

NNT : 2016SACLS559

THÈSE DE DOCTORAT
DE
L'UNIVERSITÉ PARIS-SACLAY
PRÉPARÉE À
L'UNIVERSITÉ PARIS-SUD
AU SEIN DE
INRIA SACLAY ILE-DE-FRANCE

ÉCOLE DOCTORALE N°580
Sciences et technologies de l'information et de la communication

Spécialité de doctorat : Informatique

par

M. Ruqi HUANG

Two contributions to geometric data analysis:
filamentary structures approximations,
and stability properties of functional approaches for shape comparison.

Thèse présentée et soutenue à Palaiseau, le 14 décembre 2016
Composition du jury :

| | | |
|--------------------|--------------------|--------------------------------------|
| M. Tamy BOUBEKEUR | Président | Télécom ParisTech |
| M. Pierre ALLIEZ | Rapporteur | Inria Sophia-Antipolis |
| Mme Yusu WANG | Rapporteur | The Ohio State University |
| M. Xavier GOAOC | Examineur | Université Paris-Est Marne-la-Vallée |
| Mme Pooran MEMARI | Examinatrice | CNRS |
| M. Maks OVSJANIKOV | Examineur | École Polytechnique |
| M. Boris THIBERT | Examineur | Laboratoire J. Kuntzmann, Grenoble |
| M. Frédéric CHAZAL | Directeur de thèse | Inria Saclay Île-de-France |



Abstract

Massive amounts of data are being generated, collected and processed all the time. A considerable portion of them are sampled from objects with geometric structures. Such objects can be tangible and ubiquitous in our daily life. Inferring the geometric information from the data, however, is not always an obvious task. Moreover, it's not a rare case that the underlying objects are abstract and of high dimension, where the data inference is more challenging.

This thesis studies two problems on geometric data analysis. The first one concerns metric reconstruction for filamentary structures. We in general consider a filamentary structure as a metric space being close to an underlying metric graph, which is not necessarily embedded in some Euclidean spaces. Particularly, by combining the Reeb graph and the mapper algorithm, we propose a variant of the Reeb graph, which not only faithfully recovers the metric of the filamentary structure but also allows for efficient implementation and convenient visualization of the result.

Then we focus on the problem of shape comparison. In this part, we study the stability properties of some recent and promising approaches for shape comparison, which are based on the notion of functional maps. Our results show that these approaches are stable in theory and potential for being used in more general setting such as comparing high-dimensional Riemannian manifolds.

Lastly, we propose a pipeline for implementing the functional-maps-based frameworks under our stability analysis directly on point cloud data. Though our pipeline is experimental, it undoubtedly extends the range of applications of these frameworks.

Acknowledgment

First of all, I would like to thank my adviser Frédéric Chazal for offering the opportunity to conduct the research presented in this thesis. Without his patient guidance and insightful suggestions along these years, this thesis would not have been possible. I would also like to extend my appreciation to Jian Sun and Maks Ovsjanikov for the fruitful collaborations, which respectively turn into the two major parts of the thesis.

I would like to thank Pierre Alliez and Yusu Wang, who took the time to carefully review my thesis. And I would also like to thank Tamy Boubekeur, Xavier Goaoc, Pooran Memari, Maks Ovsjanikov and Boris Thibert, who accepted to be the jury of this thesis.

I would like to express my gratitude to my colleagues in the Datashape team for the delightful working atmosphere. Especially I'd thank our team assistant Christine Biard. Her kind help makes my life as a non-French speaker much more easier than it could have been.

Last but not least, I owe my deepest gratitude to my parents for their consistent encouragement and support over the past 3 years.

Contents

| | | |
|----------|------------------------------------------------------------------------|-----------|
| 1 | Résumé | 1 |
| 2 | Introduction | 3 |
| 2.1 | Contributions | 7 |
| 3 | Approximation for Filamentary Structures Using Reeb-Type Graphs | 9 |
| 3.1 | Introduction | 9 |
| 3.1.1 | Overview | 9 |
| 3.2 | Related Works | 11 |
| 3.3 | Preliminaries | 12 |
| 3.3.1 | Metric Spaces | 12 |
| 3.3.2 | Topology | 14 |
| 3.4 | Reeb-type Graph | 15 |
| 3.4.1 | Metrics on Reeb and α -Reeb Graphs. | 16 |
| 3.5 | Recovery of Geometry | 17 |
| 3.5.1 | Bounding the Diameter M | 20 |
| 3.6 | Recovery of Topology | 23 |
| 3.6.1 | Construction of Open Cover for X | 25 |
| 3.6.2 | Construction of Open Cover for G' | 31 |
| 3.7 | Metric Reconstruction from Discrete Sampling | 32 |
| 3.7.1 | Regularity Conditions. | 33 |
| 3.7.2 | Upper-bound of $ d_G - d_H $ | 35 |
| 3.7.3 | Upper-bound of $ d_H - d_{\tilde{X},\sigma} $ | 39 |
| 3.8 | Algorithm | 43 |
| 3.9 | Experimental Results | 44 |
| 3.9.1 | Earthquake Data | 44 |
| 3.9.2 | GPS Data | 44 |
| 3.9.3 | GPS Data with Crossing | 46 |
| 4 | On Stability of Shape Difference Operators | 49 |
| 4.1 | Introduction | 49 |
| 4.1.1 | Overview | 50 |
| 4.2 | Related Work | 52 |
| 4.3 | Preliminaries | 52 |
| 4.3.1 | Differential Geometry | 53 |
| 4.3.2 | The Laplace-Beltrami Operator | 56 |
| 4.3.3 | Functional Maps | 58 |
| 4.3.4 | Map Analysis | 59 |

| | | |
|----------|----------------------------------------------------------------------------------|------------|
| 4.3.5 | Shape Difference Operators | 59 |
| 4.4 | Perturbation Model and Bounded-distortion Condition | 60 |
| 4.4.1 | Perturbation Model | 60 |
| 4.4.2 | Bounded-distortion Condition | 62 |
| 4.5 | Stability of the Shape Difference Operators | 63 |
| 4.6 | Stability of the Shape Difference Operators in a Multi-Scale Framework | 66 |
| 4.6.1 | An Alternative Collection of Multi-scale Subdomains | 66 |
| 4.6.2 | Stability with respect to the Changes in Scale | 67 |
| 4.6.3 | Stability With Respect To Perturbed Input Manifolds | 70 |
| 4.6.4 | Approximating $\ V\ _C$ | 72 |
| 4.6.5 | Analysis for the Conformal Shape Difference Operator | 73 |
| 4.7 | Experimental Results | 75 |
| 4.7.1 | Approximating $\ V\ _C$ | 76 |
| 4.7.2 | Capturing Conformal Differences | 77 |
| 4.7.3 | Stability of the Area-based Shape Difference Operators | 78 |
| 4.8 | Proofs for Theorems in Section 4.6.5 | 79 |
| 5 | functional-maps-based Frameworks on Point Cloud Data | 85 |
| 5.1 | Introduction | 85 |
| 5.2 | Pipeline for PCD | 86 |
| 5.2.1 | Rationale of Our Pipeline | 89 |
| 5.3 | Experimental Results | 90 |
| 5.3.1 | Functional Maps on PCD | 90 |
| 5.3.2 | Selection of k for k -NN | 92 |
| 5.3.3 | Reliability of PCD Setting in Multi-scale Framework | 93 |
| 5.3.4 | Analyzing Shape Collections | 94 |
| 5.3.5 | Non-uniformly Sampled Data | 95 |
| 5.4 | Limitation and Perspective | 97 |
| 5.4.1 | Limitations | 98 |
| 5.4.2 | Beyond Shapes | 98 |
| 6 | Conclusion and Perspectives | 101 |
| | Bibliography | 103 |

Résumé

Des quantités importantes de données sont générées, collectées et analysées en permanence. Dans de nombreux cas, ces données sont échantillonnées sur des objets possédant une structure géométrique spécifique. Certains de ces objets sont omniprésents dans notre vie quotidienne. L'inférence de leur structure géométrique à partir de ces données est une tâche souvent ardue. Cette tâche est d'ailleurs rendue plus difficile si les objets sous-jacents sont abstraits ou de grande dimension.

Dans cette thèse, nous nous intéressons à deux problèmes liés à l'analyse géométrique de données. Dans un premier temps, nous nous penchons sur l'inférence de métriques pour des structures filamenteuses. Les structures filamenteuses apparaissent naturellement dans les données du quotidien telles que les collections de traces GPS collectées par des véhicules sillonnant le réseau routier, les distributions d'épicentres de tremblements de terre se concentrant le long de failles géologiques ou encore les réseaux de vaisseaux sanguins. Cette liste n'étant pas exhaustive. Dans de nombreuses applications pratiques, les données apparaissent sous la forme d'une matrice de distances entre points, ce qui permet de mettre sur les données une métrique discrète. Par conséquent, il est intéressant de construire une inférence des données à partir de distances.

En général, nous représentons des structures filamenteuses par des espaces métriques que nous supposons être proches d'un graphe métrique. La méthode que nous proposons est une variante des graphes de Reeb, elle combine les graphes de Reeb et l'algorithme Mapper. Cette variante permet non seulement d'approcher de façon judicieuse la métrique associée aux structures filamenteuses, mais elle permet également d'implémenter et de visualiser le résultat aisément. En particulier, nous fournissons des garanties théoriques fortes pour cet algorithme.

Nous nous concentrons ensuite sur le problème de la comparaison de formes. Nous étudions un ensemble de méthodes récentes et prometteuses pour la comparaison de formes qui utilisent la notion d'applications fonctionnelles. Depuis les travaux précurseurs sur les applications fonctionnelles, de nombreux travaux ont suivi traitant de l'exploration de formes, de la segmentation d'images, de traitement de champs de vecteurs etc. Cependant, malgré le succès de telles méthodes, beaucoup moins d'attention a été portée à l'étude des propriétés théoriques fondamentales de ces approches fonctionnelles. Nos résultats théoriques montrent que ces approches sont stables et peuvent être utilisées dans un contexte plus général que la comparaison de formes comme par exemple la comparaison de variétés Riemanniennes de grande dimension.

Enfin, en nous reposant sur notre analyse théorique, nous proposons une généralisation des applications fonctionnelles aux données de type nuages de points. Dans le domaine de l'analyse de forme, la représentation des formes 3D par des nuages de points est sans doute la plus fréquente et primitive. Cependant, la plupart des approches fonctionnelles en traitement de données géométriques utilisent comme données des formes discrétisées par des maillages triangulaires. Nos travaux produisent des résultats raisonnables par rapport à ceux obtenus avec les maillages, malgré le fait que les nuages de points contiennent une information géométrique intrinsèque beaucoup plus pauvre que celle fournie par des maillages.

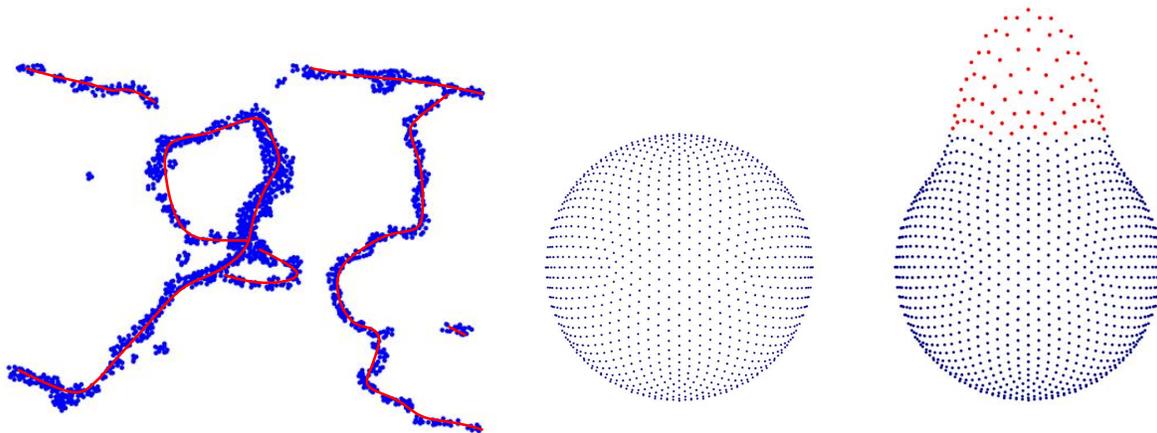
Bien que cette généralisation ne bénéficie pas des garanties théoriques, elle permet d'étendre le champ d'application des méthodes fondées sur les applications fonctionnelles. En effet, du point de vue plus général de l'analyse de données, les résultats expérimentaux que nous obtenons en étudiant les formes 3D laissent à penser que l'approche fonctionnelle appliquée à l'analyse géométrique de données plus abstraites et/ou de grande dimension a du potentiel.

