right) and mean curvature (middle), with blue to red indicating low to high values

al. (Belkin, Sun, and Wang, 2009) investigated the operator for point clouds. The Laplace-Beltrami operator eigenbasis has also been shown to be optimal for representing smooth functions with bounded variation on shapes, as it is complete, compact, orthogonal, multi-scale, and computationally efficient (Aflalo, Brezis, and Kimmel, 2015).

2.4 Implicit Embedding

In implicit embedding, instead of embedding the shape directly, functions or properties defined on the shape such as curvature, distance functions, etc. are embedded into an alternative space. Correspondence between shapes is determined through these embedded functions, rather than direct point-to-point matches. This approach offers greater robustness against non-rigid deformations, but sacrifices straightforward point-to-point interpretation. Additionally, it requires an extra step to convert the mapping from the embedded space back to the original domain.

2.4.1 Functional Maps

The concept of functional maps, introduced by Ovsjanikov et al. (Ovsjanikov, Ben-Chen, Solomon, Butscher, and Guibas, 2012), provides a framework that focuses on deriving a linear transformation between function spaces defined on the shapes. Since their introduction, functional maps have become a powerful tool in the geometry processing field, widely applied across various shape analysis tasks. Unlike traditional methods that rely on comparing physical coordinates, functional maps operate in a more abstract space defined by geometric functions. They encode shape similarities at varying levels of detail, offering a flexible and adaptable representation. By expressing these correspondences as matrices, functional maps can leverage the robust tools of numerical linear algebra for efficient manipulation and sophisticated inference tasks.

Formally, a real-valued function $f\colon \mathcal{S} \to \mathbb{R}$ defined on a shape $\mathcal{S}$ characterizes a geometric property, such as curvature at every point or the surface distance from a specific point to all other points on the shape (refer Figure 2.4). Let $\mathcal{F}(\mathcal{S}, \mathbb{R})$ represent the space of scalar-valued functions defined on shape $\mathcal{S}$. For shapes $\mathcal{S}_1$ and $\mathcal{S}_2$ with corresponding functional spaces $\mathcal{F}(\mathcal{S}_1, \mathbb{R})$ and $\mathcal{F}(\mathcal{S}_2, \mathbb{R})$ respectively, a bijective mapping $\eta : \mathcal{S}_1 \to \mathcal{S}_2$ induces a linear transformation $\eta_F\colon \mathcal{F}(\mathcal{S}_1, \mathbb{R}) \to \mathcal{F}(\mathcal{S}_2, \mathbb{R})$. This functional map provides a unique representation, allowing any function $f_1 \in \mathcal{F}(\mathcal{S}_1, \mathbb{R})$ to correspond to $f_2 \in \mathcal{F}(\mathcal{S}_2, \mathbb{R})$ through the relation $f_2 = \eta_F(f_1)$ or via $f_2 = f_1 \circ \eta^{-1}$. Despite the complexity of η, η_F operates linearly between function spaces, and understanding η_F equates to understanding η. For proofs, kindly refer (Ovsjanikov, Ben-Chen, Solomon, Butscher, and Guibas, 2012). The space $\mathcal{F}(\mathcal{S}_1, \mathbb{R})$ has a basis $\{\phi_i\}_{i\geq 1}$ such that any function $f_1 \in \mathcal{F}(\mathcal{S}_1, \mathbb{R})$ can be represented as a linear combination $f_1 = \sum_i \alpha_i \phi_i$. The coefficients α_i are determined by projecting f_1 onto each basis function ϕ_i. Since η_F is a linear mapping,

$$\eta_F\ (f_1) = \eta_F\ (\sum_i \alpha_i \phi_i) = \sum_i \alpha_i\ \eta_F(\phi_i)$$

. Simultaneously, if the function space $\mathcal{F}(\mathcal{S}_2, \mathbb{R})$ is characterized by basis $\{\psi_j\}_{j\geq 1}$, then $\eta_F(\phi_i) = \sum_j \beta_{ji}\psi_j$ for some $\{\beta_{ji}\}$ such that

$$\eta_F\ (f_1) = \sum_j \sum_i \alpha_i\ \beta_{ji}\psi_j \tag{2.3}$$

If f_1 is represented by the coefficient vector $\mathbf{a} = (\alpha_0, \alpha_1, ..., \alpha_i, ...)$ and $f_2 = \eta_F(f_1)$ by $\mathbf{b} = (b_0, b_1, ..., b_i, ...)$, then Equation (2.3) indicates that $b_j = \sum_i \beta_{ji}\alpha_i$, where β_{ji} represents the j^{th} coefficient of the i^{th} basis function from $\mathcal{F}(\mathcal{S}_1, \mathbb{R})$ mapped to $\mathcal{F}(\mathcal{S}_2, \mathbb{R})$, *i.e.*, j^{th} coefficient of $\eta_F(\phi_i)$ expressed in basis ψ_j. The coefficients β_{ji} are independent of f_1 and are determined solely by the basis and the map η. If the basis functions $\{\psi_i\}$ are orthonormal with respect to an inner product $\langle \cdot, \cdot \rangle$, then $\beta_{ji} = \langle \eta_F(\phi_i), \psi_j \rangle$, simplifies the structure of the matrix $\mathbf{C}$.

For discrete shapes $\mathcal{X}_1$ and $\mathcal{X}_2$, sampled at n points with basis matrices $\Phi, \Psi \in \mathbb{R}^{n \times k}$, the functional map η_F can be expressed as a matrix $\mathbf{C} \in \mathbb{R}^{k \times k}$ that translates coefficients between bases. For any function $\mathbf{f}_1 \in \mathcal{F}(\mathcal{X}_1, \mathbb{R})$, with coefficient vector $\mathbf{a}$, the map acts as $\eta_F(\mathbf{a}) = \mathbf{Ca}$, fully encoding the pointwise correspondence map η.

The approach focuses on adding structure to the shape and operating on that structure, rather than directly on the shapes themselves. Functional maps are effective for shape matching, as they handle high-resolution shapes efficiently by reducing the dimensionality where shape analysis occurs. This work primarily focuses on embedding deformable shapes into an alternative domain to enhance pointwise correspondence accuracy.

Chapter 3

Shape as Metric Space

3.1 Introduction

Establishing accurate correspondences between non-rigid shapes while preserving their intrinsic properties remains a fundamental challenge, given the diverse ways a shape can deform. While extrinsic properties, such as position and orientation, depend on the surrounding environment, intrinsic properties, like curvature and topology, are tied to the shape's internal structure and remain independent of the ambient space. Viewing a shape as a metric space, which is a space of points equipped with distance functions, highlights this distinction: extrinsic properties are quantified using Euclidean distances between points, whereas intrinsic properties are captured through surface-based distances, measured as the length of paths along the shape itself. As discussed in Chapter 2, isometrically deformed shapes, $\mathcal{S}_1$ and $\mathcal{S}_2$, are formally treated as distinct metric spaces. These are represented as $\left(\mathcal{S}_1, d_{\mathcal{S}_1}\right)$ and $\left(\mathcal{S}_2, d_{\mathcal{S}_2}\right)$ respectively, where $d_{\mathcal{S}}$ denotes the associated metric for each shape. A metric space is defined for a set M along with the distance function $d \colon M \times M \to \mathbb{R}$, denoted as (M, d_M) and satisfies various properties for elements $m_1, m_2, m_3 \in M$ such as

- non-negativity *i.e.*, $d_M(m_1, m_2) \geq 0$
- indistinguishability *i.e.*, $d_M(m_1, m_2) = 0$ iff $m_1 = m_2$
- symmetry *i.e.*, $d_M(m_1, m_2) = d_M(m_2, m_1)$
- triangle inequality *i.e.*, $d_M(m_1, m_3) \leq d_M(m_1, m_2) + d_M(m_2, m_3)$

The elements of a set M are referred to as points, and the metric $d_M(m_1, m_2)$ defines the distance between any two points m_1 and m_2. For example, an m-dimensional Euclidean space $\mathbb{R}^m$ can be viewed as a metric space, where the associated metric is given by the Euclidean norm, $d_{\mathbb{R}^m}(x_1, x_2) = ||x_1 - x_2||_m$, for points $x_1, x_2 \in \mathbb{R}^m$. This metric establishes a sense of proximity, categorizing points as *near* or *far*. During shape deformations, the metric structure plays a crucial role in enforcing the preservation of intrinsic geometry by ensuring that points close to each other in the original shape remain mapped to nearby points in the deformed shape. The metric space model extends the notion of pointwise distances to comparisons between entire spaces. Key distance measures include the Hausdorff distance, which compares subsets within a metric space, the Wasserstein metric, designed for comparing probability distributions, and the Gromov-Hausdorff distance, which evaluates pairwise distances between entire metric spaces (Gromov, 2007; Bronstein, Bronstein, Kimmel, Mahmoudi, and Sapiro, 2009).

Direct comparison of shapes represented as distinct metric spaces is challenging due to differences in their metric structures. While rigid shapes are often treated as subsets of a shared Euclidean space, $\mathbb{R}^3$, non-rigid shapes require embedding into a common metric space to establish meaningful correspondences. This process involves identifying point-to-point correspondences by encoding all pairwise relationships—such as distances, similarities, or affinities—among shape elements into a matrix. An embedding is then computed to preserve these relationships as faithfully as possible, even accounting for distortions introduced during the transition to the Euclidean domain (see Chapter 2). Once these canonical forms, serving as shape embeddings, are derived, alignment and closest-point techniques are employed to establish accurate pointwise correspondences between the shapes.

This chapter investigates whether enhancing the metric space can improve point-to-point correspondences between two isometrically deformed shapes. The discussion is focussed on how enriched metric structures can better capture the geometric and topological properties of shapes, facilitating more accurate correspondences. Traditionally, pairwise geodesic distances have been employed to establish deformation-invariant canonical forms. However, these methods face challenges, such as the sensitivity of geodesic distances to topological noise. To address this, shape-aware distances are utilized, offering smoother and more robust invariant representations. By integrating these refined distance measures with non-rigid alignment techniques in Euclidean space, the approach effectively improves point-to-point correspondences, even across diverse shape triangulations.

3.2 Background

Measuring the distance $d_{\mathcal{S}}(s_1, s_2)$ between two points s_1, s_2 on a surface $\mathcal{S}$ is a fundamental challenge in geometry. Various distance metrics have been developed and widely employed in shape analysis to capture surface features and underlying structures. Accurate distance measurement between surface points is crucial for establishing correspondences in deformable shapes, as it aids in understanding shape similarity despite geometric variations. These metrics not only reveal relationships among surface points but also inspire definitions of curvature, bending, and other geometric properties. Different distance metrics serve distinct purposes based on the surface characteristics and desired accuracy. Ideally, a distance metric should satisfy several criteria: it should be a true metric, locally isotropic, globally shape-aware, invariant to isometries, robust to noise and topology variations, computationally efficient, and free of parameter dependence.

3.2.1 Geodesic Distances

Distance functions are defined as metrics, with geodesic distance representing the shortest path between two points on the surface of the shape. Unlike the straight-line distance in the ambient space, geodesic distance remain invariant under isometric deformations of shape. In the early stages of developing methods for computing distances on shapes, Dijkstra's algorithm

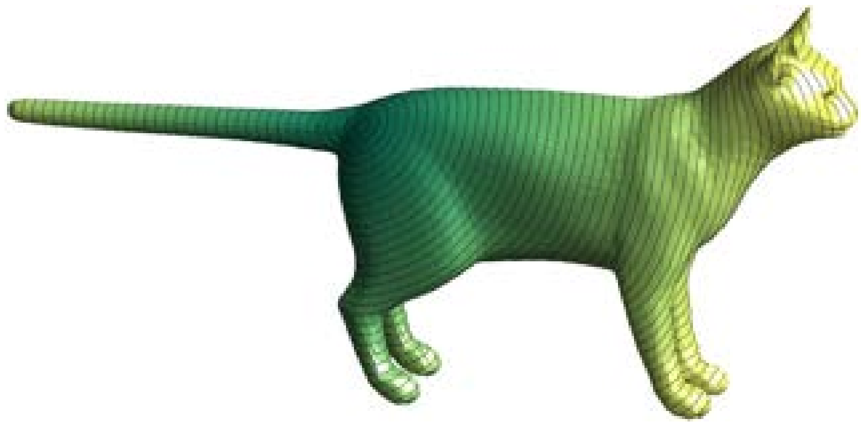

FIGURE 3.1: Geodesic distances computed from a point near the tail of a TOSCA cat mesh (Bronstein, Bronstein, and Kimmel, 2008) to all other points on the surface

(Dijkstra, 1959), from graph theory was a common choice. This method involved selecting a fixed source point and calculating the distances to all other vertices on the shape. While Dijkstra's algorithm served as a practical starting point, it had inherent limitations. As it relied on traversing mesh edges, it often failed to accurately capture the true surface distances necessary for precise shape analysis. To address these shortcomings, (Kimmel and Sethian, 1998; Sethian, 1999) introduced the widely adopted *Fast Marching* algorithm, which significantly improved the accuracy of geodesic distance computations on surfaces, by propagating the information across both the edges and within the faces of mesh triangles. This method utilizes a computationally efficient numerical approach to solve the Eikonal equation, facilitating the calculation of pairwise geodesic distances on triangulated domains. Refer to Figure 3.1 for a visual representation of geodesic distances, illustrating the measurements from a specific point near the tail of cat to all other points on the mesh, as computed using the fast marching method. The Eikonal equation is a fundamental partial differential equation used to determine geodesic distances on manifolds. It describes the propagation of wavefronts and is mathematically expressed as $\|\nabla(u)\| = 1$ where $\nabla(u)$ denotes the gradient of the distance function u from a given source point. This equation signifies that the magnitude of the gradient of the distance function is constant, implying that the distance function increases at a uniform rate as one moves along the geodesic path away from the source point.

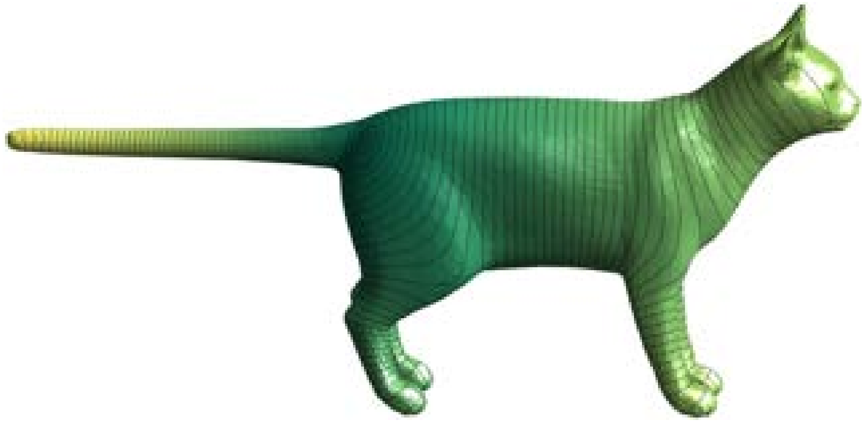

FIGURE 3.2: Biharmonic distances computed from a point near the tail of a TOSCA cat mesh (Bronstein, Bronstein, and Kimmel, 2008) to all other points on the surface

Fast Marching technique enables the rapid and reliable computation of accurate geodesic distances between any two points on the surface, essential for applications in shape analysis, computer graphics, and geospatial modeling. For a mesh with n vertices, the asymptotic computational complexity of the Fast Marching algorithm is $\mathcal{O}(n \log n)$, with a detailed exploration of these methods available in (Kimmel, 2003). While these algorithms are widely used, they have notable drawbacks, such as their inability to reuse information, their requirement for non-obtuse triangles and lack the capability for parallelization. Crane et al. (Crane, Weischedel, and Wardetzky, 2013; Crane, Weischedel, and Wardetzky, 2017) proposed the heat method to address these challenges. This method simplifies the process of distance computation by breaking it into two distinct stages: first, it determines the direction of the steepest ascent in distance, and then it reconstructs the distance itself from this directional information. For a comprehensive review of key approaches to computing geodesic paths and distances across diverse geometric domains, a recent survey by Crane et al. (Crane, Livesu, Puppo, and Qin, 2020) offers an in-depth analysis.

FIGURE 3.3: Canonical Forms derived from Biharmonic Distances for TOSCA wolf shapes (Bronstein, Bronstein, and Kimmel, 2008) via Multidimensional Scaling techniques

3.2.2 Biharmonic Distances

Biharmonic distances, introduced by Lipman et al. (Lipman, Rustamov, and Funkhouser, 2010), is a surface distance metric derived from the *green's* function of the biharmonic differential equation $\Delta_{\mathcal{S}}^2 = 0$ for shape $\mathcal{S}$. These distances effectively integrate local characteristics, such as smoothness and isotropy near the source point, with global features such as shape-awareness and resilience to topology changes, making them valuable for shape analysis and 3D surface processing. Building on the spectral properties of shape $\mathcal{S}$, the squared biharmonic distance between any two points s_1 and s_2 along the shape is defined as:

$$d_{\mathcal{S}}(s_1, s_2)^2 = \sum_{k=1}^{\infty} \frac{(\phi_k(s_1) - \phi_k(s_2))^2}{\lambda_k^2} \tag{3.1}$$

where $\phi_k(s_1)$ and $\phi_k(s_2)$ are the k^{th} eigenfunction of Laplace-Beltrami Operator evaluated at points s_1 and s_2 respectively; while λ_k represents the corresponding eigenvalue. The quadratic normalization in Equation (3.1) balances the speed of decay by being not too slow to produce a singularity and not too fast to make the distance too global. It balances the right local precision as well as appropriate global shape awareness, offering a metric that is not only smooth and shape-aware but is efficiently computable on discrete meshes.

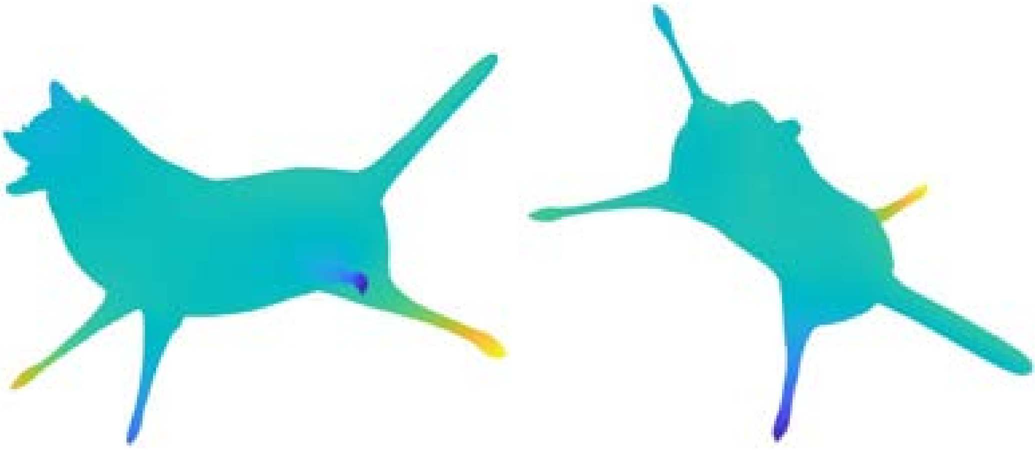

FIGURE 3.4: Canonical Forms derived from Geodesic Distances for Wolf Meshes via Multidimensional Scaling techniques

The Equation (3.1) can also be expanded as:

$$d_{\mathcal{S}}(s_1, s_2)^2 = g_{\mathcal{S}}(s_1, s_1) + g_{\mathcal{S}}(s_2, s_2) - 2g_{\mathcal{S}}(s_1, s_2)$$

where the Green's function $g_{\mathcal{S}}(s_1, s_2)$ of the biharmonic operator $\Delta^2_{\mathcal{S}}$ is defined as:

$$g_{\mathcal{S}}(s_1, s_2) = \sum_{k=1}^{\infty} \frac{\phi_k(s_1)\phi_k(s_2)}{\lambda_k^2}$$

Due to the dependence on the eigendecomposition of the Laplace-Beltrami Operator, the distance captures intrinsic properties of the shape effectively. Biharmonic distance in Equation (3.1) is approximated by considering first k eigenvectors of the LBO ordered by ascending eigenvalues, obtained by solving the generalized eigenvalue problem, as discussed in Section 2.3.2. Despite a linear error bound in $1/k$, this approximation greatly speeds up computation, which is required for high-resolution datasets. This distance metric is characterized by essential properties, including isometric invariance, computational efficiency, parameter-free operation, noise resilience, and topological insensitivity. The computational complexity of computing biharmonic distances on a mesh with n vertices and k Laplace-beltrami eigenvectors is $\mathcal{O}(kn + k)$. For a detailed demonstration of the biharmonic distance's computational efficiency and robustness to noise, see Section 5.1 of Lipman et al. (Lipman, Rustamov, and Funkhouser, 2010). Refer to Figure 3.2

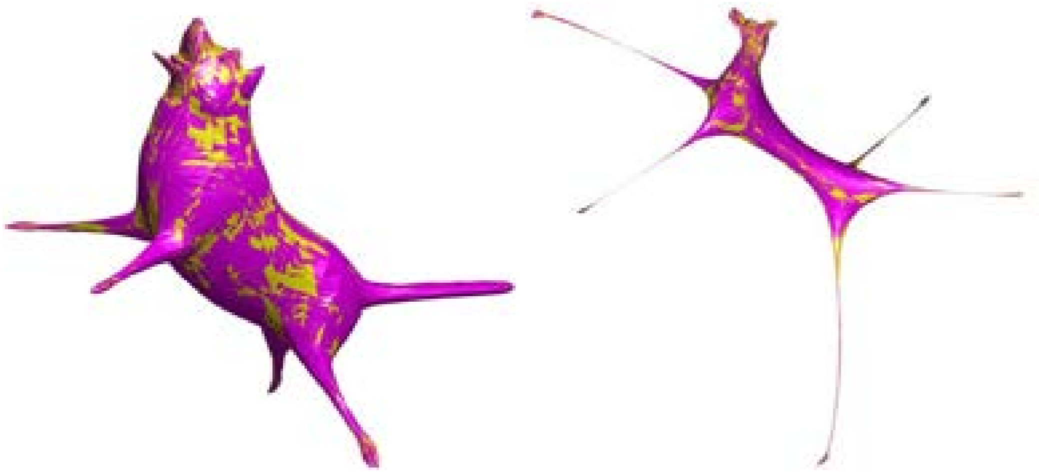

FIGURE 3.5: Non-rigidly aligned canonical forms: left shows forms using geodesic distances; right displays forms from biharmonic distances

for a visual demonstration of biharmonic distances, illustrating the measurements from a specific source point to all other points on the cat mesh. The isocontours align with the natural cross-sections of the shape, highlighting its global shape awareness.

3.3 Motivation & Contribution

Geodesic distances are computed as if ants are walking on a string, concentrating on immediate paths and ignoring the overall shape. Despite their local advantages like shortest paths and isotropy, geodesic distances poorly represent global geometry due to their shape-oblivious nature, sensitivity to noise, lack of smoothness, dependence on local segments, and strong influence from source point placement. In contrast, biharmonic distances, relying on the Laplace-Beltrami operator, characterizes the global geometry, especially for distant points, making them shape-aware and robust to noise and stable under changes in topology or tessellation.

3.4 Methodology

In this work, biharmonic distances are utilized to describe pairwise relationship between surface points of shapes to capture sharp features of shape,

FIGURE 3.6: Geodesic error visualized on the shape using heat maps: the left displays errors computed via geodesic distances, while the right illustrates errors via biharmonic distances

while also encapsulating the underlying global geometric characteristics. This relationship is used to construct an affinity matrix, which is then subjected to dimensionality reduction techniques, producing a low-dimensional embedding that effectively captures the shape's intrinsic geometry. Point-to-point correspondences between embedded shapes have traditionally depended on rigid alignment techniques, which often struggle to accurately capture complex shape variations. This work emphasizes the use of non-rigid alignment of canonical forms within the Euclidean domain, offering enhanced accuracy and flexibility in establishing robust shape correspondences.

The key contribution is the integration of biharmonic distances for determining the pose-invariant embedded forms, combined with non-rigid alignment of corresponding canonical forms, a combination that outperforms other similar existing techniques.

The methodology focuses on determining the pointwise correspondences for two wolf meshes, wolf0 and wolf1, from the TOSCA dataset (Bronstein, Bronstein, and Kimmel, 2008). To compute the mapping between shapes, the canonical forms are established for each respective shape, transforming the challenge of matching non-rigid shapes into the simpler task of matching rigid shapes. Creating these canonical forms is essential for aligning the shapes within a shared metric space. Biharmonic distances being shape-aware and smooth are employed to determine canonical forms by generating

FIGURE 3.7: Texture transfer from one wolf shape to another wolf shape: ground truth map (left), geodesic distances map (middle), and biharmonic distances map (right)

pairwise affinity matrices for both shapes.

The process begins by calculating biharmonic distances for all pairs of vertices in each mesh, utilizing the eigendecomposition of the Laplace-Beltrami operator (LBO) as detailed in Section 3.2.2. Once the biharmonic distance matrix is obtained, multi-dimensional scaling algorithms (refer Section 2.3.1) are used to identify canonical forms embedded in a Euclidean space. Specifically, the SMACOF (Scaling by MAjorizing a COmplicated Function) multidimensional scaling algorithm (F. Groenen and Velden, 2014) is applied which minimizes a stress function that measures the difference between original and embedded distances, guaranteeing stress reduction at each iteration through majorization. Incorporating Regularized Residual Embedding (RRE) enhances robustness by introducing a regularization term that penalizes large deviations, mitigating sensitivity to noise and outliers. This integration ensures smoother, more stable embeddings suitable for complex deformation. See Figure 3.3 for the canonical forms of wolf meshes, computed via pairwise biharmonic distances. For comparison, refer to Figure 3.4, which illustrates the canonical forms derived from pairwise geodesic distances. After embedding the shapes, they are carefully aligned ensuring the closest possible match between them. The conventional approach for aligning canonical forms often uses Iterative Closest Point (ICP), which is a rigid alignment technique as detailed in Section 1.2.1. In this

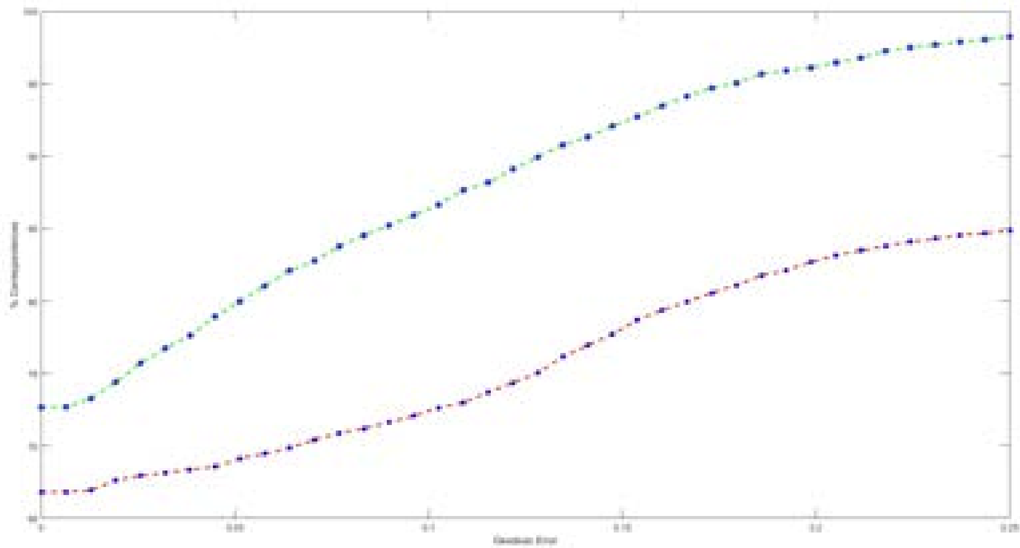

FIGURE 3.8: Point-to-point correspondence accuracy specified for wolf shape, comparing Biharmonic distances (green circles) and Geodesic distances (red squares) against increasing geodesic error on the x-axis

work, we enhance the alignment process by using a non-rigid version of ICP (Jain, Zhang, and Van Kaick, 2007) that incorporates thin-plate splines, allowing for a more accurate representation of deformations in non-rigid shapes. To align the canonical form of the wolf1 mesh with that of wolf0, we deform wolf1 while maintaining wolf0 in its fixed position, via non-rigid ICP. Refer to Algorithm 1 for a concise overview of the steps used to obtain the pointwise correspondences. Figure 3.5 illustrates the aligned canonical forms of both meshes, computed using pairwise relationship matrices based on biharmonic and geodesic distances. The aligned forms are matched using a closest-point method, optimizing an energy function based on distance metrics, which generates a point-to-point correspondence map.

Results To assess the accuracy of the point-to-point correspondence map, geodesic error is computed by summing the geodesic distances between the computed and ground-truth corresponding vertices on the wolf1 mesh for each vertex in the wolf0 mesh. Figure 3.8 visualizes the accuracy of the point-to-point correspondence map, showing the percentage of correct matches (y-axis) against increasing geodesic error (x-axis), with canonical forms computed using biharmonic and geodesic distance matrices. Figure

3.5 illustrates the geodesic error in the point-to-point correspondence map via a heat map indicating higher errors in red and zero error in black. The computational complexity of the proposed approach primarily hinges on distance calculations, making it faster with biharmonic distances than with geodesic distances. For wolf meshes, only 72.68% of vertices mapped with geodesic error under 0.1, while using biharmonic distances, it increased the accuracy to 86.65%. Refer Table 3.1 for improvement in the point-to-point correspondence accuracy for shapes like Centaur and Human from TOSCA dataset (Bronstein, Bronstein, and Kimmel, 2008).

Algorithm 1: Correspondence computation via Surface distances

Input: Two isometrically deformed shape meshes

Output: Pointwise correspondence map from one shape to another

begin

- **Step 1**: Compute surface distance matrices (Geodesic or Biharmonic) for both the meshes using LBO eigendecomposition.
- **Step 2**: Embed both meshes into a low-dimensional Euclidean space by computing canonical forms via SMACOF MDS with RRE.
- **Step 3**: Align one canonical form to other using non-rigid ICP with thin-plate splines.
- **Step 4**: Obtain pointwise correspondences by nearest-neighbor matching in the aligned canonical space.