Globalement, dans cette thèse, nous présentons une analyse théorique dans deux contextes différents, cette analyse consiste à inférer des informations à partir de données géométriques. Nous avons proposé des garanties théoriques variées, mais avons également apporté des contributions dans le domaine des applications : un nouvel algorithme de reconstruction de structure filamentaire et quelques extensions des approches fonctionnelles existantes sont proposées dans les première et seconde parties.

Introduction

Data play a role that is becoming increasingly important nowadays. Massive amounts of data are being generated, gathered and processed in various areas such as scientific research, economic activities, manufacturing productions all the time.

Perhaps data come with geometric structures are the most familiar to us. As a matter of fact, our perception system takes in and analyzes geometric data everyday. For example, looking at the scattering blue points in Figure 2.1(a), probably people would outline the skeleton of the point cloud and draw the red curves *in mind*. If the perception system is fed multiple objects like in Figure 2.1(b), then a more complicated task is accomplished—people would easily discriminate the two shapes and highlight the red bump of the second one as a prominent difference *in mind*.



(a) Outlining a shape

(b) Differentiating two shapes

Figure 2.1: (a) scattering points with a filamentary structure, which is approximated by the red curves. (b) a sphere and a deformed one, the deformation in-between can be visually captured.

However, these *geometric data inferences* are limited in several aspects. First of all, our perception system outputs only qualitative results. When differentiating shapes, we can tell where the difference is, but can't describe quantitatively how large it is. Besides, though we are able to comfortably deal with low-dimensional objects such as those depicted in Figure 2.1, we have trouble even visualizing objects of dimension more than 3, not to mention outlining or comparing. Moreover, our perception system can not

process abstract inputs directly. If the plots in Figure 2.1 are replaced with tables listing the coordinates of the points, we can hardly see anything meaningful there.

Fortunately, with the aid of computers, we are able to quantitatively process complex and abstract data. However, computers are not gifted with a perception system. It is then critical to develop methods that guide the computers to efficiently infer information from data. Approaches have been taken for data inference from various perspectives. For example, dimension reduction algorithms [Belkin 2003, Lafon 2004, Tenenbaum 2000b] attempt to project high-dimensional data into low-dimension spaces for visualization; clustering algorithms [Planck 2006, Lloyd 1982, Ester 1996, Comaniciu 2002] separate data into groups of a relatively small number; and more recent frameworks of topological data analysis [Edelsbrunner 2002, Zomorodian 2005, Carlsson 2009, Chazal 2012] enable extracting data that are not embedded in an Euclidean space.

In this thesis, we study two problems on geometric data inference respectively related to the perception tasks illustrated in Figure 2.1.

The first problem is metric reconstruction for filamentary structures. An intuitive example of data sampled from a filamentary structure is the blue point cloud in Figure 2.1(a), which *looks like* some graph, say, the red curves. We assume that the given filamentary structure X is a metric space (not necessarily embedded in some Euclidean space as the example) that is close with respect to the so-called Gromov-Hausdorff distance to an unknown metric graph¹. And our goal is to construct a metric graph that is close to X in the Gromov-Hausdorff distance. We propose a new method to address this question in a rigorous mathematical framework. This method relies on the Reeb graph and so-called Mapper algorithm, which are soon to be introduced.

As a result of our investigation to the first problem, we develop a method to construct a metric graph G that is guaranteed to be close to X . In general, The metric of X is impossible to be fully reconstructed by G . How to evaluate and visualize the parts of X which are less well approximated by the resulting graph is then a natural follow-up problem. Though there is a natural map between X and G constructed by our method, it is still a challenging problem. For instance, naive point-to-point comparison is unstable due to noisy data in practice. Motivated by this, the second problem is about comparing two geometrical objects associated by a map.

This is indeed a vast and difficult problem. Our contribution to it is different than that to the first one: instead of proposing a new model/method, we consider some recent and promising methods on comparing 3D shapes that are potentially applicable in the general setting, say, comparing Riemannian manifolds of higher dimensions. With the understandings obtained from analyzing the chosen methods, we take some experimental approach based on them to the general problem. Especially, our main effort is devoted to analyzing the *stability properties* of the chosen methods. More precisely, the goal is to understand how the outputs of the method change with respect to perturbations on the inputs and/or the parameters. Investigating the stability of a method not only verifies its robustness in theory, but also provides valuable insights for implementations.

In the following we give a brief overview of the methods involved in this thesis.

¹A metric graph is a graph endowed with a metric.

Reeb Graphs

Let f be a real-valued continuous function defined on a topological space X . We then define a relation on X with respect to f : for $x, y \in X$, we let $x \sim_f y$ if and only if $f(x) = f(y)$ and x, y are in the same connected component of the pre-image $f^{-1}(f(x))$. The *Reeb Graph* (originally proposed in [Reeb 1946]) of f is then defined as the quotient space $X \setminus \sim_f$. For the point of view of level-set of f , for any $a \in \mathbb{R}$, we decompose the subset $f^{-1}(a) = \{x \in X : f(x) = a\}$ into one or more connected components, and then collapse each component into a point in the quotient space (see Figure 2.2 for an illustration).

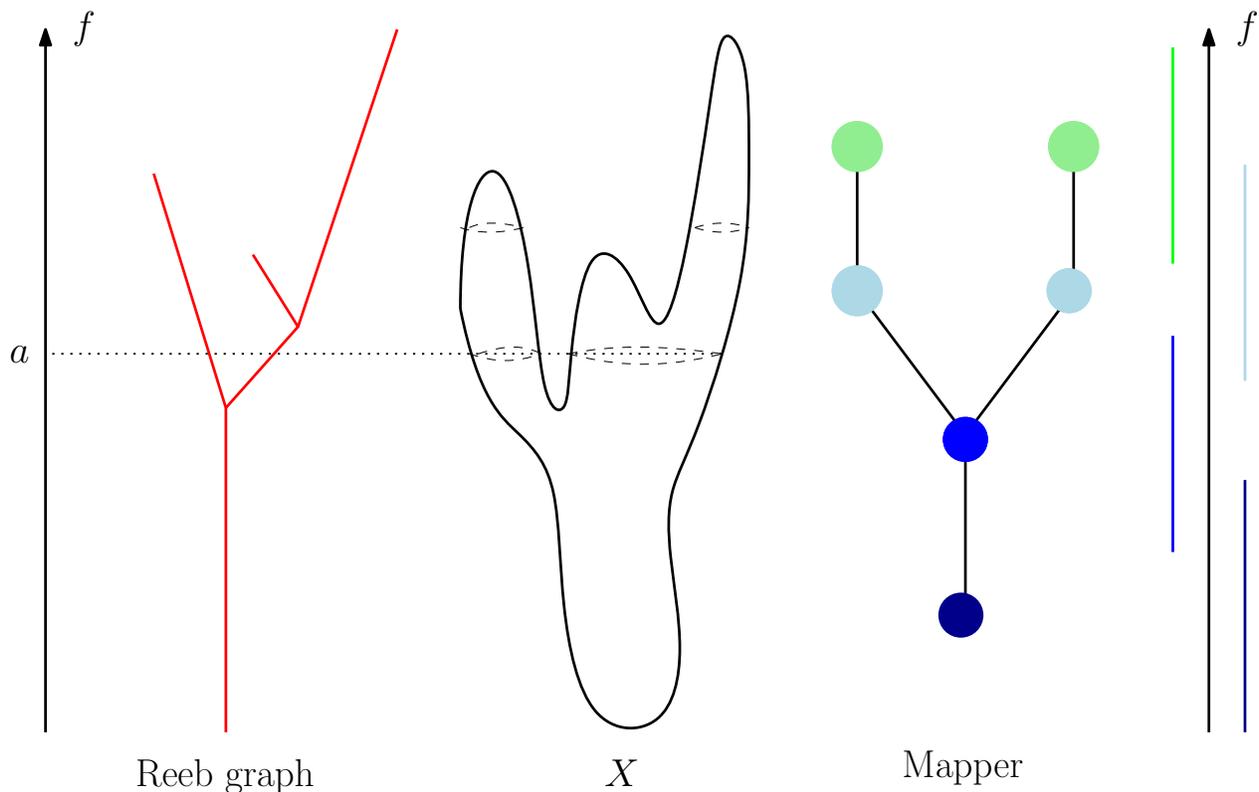


Figure 2.2: Illustrations of the Reeb graph (on the left of X) and the output of the mapper algorithm (on the right of X) with respect to the same function f . The cover used in the mapper algorithm is the 4 intervals colored differently to the right-most.

Under some regularity conditions, the Reeb graph is a 1-dimensional structure, making interpreting results convenient. Besides, the Reeb graph is defined upon a function on X instead of upon X itself. This property allows the users to view X through the lens of different functions. In the recent decades, its capability of encoding both the geometrical and the topological features of shapes is appreciated in the area of computer graphics, say, shape comparison [Hilaga 2001, Escolano 2013], skeleton extraction [Ge 2011] (see [Biasotti 2008] for a survey).

Mapper Algorithm

Many works studying variations on the Reeb graph have been proposed, such as the extend Reeb graphs [Biasotti 2000, Escolano 2013], the contour trees [Pascucci 2003, Van Kreveld 2006, Carr 2003] and so on. A generalization of the Reeb graph—the mapper algorithm [Singh 2007] also falls into this category. The mapper algorithm takes in a continuous function f on X as an input. Instead of directly collapsing the connected components of the level-set of each single point $a \in f(X)$, the algorithm considers a cover of $f(X)$, the image of f . By a cover we mean a collection of open sets $\{U_i\}_{i \in \Gamma}$ such that $f(X) \subset \bigcup_{i \in \Gamma} U_i$. For

any U_i , let $f^{-1}(U_i) = \bigcup_{\alpha \in \Lambda_i} V_i^\alpha$, where V_i^α 's are the connected components of $f^{-1}(U_i)$. In the output of the mapper algorithm, each set V_i^α is represented as a node. An edge connects two nodes if the intersection of the corresponding sets is non-empty. A triangle is added among three nodes if the three corresponding sets are intersecting, likewise a tetrahedron is created in the case of four sets intersecting. In general, a k -simplex is added to the output if there exist $k - 1$ sets have a common non-empty subset. Thus the output of the mapper algorithm is not necessarily a graph. In general it is a so-called abstract simplicial complex (see the definition in Section 3.3.2).

The mapper algorithm is more of a qualitative analysis tool for high-dimensional data set. The main aim of this algorithm is to robustly capture and easily visualize the topological structure of the input data.

Functional Maps

In [Ovsjanikov 2012], the authors propose a functional representation of a map between two shapes. Given two shapes M and N and a map $T : M \rightarrow N$, a functional map, T_F , is a pull-back induced by T . Namely, given a real-valued function $w \in \mathcal{F}(N)$, we define $T_F(w) = w \circ T \in \mathcal{F}(M)$. Therefore T_F is a map from $\mathcal{F}(N)$ to $\mathcal{F}(M)$. The first benefit of this representation is that T_F is a linear operator between the two function spaces. In fact, by definition of T_F , $T_F(\alpha f + \beta g) = (\alpha f + \beta g) \circ T = \alpha f \circ T + \beta g \circ T = \alpha T_F(f) + \beta T_F(g)$, where $\alpha, \beta \in \mathbb{R}$ and $f, g \in \mathcal{F}(N)$.