Application Texture transfer between shapes visually demonstrates the effectiveness of the point-to-point correspondence map, with the smoothness of the process serving as an indicator of its efficacy in transferring texture from the wolf0 mesh to the wolf1 mesh. Figure 3.7 demonstrates the quality of texture transfer through three distinct mappings. On the left, the ground truth serves as a reference for accurate texture alignment. The middle shows the texture transfer achieved using a correspondence map based on geodesic distance matrices, while the right illustrates the results with biharmonic

Shapes	Correspondence Accuracy with less than 0.1 geodesic error	
	Geodesic	Biharmonic
Wolf	72.68%	86.65%
Centaur	17.38%	39.2%
Human	46.90%	56.19%

TABLE 3.1: Correspondence accuracy (<0.1 geodesic error) on TOSCA shapes (Bronstein, Bronstein, and Kimmel, 2008) when canonical forms are built with biharmonic vs. geodesic distances

distance-based correspondence. Notably, biharmonic distances enhance the accuracy of texture transfer, especially in the facial region, where the texture aligns more closely with key features. This suggests that biharmonic distances are more effective in preserving fine details in texture mapping than geodesic distances.

3.5 Discussions and Future Scope

This work discusses the impact of shape-aware biharmonic distances in accurately capturing shape geometry for shape correspondence. The incorporation of a non-rigid version of the ICP algorithm to align canonical forms obtained via pairwise biharmonic distances leads to more effective point-to-point correspondences by better capturing deformation information. As shown in Figure 3.7, canonical forms generated via biharmonic distances more effectively represents the facial features of the wolf mesh compared to those generated via geodesic distances. This highlights the superiority of biharmonic distances in shape matching by enhancing the balancing of fine local details and global geometric characteristics of the shape.

However, while biharmonic distances enhances pointwise mapping between shapes, they may smooth out high-frequency details due to their reliance on Laplace-Beltrami operators. Future exploration could focus on integrating Hamiltonian operators with intrinsic potentials to better capture high-frequency features in their spectra, enhancing the accuracy of point-to-point correspondence maps.

Chapter 4

Localization via Functional Maps

4.1 Introduction

Shape matching has largely concentrated on capturing the global geometric properties of a shape, like distinguishing between an SUV and a sedan based on overall design. However, finer-grained shape characteristics are defined by local features, including hood curvature, door handle contours, and headlight details. Addressing complex non-rigid shape deformations along with the integration of detailed shape features into analysis, requires robust and sophisticated approaches. One such approach, introduced by Ovsjanikov et al. (Ovsjanikov, Ben-Chen, Solomon, Butscher, and Guibas, 2012), is the functional map framework, which transforms point-wise correspondences into a linear relationship between function spaces defined on shapes. It focusses on mapping functions rather than directly matching points on shape. This framework aims at reducing computational complexity, by utilizing a compact set of basis functions defined for each shape respectively. By projecting shape features onto the basis functions, this coordinate-free method enhances both analytical efficiency and the capability to capture better geometric properties, improving the accuracy of point-to-point correspondences. Moreover, without explicitly computing correspondences, functional maps allow for the seamless transfer of geometric information, such as curvature and deformation, between shapes, making it valuable for learning-based techniques (Marin, Rakotosaona, Melzi, and Ovsjanikov, 2020) as well. Although the point-to-point mapping between surfaces might

FIGURE 4.1: Hand reconstruction with Laplacian (left) and Hamiltonian (right) basis functions

be complex, functional maps act linearly when considered between function spaces. Refer Section 2.4.1 for an in-depth discussion on functional maps.

The functional map framework converts the task of determining point-to-point correspondences between non-rigid shapes into a more tractable problem by computing a linear relationship between their function spaces. First, basis functions are computed for both shapes, and key descriptors or features are represented as functions on each shape. Then, a compact low-rank mapping matrix is built to align these functional representations across the shapes using the appropriate basis. This linear map is then efficiently translated into point-to-point correspondences, offering an automatic, user-independent solution that simplifies the non-rigid shape matching problem. The focus of this chapter is to investigate whether incorporating sharp geometric features of the shapes into this framework can further enhance the accuracy of point-to-point correspondences, leveraging both the functional map framework and efficient computational methods.

Selecting the appropriate basis is essential within the functional map framework, as it directly influences the accuracy of point-to-point correspondences. High-frequency basis functions capture fine details like sharp edges, while functions with localized support offer precise control over specific shape regions. A careful basis choice enhances geometric feature representation along with improving correspondence mapping. Laplace-Beltrami eigenfunctions serve as a foundational basis function in shape analysis due to their effectiveness in representing global geometric properties at multiple levels. To better capture high-frequency details, Neumann et al. (Neumann, Varanasi, Theobalt, Magnor, and Wacker, 2014) proposed *compressed manifold modes*, a sparse basis with localized support, invariant to sign flips and reordering. Melzi et al. (Melzi, Rodolà, Castellani, and Bronstein, 2018) introduced *localized manifold harmonics* (LMH), which allow explicit control over the region of localization. However, applying these harmonics becomes increasingly complex when handling multiple meshes simultaneously, as in tasks like shape correspondence and pose transfer. While Kovnatsky et. al. (Kovnatsky, Bronstein, Bronstein, Glashoff, and Kimmel, 2013) and Eynard et. al. (Eynard, Kovnatsky, Bronstein, Glashoff, and Bronstein, 2015) account for multiple meshes and derive bases that incorporate the geometric information across all meshes, they struggle to capture the high-frequency details essential for accurately representing complex shapes. This highlights the necessity for an approach that effectively integrates geometric context in multiple meshes while preserving critical high-frequency features.

4.2 Background

For shapes $\mathcal{S}_1$ and $\mathcal{S}_2$, a bijection is required for a point-to-point map $\eta\colon \mathcal{S}_1 \to \mathcal{S}_2$ when both shapes have the same number of points. The mathematical foundations for computing functional maps are outlined in Section 2.4.1, while the next section addresses the remaining components needed to complete the framework for deriving pointwise correspondences from functional maps.

4.2.1 Correspondence via functional maps

Functional map framework defines spaces of scalar-valued functions $\mathcal{F}(\mathcal{S}_1, \mathbb{R})$ and $\mathcal{F}(\mathcal{S}_2, \mathbb{R})$ for isometrically deformed shapes $\mathcal{S}_1$ and $\mathcal{S}_2$, aiming to compute a linear functional mapping $\mathcal{T}_F\colon \mathcal{F}(\mathcal{S}_1, \mathbb{R}) \to \mathcal{F}(\mathcal{S}_2, \mathbb{R})$. This mapping associates values from functions $f\colon \mathcal{S}_1 \to \mathbb{R}$ and $g\colon \mathcal{S}_2 \to \mathbb{R}$ and can be represented as a matrix $\mathbb{C} \in \mathbb{R}^{k_1 \times k_2}$, with bases $\phi_{\mathcal{S}_1}$ and $\phi_{\mathcal{S}_2}$ containing k_1 and k_2 elements, respectively. The framework follows a structured pipeline designed to compute shape correspondences effectively. First, it calculates invariant feature descriptors **F** and **G** for each shape, ensuring that these descriptors account for isometric deformations. Next, basis functions $\phi_{\mathcal{S}1}$ and $\phi_{\mathcal{S}_2}$ are selected for both shapes. Function preservation constraints are then established by projecting the feature descriptors onto their respective bases, resulting in matrices **A** and **B**. Additional constraints, such as operator commutativity and regularization, are incorporated to enhance the model's performance. The optimal functional map **C** is determined by minimizing the following energy function:

$$E(\mathbf{C}) = ||\mathbf{CA} - \mathbf{B}||^2 + \|\mathcal{S}_F^{\mathcal{S}_2}\mathbf{C} - \mathbf{C}\mathcal{S}_F^{\mathcal{S}_1}\|^2$$

where $\mathcal{S}_F^{\mathcal{S}_1}$ and $\mathcal{S}_F^{\mathcal{S}_2}$ are linear operators, chosen according to the deformation involved. Second term in the above equation describes that say if $\mathcal{S}$ is a symmetry operator, then it should commute with the functional map. Finally, the computed map **C** is further refined, and point-to-point correspondences are derived using an iterative closest point (ICP) algorithm, applied in the embedded functional space rather than the traditional Euclidean space. By incorporating functional maps, which efficiently reduce dimensionality for high-resolution shape analysis, the method focuses on adding structure to shapes rather than directly manipulating them, enhancing the efficacy of shape matching.

4.2.2 Hamiltonian operator

On the shape $\mathcal{S}$, given a scalar potential function $v : \mathcal{S} \to \mathbb{R}_+$ Hamiltonian operator $\mathcal{H}_{\mathcal{S}}$ is described as the extension of Laplace-Beltrami operator $\Delta_{\mathcal{S}}$ operating on $f \in \mathcal{S}$ as $\mathcal{H}_{\mathcal{S}}(f) = \Delta_{\mathcal{S}}(f) + \mu\, v(f)$ with parameter $\mu \in \mathbb{R}$. In the discrete setting, for triangle mesh $\mathcal{X}$, Hamiltonian operator can be specified via LBO's cotangent weight matrix **W** and **D** respectively and

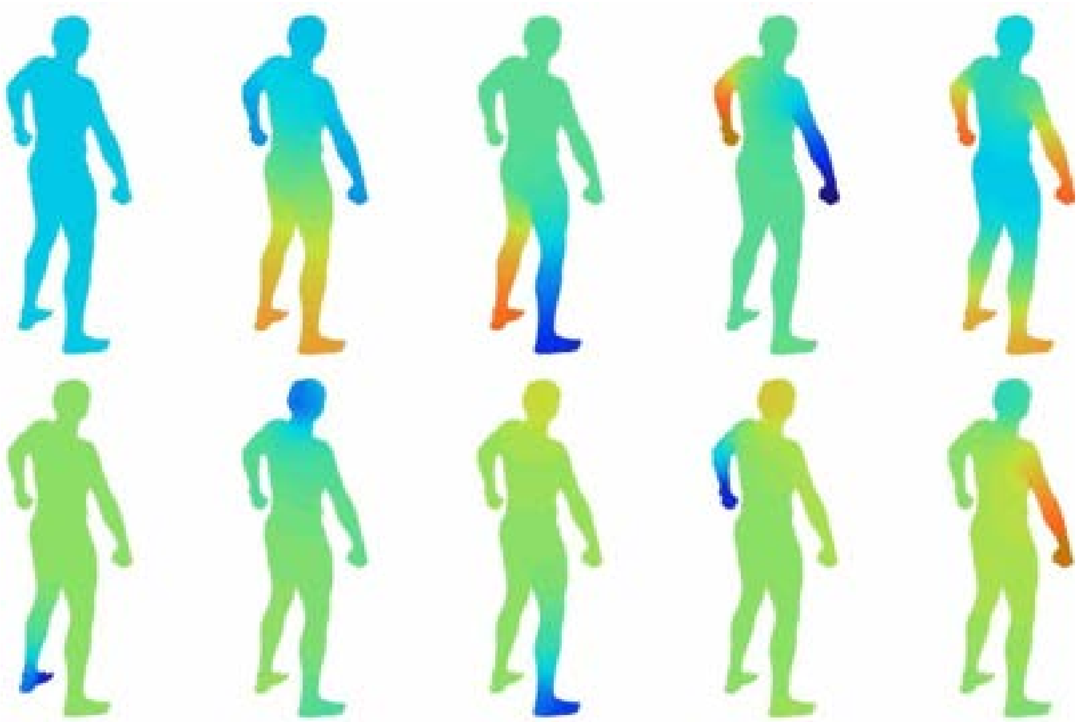

FIGURE 4.2: Visualizing the spectral properties of Laplace-Beltrami eigenfunctions (top) and Hamiltonian eigenfunctions (bottom) ordered by ascending eigenvalues

the potential function v is applied diagonally (refer Section 2.3.2). Hamiltonian operator, a sum of two self-adjoint operators is also self-adjoint and hence admits an eigendecomposition with generalized eigenvalue problem $(\mathbf{W} + \mu\, \mathbf{D}\, diag(v))\, \Phi = \mathbf{D}\Phi\Lambda$ which can numerically be posed as the respective optimization problem as:

$$\arg\min_{\Phi} \; tr(\Phi^T \mathbf{W} \Phi) + \mu\, tr(\Phi^T diag(v)\, \Phi) \;\; s.t.\; \Phi^T \mathbf{D} \Phi = I$$

with parameter μ controlling the trade-off between global and local support of Hamiltonian eigenbasis.

4.3 Motivation & Contribution

The functional map framework offers inherent flexibility, allowing for enhancements at any stage of the pipeline (refer Section 4.2.1). Despite various attempts to improve functional maps, the selection of an appropriate basis

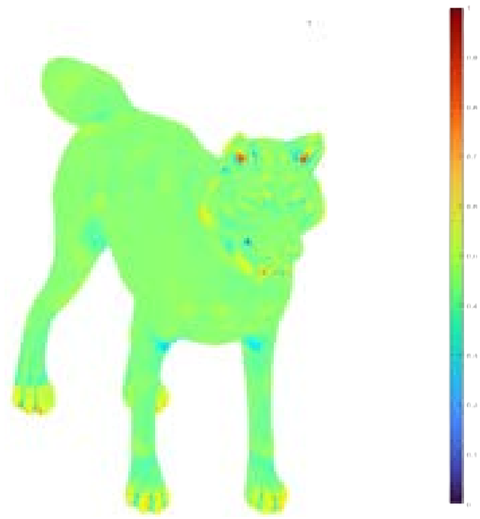

FIGURE 4.3: Visualizing Gaussian Curvature on a wolf shape from TOSCA dataset (Bronstein, Bronstein, and Kimmel, 2008)

remains crucial, as it defines the analysis domain and must effectively reduce representation complexity while ensuring stability and compactness. While Laplace-Beltrami eigenfunctions (Levy, 2006) are commonly used due to their multi-scale properties and invariance to isometric deformations, their tendency to smoothen out sharp features and sensitivity to topological changes limit their ability to capture the localized details of complex shapes. This work proposes a basis that better captures shape geometry by emphasizing sharp features while maintaining computational efficiency of LBO eigenbasis. Representing the discrete potential function as a diagonal matrix ensures no additional computational cost compared to the Laplace-Beltrami basis (refer Section 4.2.2). Utilizing the Hamiltonian operator's potential function to localize regions enables its eigenfunctions to capture high-frequency information while preserving the characteristics of Laplace-Beltrami eigenfunctions.