This linearity property indicates that T_F admits a (potentially infinite) matrix representation. Suppose that $\{\varphi_i^N\}$, $\{\psi_j^M\}$ form a basis of $\mathcal{F}(N)$ and of $\mathcal{F}(M)$ respectively. Let $w = \sum_{i=1} a_i \varphi_i^N$ and accordingly

$T_F(w) = \sum_{j=1} b_j \psi_j^M$, then there exists a unique matrix C_T such that $C_T \mathbf{a} = \mathbf{b}$ holds for any w , where $\mathbf{a} = (a_1, a_2, \dots, a_n, \dots)$ and $\mathbf{b} = (b_1, b_2, \dots, b_n, \dots)$. We then call C_T the matrix representation of T_F with respect to basis $\{\varphi_i^N\}$ and $\{\psi_j^M\}$.

Especially, the authors take the first k_M and k_N eigenfunctions of the Laplace-Beltrami operators on M and N respectively as approximations to the full bases of the function spaces, and obtain a compact matrix representation of C_T which is of dimension $k_M \times k_N$.

Geometrically, truncating the eigenbasis in this way amounts to putting a low-frequency filter on the function space. In discrete setting, M and N are usually represented by m and n discrete points. Thus the function spaces are respectively vector spaces of dimension m and n , and T_F is a $m \times n$ matrix. Using above scheme, T_F is then approximated by a $k_M \times k_N$ matrix C_T written in the truncated bases. It has been validated empirically that for a pair of shapes with $m, n \sim 10^4$, $k_M, k_N \sim 10^2$ are sufficient to obtain a

reasonable approximation.

Shape Difference Operators

Based on the functional map, in [Rustamov 2013] the authors propose a framework to encode the differences between shape with so-called shape difference operators. Let $\langle \cdot, \cdot \rangle_{\mathcal{F}(M)}$ and $\langle \cdot, \cdot \rangle_{\mathcal{F}(N)}$ be two inner products on the function spaces $\mathcal{F}(M)$ and $\mathcal{F}(N)$ respectively. In general, the equality $\langle f, g \rangle_{\mathcal{F}(N)} = \langle T_F(f), T_F(g) \rangle_{\mathcal{F}(M)}$ does not hold for all $f, g \in \mathcal{F}(N)$. The key observation of [Rustamov 2013] is that by the Riesz representation theorem and the linearity of T_F , there always exist a self-adjoint linear operator on the function space $\mathcal{F}(N)$, S such that $\langle f, S(g) \rangle_{\mathcal{F}(N)} = \langle T_F(f), T_F(g) \rangle_{\mathcal{F}(M)}$.

This operator compensating the differences between the two inner products in some sense also captures the differences between the shapes induced by T_F (and equivalently T), and is called the shape difference operator. Particularly, in [Rustamov 2013], for each of the shapes M and N , two typical function spaces and the associated inner products are discussed, resulting in two types of shape difference operators that capture differences between M and N from separated aspects.

On the other hand, it is also interesting to study the *gap* between $\langle f, g \rangle_{\mathcal{F}(N)}$ and $\langle T_F(f), T_F(g) \rangle_{\mathcal{F}(M)}$. The framework of [Ovsjanikov 2013] constructs a functional to measure the gap at any function f : $F(f) = \frac{\langle T_F(f), T_F(f) \rangle_{\mathcal{F}(M)}}{\langle f, f \rangle_{\mathcal{F}(N)}}$ for a specified inner product, which is one of the two considered in [Rustamov 2013]. By maximizing the functional, the authors identify a function at which the gap is the largest. Moreover, a multi-scale framework is presented there by putting an extra constraint on the maximization. This constraint is related to the truncated function basis mentioned before. Particularly, the authors force f to be spanned by the first k eigenfunctions of the Laplace-Beltrami operator on N and generate a collection of maximizers corresponding to different k 's.

2.1 Contributions

This thesis consists of two major parts: the first part (chapter 2) investigates the problem of metric reconstruction for filamentary structures and the second part (chapter 3 and chapter 4) studies the stability properties of the shape difference operators and presents an empirical method to perform the functional-map-based frameworks directly on point cloud data.

Chapter 2: Particularly, given a geodesic space² (X, d_X) , we study the Reeb graph of the distance function to a point. A distance function to $x \in X$ is defined as $d(y) = d_X(x, y)$, which is the distance in X from x to y . The same Reeb graph has been studied in [Ge 2011] as a tool for capturing a 1-dimension skeleton of some general data. We further define a metric on the Reeb graph of the distance function so that it's as well a metric graph. We prove that this metric graph is well-approximating X in the sense of the Gromov-Hausdorff distance.

From the practical point of view, data usually come as discrete points sampled from the underlying metric space X . We incorporate the idea of the mapper algorithm to construct a variant of the Reeb graph called α -Reeb graph, with which it is more convenient to deal with the discrete sampling points in practice. And

²see Section 3.3.1 for the definition

theoretically, the α -Reeb graph not only enjoys similar guarantees on the metric approximation as the Reeb graph, but also comes with certain topological guarantees.

Overall, we conclude that the α -Reeb graph is a reliable approximation of X in both geometrical and topological sense. The above are results of a collaboration with F. Chazal and J. Sun, which have been published in *Discrete and Computational Geometry* [Chazal 2015].

Our framework assumes that the input X is already a (discrete) metric space. In many cases, the raw data are discrete isolated points sampled from some underlying metric space X . It's then appealing to study how accurately we can recover the metric of X with the sampling points. In this chapter we study in a special case where X is a metric graph embedded in an Euclidean space. We assure that under some technique conditions on X and the sampling density, the metric we recover from the sampling points is well-approximating the ground truth.

Chapter 3: As mentioned above, the second problem considered in this thesis is about comparing geometrical objects associated by a map. We notice that similar approaches have been made in the area of computer graphics. Particularly, the map-based shape analysis frameworks based on the functional map seem promising approaches to this problem. Though the original formulations are proposed for shapes, i.e., 2-dimensional Riemannian manifolds, they can naturally be defined for Riemannian manifolds of arbitrary dimension.

In this chapter, we concentrate on stability analysis of two frameworks—the shape difference operators and the map analysis and visualization from [Ovsjanikov 2013]. More generally, we assume that the inputs are two n -dimensional Riemannian manifolds M, N and a map $T : M \rightarrow N$. We verify two type of stability properties: one is with respect to perturbations on the input manifolds and the other is with respect to the changing scale, which is peculiar to the latter framework.

The results of this chapter are obtained in collaboration with F. Chazal and M. Ovsjanikov. The manuscript is soon to be submitted.

Chapter 4: The functional-map-based frameworks are usually proposed to analyze 3D shapes, i.e., 2-dimensional Riemannian manifold embedded in \mathbb{R}^3 . The 3D shapes can be nicely approximated and represented by polygon meshes. Implementing these frameworks on meshes are obvious and efficient. In a more general setting, where the input are Riemannian manifolds of dimension higher than 2 or embedded in a Euclidean of dimension higher than 3, the implementation is not straightforward any more.

In this chapter, we propose a pipeline for dealing with data in the most primitive form—point cloud data (PCD). The idea is to construct counterparts in the PCD setting of the ingredients necessary in the mesh setting. To test our method, we perform the functional-map-based frameworks with our pipeline on point clouds sampled from 3D shapes. The results are compared to the ones obtained from the mesh setting. Empirically, the pipeline works well with the test data and shows robustness with respect to noisy data.

At the end, we emphasize that the work in this chapter is exploratory and experimental. The experimental results suggest that this direction is worth further exploration and requires corresponding theoretical analyses.

The work presented in this chapter is obtained in collaboration with F. Chazal and M. Ovsjanikov. The manuscript is in preparation.

Approximation for Filamentary Structures Using Reeb-Type Graphs

3.1 Introduction

With the advance of sensor technology, computing power and the Internet, massive amounts of geometric data are being generated and collected in various areas of science, engineering and business. As they are becoming widely available, there is a real need to analyze and visualize these large scale geometric data to extract useful information out of them. In many cases these data are not embedded in Euclidean spaces and come as (finite) sets of points with pairwise distance information, i.e. (discrete) metric spaces. A large amount of research has been done on dimensionality reduction, manifold learning and geometric inference for data embedded in (possibly high dimensional) Euclidean spaces and assumed to be concentrated around low dimensional manifolds [Belkin 2003, Lafon 2004, Tenenbaum 2000b]. However, the assumption of data lying on a manifold may fail in many applications. In addition, the strategy of representing data by points in an Euclidean space may introduce large metric distortions as the data may lie in highly curved spaces, instead of in flat Euclidean space raising many difficulties in the analysis of metric data. In the past decade, with the development of topological methods in data analysis, new theories such as topological persistence (see, for example, [Edelsbrunner 2002, Zomorodian 2005, Carlsson 2009, Chazal 2012]) and new tools such as the Mapper algorithm [Singh 2007] have given rise to new algorithms to extract and visualize geometric and topological information from metric data without the need of an embedding into an Euclidean space. In this chapter we focus on a simple but important setting where the underlying geometric structure approximating the data can be seen as a branching filamentary structure i.e., more precisely, as a *metric graph* which is a topological graph endowed with a length assigned to each edge. Such structures appear naturally in various real-world data such as collections of GPS traces collected by vehicles on a road network, earthquakes distributions that concentrate around geological faults, distributions of galaxies in the universe, networks of blood vessels in anatomy or hydrographic networks in geography just to name a few. It is thus appealing to try to capture such filamentary structures and to approximate the data by metric graphs that will summarize the metric and allow convenient visualization.

3.1.1 Overview

In this chapter we address the metric reconstruction problem for filamentary structures. The input of our method and algorithm is a metric space (X, d_X) that is assumed to be close with respect to the so-called Gromov-Hausdorff distance d_{GH} to a much simpler, but unknown, metric graph $(G', d_{G'})$. Our algorithm outputs a metric graph (G, d_G) that is proven to be close to $(G', d_{G'})$ in both geometry and topology. Our

approach relies on the notion of Reeb graph (and some variants of it introduced in Section 3.4) and our main theoretical results are stated in the following two theorems.

Theorem 3.5 [Recovery of Geometry]. *Let (X, d_X) be a compact connected geodesic space, let $r \in X$ be a fixed base point such that the metric Reeb graph (G, d_G) of the function $d = d_X(r, \cdot) : X \rightarrow \mathbb{R}$ is a finite graph. If for a given $\varepsilon > 0$ there exists a finite metric graph $(G', d_{G'})$ such that $d_{GH}(X, G') < \varepsilon$ then we have*

$$d_{GH}(X, G) < 2(\beta_1(G) + 1)(17 + 8N_{E, G'}(8\varepsilon))\varepsilon$$

where $N_{E, G'}(8\varepsilon)$ is the number of edges of G' of length at most 8ε and $\beta_1(G)$ is the first Betti number of G , i.e. the number of edges to remove from G to get a spanning tree. In particular if X is at distance less than ε from a metric graph with shortest edge larger than 8ε then $d_{GH}(X, G) < 34(\beta_1(G) + 1)\varepsilon$.

Note that $\beta_1(G) \leq \beta_1(X)$ and thus $d_{GH}(X, G)$ is upper bounded by the quantities depending only on the input X .