Gaussian curvature (K), which fully characterizes shape geometry (Alexa, 2002) as the product of principal curvatures $K = k_1 * k_2$, highlights regions with positive and negative curvatures, indicating areas with sharp features (refer Figure 4.3, where positive curvature regions are marked in yellow, such as paws, and negative ones in red, such as the inside of the ear). This chapter

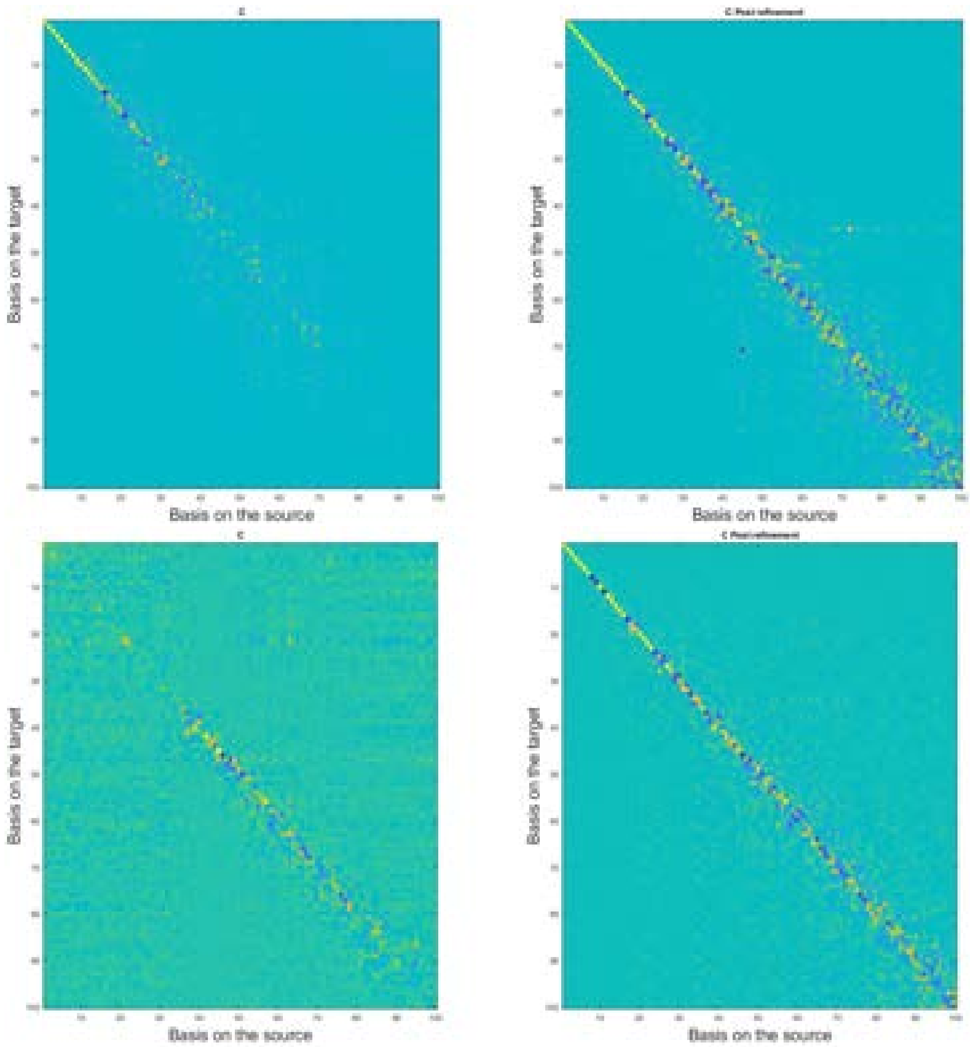

FIGURE 4.4: Comparing Laplace-Beltrami eigenbasis (top) and Hamiltonian bases for functional maps (bottom); pre-refinement (left) post-refinement (right)

advocates for using Gaussian curvature as a potential function for determining the Hamiltonian basis, effectively capturing shape geometry and enhancing point-to-point correspondences without additional costs. When a shape is projected onto the first few basis functions and reconstructed from those projections, Figure 4.1 depicts the comparison, demonstrating that Hamiltonian bases preserve sharp features, such as fingers, while Laplace-Beltrami bases smoothen them out. Unlike Laplace-Beltrami bases, Hamiltonian bases are better at capturing and representing sharp features of shape, resulting in more accurate reconstructions, even when only a limited number of basis functions are used.

FIGURE 4.5: Geodesic error heat maps: Hamiltonian (left) vs Laplace-Beltrami (right) for wolf mesh

4.4 Methodology

Given two wolf meshes in different poses from TOSCA dataset, the functional map framework is applied to compute vertex-to-vertex correspondences. The technique involves computing feature descriptors that are invariant to isometric deformations such as heat kernel signatures (HKS) (Sun, Ovsjanikov, and Guibas, 2009) at varied timestamps and wave kernel signatures (WKS) (Aubry, Schlickewei, and Cremers, 2011) at varied energy values, describing the functional spaces for both the respective meshes. Basis functions for each shape are identified by eigendecomposing the Hamiltonian operator, using Gaussian curvature as the potential function. The first 100 Hamiltonian eigenfunctions, ordered by increasing eigenvalues, form the shape basis. Feature descriptors are projected onto the Hamiltonian basis, and a gradient descent algorithm is employed to minimize the energy functional while adhering to the Laplacian commutativity constraint, thereby computing an optimal functional map. Refer Section 2.4.1 for details regarding optimization energy needed to determine functional maps. The functional map is then refined via ICP algorithm as discussed in Section 4.2.1, in order to obtain point-to-point correspondence map between both the shapes. Refer to Algorithm 2 for a concise overview of the steps used to obtain the pointwise correspondences.

Algorithm 2: Functional correspondence via Hamiltonian eigenbasis

Input: Two isometrically deformed shape meshes

Output: Pointwise correspondence map from one shape to another

begin

Step 1: Compute isometry-invariant descriptors on both meshes:

- Heat Kernel Signature (HKS) at multiple time stamps
- Wave Kernel Signature (WKS) at multiple energy values

Step 2: Construct the Hamiltonian basis:

- Use Gaussian curvature as the potential function
- Eigendecompose the Hamiltonian operator and keep the 100 lowest-eigenvalue eigenfunctions

Step 3: Project the descriptors (HKS/WKS) onto the Hamiltonian basis for each mesh

Step 4: Estimate the functional map via gradient descent, minimizing functional-map energy with Laplacian-commutativity enforced

Step 5: Refine the functional map via ICP in functional map space to recover vertex-to-vertex correspondence

Refer Figure 4.4 (top) for illustration of the structure of the functional map matrix with 100 basis functions, both before and after ICP refinement, using the Laplace-Beltrami basis, while the bottom row showcases the results using the Hamiltonian basis. The Laplace-Beltrami captures low-frequency features, whereas the Hamiltonian picks up high-frequency features as depicted in the 1^{st} column of the same figure highlighting the functional map obtained before refinement. The accuracy of the point-to-point correspondence map is verified by calculating the geodesic error, which sums the distances between computed and ground-truth mappings for each vertex on the source mesh. The proposed basis outperforms Laplace-Beltrami basis and

Basis Type	**Accuracy with less than 0.1 geodesic error**		
	Wolf	Human	Centaur
LBO Eigenfunctions	84.7%	64.09%	35.81%
CMM	14.73%	12.6%	10.33%
Proposed basis	98.6%	72.23%	78.94%

TABLE 4.1: Diverse bases (LBO (Ovsjanikov, Ben-Chen, Solomon, Butscher, and Guibas, 2012); CMM (Neumann, Varanasi, Theobalt, Magnor, and Wacker, 2014); Proposed basis (Bindal and Kamat, 2023)) define the functional space used for computing point-to-point correspondences

compressed manifold modes in terms of point-to-point correspondence accuracy, especially with minimal error allowed (refer Table 4.1). It achieves similar exact error to Laplace-Beltrami basis but with better overall performance. Figure 4.5 visualizes the error as a heat map, showing that the geodesic error for the proposed basis is localized around the tail, while for the Laplace-Beltrami basis, it is more dispersed, particularly near sharp features such as the paws. Figure 4.6 shows that the proposed basis consistently delivers higher point-to-point accuracy compared to other basis functions, highlighting its enhanced precision in correspondence mapping. Texture transfer is used to further illustrate the effectiveness of the proposed basis in Figure 4.7. When texture is mapped from a source shape to a target shape using the computed correspondences, the results reveal that maps derived from the Laplace-Beltrami Operator (LBO) basis exhibit significant defects, particularly around the paws. In contrast, the Hamiltonian basis accurately preserves sharp details, closely aligning with the ground-truth map and demonstrating superior correspondence quality.

4.5 Discussions and Future Scope

This chapter investigates the impact of using an intrinsic potential function, specifically Gaussian curvature, to define the Hamiltonian operator within the functional map framework. By modulating manifold harmonics with this potential, the Hamiltonian basis captures sharper features than the Laplace-Beltrami basis, enhancing the accuracy of point-to-point correspondences.

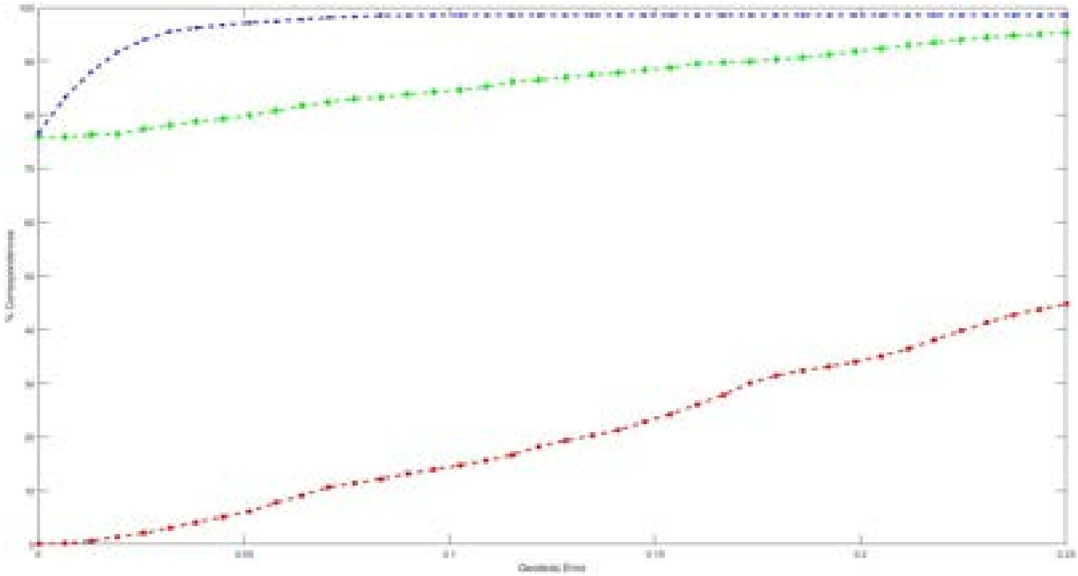

FIGURE 4.6: Point-to-point correspondence accuracy on wolf mesh using different bases via functional maps: Laplace-Beltrami (green), Hamiltonian (blue), compressed manifold modes (red)

While compressed manifold modes aim to capture high-frequency details, their sparsity and lack of global shape information limit their effectiveness for non-rigid shape matching (refer Figure 4.6). The proposed Hamiltonian basis resolves these limitations by balancing the representation of both global structure and sharp features, leading to improved performance. Empirical testing on multiple shape pairs demonstrates that incorporating sharp feature information into the basis significantly enhances shape's geometric representation (Bindal and Kamat, 2023). In Figure 4.5, which visualizes the error heat maps obtained using the Hamiltonian and Laplace–Beltrami bases, the Hamiltonian-based reconstruction exhibits localized error in the tail region. This localization can be attributed either to limited sensitivity of the Hamiltonian eigenfunctions in the high-potential regions of the domain (like tail), or it may attribute more fundamentally to the fact that the Hamiltonian basis is intentionally biased to represent signal content in low-potential regions, thereby de-emphasizing details outside those regions.

A key limitation of the proposed approach is that the potential function used for computing Hamiltonian eigenfunctions is fixed prior to analysis, as in the case of Gaussian curvature or predefined step functions. This rigidity

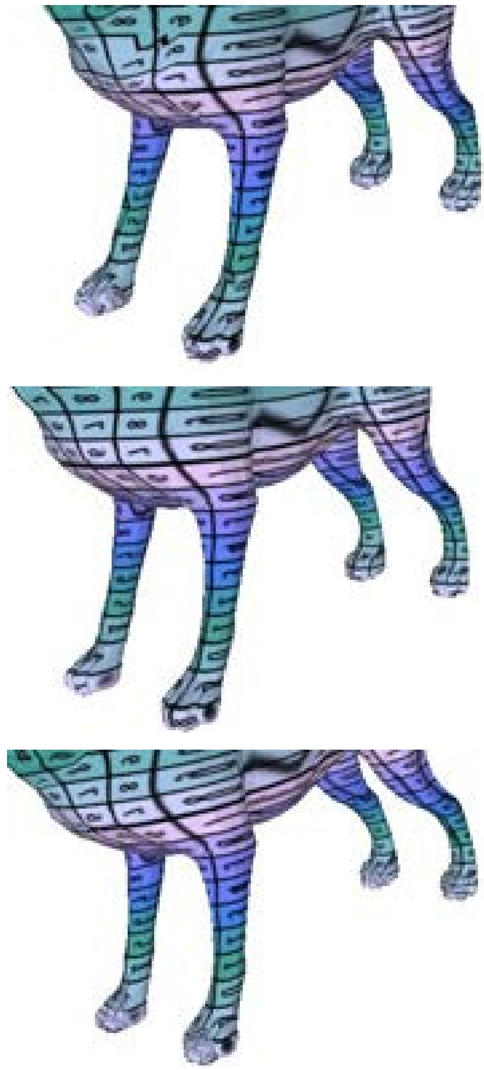

FIGURE 4.7: Texture mapping: Hamiltonian (top), ground truth (middle), Laplace-Beltrami (bottom) for wolf0 to wolf1

may contribute to the observed error localization, as flat regions are not effectively captured. A promising avenue for future work is the development of an adaptive potential function capable of dynamically identifying critical regions across shapes, which could further improve the precision of point-to-point correspondences.

Chapter 5

Bases for Functional Spaces

5.1 Introduction

Recent years have witnessed significant advancements in functional map estimation and point-to-point correspondences through various regularization techniques. Improvements include managing the unknown ordering of descriptor constraints (Pokrass, Bronstein, Bronstein, Sprechmann, and Sapiro, 2013), imposing geometric structure on the functional map matrix by treating rows and columns as shape-specific functions (Kovnatsky, Bronstein, Bresson, and Vandergheynst, 2015), treating descriptor functions as linear operators to refine pointwise mapping (Nogneng and Ovsjanikov, 2017), and introducing a bounded Laplacian operator to preserve structural integrity in functional maps (Ren, Panine, Wonka, and Ovsjanikov, 2019). ZoomOut (Melzi, Ren, Rodolà, Sharma, Wonka, and Ovsjanikov, 2019) is an effective end-to-end technique that iteratively expands the functional basis, optimizing functional maps and pointwise mappings subsequently. Rodola et al. (Rodolà, Cosmo, Bronstein, Torsello, and Cremers, 2017) adapted functional maps for partial shape matching by imposing specific structural constraints. To improve resilience to mesh tessellation variations and enhance correspondence precision, Ezuz et al. (Ezuz and Ben-Chen, 2017) proposed a map deblurring and denoising method. For the past decade, continuous optimization techniques such as ADMM and MADMM have been prevalent in optimal basis and functional map optimization. However, Ren et al. (Ren, Melzi, Wonka, and Ovsjanikov, 2021) introduced a promising discrete solver tailored for energies derived from functional maps.

This chapter builds upon the functional map framework, as introduced in Section 4.2.1. The framework offers a robust and systematic process, starting with the determination of feature descriptors for each shape. Next, a suitable basis for the function spaces is selected to facilitate the mapping. The process concludes with the estimation of the functional map, achieved by solving an optimization problem that preserves the descriptor functions while regularizing the map to maintain structural accuracy. Central to this framework is the choice of basis functions, which provide the structure essential for embedding shapes computationally, enhancing efficiency, and minimizing storage requirements, especially for high-resolution shapes. The accuracy of point-to-point correspondences within the functional map framework relies greatly on the choice of basis functions.

Despite the broad applicability of functional maps in shape correspondence, selecting the optimal basis function remains an open challenge due to its significant impact on correspondence accuracy. Over the years, researchers have proposed various basis functions, each with unique strengths and limitations. Recently, Cammarasana and Patané (Cammarasana and Patané, 2021) explored several basis functions, focusing solely on Laplacian-based approaches for shape analysis tasks. This chapter reviews various basis functions for non-rigid shape correspondence, covering their principles, limitations, and practical implications. Researchers, students, and practitioners seeking guidance on selecting the most suitable basis functions for shape matching are directed to (Bindal and Kamat, 2024) for a comprehensive reference.

5.2 Diverse Basis functions

Effective basis functions for shape correspondence must accurately capture key geometric details while ensuring computational efficiency and robustness to deformations. They should exhibit *completeness* to represent all shape variations, *compactness* for efficient computation with minimal basis functions, and *stability* under isometric deformations should span the same function space. Additional properties like *orthogonality* simplify processing by making basis functions linearly independent, while *multi-scale adaptability* allows capturing both fine details and coarse structures. *Smoothness* ensures continuity, reducing sensitivity to small shape perturbations. Moreover, *robustness* to topological changes enables handling significant structural differences, and *locality* supports efficient processing and feature capture by focusing on local geometry. Together, these properties enhance the accuracy, versatility, and efficiency of shape correspondence across diverse applications. The following sections describe the individual principles underlying various basis functions, while Section 5.3 provides a comparative analysis based on the properties satisfied by each basis function.

5.2.1 Laplace-Beltrami Operator Eigenfunctions

The Laplace-Beltrami operator eigenbasis, analogous to the Fourier basis in Euclidean spaces, offers an optimal and efficient representation for smooth functions on shapes. The Laplace-Beltrami operator $\Delta_{\mathcal{S}}$, as discussed in Section 2.3.2 is intrinsic to the shape $\mathcal{S}$ itself *i.e.*, it is independent of any coordinate system or representation format, such as meshes or point clouds, allowing it to capture core geometric features of the shape. Moreover, it is theoretically proven to be the optimal basis for representing smooth functions with bounded variation (Aflalo, Brezis, and Kimmel, 2015), providing completeness, compactness, orthogonality, multi-scale characteristics, and high computational efficiency. Although constructed locally, it captures global shape information and is widely used in geometry processing for isometrically deformed shapes. However, even minor deviations in isometric deformations can disrupt consistency in low-frequency bases. LBO eigenbasis captures global shape aspects, resists local topological changes leading to numerical instability at higher frequencies (Zhang, Van Kaick, and Dyer, 2010). Learning-based methods recognize these challenges and aim to improve basis for shape correspondence, despite the LBO's extensive application (Litany, Remez, Rodola, Bronstein, and Bronstein, 2017). Refer to Figure 5.1(a) for the visualization of the first few LBO basis functions on human shape from the SHREC dataset, where red represents high values and blue indicates low values.

5.2.2 Coupled-quasi Harmonic basis

When computing shape correspondences between shapes that undergo not just isometry but also stretching and other deformations, consistency among the shape bases becomes essential. To address this, the coupled-quasi harmonic basis (CQHB) was introduced in (Kovnatsky, Bronstein, Bronstein, Glashoff, and Kimmel, 2013), allowing for stretching and matching of non-isometric shapes, such as semantically aligning a human shape with that of a gorilla. It seeks to identify coupled bases that can simultaneously diagonalize the Laplace-Beltrami operators of both shapes, focusing on finding basis functions that are similar, ensuring an approximate relation through deformation. This is achieved by solving a non-linear optimization problem with orthogonality constraints, minimizing the off-diagonal elements of the diagonalized matrices while also reducing the difference between the descriptor functions of the two shapes. These descriptor functions remain invariant to isometric deformations and ensure coupling between the bases. Although the CQHB addresses non-isometric shape deformations more effectively than the LBO basis, its theoretical framework is still evolving for improved computational efficiency. By jointly considering shapes, this approach highlights sharp features, enhancing localization.

5.2.3 Compressed Manifold Modes

The global support of the LBO eigenbasis hinders the interpretation of local shape features and makes it particularly sensitive to topological noise, including holes and occlusions. Compressed Manifold Modes (CMM) enhances the traditional Laplace-Beltrami Operator (LBO) eigenbasis by incorporating local support through an l_1-regularization term that promotes sparsity in the basis functions (Neumann, Varanasi, Theobalt, Magnor, and Wacker, 2014) . This approach enables the shape to be treated as a collection of parts rather than a global structure. The optimization problem is efficiently solved using the Manifold Alternating Direction Method of Multipliers (MADMM), optimizing the basis on the Stiefel manifold of orthogonal matrices. CMMs (refer Figure 5.1(b)) are robust to noise and deformation, outperforming LBO eigenbasis, especially in partial shape matching. Haas et al. (Haas, Baskurt, Dupont, and Denis, 2017) optimized CMM computation for better shape coverage. However, CMMs remain computationally expensive and lack a multi-scale representation.

5.2.4 Localized Manifold Harmonics

Although CMM provides local support, it lacks explicit localization control and does not ensure a global solution. In (Melzi, Rodolà, Castellani, and Bronstein, 2018), Melzi et al. introduced *Localized Manifold Harmonics* (LMH), which allows controlled localization by specifying a region of interest on the surface and expands the LBO eigenbasis via an incremental process that maintains orthogonality with existing basis functions. CMMs are obtained by spectrally decomposing a modified LBO, incorporating a regularization term to induce localized support and orthogonality to a precomputed LBO basis. LMH is isometry-invariant, smooth, captures sharp details, and enhances LBO eigenbasis with local information. However, manually specifying regions for the shape to improve correspondences, such as areas with high reconstruction error (refer Figure 5.1(c)), remains a challenge.

5.2.5 Hamiltonian operator Basis

LBO eigenbasis struggles with matching complex shapes due to its global sensitivity. Hamiltonian operator (HO), introduced in (Choukroun, Shtern, Bronstein, and Kimmel, 2018), leverages quantum mechanics to modify manifold harmonics, enabling the incorporation of sharp features and efficient spectral analysis. The Hamiltonian operator, a modified Laplace-Beltrami operator, incorporates a potential function to control the localization of eigenfunctions, balancing global and local shape information. Hamiltonian operators, using step potential functions, can concentrate harmonics within specific regions of a manifold (refer Figure 5.1(d)). This property has been exploited for partial shape matching. Rampini et al. (Rampini, Tallini, Ovsjanikov, Bronstein, and Rodola, 2019) aligned Hamiltonian eigenvalues, while Postolache et al. (Postolache, Fumero, Cosmo, and Rodolà, 2020) utilized Hamiltonian eigenfunctions within the functional map framework to achieve state-of-the-art results in partial shape matching.

5.2.6 Coordinate Manifold Harmonics

Traditionally, the dimensionality of basis functions for preserving shape details is set empirically. However, by adding just three supplementary functions to standard manifold harmonics, Coordinate Manifold Harmonics (CMH) (Marin, Melzi, Musoni, Bardon, Tarini, and Castellani, 2019) significantly enhances local isometry preservation, capturing both intrinsic and extrinsic shape information (see Figure 5.1(e)). To incorporate extrinsic information into shape analysis, three basis functions are added to standard LBO eigenfunctions, capturing orthogonal projection of the representation error from the low-pass filter of all the 3D coordinates of the shape's vertices. CMH, primarily used for tessellation transfer and extrinsic information inclusion, requires similar poses across shapes due to its vertex-coordinate dependency (Melzi, Marin, Musoni, Bardon, Tarini, and Castellani, 2020).

5.2.7 Data-driven Basis

The LBO focuses on low-frequency eigenfunctions for smoothness, while Azencot et. al. in (Azencot and Lai, 2021) introduced an adaptive basis to enhance feature matching. The unified framework combines functional map

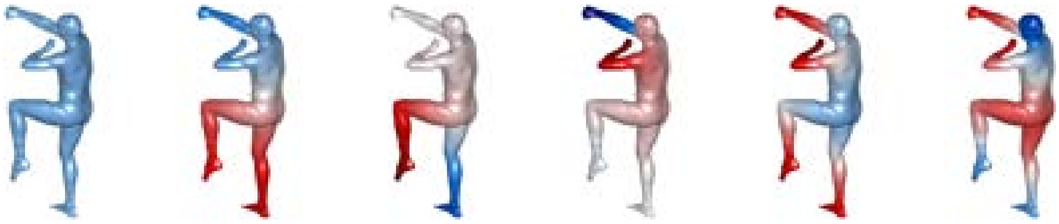

(a) Global Laplace-Beltrami eigenfunctions characterize the overall shape

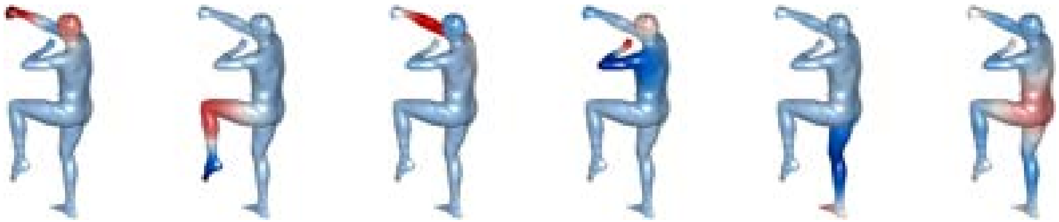

(b) Compressed Manifold Modes capture localized shape features

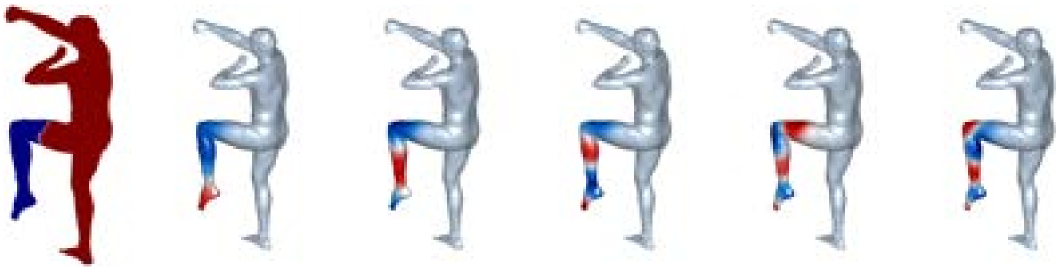

(c) Manifold Harmonics localized to the left leg of a human shape, region defined explicitly

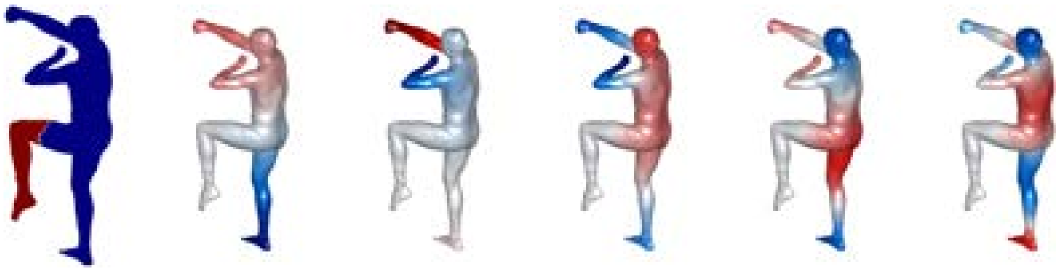

(d) Hamiltonian eigenfunctions with localized support, defined by step potential functions (red: high potential, blue: low potential)

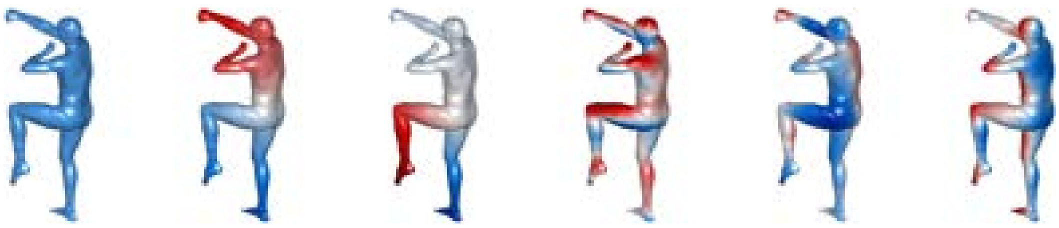

(e) Coordinate Manifold Harmonics, with the first three figures showing intrinsic shape information and the last three showing extrinsic coordinate information

FIGURE 5.1: Comparing the localization characteristics of different basis functions on a human shape from the SHREC 2010 dataset (Bronstein, Bronstein, Castellani, Falcidieno, Fusiello, Godil, Guibas, Kokkinos, Lian, Ovsjanikov, Patané, Spagnuolo, and Toldo, 2010)

computation with basis design to align descriptor functions between two shapes through a complex non-linear optimization problem. This process is simplified using Proper Orthogonal Decomposition (POD) modes, improving the representation of high-frequency signals. Regularization terms promote smoothness and consistency, with optimization achieved via ADMM. However, the accuracy of shape correspondence relies on the choice and size of the subspace; larger subspaces offer flexibility but require stronger regularization to avoid poor solutions. Marin et al. (Marin, Rakotosaona, Melzi, and Ovsjanikov, 2020) noted that bases learned through machine learning improve accuracy in challenging scenarios, but they do not exploit the underlying geometric structure of the shapes.

5.2.8 Landmark Adapted Basis

To transfer textures accurately, artists use landmark points that must be preserved and expanded into a dense correspondence map, even with non-rigid deformations. While functional maps can incorporate these landmarks, they may not ensure precise preservation. Panine et. al. in (Panine, Kirgo, and Ovsjanikov, 2022) introduced landmark-adapted basis (LAB) functions, combining Dirichlet-Laplacian and Dirichlet-Steklov eigenproblems to ensure exact landmark preservation and support conformal mapping. Using the cotangent Laplacian and a non-lumped Steklov mass matrix on the boundary, LAB improves correspondence accuracy for non-isometric shapes. This descriptor-free approach is efficient, robust, and handles significant mesh variations and noise effectively.

5.2.9 Elastic Basis

Computing correspondences for shapes undergoing non-isometric deformations, especially for aligning extrinsic features, is challenging. To address this, Hartwig et al. in (Hartwig, Sassen, Azencot, Rumpf, and Ben-Chen, 2023) introduced a non-orthogonal crease-aware elastic basis, derived from the Hessian of elastic thin shell energy, which captures bending and stretching properties based on shell thickness. In contrast to the previous orthogonal bases utilized in functional maps, the non-orthogonal nature of this elastic basis introduces challenges for standard computations, leading the authors to modify the functional map framework to achieve effective point-to-point correspondences. For optimal results, the elastic basis necessitates the use of regular triangular meshes.

5.3 Comparative Analysis of Bases Functions

Table 5.1 summarizes the properties satisfied by different basis functions employed in the functional map framework since its inception for addressing shape matching. The Laplace-Beltrami Operator (LBO) eigenbasis, a popular choice for shape analysis, is a well-established basis which is invariant to isometric deformations. However, it struggles with non-isometric deformations, is inconsistent in aligning across different shapes and poorly represents high-frequency signals. Coupled-Quasi Harmonic Basis (CQHB) addresses consistency issues but is computationally expensive. While LBO eigenfunctions capture global information, they lack localization, hindering partial shape matching and requiring large number of basis functions for detailed representation. Compressed Manifold Modes (CMMs) and Localized Manifold Harmonics (LMH) offer localized support, improving shape interpretation and handling topological noise. LMH, in particular, excels at capturing sharp details with fewer basis functions compared to LBO. However, LMH's reliance on precomputed LBO basis and region of support information requirement limits its popularity. Hamiltonian basis, derived from the vibrational modes of a shape, provides guided localization and is well-suited for partial shape matching. Coordinate Manifold Harmonics (CMH) incorporate extrinsic information, making them valuable for matching shapes with similar intrinsic structure but different spatial positioning. While CQHB and data-driven basis both consider multiple shapes, CQHB prioritizes smooth basis functions, whereas data-driven basis does not. Elastic basis, designed for non-isometric deformations, aligns creases and folds accurately, outperforming conventional orthogonal basis. Landmark Adapted Basis (LAB) precisely preserves landmark correspondences, making it ideal for texture transfer applications. Indicator basis excels at transferring non-smooth functions, a challenging task for other smooth operator basis.