Theorem 3.7 [Recovery of Topology]. *Let (X, d_X) be a compact connected path metric space and $(G', d_{G'})$ is a metric graph so that $d_{GH}(X, G') < \varepsilon$. Let $r \in X$, $\alpha > 60\varepsilon$ and $\mathcal{I} = \{[0, 2\alpha), (i\alpha, (i + 2)\alpha) | 1 \leq i \leq m\}$ covers the segment $[0, \text{Diam}(X)]$ such that the 2α -Reeb graph G associated to \mathcal{I} and the function $d = d_X(r, \cdot) : X \rightarrow \mathbb{R}$ is a finite graph. If no edges of G' are shorter than L and no loops of G' are shorter than $2L$ with $L \geq 32\alpha + 9\varepsilon$, then we have G and G' are homotopy equivalent.*

Approximating the Reeb graph (G, d_G) from a neighborhood graph is usually not obvious. If we compute the Reeb graph of the distance function to a given point defined on the neighborhood graph we obtain the neighborhood graph itself and do not achieve our goal of representing the input data by a simple graph. See Table 3.1. It is then appealing to build a two dimensional complex having the neighborhood graph as 1-dimensional skeleton and use the algorithm of [Harvey 2010, Parsa 2013] to compute the Reeb graph of the distance to the root point. Unfortunately adding triangles to the neighborhood graph may widely change the metric between the data points on the resulting complex and significantly increase the complexity of the algorithm. We overcome this issue by introducing a variant of the Reeb graph, the α -Reeb graph, inspired from [Singh 2007] and related to the recently introduced notion of graph induced complex [Dey 2013a], that is easier to compute than the Reeb graph but also comes with approximation guarantees (see Theorem 3.6). As a consequence our algorithm runs in almost linear time (see Section 3.8).

Raw data usually do not come as geodesic spaces. They are given as discrete sets of points (and thus not connected metric spaces) sampled from the underlying space (X, d_X) . Moreover in many cases only distances between nearby points are known. A geodesic space (see Section 3.3.1 for a definition of geodesic space) can then be obtained from these raw data as a neighborhood graph where nearby points are connected by edges whose length is equal to their pairwise distance. The shortest path distance in this graph is then used as the metric. In our experiments we use this new metric as the input of our algorithm. Particularly, in Section 3.7 we study a special case that the underlying metric space (X, d_X) is a metric graph embedded in a d -dimensional Euclidean space \mathbb{R}^d and data are sampled from (X, d_X) . Our result shows that the metric of a certain neighboring graph built on the sampling points well approximates d_X under proper conditions on X and the sampling density.

3.2 Related Works

Approximation of data by 1-dimensional geometric structures has been considered by different communities. In statistics, several approaches have been proposed to address the problem of detection and extraction of filamentary structures in point cloud data. For example Arial-Castro et al [Arias-Castro 2006] use multi-scale anisotropic strips to detect linear structure while [Genovese 2009, Genovese 2012] and more recently [Genovese 2014] base their approach upon density gradient descents or medial axis techniques. These methods apply to data corrupted by outliers embedded in Euclidean spaces and focus on the inference of individual filaments without focus on the global geometric structure of the filaments network.

In computational geometry, the curve reconstruction problem from points sampled on a curve in an Euclidean space has been extensively studied and several efficient algorithms have been proposed [Amenta 1998, Dey 2000, Dey 2001]. Unfortunately, these methods restrict to the case of simple embedded curves (without singularities or self-intersections) and hardly extend to the case of topological graphs. In a more intrinsic setting where data come as finite abstract metric spaces, [Aanjaneya 2012] propose an algorithm that outputs a topologically correct (up to a homeomorphism) reconstruction of the approximated graph. However this algorithm requires some tedious parameters tuning and relies on quite restrictive sampling assumptions. When these conditions are not satisfied, the algorithm may fail and not even outputs a graph. Compared to the algorithm of [Aanjaneya 2012], our algorithm not only comes with metric guarantees but also whatever the input data is, it always outputs a metric graph and does not require the user to choose any parameters. Closely related to our approach is the data skeletonization algorithm proposed in [Ge 2011] that computes the Reeb graph of an approximation of the distance function to a root point on a 2-dimensional complex built on top of the data whose size might be significantly larger than a neighboring graph. The algorithm of [Ge 2011] also always outputs a graph but it does not come with metric guarantees. Recently, Bauer, Ge and Wang [Bauer 2014] define a metric based on the function for Reeb graph and show it is stable under Gromov-Hausdorff distance. The implementation of our algorithm relies on the Mapper algorithm [Singh 2007], that provides a way to visualize data sets endowed with a real valued function as a graph, where the considered function is the distance to the chosen root point. However, unlike the general Mapper algorithm, our methods provides an upper bound on the Gromov-Hausdorff distance between the reconstructed graph and the underlying space from which the data points have been sampled.

In theoretical computer science, there is much of work on approximating metric spaces using trees [Bădoiu 2007, Abraham 2007, Chepoi 2008] or distribution of trees [Dhamdhere 2006, Fakcharoenphol 2004] where the trees are often constructed as spanning trees possibly with Steiner points. Our approach is different as our reconstructed graph or tree is a quotient space of the original metric space where the metric only gets contracted (see Proposition 3.2). Finally we remark that the recovery of filament structure is also studied in various applied settings, including road networks [Chen 2010, Tupin 1998], galaxies distributions [Choi 2010].

3.3 Preliminaries

3.3.1 Metric Spaces

A metric space is a pair (X, d_X) where X is a set and $d_X : X \times X \rightarrow \mathbb{R}$ is a non-negative map such that for any $x, y, z \in X$, $d_X(x, y) = 0$ if and only if $x = y$, $d_X(x, y) = d_X(y, x)$ and $d_X(x, z) \leq d_X(x, y) + d_X(y, z)$.

A continuous path (or curve) σ , is a continuous map from I , a real interval, to X . A path is called *simple* if it is not self-intersecting, or equivalently σ is an injective map. The length of a path can be induced by the metric space (X, d_X)

Definition 3.1 Let $I = [a, b]$ be an interval. A partition of I is a finite sequence of point in I such that $a = t_0 \leq t_1 \leq t_2 \leq \dots \leq t_k = b$. Let Γ be the set of all partitions of I . Then the length of σ , $L(\sigma)$ is defined as

$$L(\sigma) = \sup_{\{t_0, t_1, \dots, t_k\} \subset \Gamma} \sum_{i=1}^k d_X(\sigma(t_{i-1}), \sigma(t_i))$$

We call (X, d_X) a path metric space if the distance between any pair of points is equal to the infimum of the length of the continuous curves joining them. Hereinafter we only consider *compact path metric spaces*, which are *geodesic* according to the following theorem (see also [Gromov 2002]).

Theorem 3.1 [Hopf-Rinow Theorem] If (X, d_X) is a complete, locally compact path metric space, then:

- Closed balls are compact, or, equivalently, each bounded, closed domain is compact.
- Each pair of points can be joined by a minimizing geodesic.

A non-geodesic metric space is illustrated in Figure 3.1.

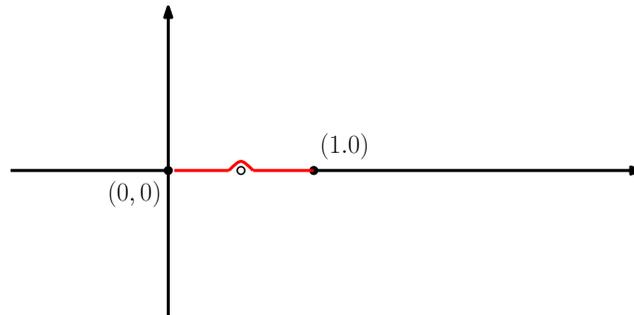


Figure 3.1: A metric space that is *not* geodesic: let $X = \mathbb{R}^2 \setminus (\frac{1}{2}, 0)$, and d_X be the euclidean metric. The distance between point $(0, 0)$ and $(1, 0)$ is 1, however there does not exist a continuous curve joining them of lenth 1.

Two compact metric spaces (X, d_X) and (Y, d_Y) are *isometric* if there exists a bijection $\phi : X \rightarrow Y$ that preserves the distances, i.e., for any $x, x' \in X$, $d_Y(\phi(x), \phi(x')) = d_X(x, x')$. The set of isometry classes of

compact metric spaces can be endowed with the Gromov-Hausdorff distance that can be defined using the following notion of correspondence (see [Burago 2001]).

Definition 3.2 Let (X, d_X) and (Y, d_Y) be two compact metric spaces. Given $\varepsilon > 0$, an ε -correspondence between (X, d_X) and (Y, d_Y) is a subset $C \subset X \times Y$ such that: i) for any $x \in X$ there exists $y \in Y$ such that $(x, y) \in C$; ii) for any $y \in Y$ there exists $x \in X$ such that $(x, y) \in C$; iii) for any $(x, y), (x', y') \in C$, $|d_X(x, x') - d_Y(y, y')| \leq \varepsilon$.

Definition 3.3 The Gromov-Hausdorff distance between two compact metric spaces (X, d_X) and (Y, d_Y) is defined by

$$d_{GH}(X, Y) = \frac{1}{2} \inf \{ \varepsilon \geq 0 : \text{there exists an } \varepsilon\text{-correspondence between } X \text{ and } Y \}$$

It follows obviously from the definitions that two metric spaces are isometric if and only if the Gromov-Hausdorff distance between them is zero.

Neighboring Graphs on Point Clouds in \mathbb{R}^d In many applications, many metric spaces are isometrically embedded in \mathbb{R}^d . Sampling from an embedded metric space results in a collection of points $\tilde{X} = \{x_1, x_2, \dots, x_n\} \subset \mathbb{R}^d$. No information about the underlying metric space can be inferred from solely the isolated points, thus a first step towards processing \tilde{X} is to build connections between points. Usually, we build a neighboring graph on top of \tilde{X} by connecting points being close to each other in a certain criterion. Now that \tilde{X} is a subset of \mathbb{R}^d , the Euclidean distance is available for measuring closeness between points. Hereinafter we denote the Euclidean distance between x and y by $\|x - y\|$. Based on the Euclidean distances, two typical neighboring graph constructions are defined as the following

Definition 3.4 Given an integer k , we denote the set of the nearest k neighborhoods of x_i among $\tilde{X} \setminus x_i$ by $N(x_i, k)$. The k -nearest graph on top of \tilde{X} is then constructed by connecting x_i to x_j if and only if $x_i \in N(x_j, k)$ or $x_j \in N(x_i, k)$.

Additionally, the mutual k -nearest graph is constructed by connecting x_i to x_j if and only if $x_i \in N(x_j, k)$ and $x_j \in N(x_i, k)$.

Definition 3.5 Given a positive constant δ , the δ -neighborhood graph of \tilde{X} is constructed by connecting x_i to x_j if and only if $\|x_i - x_j\| \leq \delta$.

There certainly exists many other graph constructions for a point cloud. Starting with a graph $R = (\tilde{X}, E)$ embedded in \mathbb{R}^d , a metric on \tilde{X} is induced by the graph and we denote it by d_R , whose definition is actually a discrete version of the definition of curve length 3.1.

Definition 3.6 Given a pair of points $x, y \in \tilde{X} \subset \mathbb{R}^d$, let $P = (x_1, x_2, \dots, x_i)$ be a path in $R = (\tilde{X}, E)$ such that $x_1 = x, x_i = y$ and $(x_j, x_{j+1}) \in E, \forall 1 \leq j \leq i-1$.

A metric $d_R : \tilde{X} \times \tilde{X} \rightarrow [0, +\infty)$ is then defined as:

$$d_R(x, y) = \inf_{P \in \Gamma} \sum_{j=1}^{i-1} \|x_j - x_{j+1}\|$$

where Γ is the set of all possible paths connecting x to y in R .

It follows from the above definition and the triangle inequality that $d_R(x, y) \geq \|x - y\|$ for any $x, y \in \tilde{X}$.

Metric Graphs A graph $G = (V, E)$ is obtained by taking a finite set of vertices, V , and joining some of them by edges, which are elements of the edge set E . A metric segment of length a ($a > 0$) is a metric space isometric to an interval $[0, a] \subset \mathbb{R}$. A metric graph is then a graph whose edges are metric segments.

Equivalently, a metric graph is also seen as a graph with a length assigned to each of its edges. Such a length assignment naturally induces a metric on G , we then denote the metric by $d_G : G \times G \rightarrow \mathbb{R}$. For any two points $g_1, g_2 \in G$ (they are not necessarily elements of V), $d_G(g_1, g_2)$ is the length of the shortest path between them along edges of G .

As we will mention soon, the goal of this part of our work is to approximate filamentary metric spaces with certain graphs. Especially, we consider filamentary metric spaces that are close to metric graphs in Gromov-Hausdorff distance (see Definition 3.3).

3.3.2 Topology

After introducing the geometrical notions, we give a brief review of topological notions that will be involved in Section 3.6.

It is well-known that a metric on a space induces topological properties such as open sets. In fact, let $B(x, r)$ be the open ball centered at $x \in X$ with radius r , i.e., $B(x, r) = \{y : d_X(x, y) < r\}$. The set of all the open balls in X forms a basis of a topology on space X . Denote the topology by τ_X , then (X, τ_X) is a topological space induced by the metric d_X .

Homotopy Equivalence We first introduce the homotopic relationship between functions from a topological space (X, τ_X) to (Y, τ_Y) .

Definition 3.7 Two continuous functions f_1, f_2 from X to Y are homotopic if and only if there exists a continuous function $H : X \times [0, 1] \rightarrow Y$, such that $H(x, 0) = f_1(x)$ and $H(x, 1) = f_2(x)$

Then we define homotopy equivalence between topological spaces.

Definition 3.8 Two topological spaces X and Y are homotopy equivalent if and only if there exists two continuous functions $f : X \rightarrow Y$ and $g : Y \rightarrow X$ such that $f \circ g$ (resp., $g \circ f$) is homotopic to id_Y (resp., id_X), which is the identity map.

Especially, if a topological space X is homotopy equivalent to a one-point space $Y = \{y\}$, then we call X a contractible space. From the point view of homotopy, cycles are the only non-trivial topological features in a graph. In fact, one can classify graphs with the number of circles. A simple illustration of the following theorem is in Figure 3.2.

Theorem 3.2 [Homotopy Classification of graphs] Any connected graph $G = (V, E)$ is homotopy equivalent to a wedge of $|E| - |V| + 1$ circles, where $|V|$ (resp., $|E|$) is the cardinality of V (resp., E).

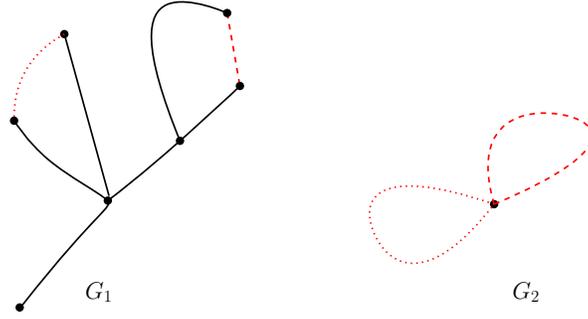


Figure 3.2: The black part of G_1 is contractible to the black point in G_2 . After adding two red edges, G_1 is no longer contractible and is homotopy equivalent to a wedge of 2 cycles as G_2 .

The number of cycles of a graph, $|E| - |V| + 1$, is also known as the first betti number $\beta_1(G)$ of G , i.e., the rank of the first homology group of G . Roughly speaking, one need to remove at least $\beta_1(G)$ edges from G to obtain an acyclic graph, i.e., a tree.

Abstract Simplicial Complex An abstract simplicial complex Δ is a family of non-empty sets that satisfies if $A \in \Delta$ and $B \subset A$, then $B \in \Delta$.

Covering and Nerve Let $\mathcal{V} = \{V_j\}_{j \in J}$ be a finite open covering of space X , meaning that $X = \bigcup_{j \in J} V_j$ and J is a finite indexing set.

The *nerve* of the covering \mathcal{V} is a simplicial complex $N(\mathcal{V})$ whose vertex set is the indexing set J , and vertices $\{j_1, j_2, \dots, j_{k+1}\}$ span a k -simplex in $N(\mathcal{V})$ if and only if $V_{j_1} \cap V_{j_2} \cap \dots \cap V_{j_{k+1}} \neq \emptyset$.

The following Nerve lemma asserts that the homotopy type of X is possibly recovered by a *good* cover.

Lemma 3.1 [Nerve Lemma] *Let $\mathcal{V} = \{V_j\}_{j \in J}$ be a finite open covering of X . If the intersection of any subset of \mathcal{V} is either empty or contractible, then the nerve of \mathcal{V} , $N(\mathcal{V})$, is homotopy equivalent to the union of X .*

3.4 Reeb-type Graph

In this section, we describe a construction to build a Reeb-type graph for approximating a metric space (X, d_X) . Let (X, d_X) be a compact geodesic space and let $r \in X$ be a fixed base point. Let $d : X \rightarrow \mathbb{R}$ be the distance function to r , i.e., $d(x) = d_X(r, x)$. The set of points in X with distance to r equal to $d(x)$ is then the pre-image of $d(x)$, i.e., $d^{-1}(d(x))$.

The Reeb graph. Define relation $x \sim y$ if and only if $d(x) = d(y)$ and x, y are in the same path connected component of $d^{-1}(d(x))$. This relation is an equivalence relation. The quotient space $G = X / \sim$ is called the *Reeb graph* of d .

We denote by $\pi : X \rightarrow G$ the quotient map. Notice that π is continuous and as X is path connected, G is path connected. The function d induces a function $d_* : G \rightarrow \mathbb{R}_+$ that satisfies $d = d_* \circ \pi$. A relation on

G is defined by: for any $g, g' \in G$, $g \leq_G g'$ if and only if $d_*(g) \leq d_*(g')$ and there exist a continuous path γ in G connecting g to g' such that $d \circ \gamma$ is non decreasing, makes G a partially ordered set.

The α -Reeb graphs. Computing or approximating the Reeb graph of (X, d) from a finite set of point sampled on X is usually a difficult task. To overcome this issue we also consider a variant of the Reeb graph that shares very similar properties to the Reeb graph. This variant is closely related to the graph constructed by the Mapper algorithm introduced in [Singh 2007] making its computation much easier than the Reeb graph (see Section 3.8).

Let $\alpha > 0$ and let $\mathcal{I} = \{I_i\}$ be a covering of the range of d by open intervals of length at most α . The transitive closure of the relation $x \sim_\alpha y$ if and only if $d(x) = d(y)$ and x, y are in the same path connected component of $d^{-1}(I_i)$ for some interval $I_i \in \mathcal{I}$ is an equivalence relation that is also denoted by \sim_α . The quotient space $G_\alpha = X / \sim_\alpha$ is called the α -Reeb graph¹ of d .

We denote by $\pi : X \rightarrow G_\alpha$ the quotient map. Notice that π is continuous and as X is path connected, G_α is path connected. The function d induces a function $d_* : G_\alpha \rightarrow \mathbb{R}_+$ that satisfies $d = d_* \circ \pi$. The relation on G_α is defined by: for any $g, g' \in G_\alpha$, $g \leq_{G_\alpha} g'$ if and only if $d_*(g) \leq d_*(g')$ and there exist a continuous path γ in G_α connecting g to g' such that $d \circ \gamma$ is non decreasing, makes G_α a partially ordered set.

Notice that without making assumptions on X and d , in general G and G_α are not finite graphs. However when the number of path connected components of the level sets of d is finite and changes only a finite number of times then the Reeb graph turns out to be a finite directed acyclic graph. Similarly, when the covering of X by the connected components of $d^{-1}(I_i), I_i \in \mathcal{I}$ is finite, the α -Reeb graph also turns out to be a finite directed acyclic graph. This happens in most applications and for example when (X, d_X) is a finite simplicial complex or a compact semialgebraic (or more generally a compact subanalytic space) with d being semi-algebraic (or subanalytic).

All the results and proofs presented in Section 3.5 are exactly the same for the Reeb and the α -Reeb graphs. In the following paragraph and in Section 3.5, G denotes indifferently the Reeb graph or an α -Reeb graph for some $\alpha > 0$. We also always assume that X and d (and α and \mathcal{I}) are such that G is a finite graph.

3.4.1 Metrics on Reeb and α -Reeb Graphs.

Let us define the set of vertices V of G as the union of the set of points of degree not equal to 2 with the set of local maxima of d_* over G , and the base point $\pi(r)$. The set of edges E of G is then the set of the connected components of the complement of V . Notice that $\pi(r)$ is the only local (and global) minimum of d_* : since X is path connected, for any $x \in X$ there exists a geodesic γ joining r to x along which d is increasing; d_* is thus also increasing along the continuous curve $\pi(\gamma)$, so $\pi(x)$ cannot be a local minimum of d_* . As a consequence d_* is monotonic along the edges of G . We can thus assign an orientation to each edge: if $e = [p, q] \in G$ is such that $d_*(p) < d_*(q)$ then the positive orientation of e is the one pointing from p to q . Finally, we assign a metric to G . Each edge $e \in E$ is homeomorphic to an interval to which we assign a length equal to the absolute difference of the function d_* at two endpoints. The distance between two points p, p' of e is then $|d_*(p) - d_*(p')|$. This makes G a metric graph (G, d_G) isometric to the quotient space of

¹strictly speaking we should call it the α -Reeb graph associated to the covering \mathcal{I} but we assume in the sequel that some covering \mathcal{I} has been chosen and we omit it in notations

the union of the intervals isometric to the edges by identifying the endpoints if they correspond to the same vertex in G . Note that d_* is continuous in (G, d_G) and for any $p \in G$, $d_*(p) = d_G(\pi(r), p)$. Indeed this is a consequence of the following lemma.

Lemma 3.2 *If δ is a path joining two points $p, p' \in G$ such that $d_* \circ \delta$ is strictly increasing then δ is a shortest path between p and p' and $d_G(p, p') = d_*(p') - d_*(p)$.*

Proof: As $d_* \circ \delta$ is strictly increasing, when δ enters an edge e by one of its end points, either it exits at the other end point or it stops at p' if $p' \in e$. Moreover δ cannot go through a given edge more than one time. As a consequence δ can be decomposed in a finite sequence of pieces $e_0 = [p, p_1], e_1 = [p_1, p_2], \dots, e_{n-1} = [p_{n-1}, p_n], e_n = [p_n, p']$ where e_0 and e_n are the segments joining p and p' to one of the endpoint of the edges that contain them and e_1, \dots, e_{n-1} are edges. So, the length of δ is equal to $(d_*(p_1) - d_*(p)) + (d_*(p_2) - d_*(p_1)) + \dots + (d_*(p') - d_*(p_n)) = d_*(p') - d_*(p)$ and $d_G(p, p') \leq d_*(p') - d_*(p)$.

Similarly any simple path joining p to p' can be decomposed in a finite sequence of pieces $e'_0 = [p, p'_1], e'_1 = [p'_1, p'_2], \dots, e'_{k-1} = [p'_{k-1}, p'_k], e'_k = [p'_k, p']$ where e'_0 and e'_k are the segments joining p and p' to one of the endpoint of the edges that contain them, and e'_1, \dots, e'_{k-1} are edges. Now, as we do not know that d_* is increasing along this path, its length is thus equal to $|d_*(p'_1) - d_*(p)| + |d_*(p'_2) - d_*(p'_1)| + \dots + |d_*(p') - d_*(p'_k)| \geq d_*(p') - d_*(p)$. So, $d_G(p, p') \geq d_*(p') - d_*(p)$. \square

3.5 Recovery of Geometry

The goal of this section is to provide an upper bound of the Gromov-Hausdorff distance between X and G , and we conclude our result as the following theorem:

Theorem 3.5 [Recovery of Geometry]. *Let (X, d_X) be a compact connected geodesic space, let $r \in X$ be a fixed base point such that the metric Reeb graph (G, d_G) of the function $d = d_X(r, \cdot) : X \rightarrow \mathbb{R}$ is a finite graph. If for a given $\varepsilon > 0$ there exists a finite metric graph $(G', d_{G'})$ such that $d_{GH}(X, G') < \varepsilon$ then we have*

$$d_{GH}(X, G) < 2(\beta_1(G) + 1)(17 + 8N_{E, G'}(8\varepsilon))\varepsilon$$

where $N_{E, G'}(8\varepsilon)$ is the number of edges of G' of length at most 8ε and $\beta_1(G)$ is the first Betti number of G , i.e. the number of edges to remove from G to get a spanning tree. In particular if X is at distance less than ε from a metric graph with shortest edge larger than 8ε then $d_{GH}(X, G) < 34(\beta_1(G) + 1)\varepsilon$.