	LBO basis	CMM	CQHB	LMH	CMH	HO basis	Data-driven basis	Elastic Basis	LAB
Completeness	✓	✓		✓	✓	✓			✓
Compactness	✓			✓	✓	✓		✓	✓
Stability	✓	✓	✓	✓	✓	✓			✓
Orthogonality	✓	✓	✓	✓	✓	✓	×	×	✓*
Multi-scale Property	✓	×	✓	✓	✓	✓	×	✓	✓
Smoothness	✓	✓	✓	✓	✓	✓	✓*	✓	✓
Computational Efficiency	✓	×	×	✓	✓	✓	✓*	×	✓*
Numerical stability	✓*	✓	✓	✓	✓	✓	✓*	✓*	✓*
Consistency for Isometric deformations	✓	✓	✓	✓	✓	✓	✓*	×	×
Consistency for non-Isometric deformations	×	×	✓	×	×	×	✓*	✓	✓
Local Support	×	✓	×	✓	×	✓*	✓*	×	✓*
Global Support	✓	×	✓	✓	✓	✓*	✓	✓	✓*
Robustness to topological Noise	×	✓	✓	✓	×	✓	✓	×	✓
Shared-basis	×	×	✓	×	×	×	✓	×	×

TABLE 5.1: Diverse basis functions in non-rigid shape correspondence with distinct properties (✓: fully satisfied, ✓*: conditionally satisfied, ×: not satisfied, " " (space): not studied in literature)

The diversity of functional bases provides unique advantages for non-rigid shape correspondence, with each basis offering trade-offs in consistency, localization, and computational efficiency. While the LBO remains foundational, alternative bases like CQHB, LMH, and LAB extend its capabilities to better handle non-isometric deformations and complex, localized shape details. Refer to Table 5.1 for a summary of various functional bases used in non-rigid shape correspondence, highlighting their adherence to specific properties. For a visual comparison of localization across basis functions, refer to Figure 5.1. LBO eigenfunctions, lacking local support, are shown in Figure 5.1(a). CMMs (Figure 5.1(b)) capture specific local regions like hands and legs, aiding partial shape matching. CMH (Figure 5.1(e)) reflects both intrinsic (left) and extrinsic spatial information (right). LMH (Figure 5.1(c)) localizes information to designated regions (blue area). In Figure 5.1(d), the Hamiltonian operator basis highlights low-potential regions, avoiding high-potential areas (red).

5.4 Discussions and Future Scope

Research on point-to-point correspondences for non-rigid shapes through functional maps highlights several promising directions. Challenges persist in defining functional bases for partial and non-isometric shapes, determining the optimal number of basis functions, and exploring the behavior of alternative operators. Expanding the use of the Hamiltonian operator beyond step potential and further adapting functional maps for non-orthogonal bases like the Elastic basis could yield improvements. Additionally, identifying descriptors that capture both local and global geometry in a stable manner, optimizing landmark clustering, and exploring basis functions for diverse shape representations beyond triangle meshes are encouraging avenues for future exploration. While no single basis meets all challenges, recent advancements emphasize localization, extrinsic feature integration, and practical considerations, marking a shift from traditional orthogonality and smoothness characterization.

Chapter 6

Conclusion

This thesis explores techniques for achieving robust point-to-point correspondence between isometrically deformable non-rigid shapes, emphasizing the preservation of sharp geometric details. By investigating the impact of selecting intrinsic potential functions, particularly Gaussian curvature, in defining the Hamiltonian operator, it was established that sharper geometric features could be effectively captured in the basis functions used within the functional map framework. This enhancement improved point-to-point shape correspondence by addressing the Laplace-Beltrami operator's limitations in capturing high-frequency details and global shape information in non-rigid settings. The proposed Hamiltonian basis effectively combined sharp feature detection with global shape representation, achieving superior performance in correspondence tasks. However, further exploration is needed to enhance its adaptability for better handling of flat regions.

The integration of shape-aware biharmonic distances to derive canonical forms, the pose-invariant representations of shapes in a common metric space, alongside with the non-rigid Iterative Closest Point (ICP) algorithm for aligning these canonical forms, resulted in enhanced point-to-point correspondences. The non-rigid alignment approach using thin-plate splines offered a robust representation of deformation characteristics, surpassing the traditional ICP algorithm, which is designed for rigid transformations. Biharmonic distances also effectively captured surface variations as opposed to Geodesic distances. The improvements were particularly evident in applications such as texture transfer, where the method effectively preserved local and global shape fidelity, demonstrating its utility in tasks requiring detailed deformation modeling. However, the reliance of Biharmonic distances on Laplace-Beltrami operator for its computation resulted in the smoothing of high-frequency details. Incorporating Hamiltonian operators with intrinsic potentials could further enhance the accuracy of correspondence maps by preserving finer geometric details.

The thesis also presents a comprehensive study of various basis functions in the functional map framework, examining their effectiveness in representing non-rigidly deformed shapes. Basis functions are crucial for shape correspondence, enabling efficient mappings that capture global and local geometric properties, minimizing ambiguities and ensuring robust matching. While each basis function offers specific advantages for different deformation types, no single basis is universally ideal. The work highlights the trade-offs between computational efficiency, theoretical rigor, and the ability to capture both global and local shape features. The Laplace-Beltrami eigenbasis excels in isometric deformations, while alternatives like the Coupled-Quasi Harmonic Basis (CQHB) and Elastic Basis address more complex non-isometric scenarios. Localized bases such as Compressed Manifold Modes, along with novel approaches integrating extrinsic features like Coordinate Manifold Harmonics, highlights the evolving focus on handling real-world challenges like partial matching and artistic attribute preservation. The emerging focus on localization and integration of extrinsic features marks a shift from traditional smooth and orthogonal bases to those that better handle complex, real-world shape variations.

In conclusion, this thesis highlights the potential of advanced basis functions, such as Hamiltonian eigenfunctions and improved distance metric like biharmonic distances, in enhancing the accuracy of shape correspondence. However, further exploration of adaptive models and hybrid methodologies is essential to fully leverage these techniques, addressing the complexities of non-rigid shapes and enabling more precise and efficient matching for a wide range of applications.

6.1 Future Scope

The thesis explores several promising avenues for future research focused on enhancing pointwise correspondence accuracy, robustness, and efficiency in non-rigid shape analysis. Future work can refine current methods, including more reliable error localization, while improving scalability. Key directions include:

6.1.1 Enhancing Basis Functions and Operators

- **Exploring Alternative Operators:** Investigate eigenfunctions of alternative operators beyond the Laplace-Beltrami operator and Hamiltonian operator to provide novel geometric insights.
- **Developing Adaptive Potential Functions:** Develop adaptive potential functions for Hamiltonian operator, capable of dynamically identifying critical geometric regions. This could address the limitation that the potential function is currently fixed prior to analysis (e.g., Gaussian curvature) and contributes to error localization in flat areas.
- **Bases for Partial and Non-Isometric Matching** Design bases tailored to partial shapes and non-isometric matching to handle complex, real-world deformations.

6.1.2 Optimising the Functional Map Framework

- **Accommodating Non-Orthogonal Bases:** Extend functional maps to accommodate non-orthogonal bases like Elastic bases or a shape-adaptive non-orthogonal bases, to enhance flexibility in representing shape variations.
- **Determining Optimal Basis Number:** Develop methods for determining the optimal number of basis functions used in functional map calculations, to balance accuracy and computational cost.

- **Designing Robust Descriptor Functions:** Designing multi-scale shape descriptors that capture both local detail and global structure across diverse deformation types, can improve correspondence stability compared to purely local (e.g., point curvature) or purely global (e.g., eigenvalues) descriptors.

- **Enhancing post-map refinement** After computing functional maps, develop stronger optimization strategies to convert them into more accurate and stable pointwise correspondences.

6.1.3 Real-world Shape Correspondence

- **Robustness to complex deformations** Improve handling of deformations such as stretching while preserving high-frequency details alongside global geometric consistency for tougher real-world datasets, going beyond triangle meshes.

- **Data-driven geometric consistency** Explore hybrid approaches that combine geometric pipelines (e.g., functional maps) with neural predictions to learn better descriptors and improve correspondence robustness on noisy, diverse real-world shapes.

Bibliography

Euclid (300 BC). *Elements*.

Dijkstra, E. W. (Dec. 1959). "A note on two problems in connexion with graphs". In: *Numerische Mathematik* 1.1, 269–271.
URL: https://doi.org/10.1007/BF01386390.

Kendall, David G (1977). "The Diffusion of Shape". In: *Advances in Applied Probability* 9.3, pp. 428–430.
URL: http://www.jstor.org/stable/1426091.

Levoy, Marc and Turner Whitted (1985). *The use of points as a display primitive*. Tech. rep. 85-022, University of North Carolina at Chapel Hill, NC.
URL: https://graphics.stanford.edu/papers/points.

Fischler, Martin A. and Robert C. Bolles (1987). "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography". In: *Readings in Computer Vision*. Ed. by Martin A. Fischler and Oscar Firschein. San Francisco (CA): Morgan Kaufmann, pp. 726–740.
URL: https://www.sciencedirect.com/science/article/pii/B9780080515816500702.

Schwartz, E.L., A. Shaw, and E. Wolfson (1989). "A numerical solution to the generalized mapmaker's problem: flattening nonconvex polyhedral surfaces". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 11.9, pp. 1005–1008.
URL: https://doi.org/10.1109/34.35506.

Cox, Trevor F. and Michael A. A. Cox (Jan. 1991). "Multidimensional scaling on a sphere". In: *Communications in Statistics - Theory and Methods* 20.9, pp. 2943–2953.
URL: http://doi.org/10.1080/03610929108830679.

Besl, P.J. and Neil D. McKay (1992). "A method for registration of 3-D shapes". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14.2, pp. 239–256.
URL: https://doi.org/10.1109/34.121791.

Chung, Fan (Dec. 1997). *Spectral Graph Theory*. Vol. 92. American Mathematical Society.
URL: http://doi.org/10.1090/cbms/092.

Kimmel, Ron and James A Sethian (1998). "Computing geodesic paths on manifolds". In: *Proceedings of the national academy of Sciences* 95.15, pp. 8431–8435.
URL: https://www.pnas.org/doi/abs/10.1073/pnas.95.15.8431.

Sethian, James A. (1999). *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science*. Second. Cambridge Monograph on Applied and Computational Mathematics. Cambridge Univeristy Press.
URL: https://math.berkeley.edu/~sethian/Books/hold_sethian_book.pdf.

Karni, Zachi and Craig Gotsman (2000). "Spectral compression of mesh geometry". In: *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '00. USA: ACM Press/Addison-Wesley Publishing Co., 279–286.
URL: https://doi.org/10.1145/344779.344924.

Zorin, Denis, Peter Schröder, Tony DeRose, Leif Kobbelt, Adi Levin, and Wim Sweldens (2000). "Subdivision for modeling and animation". In: *ACM SIGGRAPH 2000 Course Notes*.
URL: https://mrl.cs.nyu.edu/publications/subdiv-course2000/.

Rusinkiewicz, S. and M. Levoy (2001). "Efficient variants of the ICP algorithm". In: *Proceedings of the Third International Conference on 3-D Digital Imaging and Modeling*, pp. 145–152.
URL: https://doi.org/10.1109/IM.2001.924423.

Veltkamp, Remco C. and Michiel Hagedoorn (2001). "State of the Art in Shape Matching". In: *Principles of Visual Information Retrieval*. Ed. by

Michael S. Lew. London: Springer London, pp. 87–119.
URL: https://doi.org/10.1007/978-1-4471-3702-3_4.

Alexa, Marc (2002). "Recent advances in mesh morphing". In: *Computer graphics forum*. Vol. 21. 2. Wiley Online Library, pp. 173–198.
URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-8659.00575.

Linial, Nathan (2002). "Finite metric spaces: combinatorics, geometry and algorithms". In: *Proceedings of the Eighteenth Annual Symposium on Computational Geometry*. SCG '02. Barcelona, Spain: Association for Computing Machinery, p. 63.
URL: https://doi.org/10.1145/513400.513441.

Zigelman, G., R. Kimmel, and N. Kiryati (2002). "Texture mapping using surface flattening via multidimensional scaling". In: *IEEE Transactions on Visualization and Computer Graphics* 8.2, pp. 198–207.
URL: http://doi.org/10.1109/2945.998671.

Alexa, Marc, Johannes Behr, Daniel Cohen-Or, Shachar Fleishman, David Levin, and Claudio T. Silva (2003). "Computing and rendering point set surfaces". In: *IEEE Transactions on Visualization and Computer Graphics* 9.1, pp. 3–15.
URL: https://doi.org/10.1109/TVCG.2003.1175093.

Belkin, Mikhail and Partha Niyogi (2003). "Laplacian Eigenmaps for Dimensionality Reduction and Data Representation". In: *Neural computation* 15.6, pp. 1373–1396.
URL: https://doi.org/10.1162/089976603321780317.

Elad, Asi and Ron Kimmel (2003). "On bending invariant signatures for surfaces". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25.10, pp. 1285–1295.
URL: https://doi.org/10.1109/TPAMI.2003.1233902.

Kimmel, Ron (2003). *Numerical Geometry of Images: Theory, Algorithms, and Applications*. Springer New York, NY, pp. XIII, 209.
URL: https://doi.org/10.1007/978-0-387-21637-9.

Meyer, Mark, Mathieu Desbrun, Peter Schröder, and Alan H. Barr (2003). "Discrete Differential-Geometry Operators for Triangulated 2-Manifolds". In: *Visualization and Mathematics III, Springer Berlin Heidelberg*, pp. 35–57. URL: `https://doi.org/10.1007/978-3-662-05105-4_2`.

Sumner, Robert W and Jovan Popović (Aug. 2004). "Deformation transfer for triangle meshes". In: *ACM Transactions on Graphics (TOG)* 23.3, pp. 399–405. URL: `https://doi.org/10.1145/1015706.1015736`.

Berg, A.C., T.L. Berg, and J. Malik (2005). "Shape matching and object recognition using low distortion correspondences". In: *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*. Vol. 1, pp. 26–33. URL: `https://doi.org/10.1109/cvpr.2005.320`.

Borg, Ingwer and Patrick JF Groenen (2005). *Modern Multidimensional Scaling: Theory and applications*. Springer Series in Statistics. Springer New York, NY. URL: `https://doi.org/10.1007/0-387-28981-X`.

Dinh, Huong Quynh, Anthony Yezzi, and Greg Turk (Apr. 2005). "Texture transfer during shape transformation". In: *ACM Transactions on Graphics (ToG)* 24.2, pp. 289–310. URL: `https://doi.org/10.1145/1061347.1061353`.

Gelfand, Natasha, Niloy J Mitra, Leonidas J Guibas, and Helmut Pottmann (2005a). "Robust global registration". In: *SGP '05: Proceedings of the third Eurographics Symposium on Geometry processing*. Vol. 2. 3. Vienna, Austria, pp. 197–206. URL: `http://vecg.cs.ucl.ac.uk/Projects/SmartGeometry/global_registration/paper_docs/global_registration_sgp_05.pdf`.

Gelfand, Natasha, Niloy J. Mitra, Leonidas J. Guibas, and Helmut Pottmann (2005b). "Robust global registration". In: *Proceedings of the Third Eurographics Symposium on Geometry Processing*. Vol. 2. SGP '05 3. Vienna, Austria: Eurographics Association, 197–es.