We decompose our proofs into two parts: the first one (Theorem 3.3) asserts that the Gromov-Hausdorff distance between X and G only depends on the first Betti number $\beta_1(G)$ of G and the maximal diameter M of the level sets of π ; and the rest part that estimates an upper bound of M is given in Section 3.5.1

Theorem 3.3 $d_{GH}(X, G) < (\beta_1(G) + 1)M$ where $d_{GH}(X, G)$ is the Gromov-Hausdorff distance between X and G , $\beta_1(G)$ is the first Betti number of G and $M = \sup_{p \in G} \{\text{diam}(\pi^{-1}(p))\}$ is the supremum of the diameters of in the level sets of π .

Remark that as $\beta_1(G) \leq \beta_1(X)$, from the above theorem, $d_{GH}(X, G)$ is upper bounded by the quantities depending only on the input X . The proof of Theorem 3.3 is deduced from two propositions comparing the distances between pairs of points $x, y \in X$ and their images $\pi(x), \pi(y) \in G$ whose proofs rely on the notion

of merging vertex. A vertex $v \in V$ is called a *merging vertex* if it is the end point of at least two edges e_1 and e_2 that are pointing to it according to the orientation defined in Section 3.4. Geometrically this means that there are at least two distinct connected components of $\pi^{-1}(d_*^{-1}(d_*(v) - \varepsilon))$ that accumulate to $\pi^{-1}(v)$ as $\varepsilon > 0$ goes to 0. The set of merging vertices is denoted by V_m . We have

Lemma 3.3 *The cardinality of V_m is at most $\beta_1(G)$ where $\beta_1(G)$ is the rank of the first homology group of G .*

Proof: The result follows from classical persistence homology theory [Edelsbrunner 2010]. First remark that, as $\pi(r)$ is the only local minimum of d_* , the sublevel sets of the function $d_* : G \rightarrow \mathbb{R}_+$ are all path connected. Indeed if $\pi(x), \pi(y) \in G$ are in the same sublevel set $d_*^{-1}([0, \alpha])$, $\alpha > 0$, then the images by π of the shortest paths in X connecting x to r and y to r are contained in $d_*^{-1}([0, \alpha])$ and their union is a continuous path joining $\pi(x)$ to $\pi(y)$. As a consequence, the 0-dimensional persistence of d_* is trivial. So as we increase the α value, no merging vertices serve as connecting two different connected components. Thus, each merging vertex in V_m creates at least a cycle that never dies as G is one dimensional and does not contain any 2-dimensional simplex. Thus $|V_m| \leq \beta_1(G)$. \square

The following lemma shows that a shortest path in G is the projection of a shortest path in X as long as it does not meet a merging vertex and allow to prove proposition 3.1 below.

Lemma 3.4 *Let $p, p' \in G$ and let $\delta : [d_*(p), d_*(p')] \rightarrow G$ be a strictly increasing path going from p to p' that does not contain any point of V_m in its interior. Then for any $x' \in \pi^{-1}(p') \cap \text{cl}(\pi^{-1}(\delta(d_*(p), d_*(p'))))$ where $\text{cl}(\cdot)$ denotes the closure, there exists a shortest path γ connecting a point x of $\pi^{-1}(p)$ to x' such that $\pi(\gamma) = \delta$ and $d_X(x, x') = d(x') - d(x) = d_*(p') - d_*(p) = d_G(p, p')$.*

Proof: First assume that p' is not a merging point. Let $\gamma_0 : [0, d(x')] \rightarrow X$ be any shortest path between r and x' and let γ be the restriction of γ_0 to $[d_*(p), d(x')] = [d_*(p), d_*(p')]$. If the infimum t_0 of the set $I = \{t \in [d_*(p), d_*(p')] : \pi(\gamma(t)) \in \delta, \forall t' \geq t\}$ is larger than $d_*(p)$, then there exists an increasing sequence (t_n) that converges to t_0 such that $\pi(\gamma(t_n)) \notin \delta$. As a consequence $\delta(t_0)$ is a merging point; a contradiction. So $t_0 = d_*(p)$ and $\gamma(d_*(p))$ intersects $\pi^{-1}(p)$ at a point x .

Now if p' is a merging point, as x' is chosen in the closure of $\pi^{-1}(\delta(d_*(p), d_*(p')))$, for any sufficiently large $n \in \mathbb{N}$ one can consider a sequence of points $x'_n \in \pi^{-1}(\delta(d_*(p') - 1/n))$ that converges to x' and apply the first case to get a sequence of shortest path γ_n from a point $x_n \in \pi^{-1}(p)$ and x'_n . Then applying Arzelà-Ascoli's theorem (see [Dieudonné 1960] 7.5) we can extract from γ_n a sequence of points converging to a shortest path γ between a point $x \in \pi^{-1}(p)$ and x' .

To conclude the proof, notice that from Lemma 3.2 we have $d_G(p, p') = d_*(p') - d_*(p) = d(x') - d(x)$. Since γ is the restriction of a shortest path from r to x we also have $d_X(x, x') = d(x') - d(x)$. \square

Notice that from Lemma 3.2, δ is a shortest path and the parametrization by the interval $[d_*(p), d_*(p')]$ can be chosen to be an isometric embedding.

Proposition 3.1 *For any $x, y \in X$ we have*

$$d_X(x, y) \leq d_G(\pi(x), \pi(y)) + 2(\beta_1(G) + 1)M$$

where $M = \sup_{p \in G} \{\text{diam}(\pi^{-1}(p))\}$ and $\beta_1(G)$ is the first Betti number of G .

Proof: Let δ be a shortest path between $\pi(x)$ and $\pi(y)$. Remark that except at the points $\pi(x)$ and $\pi(y)$ the local maxima of the restriction of d_* to δ are in V_m . Indeed as δ is a shortest path it has to be simple, so if $p \in \delta$ is a local maximum then p has to be a vertex and δ has to pass through two edges having p as end point and pointing to p according to the orientation defined in 3.4. So p is a merging point.

Since δ is simple and V_m is finite, δ can be decomposed in at most $|V_m| + 1$ connected paths along the interior of which the restriction of d_* does not have any local maxima. So along each of these connected paths the restriction of d_* can have at most one local minimum. As a consequence, δ can be decomposed in a finite number of continuous paths $\delta_1, \delta_2, \dots, \delta_k$ with $k \leq 2(|V_m| + 1)$, such that the restriction of d_* to each of these path is strictly monotonic. For any $i \in \{1, \dots, k\}$ let p_i and p_{i+1} the end points of δ_i with $p_1 = \pi(x)$ and $p_{k+1} = \pi(y)$. We can apply Lemma 3.4 to each δ_i to get a shortest path γ_i in X between a point $x_i \in \pi^{-1}(p_i)$ and a point in $y_{i+1} \in \pi^{-1}(p_{i+1})$ such that $\pi(\gamma_i) = \delta_i$ and $d_X(x_i, y_{i+1}) = d_G(p_i, p_{i+1})$. The sum of the lengths of the paths γ_i is equal to the sum of the lengths of the path δ_i which is itself equal to $d_G(\pi(x), \pi(y))$. Now for any $i \in \{1, \dots, k\}$, since $\pi(x_i) = \pi(y_i)$ we have $d_X(x_i, y_i) \leq M$ and x_i and y_i can be connected by a path of length at most M (x_1 is connected to x and y_{k+1} is connected to y). Gluing these paths to the paths γ_i gives a continuous path from x to y whose length is at most $d_G(\pi(x), \pi(y)) + kM \leq d_G(\pi(x), \pi(y)) + 2(|V_m| + 1)M$. Since from Lemma 3.3, $|V_m| \leq \beta_1(G)$, we finally get that $d_X(x, y) \leq d_G(\pi(x), \pi(y)) + 2(\beta(G) + 1)M$. \square

Proposition 3.2 *The map $\pi : X \rightarrow G$ is 1-Lipschitz: for any $x, y \in X$ we have*

$$d_G(\pi(x), \pi(y)) \leq d_X(x, y).$$

Proof: Let $x, y \in X$ and let $\gamma : I \rightarrow X$ be a shortest path from x to y in X where $I \subset \mathbb{R}$ is a closed interval. The path $\pi(\gamma)$ connects $\pi(x)$ and $\pi(y)$ in G .

We first claim that there exists a continuous path Γ contained in $\pi(\gamma)$ connecting $\pi(x)$ and $\pi(y)$ that intersects each vertex of G at most one time. The path Γ can be defined by iteration in the following way. Let $v_1, \dots, v_n \in V$ be the vertices of G that are contained in $\pi(\gamma) \setminus \{\pi(x), \pi(y)\}$ and let $\Gamma_0 = \pi(\gamma) : J_0 \rightarrow G, J_0 = I$. For $i = 1, \dots, n$ let $t_i^- = \inf\{t : \Gamma_{i-1}(t) = v_i\}$ and $t_i^+ = \sup\{t : \Gamma_{i-1}(t) = v_i\}$ and define Γ_i as the restriction of Γ_{i-1} to $J_i = J_{i-1} \setminus (t_i^-, t_i^+)$. The path Γ_i is a connected continuous path (although J_i is a disjoint union of intervals) that intersects the vertices v_1, v_2, \dots, v_i at most one time. We then define $\Gamma = \Gamma_n : J = J_n \rightarrow G$ where $J \subset I$ is a finite union of closed intervals. Notice that Γ is the image by π of the restriction of γ to J and that $\Gamma(t) \in \{v_1, \dots, v_n\}$ only if t is one of the endpoints of the closed intervals defining J .

Now, for each connected component $[t, t']$ of J , $\gamma([t, t'])$ is contained in $\pi^{-1}(e)$ where e is the edge of G containing $\Gamma([t, t'])$. As a consequence,

$$\begin{aligned} d_G(\pi(\gamma)(t), \pi(\gamma)(t')) &= |d_*(\pi(\gamma)(t)) - d_*(\pi(\gamma)(t'))| \\ &= |d(\gamma(t)) - d(\gamma(t'))|. \end{aligned}$$

Recalling that $d(\gamma(t)) = d_X(x, \gamma(t))$ and $d(\gamma(t')) = d_X(x, \gamma(t'))$ and using the triangle inequality we get that $|d(\gamma(t)) - d(\gamma(t'))| \leq d_X(\gamma(t), \gamma(t'))$. To conclude the proof, since γ is a geodesic path we just need

to sum up the previous inequality over all connected components of J :

$$\begin{aligned} d_X(x, y) &\geq \sum_{[t, t'] \in cc(J)} d_X(\gamma(t), \gamma(t')) \\ &\geq \sum_{[t, t'] \in cc(J)} d_G(\pi(\gamma)(t), \pi(\gamma)(t')) \geq d_G(\pi(x), \pi(y)) \end{aligned}$$

where $cc(J)$ is the set of connected components of J . \square

The proof of Theorem 3.3 now easily follows from Propositions 3.1 and 3.2.

Proof: (of Theorem 3.3) Consider the set $C = \{(x, \pi(x)) : x \in X\} \subset X \times G$. As π is surjective this is a correspondence between X and G . It follows from Propositions 3.1 and 3.2 that for any $(x, \pi(x)), (y, \pi(y)) \in C$,

$$|d_X(x, y) - d_G(\pi(x), \pi(y))| \leq 2(\beta_1(G) + 1)M$$

So C is a $2(\beta_1(G) + 1)M$ -correspondence and $d_{GH}(X, G) \leq (\beta_1(G) + 1)M$. \square

3.5.1 Bounding the Diameter M

The two following lemmas allow to bound M , the diameter of the level sets of π .

Lemma 3.5 *Let (G, d_G) be a connected finite metric graph and let $r \in G$. We denote by $d_r = d_G(r, \cdot) : G \rightarrow [0, +\infty)$ the distance to r . For any edge $e \subset G$, the restriction of d_r to e is either strictly monotonic or has only one local maximum. Moreover the length $l = l(e)$ of e is upper bounded by two times the difference between the maximum and the minimum of d_r restricted to e .*

Proof: Let l be the length of E and let $t \mapsto e(t)$, $t \in [0, l]$, be an arc length parametrization of E . Since E is an edge of G , for $t \in [0, l]$ any shortest geodesic γ_t joining r to $e(t)$ must contain either $x_1 = e(0)$ or $x_2 = e(l)$. If it contains x_1 then for any $t' < t$ the restriction of γ_t between r and $e(t')$ is a shortest geodesic containing x_1 and if it contains x_2 then for any $t' > t$ the restriction of γ_t between r and $e(t')$ is a shortest geodesic containing x_2 . Moreover in both cases, the function d_r is strictly monotonic along γ . As a consequence, the set $I_1 = \{t \in [0, l] : \text{a shortest geodesic joining } r \text{ to } e(t) \text{ contains } x_1\}$ is a closed interval containing 0. Similarly the set $I_2 = \{t \in [0, l] : \text{a shortest geodesic joining } r \text{ to } e(t) \text{ contains } x_2\}$ is a closed interval containing l and $[0, l] = I_1 \cup I_2$. Moreover d_r is strictly monotonic on $e(I_1)$ and on $e(I_2)$. As a consequence $I_1 \cap I_2$ is reduced to a single point t_0 that has to be the unique local maximum of d_r restricted to E .

The second part of the lemma follows easily from the previous proof: the minimum of d_r restricted to E is attained either at x_1 or x_2 and $d_r(e(t_0)) = d_r(x_1) + t_0 = d_r(x_2) + l - t_0$ is the maximum of d_r restricted to E . We thus obtain that $2t_0 = l + (d_r(x_2) - d_r(x_1))$. As a consequence if $d_r(x_1) \leq d_r(x_2)$ then $l/2 \leq t_0 = d_r(e(t_0)) - d_r(x_1)$; similarly if $d_r(x_1) \geq d_r(x_2)$ then $l/2 \leq l - t_0 = d_r(e(t_0)) - d_r(x_2)$. \square

Proposition 3.3 *Let (G, d_G) be a connected finite metric graph and let $r \in G$. For $\alpha > 0$ we denote by $N_E(\alpha)$ the number of edges of G of length at most α . For any $d > 0$ and any connected component B of the set $B_{d, \alpha} = \{x \in G : d - \alpha \leq d_G(r, x) \leq d + \alpha\}$ we have*

$$\text{diam}(B) \leq 4(2 + N_E(4\alpha))\alpha$$

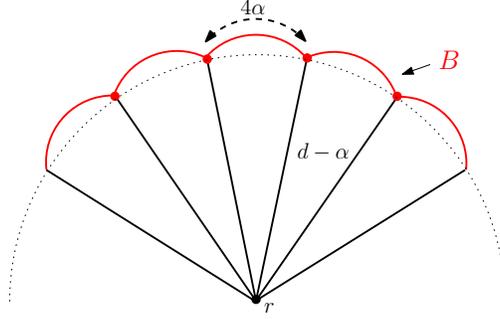


Figure 3.3: Tightness of the bound in Lemma 3.3: there are 3 edges of length at most 4α and the diameter of B is equal to 20α . The range of the distances from r to the points on the red curve is $[d - \alpha, d + \alpha]$.

Proof: Let $x, y \in B$ and let $t \mapsto \gamma(t) \in B$ be a continuous path joining x to y in B . Let E be an edge of G that does not contain x or y and with end points x_1, x_2 such that γ intersects the interior of E . Then $\gamma^{-1}(E)$ is a disjoint union of closed intervals of the form $I = [t, t']$ where $\gamma(t)$ and $\gamma(t')$ belong to the set $\{x_1, x_2\}$. If $\gamma(t) = \gamma(t')$ we can remove the part of γ between t and t' and still get a continuous path between x and y . So without loss of generality we can assume that if γ intersects the interior of E , then E is contained in γ . Using the same argument as previously we can also assume that if γ goes across E , it only does it one time, i.e. $\gamma^{-1}(E)$ is reduced to only one interval. As a consequence, γ can be decomposed in a sequence $[x, v_0], E_1, E_2, \dots, E_k, [v_k, y]$ where $[x, v_0]$ and $[v_k, y]$ are pieces of edges containing x and y respectively and $E_1 = [v_0, v_1], E_2 = [v_1, v_2], \dots, E_k = [v_{k-1}, v_k]$ are pairwise distinct edges of G contained in B . It follows from Lemma 3.5 that the lengths of the edges E_1, \dots, E_k and of $[x, v_0]$ and $[v_k, y]$ are upper bounded by 4α . As a consequence the length of γ is upper bounded by $4(k+2)\alpha$ which is itself upper bounded by $4(N_E(4\alpha)+2)\alpha$ since the edges E_1, \dots, E_k are pairwise distinct. It follows that $d_G(x, y) \leq 4(N_E(4\alpha)+2)\alpha$. \square

The example of Figure 3.3 shows that the bound of Lemma 3.3 is tight.

Lemma 3.6 *Let X and Y be compact geodesic metric spaces and $C \subset X \times Y$ be an ε' -correspondence between them. Assume $(x_0, y_0) \in C(X, Y)$, we define functions $d_{x_0}(\cdot) = d_X(x_0, \cdot)$ and $d_{y_0}(\cdot) = d_Y(y_0, \cdot)$ in X and Y respectively. Then for any path γ_x in X connecting $x_a, x_b \in X$, we can find a path γ_y in Y such that its end points, y_a, y_b , are corresponding to x_a, x_b . And further more:*

$$[\min_{y \in \gamma_y} d_{y_0}(y), \max_{y \in \gamma_y} d_{y_0}(y)] \subset [\min_{x \in \gamma_x} d_{x_0}(x) - 2\varepsilon', \max_{x \in \gamma_x} d_{x_0}(x) + 2\varepsilon']$$

Proof: Let $\varepsilon > \varepsilon' > 0$ and u, l be the maximum and minimum of d_{x_0} restricted to γ_x . Since C is an ε' -correspondence for any $x \in \gamma_x$ there exists a point $(x, y) \in C$ such that $d_{x_0}(x) - \varepsilon' \leq d_{y_0}(y) \leq d_{x_0}(x) + \varepsilon'$. As illustrated in Figure 3.4, the set of points y obtained in this way is not necessarily a continuous path from y_a to y_b . However one can consider a finite sequence $x_1 = x_a, x_2, \dots, x_n = x_b$ of points in γ_x such that for any $i = 1, \dots, n-1$ we have $d_X(x_i, x_{i+1}) < \varepsilon - \varepsilon'$. If $(x_i, y_i) \in C$ then we have $d_Y(y_i, y_{i+1}) < \varepsilon - \varepsilon' + \varepsilon' = \varepsilon$. As a consequence, since $l - \varepsilon < l - \varepsilon' < d_{y_0}(y_i) < u + \varepsilon' < u + \varepsilon$ the shortest geodesic connecting y_i to y_{i+1} in G remains in the set $d_{y_0}^{-1}([l - 2\varepsilon, u + 2\varepsilon])$ and connecting these geodesics for all

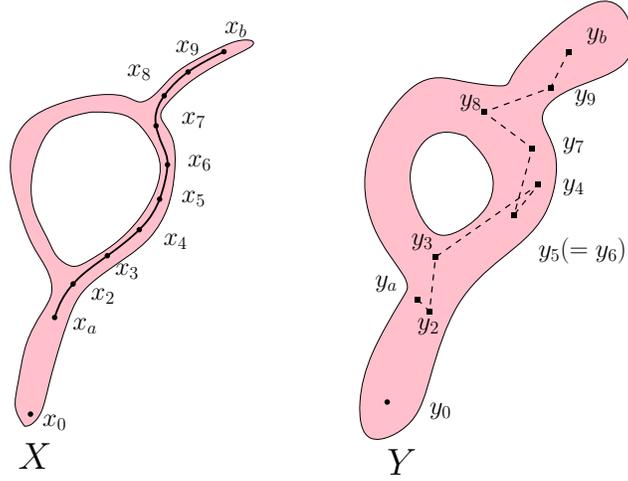


Figure 3.4: The path correspondence between two metric spaces X and Y .

$i = 1, \dots, n - 1$ we get a continuous path from y_a to y_b in $d_r^{-1}([l - 2\varepsilon, u + 2\varepsilon])$. Now decreasing ε to ε' , we finish the construction. \square

As a corollary, we have the following theorem.

Theorem 3.4 *Let (G, d_G) be a connected finite metric graph and let (X, d_X) be a compact geodesic metric space such that $d_{GH}(X, G) < \varepsilon$ for some $\varepsilon > 0$. Let $x_0 \in X$ be a fixed point and let $d_{x_0} = d_X(x_0, \cdot) : X \rightarrow [0, +\infty)$ be the distance function to x_0 . Then for $d \geq \alpha \geq 0$ the diameter of any connected component L of $d_{x_0}^{-1}([d - \alpha, d + \alpha])$ satisfies*

$$\text{diam}(L) \leq 4(2 + N_E(4(\alpha + 2\varepsilon)))(\alpha + 2\varepsilon) + \varepsilon$$

where $N_E(4(\alpha + 2\varepsilon))$ is the number of edges of G of length at most $4(\alpha + 2\varepsilon)$. In particular if $\alpha = 0$ and 8ε is smaller than the length of the shortest edge of G then $\text{diam}(L) < 17\varepsilon$.

Proof: Let $\varepsilon' > 0$ be such that $d_{GH}(X, G) < \varepsilon' < \varepsilon$. Let $C \subset X \times G$ be an ε' -correspondence between X and G and $(x_0, r) \in C$. We denote by $d_r = d_G(r, \cdot) : G \rightarrow [0, +\infty)$ the distance function to r in G . Let $x_a, x_b \in L$ and let $(x_a, y_a), (x_b, y_b) \in C$. There exists a continuous path $\gamma \subseteq L$ joining x_a to x_b . Following lemma 3.6, we get a continuous path from y_a to y_b in $d_r^{-1}([d - \alpha - 2\varepsilon', d + \alpha + 2\varepsilon'])$. It then follows from Proposition 3.3 that $d_G(y_a, y_b) \leq 4(2 + N_E(4(\alpha + 2\varepsilon)))(\alpha + 2\varepsilon)$ and since C is an ε' -correspondence (and so an ε -correspondence), $d_X(x_a, x_b) < 4(2 + N_E(4(\alpha + 2\varepsilon)))(\alpha + 2\varepsilon) + \varepsilon$. \square

From Theorems 3.4 and 3.3 we obtain the following results for the Reeb and α -Reeb graphs.