Mémoli, Facundo and Guillermo Sapiro (June 2005). "A Theoretical and Computational Framework for Isometry Invariant Recognition of Point Cloud Data". In: *Found. Comput. Math.* 5.3, 313–347.
URL: `http://doi.org/10.1007/s10208-004-0145-y`.

Biasotti, Silvia, Simone Marini, Michela Spagnuolo, and Bianca Falcidieno (2006). "Sub-part correspondence by structural descriptors of 3D shapes". In: *Computer-Aided Design* 38.9. Shape Similarity Detection and Search for CAD/CAE Applications, pp. 1002–1019.
URL: `https://www.sciencedirect.com/science/article/pii/S0010448506001345`.

Bronstein, Alexander M., Michael M. Bronstein, and Ron Kimmel (Jan. 2006). "Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching". In: *Proceedings of the National Academy of Sciences* 103.5, 1168–1172.
URL: `https://doi.org/10.1073/pnas.0508601103`.

Jain, Varun and Hao Zhang (2006). "Robust 3D Shape Correspondence in the Spectral Domain". In: *Proceedings of the IEEE International Conference on Shape Modeling and Applications 2006*. SMI '06. USA: IEEE Computer Society, p. 19.
URL: `http://doi.org/10.1109/smi.2006.31`.

Jones, M.W., J.A. Baerentzen, and M. Sramek (July 2006). "3D distance fields: a survey of techniques and applications". In: *IEEE Transactions on Visualization and Computer Graphics* 12.4, pp. 581–599.
URL: `http://doi.org/10.1109/tvcg.2006.56`.

Levy, B. (2006). "Laplace-Beltrami Eigenfunctions Towards an Algorithm That "Understands" Geometry". In: *IEEE International Conference on Shape Modeling and Applications 2006 (SMI'06)*. Institute of Electrical and Electronics Engineers (IEEE) Computer Soceity, pp. 13–13.
URL: `https://doi.org/10.1109/smi.2006.21`.

Reuter, Martin, Franz-Erich Wolter, and Niklas Peinecke (2006). "Laplace–Beltrami spectra as 'Shape-DNA'of surfaces and solids". In: *Computer-Aided Design* 38.4, pp. 342–366.

URL: https://www.sciencedirect.com/science/article/pii/S0010448505001867.

Bronstein, Alexander M., Michael M. Bronstein, and Ron Kimmel (2007). "Rock, Paper, and Scissors: extrinsic vs. intrinsic similarity of non-rigid shapes". In: *IEEE 11th International Conference on Computer Vision (ICCV)*. IEEE, pp. 1–6.
URL: https://doi.org/10.1109/ICCV.2007.4409076.

Gromov, Mikhail (2007). *Metric Structures for Riemannian and Non-Riemannian Spaces*. Birkhäuser Boston, pp. XX,586.
URL: https://doi.org/10.1007/978-0-8176-4583-0.

Jain, Varun, Hao Zhang, and Oliver Van Kaick (2007). "Non-rigid spectral correspondence of triangle meshes". In: *International Journal of Shape Modeling* 13.01, pp. 101–124.
URL: https://doi.org/10.1142/S0218654307000968.

Mémoli, Facundo (Jan. 2007). "On the use of Gromov-Hausdorff Distances for Shape Comparison". In: *Eurographics Symposium on Point-Based Graphics*. The Eurographics Association, pp. 81–90.
URL: https://doi.org/10.2312/SPBG/SPBG07/081-090.

Rustamov, Raif M. (2007). "Laplace-Beltrami Eigenfunctions for Deformation Invariant Shape Representation". In: *SGP '07: Proceedings of the Fifth Eurographics Symposium on Geometry Processing*. Ed. by Alexander Belyaev and Michael Garland. The Eurographics Association, 225–233.
URL: https://doi.org/10.2312/SGP/SGP07/225-233.

Wardetzky, Max, Saurabh Mathur, Felix Kaelberer, and Eitan Grinspun (2007). "Discrete Laplace operators: No free lunch". In: *Geometry Processing*. The Eurographics Association.
URL: https://doi.org/10.2312/SGP/SGP07/033-037.

Bronstein, Alexander, Michael Bronstein, and Ron Kimmel (Oct. 2008). *Numerical Geometry of Non-Rigid Shapes*. 1st ed. Monographs in Computer Science. Springer New York, NY, p. 346.
URL: https://doi.org/10.1007/978-0-387-73301-2.

Castellani, Umberto, Marco Cristani, Simone Fantoni, and Vittorio Murino (2008). "Sparse points matching by combining 3D mesh saliency with statistical descriptors". In: *Computer Graphics Forum*. Vol. 27. 2. Wiley Online Library, pp. 643–652.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2008.01162.x`.

Chang, Will and Matthias Zwicker (2008). "Automatic Registration for Articulated Shapes". In: *Computer Graphics Forum* 27.5, pp. 1459–1468.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2008.01286.x`.

Memoli, Facundo (June 2008). "Gromov-Hausdorff distances in Euclidean spaces". In: *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPR)*. Institute of Electrical and Electronics Engineers (IEEE), pp. 1–8.
URL: `http://doi.org/10.1109/cvprw.2008.4563074`.

Shamir, Ariel (2008). "A survey on Mesh Segmentation Techniques". In: *Computer Graphics Forum* 27.6, pp. 1539–1556.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2007.01103.x`.

Tangelder, Johan W. H. and Remco C. Veltkamp (Dec. 2008). "A survey of content based 3D shape retrieval methods". In: *Multimedia Tools and Applications* 39.3, pp. 441–471.
URL: `http://doi.org/10.1007/s11042-007-0181-0`.

Taylor Chris, Davies Rhodri and Twining Carole (2008). Springer-Verlag London.
URL: `https://doi.org/10.1007/978-1-84800-138-1`.

Vallet, Bruno and Bruno Lévy (2008). "Spectral Geometry Processing with Manifold Harmonics". In: *Computer Graphics Forum* 27.2, pp. 251–260.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2008.01122.x`.

Zass, Ron and Amnon Shashua (2008). “Probabilistic graph and hypergraph matching”. In: *2008 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1–8.
URL: `https://doi.org/10.1109/cvpr.2008.4587500`.

Zhang, H., A. Sheffer, D. Cohen-Or, Q. Zhou, O. Van Kaick, and A. Tagliasacchi (2008). “Deformation-Driven Shape Correspondence”. In: *Computer Graphics Forum* 27.5, pp. 1431–1439.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2008.01283.x`.

Belkin, Mikhail, Jian Sun, and Yusu Wang (2009). “Constructing Laplace Operator from Point Clouds in R^d”. In: *Proceedings of the 2009 Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SODA ’09. New York, New York: Society for Industrial and Applied Mathematics, pp. 1031–1040.
URL: `https://epubs.siam.org/doi/abs/10.1137/1.9781611973068.112`.

Bronstein, Alexander M, Michael M Bronstein, Alfred M Bruckstein, and Ron Kimmel (2009). “Partial similarity of objects, or how to compare a centaur to a horse”. In: *International Journal of Computer Vision* 84, pp. 163–183.
URL: `https://doi.org/10.1007/s11263-008-0147-3`.

Bronstein, Alexander M, Michael M Bronstein, and Ron Kimmel (2009). “Topology-Invariant Similarity of Nonrigid Shapes”. In: *International journal of computer vision* 81, pp. 281–301.
URL: `https://doi.org/10.1007/s11263-008-0172-2`.

Bronstein, Alexander M., Michael M. Bronstein, Ron Kimmel, Mona Mahmoudi, and Guillermo Sapiro (Oct. 2009). “A Gromov-Hausdorff Framework with Diffusion Geometry for Topologically-Robust Non-rigid Shape Matching”. In: *International Journal of Computer Vision* 89.2–3, 266–286.
URL: `http://doi.org/10.1007/s11263-009-0301-6`.

Hasler, N., C. Stoll, M. Sunkel, B. Rosenhahn, and H.-P. Seidel (2009). “A Statistical Model of Human Pose and Body Shape”. In: *Computer Graphics Forum* 28.2, pp. 337–346.

URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2009.01373.x.

Lipman, Yaron and Thomas Funkhouser (2009). "Möbius voting for surface correspondence". In: *ACM SIGGRAPH 2009 Papers*. Vol. 28. SIGGRAPH '09 3. New Orleans, Louisiana: Association for Computing Machinery, pp. 1–12.
URL: https://doi.org/10.1145/1576246.1531378.

Liu, Rong, Hao Zhang, Ariel Shamir, and Daniel Cohen-Or (2009). "A Part-aware Surface Metric for Shape Analysis". In: *Computer Graphics Forum* 28.2, pp. 397–406.
URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2009.01379.x.

Sun, Jian, Maks Ovsjanikov, and Leonidas Guibas (2009). "A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion". In: *Computer Graphics Forum* 28.5, pp. 1383–1392.
URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2009.01515.x.

Botsch, Mario, Leif Kobbelt, Mark Pauly, Pierre Alliez, and Bruno Levy (2010). *Polygon Mesh Processing*. 1st ed. A K Peters/CRC Press.
URL: https://doi.org/10.1201/b10688.

Bronstein, A. M., M. M. Bronstein, U. Castellani, B. Falcidieno, A. Fusiello, A. Godil, L. J. Guibas, I. Kokkinos, Z. Lian, M. Ovsjanikov, G. Patané, M. Spagnuolo, and R. Toldo (2010). "SHREC'10 track: robust shape retrieval". In: *Proceedings of the 3rd Eurographics Conference on 3D Object Retrieval*. 3DOR '10. Norrköping, Sweden: Eurographics Association, 71–78.

Bronstein, Michael M. and Iasonas Kokkinos (2010). "Scale-invariant heat kernel signatures for non-rigid shape recognition". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, pp. 1704–1711.
URL: https://doi.org/10.1109/CVPR.2010.5539838.

Lévy, Bruno and Hao (Richard) Zhang (2010). "Spectral mesh processing". In: *ACM SIGGRAPH 2010 Courses*. SIGGRAPH '10. Los Angeles, California:

Association for Computing Machinery.
URL: `https://doi.org/10.1145/1837101.1837109`.

Lipman, Yaron, Raif M. Rustamov, and Thomas A. Funkhouser (2010). "Biharmonic distance". In: *ACM Transactions on Graphics (TOG)* 29.3, pp. 0730–0301.
URL: `https://doi.org/10.1145/1805964.1805971`.

Xu, Kai, Honghua Li, Hao Zhang, Daniel Cohen-Or, Yueshan Xiong, and Zhi-Quan Cheng (2010). "Style-content separation by anisotropic part scales". In: *ACM Transactions on Graphics (ToG)*. SIGGRAPH ASIA '10 29.6, pp. 0730–0301.
URL: `https://doi.org/10.1145/1866158.1866206`.

Zhang, H., O. Van Kaick, and R. Dyer (2010). "Spectral Mesh Processing". In: *Computer Graphics Forum* 29.6, pp. 1865–1894.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2010.01655.x`.

Aubry, Mathieu, Ulrich Schlickewei, and Daniel Cremers (2011). "The wave kernel signature: A quantum mechanical approach to shape analysis". In: *IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pp. 1626–1633.
URL: `https://doi.org/10.1109/ICCVW.2011.6130444`.

Kaick, Oliver van, Hao Zhang, Ghassan Hamarneh, and Daniel Cohen-Or (2011). "A Survey on Shape Correspondence". In: *Computer Graphics Forum* 30.6, pp. 1681–1707.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2011.01884.x`.

Kim, Vladimir G., Yaron Lipman, and Thomas Funkhouser (2011). "Blended intrinsic maps". In: *ACM SIGGRAPH 2011 Papers*. Vol. 30. SIGGRAPH '11 4. Vancouver, British Columbia, Canada: Association for Computing Machinery, pp. 1–12.
URL: `https://doi.org/10.1145/1964921.1964974`.

Ovsjanikov, Maks, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas (July 2012). "Functional maps: a flexible representation

of maps between shapes". In: *ACM Transactions on Graphics (TOG)* 31.4, pp. 1–11.
URL: `https://doi.org/10.1145/2185520.2185526`.

Raviv, Dan, Anastasia Dubrovina, and Ron Kimmel (2012). "Hierarchical Matching of Non-rigid Shapes". In: *Scale Space and Variational Methods in Computer Vision: Third International Conference, SSVM 2011, Ein-Gedi, Israel*. Springer Berlin Heidelberg, 604–615.
URL: `https://doi.org/10.1007/978-3-642-24785-9_51`.

Berger, Matthew, Joshua A Levine, Luis Gustavo Nonato, Gabriel Taubin, and Claudio T Silva (Apr. 2013). "A benchmark for surface reconstruction". In: *ACM Transactions on Graphics (TOG)* 32.2, pp. 1–17.
URL: `https://doi.org/10.1145/2451236.2451246`.

Crane, Keenan, Clarisse Weischedel, and Max Wardetzky (Oct. 2013). "Geodesics in heat: A new approach to computing distance based on heat flow". In: *ACM Transactions on Graphics (TOG)* 32.5, pp. 1–11.
URL: `https://doi.org/10.1145/2516971.2516977`.

Elghoul, Esma and Anne Verroust-Blondet (2013). "A Segmentation Transfer Method for Articulated Models". In: *Eurographics 2013 - Short Papers*. Ed. by M.-A. Otaduy and O. Sorkine. The Eurographics Association.
URL: `https://doi.org/10.2312/conf/EG2013/short/017-020`.

Kovnatsky, Artiom, Michael M Bronstein, Alexander M Bronstein, Klaus Glashoff, and Ron Kimmel (2013). "Coupled quasi-harmonic bases". In: *Computer Graphics Forum* 32, pp. 439–448.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12064`.

Litman, Roee and Alexander M Bronstein (2013). "Learning spectral descriptors for deformable shape correspondence". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36.1, pp. 171–180.

Mitra, Niloy J., Mark Pauly, Michael Wand, and Duygu Ceylan (2013). "Symmetry in 3D Geometry: Extraction and Applications". In: *Computer Graphics Forum* 32.6, pp. 1–23.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12010`.

Petronetto, F., A. Paiva, E. S. Helou, D. E. Stewart, and L. G. Nonato (Sept. 2013). "Mesh-Free Discrete Laplace-Beltrami Operator". In: *Computer Graphics Forum* 32.6, 214–226.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12086`.

Pokrass, J., A. M. Bronstein, M. M. Bronstein, P. Sprechmann, and G. Sapiro (2013). "Sparse Modeling of Intrinsic Correspondences". In: *Computer Graphics Forum* 32, pp. 459–468.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12066`.

Tam, Gary K., Zhi-Quan Cheng, Yu-Kun Lai, Frank Langbein, Yonghuai Liu, A. David Marshall, Ralph Martin, Xianfang Sun, and Paul Rosin (July 2013). "Registration of 3D Point Clouds and Meshes: A Survey from Rigid to Nonrigid". In: *IEEE Transactions on Visualization and Computer Graphics* 19.7, pp. 1199–1217.
URL: `https://doi.org/10.1109/TVCG.2012.310`.

Wenger, Rephael (June 2013). In: *Isosurfaces: Geometry, Topology, and Algorithms (1st ed.)* A K Peters/CRC Press, 369–392.
URL: `http://doi.org/10.1201/b15025-15`.

F. Groenen, Patrick J. and Michel Velden (2014). "Multidimensional Scaling". In: *Wiley StatsRef: Statistics Reference Online*. John Wiley & Sons, Ltd.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118445112.stat06476`.

Li, Xin and S. S. Iyengar (Dec. 2014). "On computing mapping of 3d objects: A survey". In: *ACM Computing Surveys (CSUR)* 47.2, pp. 1–45.
URL: `https://doi.org/10.1145/2668020`.

Neumann, T., K. Varanasi, C. Theobalt, M. Magnor, and M. Wacker (2014). "Compressed Manifold Modes for Mesh Processing". In: *Computer Graphics Forum* 33.5, pp. 35–44.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12429`.

Rodola, Emanuele, Samuel Rota Bulo, Thomas Windheuser, Matthias Vestner, and Daniel Cremers (June 2014). "Dense Non-Rigid Shape Correspondence Using Random Forests". In: *2014 IEEE Conference on Computer*

Vision and Pattern Recognition (CVPR). Institute of Electrical and Electronics Engineers (IEEE) Computer Soceity, 4177–4184.
URL: `http://doi.org/10.1109/cvpr.2014.532`.

Solomon, Justin, Keegan Crane, and Etienne Vouga (2014). "Laplace-Beltrami: The Swiss army knife of geometry processing". In: *Symposium on Geometry Processing Graduate School (Cardiff, UK, 2014)*. Vol. 2.

Aflalo, Yonathan, Haim Brezis, and Ron Kimmel (2015). "On the Optimality of Shape and Data Representation in the Spectral Domain". In: *SIAM Journal on Imaging Sciences* 8.2, pp. 1141–1160.
URL: `https://doi.org/10.1137/140977680`.

Corman, Étienne, Maks Ovsjanikov, and Antonin Chambolle (2015). "Supervised Descriptor Learning for Non-Rigid Shape Matching". In: *Computer Vision - ECCV 2014 Workshops*. Lecture Notes in Computer Science. Cham: Springer International Publishing, pp. 283–298.

Eynard, Davide, Artiom Kovnatsky, Michael M Bronstein, Klaus Glashoff, and Alexander M Bronstein (2015). "Multimodal Manifold Analysis by Simultaneous Diagonalization of Laplacians". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 37.12, pp. 2505–2517.
URL: `https://doi.org/10.1109/TPAMI.2015.2408348`.

Gasparetto, Andrea, Giorgia Minello, and Andrea Torsello (2015). "Non-parametric Spectral Model for Shape Retrieval". In: *2015 International Conference on 3D Vision*, pp. 344–352.
URL: `https://doi.org/10.1109/3dv.2015.46`.

Kovnatsky, Artiom, Michael M. Bronstein, Xavier Bresson, and Pierre Vandergheynst (2015). "Functional correspondence by matrix completion". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 905–914.
URL: `https://doi.org/10.1109/CVPR.2015.7298692`.

Maglo, Adrien, Guillaume Lavoué, Florent Dupont, and Céline Hudelot (Feb. 2015). "3d mesh compression: Survey, comparisons, and emerging trends". In: *ACM Computing Surveys (CSUR)* 47.3, pp. 1–41.
URL: `https://doi.org/10.1145/2693443`.

Shtern, Alon and Ron Kimmel (2015). "Spectral gradient fields embedding for nonrigid shape matching". In: *Computer Vision and Image Understanding* 140, pp. 21–29.
URL: https://www.sciencedirect.com/science/article/pii/S1077314215000387.

Do Carmo, Manfredo P (2016). *Differential Geometry of Curves and Surfaces: Revised and Updated Second Edition*. Dover Publications.

Masci, Jonathan, Emanuele Rodolà, Davide Boscaini, Michael M. Bronstein, and Hao Li (2016). "Geometric deep learning". In: *SIGGRAPH ASIA 2016 Courses*. SA '16. Macau: Association for Computing Machinery.
URL: https://doi.org/10.1145/2988458.2988485.

Xu, Kai, Vladimir G. Kim, Qixing Huang, Niloy Mitra, and Evangelos Kalogerakis (2016). "Data-Driven Shape Analysis and Processing". In: *SIGGRAPH ASIA 2016 Courses*. SA '16. New York, NY, USA: Association for Computing Machinery.
URL: https://doi.org/10.1145/2988458.2988473.

Crane, Keenan, Clarisse Weischedel, and Max Wardetzky (Oct. 2017). "The heat method for distance computation". In: *Communications of the ACM* 60.11, 90–99.
URL: https://doi.org/10.1145/3131280.

Ezuz, Danielle and Mirela Ben-Chen (2017). "Deblurring and denoising of maps between shapes". In: *Computer Graphics Forum* 36.5, pp. 165–174.
URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13254.

Haas, Sylvain, Atilla Baskurt, Florent Dupont, and Florence Denis (2017). "A framework based on compressed manifold modes for robust local spectral analysis". In: *Proceedings of the Eurographics Workshop on 3D Object Retrieval (3DOR'17)*. Lyon, France, 63–66.
URL: https://doi.org/10.2312/3dor.20171054.

Litany, Or, Tal Remez, Emanuele Rodola, Alex Bronstein, and Michael Bronstein (2017). "Deep functional maps: Structured prediction for dense shape correspondence". In: *Proceedings of the IEEE international conference*

on computer vision (ICCV), pp. 5660–5668.
URL: `https://doi.org/10.1109/iccv.2017.603`.

Nogneng, Dorian and Maks Ovsjanikov (2017). "Informative Descriptor Preservation via Commutativity for Shape Matching". In: *Computer Graphics Forum* 36.2, pp. 259–267.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13124`.

Rodolà, Emanuele, Luca Cosmo, Michael M Bronstein, Andrea Torsello, and Daniel Cremers (2017). "Partial functional correspondence". In: *Computer Graphics Forum* 36.1, pp. 222–236.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12797`.

Choukroun, Yoni, Alon Shtern, Alex M Bronstein, and Ron Kimmel (2018). "Hamiltonian Operator for Spectral Shape Analysis". In: *IEEE Transactions on Visualization and Computer Graphics* 26.2, pp. 1320–1331.
URL: `https://doi.org/10.1109/TVCG.2018.2867513`.

Hu, Ruizhen, Cheng Wen, Oliver Van Kaick, Luanmin Chen, Di Lin, Daniel Cohen-Or, and Hui Huang (Dec. 2018). "Semantic object reconstruction via casual handheld scanning". In: *ACM Transactions on Graphics (TOG)* 37.6.
URL: `https://doi.org/10.1145/3272127.3275024`.

Melzi, Simone, Emanuele Rodolà, Umberto Castellani, and Michael M Bronstein (2018). "Localized Manifold Harmonics for Spectral Shape Analysis". In: *Computer Graphics Forum* 37.6, pp. 20–34.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13309`.

Dyke, Roberto M., Chris Stride, Yu-Kun Lai, and Paul L. Rosin (2019). "Shrec'19: Shape Correspondence with Isometric and Non-Isometric Deformations". In: *Proceedings of the Eurographics Workshop on 3D Object Retrieval (3DOR'19), Italy, Cardiff University. Dataset*, p. 9.

Halimi, Oshri, Or Litany, Emanuele Rodolà Rodolà, Alex M. Bronstein, and Ron Kimmel (2019). "Unsupervised Learning of Dense Shape Correspondence". In: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4365–4374.
URL: `https://doi.org/10.1109/cvpr.2019.00450`.

Knoverek, Catherine R, Gaya K Amarasinghe, and Gregory R Bowman (Apr. 2019). “Advanced Methods for Accessing Protein Shape-Shifting Present New Therapeutic Opportunities”. In: *Trends in biochemical sciences* 44.4, pp. 351–364.
URL: http://doi.org/10.1016/j.tibs.2018.11.007.

Marin, Riccardo, Simone Melzi, Pietro Musoni, Filippo Bardon, Marco Tarini, and Umberto Castellani (2019). “CMH: Coordinates Manifold Harmonics for Functional Remeshing”. In: *Eurographics Workshop on 3D Object Retrieval*. Ed. by Silvia Biasotti, Guillaume Lavoué, and Remco Veltkamp. The Eurographics Association, pp. 63–70.
URL: https://doi.org/10.2312/3dor.20191063.

Melzi, Simone, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, and Maks Ovsjanikov (2019). “ZoomOut: Spectral Upsampling for Efficient Shape Correspondence”. In: *ACM Transaction on Graphics (TOG)* 38.6, pp. 1–14.
URL: https://doi.org/10.1145/3355089.3356524.

Rampini, Arianna, Irene Tallini, Maks Ovsjanikov, Alex M. Bronstein, and Emanuele Rodola (Sept. 2019). “ Correspondence-Free Region Localization for Partial Shape Similarity via Hamiltonian Spectrum Alignment ”. In: *2019 International Conference on 3D Vision (3DV)*. Los Alamitos, CA, USA: IEEE Computer Society, pp. 37–46.
URL: https://doi.ieeecomputersociety.org/10.1109/3DV.2019.00014.

Ren, Jing, Mikhail Panine, Peter Wonka, and Maks Ovsjanikov (2019). “Structured Regularization of Functional Map Computations”. In: *Computer Graphics Forum* 38.5, pp. 39–53.
URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13788.

Crane, Keenan, Marco Livesu, Enrico Puppo, and Yipeng Qin (July 2020). “A Survey of Algorithms for Geodesic Paths and Distances”. In: *arXiv e-prints*, arXiv:2007.10430, arXiv:2007.10430. arXiv: 2007.10430 [cs.GR].
URL: https://arxiv.org/pdf/2007.10430.

Donati, Nicolas, Abhishek Sharma, and Maks Ovsjanikov (2020). "Deep Geometric Functional Maps: Robust Feature Learning for Shape Correspondence". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8589–8598.
URL: `https://doi.org/10.1109/cvpr42600.2020.00862`.

Marin, Riccardo, Marie-Julie Rakotosaona, Simone Melzi, and Maks Ovsjanikov (2020). "Correspondence learning via linearly-invariant embedding". In: *Proceedings of the 34th International Conference on Neural Information Processing Systems*. NIPS '20. Vancouver, BC, Canada: Curran Associates Inc.
URL: `https://dl.acm.org/doi/10.5555/3495724.3495860`.

Melzi, Simone, Riccardo Marin, Pietro Musoni, Filippo Bardon, Marco Tarini, and Umberto Castellani (2020). "Intrinsic/extrinsic embedding for functional remeshing of 3D shapes". In: *Computers & Graphics* 88, pp. 1–12.
URL: `https://www.sciencedirect.com/science/article/pii/S0097849320300194`.

Postolache, Emilian, Marco Fumero, Luca Cosmo, and Emanuele Rodolà (2020). "A parametric analysis of discrete Hamiltonian functional maps". In: *Computer Graphics Forum* 39.5, pp. 103–118.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14072`.

Sahillioğlu, Yusuf (2020). "Recent advances in shape correspondence". In: *The Visual Computer* 36.8, pp. 1705–1721.

Saiti, Evdokia and Theoharis Theoharis (2020). "An application independent review of multimodal 3D registration methods". In: *Computers & Graphics* 91, pp. 153–178.
URL: `https://www.sciencedirect.com/science/article/pii/S009784932030114X`.

Sawhney, Rohan and Keenan Crane (Aug. 2020). "Monte Carlo geometry processing: a grid-free approach to PDE-based methods on volumetric domains". In: *ACM Transaction on Graphics* 39.4.
URL: `http://doi.org/10.1145/3386569.3392374`.

Sharp, Nicholas and Keenan Crane (2020). "A Laplacian for Nonmanifold Triangle Meshes". In: *Computer Graphics Forum* 39.5, pp. 69–80.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14069`.

Azencot, Omri and Rongjie Lai (2021). "A Data-Driven Approach to Functional Map Construction and Bases Pursuit". In: *Computer Graphics Forum* 40.5, pp. 97–110.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14360`.

Cammarasana, Simone and Giuseppe Patané (2021). "Localised and shape-aware functions for spectral geometry processing and shape analysis: A survey & perspectives". In: *Computers & Graphics* 97, pp. 1–18.
URL: `https://www.sciencedirect.com/science/article/pii/S009784932100042X`.

Liu, Shan, Min Zhang, Pranav Kadam, and C.-C. Jay Kuo (2021). *3D Point Cloud Analysis: Traditional, Deep Learning, and Explainable Machine Learning Methods*. Springer International Publishing, pp. XIV, 146.
URL: `http://doi.org/10.1007/978-3-030-89180-0`.

Ren, Jing, Simone Melzi, Peter Wonka, and Maks Ovsjanikov (Aug. 2021). "Discrete Optimization for Shape Matching". In: *Computer Graphics Forum* 40.5, pp. 81–96.
URL: `http://doi.org/10.1111/cgf.14359`.

Dammann, Maximilian Peter, Wolfgang Steger, and Ralph Stelzer (Jan. 2022). "Automated and Adaptive Geometry Preparation for AR/VR-Applications". In: *Journal of Computing and Information Science in Engineering* 22.3, p. 031010.
URL: `http://doi.org/10.1115/1.4053327`.

Deng, Bailin, Yuxin Yao, Roberto M. Dyke, and Juyong Zhang (2022). "A Survey of Non-Rigid 3D Registration". In: *Computer Graphics Forum* 41.2, pp. 559–589.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14502`.

Panine, Mikhail, Maxime Kirgo, and Maks Ovsjanikov (2022). "Non-Isometric Shape Matching via Functional Maps on Landmark-Adapted Bases". In:

Computer Graphics Forum 41.6, pp. 394–417.
URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14579`.

Zhang, Juyong, Yuxin Yao, and Bailin Deng (2022). “Fast and Robust Iterative Closest Point”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44.7, pp. 3450–3466.
URL: `https://doi.org/10.1109/TPAMI.2021.3054619`.

Abdelreheem, Ahmed, Abdelrahman Eldesokey, Maks Ovsjanikov, and Peter Wonka (2023). “Zero-Shot 3D Shape Correspondence”. In: *SIGGRAPH Asia 2023 Conference Papers*. SA ’23. Association for Computing Machinery.
URL: `https://doi.org/10.1145/3610548.3618228`.

Bindal, Manika and Venkatesh Kamat (July 2023). “Detail Preserving Nonrigid Shape Correspondences”. In: *Computer Science Research Notes*. Vol. 3301. WSCG 2023. World Society for Computer Graphics , University of West Bohemia, Czech Republic, 323–330.
URL: `http://doi.org/10.24132/csrn.3301.36`.

Bunge, A. and M. Botsch (2023). “A Survey on Discrete Laplacians for General Polygonal Meshes”. In: *Computer Graphics Forum* 42.2, pp. 521–544.
URL: `http://doi.org/10.1111/cgf.14777`.

Deng, Xin, WenYu Zhang, Qing Ding, and XinMing Zhang (June 2023). “PointVector: A Vector Representation In Point Cloud Analysis”. In: *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 9455–9465.
URL: `http://doi.org/10.1109/cvpr52729.2023.00912`.

Hartwig, Florine, Josua Sassen, Omri Azencot, Martin Rumpf, and Mirela Ben-Chen (2023). “An Elastic Basis for Spectral Shape Correspondence”. In: *ACM SIGGRAPH 2023 Conference Proceedings*. SIGGRAPH ’23. Los Angeles, CA, USA: Association for Computing Machinery, pp. 1–11.
URL: `https://doi.org/10.1145/3588432.3591518`.

Bindal, Manika and Venkatesh Kamat (Aug. 2024). "Diverse Bases for Functional Spaces for Non-Rigid Shape Correspondence". In: *Ingénierie des systèmes d information* 29.4, 1405–1422.
URL: `http://doi.org/10.18280/isi.290415`.