Theorem 3.5 *Let (X, d_X) be a compact connected path metric space, let $r \in X$ be a fixed base point such that the metric Reeb graph (G, d_G) of the function $d = d_X(r, \cdot) : X \rightarrow \mathbb{R}$ is a finite graph. If for a given $\varepsilon > 0$ there exists a finite metric graph $(G', d_{G'})$ such that $d_{GH}(X, G') < \varepsilon$ then we have*

$$d_{GH}(X, G) < (\beta_1(G) + 1)(17 + 8N_{E,G'}(8\varepsilon))\varepsilon$$

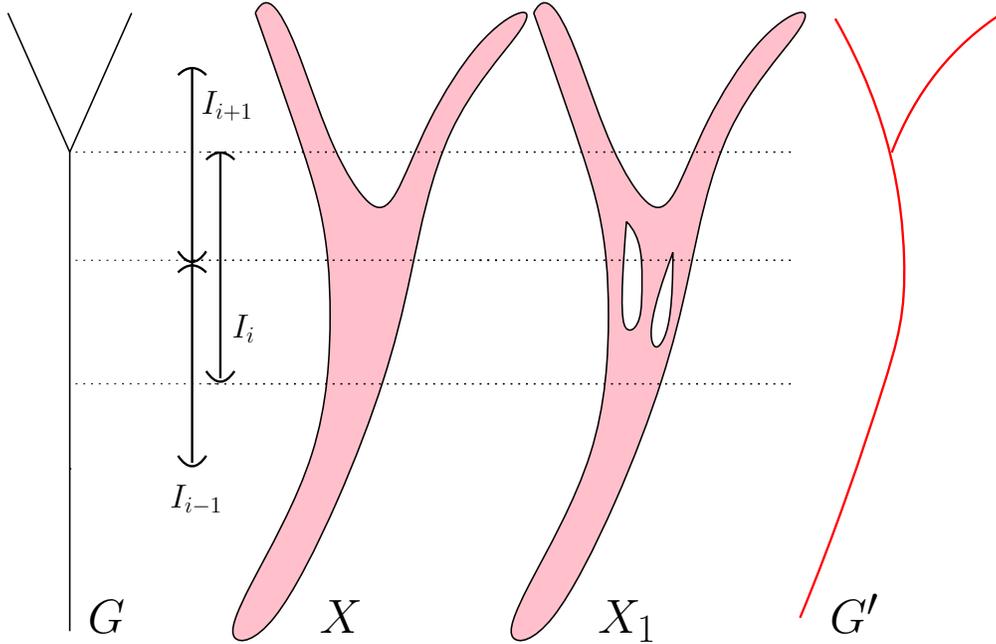


Figure 3.5: X and X_1 are metric spaces that are both close to the metric graph G' with respect to the Gromov-Hausdorff distance. The α -Reeb graph, G , captures the stable topological features and ignores the unstable ones like the small holes in X_1 .

where $N_{E,G'}(8\varepsilon)$ is the number of edges of G' of length at most 8ε . In particular if X is at distance less than ε from a metric graph with shortest edge length larger than 8ε then $d_{GH}(X, G) < 17(\beta_1(G) + 1)\varepsilon$.

Theorem 3.6 Let (X, d_X) be a compact connected path metric space. Let $r \in X$, $\alpha > 0$ and \mathcal{I} be a finite covering of the segment $[0, \text{Diam}(X)]$ by open intervals of length at most α such that the α -Reeb graph G_α associated to \mathcal{I} and the function $d = d_X(r, \cdot) : X \rightarrow \mathbb{R}$ is a finite graph. If for a given $\varepsilon > 0$ there exists a finite metric graph $(G', d_{G'})$ such that $d_{GH}(X, G') < \varepsilon$ then we have

$$d_{GH}(X, G_\alpha) < (\beta_1(G_\alpha) + 1)(4(2 + N_{E,G'}(4(\alpha + 2\varepsilon)))(\alpha + 2\varepsilon) + \varepsilon)$$

where $N_{E,G'}(4(\alpha + 2\varepsilon))$ is the number of edges of G' of length at most $4(\alpha + 2\varepsilon)$. In particular if X is at distance less than ε from a metric graph with shortest edge length larger than $4(\alpha + 2\varepsilon)$ then $d_{GH}(X, G_\alpha) < (\beta_1(G_\alpha) + 1)(8\alpha + 17\varepsilon)$.

3.6 Recovery of Topology

In this section, we consider the the α -Reeb graph G of a distance function $d : X \rightarrow \mathbb{R}$, and show the following theorem which asserts that G recovers some topology of X .

Theorem 3.7 *Let (X, d_X) be a compact connected path metric space and $(G', d_{G'})$ is a metric graph so that $d_{GH}(X, G') < \varepsilon$. Let $r \in X$, $\alpha > 60\varepsilon$ and $\mathcal{I} = \{[0, 2\alpha), (i\alpha, (i+2)\alpha) | 1 \leq i \leq m\}$ covers the segment $[0, \text{Diam}(X)]$ such that the α -Reeb graph G associated to \mathcal{I} and the function $d = d_X(r, \cdot) : X \rightarrow \mathbb{R}$ is a finite graph. If no edges of G' are shorter than L and no loops' lengths of G' are shorter than $2L$ with $L \geq 32\alpha + 9\varepsilon$, then we have G and G' are homotopy equivalent.*

First we note that under the assumption of $d_{GH}(X, G') \leq \varepsilon$, X and G' are not necessarily homotopy equivalent. For example, in Figure 3.5 we assume that $d_{GH}(X, G') \leq \varepsilon$, then $d_{GH}(X_1, G') \leq \varepsilon$ as long as the perimeters of the two small holes in X_1 are both less than ε . Nevertheless, X and X_1 are not homotopy equivalent since X is contractible while X_1 is not.

On the other hand, we usually care less about smaller topological features such as the holes in X_1 , since they can be mixed with noisy features introduced in the real data acquisition. One of the important properties of the mapper algorithm is that it provides a way to extract robust (opposite to the small unstable) topological features in the data. Inheriting this property, the α -Reeb graph G in Figure 3.5 manages to capture the homotopy type of X and negates the two small loops of X_1 .

The above observations motivate our setting in this section, we assume that the sizes of the topological features, i.e., the length of the shortest edges and circles, of the underlying G' is bounded from below and that X is allowed to have unstable small features as long as $d_{GH}(X, G') \leq \varepsilon$. And then we prove that the α -Reeb graph is capable of capturing the topology of G' .

Our strategy of proving Theorem 3.7 is to construct some open covers for X and G' and relate the α -Reeb graph G and the graph G' to the nerves of the open covers. Specifically, we construct an initial open cover \mathcal{V}_0 of X whose nerve $N(\mathcal{V}_0)$ is homotopy equivalent to G . Then we obtain a new open cover $\tilde{\mathcal{V}}$ of X by merging certain elements in \mathcal{V}_0 while preserving the homotopy type of the nerve of the open cover, i.e., $N(\mathcal{V}_0)$ and $N(\tilde{\mathcal{V}})$ are homotopy equivalent. Based on the open cover $\tilde{\mathcal{V}}$, we construct an open cover $\tilde{\mathcal{U}}$ for G' whose nerve $N(\tilde{\mathcal{U}})$ is isomorphic to $N(\tilde{\mathcal{V}})$ as graphs and at the same time is homotopy equivalent to G' . In the following, we describe the constructions of the above open covers for X and G' and show the above claimed relations between them.

Since $d_{GH}(X, G') < \varepsilon$, there exists an ε -correspondence between the two spaces, denoted $C(X, G')$. For any subset $V \subset X$, denote $C(V) = \{g' : (x, g') \in C(X, G'), x \in V\}$, and similarly for any subset $U \subset G'$, denote $C(U) = \{x : (x, g') \in C(X, G'), g' \in U\}$. We call $C(V)$ and $C(U)$ are the correspondence set of V and U respectively under $C(X, G')$. Recall that $r \in X$ is the root point. Choose a point $g_r \in C(r)$ and define a distance function $b : G' \rightarrow \mathbb{R}$ by $b(g) = d_{G'}(g_r, g)$. Let $N = \{g_{n_1}, g_{n_2}, \dots, g_{n_p}\}$ be the vertices of G' , i.e., N is the set of vertices whose degree is not equal to two. From the hypotheses of the above theorem, the distance between any pair of vertices g_{n_i}, g_{n_j} with $i \neq j$ is larger than L . For convenience, we also add into the vertices of G' the remaining local maximal/minimal points of the distance function b , which we denote using $M = \{g_{m_1}, \dots, g_{m_q}\}$. Note any newly added vertex $g_{m_i} \in M$ is of degree two. We call the graph G' before adding the vertices in M the original G' , and the edges in the original G' the original edges of G' . An original edge of G' contains at most one vertex in M and thus can be split into at most two edges in G' .

3.6.1 Construction of Open Cover for X

We start with the following open cover of X . For each $I_k \in \mathcal{I}$, denote $V_k = d^{-1}(I_k)$. V_k may have several connected components, which can be listed in an arbitrary order. Denote V_k^l the l -th connected component of V_k . Then $\mathcal{V}_0 = \{V_k^l\}_{k,l}$ is an open cover of X . Since at most two elements in \mathcal{I} are overlapped, the nerve of \mathcal{V}_0 , denoted $N(\mathcal{V}_0)$, is a graph. The following lemma states that any loop in the nerve $N(\mathcal{V}_0)$ is large, which is useful for the proof of Theorem 3.7. We say an open set $V_{k_1}^{l_1} \in \mathcal{V}_0$ is lower than the open set $V_{k_2}^{l_2} \in \mathcal{V}_0$ if $k_1 < k_2$ and is higher than $V_{k_2}^{l_2}$ if $k_2 > k_1$.

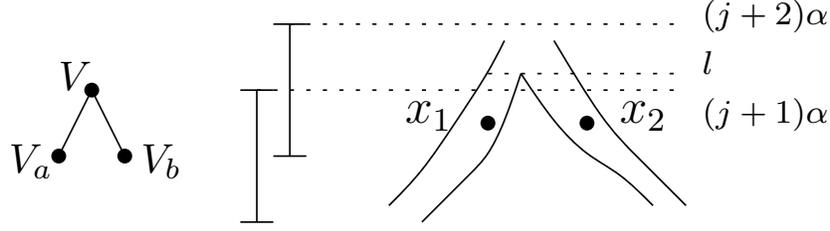
Lemma 3.7 *Let V_k^l and V_j^i are the lowest vertex and the highest vertex of a loop respectively in the nerve $N(\mathcal{V}_0)$. Then under the hypotheses of Theorem 3.7, we have $j - k \geq 15$.*

Proof: First notice that $j > k$. Let $x_1 \in V_k^l \cap d^{-1}(k\alpha, (k+1)\alpha)$ and $x_2 \in V_j^i \cap d^{-1}((j+1)\alpha, (j+2)\alpha)$. From the hypotheses of the lemma, there are two different paths β_1, β_2 connecting x_1 to x_2 so that $\beta_1 \cap d^{-1}((k+1)\alpha, (j+1)\alpha)$ and $\beta_2 \cap d^{-1}((k+1)\alpha, (j+1)\alpha)$ are in the different connected components of $d^{-1}((k+1)\alpha, (j+1)\alpha)$. Choose $g_i \in G'$ from $C(x_i)$ for $i = 1, 2$. Following Lemma 3.6, the path β_i in X for $i = 1, 2$ traces out a simple path γ_i in G' connecting g_1 to g_2 so that γ_i lies in $b^{-1}(k\alpha - \varepsilon, (j+2)\alpha + \varepsilon)$. One can verify that γ_1 and γ_2 are two different paths due to the fact that β_1 and β_2 pass through different connected components of $d^{-1}((k+1)\alpha, (j+1)\alpha)$ and thus form a loop in G' , denoted γ . We have $b(\gamma) \subset (k\alpha - \varepsilon, (j+2)\alpha + \varepsilon)$.

We claim the range of the function b restricted to any loop, in particular β , covers an interval with the length at least $\frac{L}{2}$. If the claim holds, then we have $(j - k + 2)\alpha + 2\varepsilon \geq \frac{L}{2}$, which implies $j - k \geq 15$ from the hypothesis $L \geq 32\alpha + 9\varepsilon$ of Theorem 3.7. Indeed, if β contains at least two vertices in N , then it is obvious that the range of the function b restricted to γ covers an interval the length at least $\frac{L}{2}$ as any original edge of G' is longer than L . Now consider the case where γ contains one vertex in N , say g_a . If γ does not contain g_r , then there is exactly one local maximum on γ , say g_b . If γ contains g_r , let $g_b = g_r$. The removal of g_a and g_b cuts γ into two pieces. Along either piece, the function b has at most one local maximum. As the length of γ is longer than $2L$. We have $b(\gamma)$ covers an interval with length longer than $L/2$. Finally, if β contains no vertex in N , then G' is a single loop γ and the claim obviously holds. \square

In the following, we modify this open cover by merging while preserving the homotopy type of its nerve. The main purpose of the merging operation is make it easy to relate the open cover of X to the open cover of G' as constructed in Section 3.6.2. The merging operation is done in two steps.

For any vertex $g \in M \cup N$ of G' , we construct a connected open set $V(g)$ as the union of a subset of the open cover \mathcal{V}_0 as follows. If $b(g) \geq \frac{\alpha}{2}$, then there exists a unique positive integer k' s.t. $k' \frac{\alpha}{2} \leq b(g) < (k' + 1) \frac{\alpha}{2}$. Let $k = \lfloor \frac{k'+1}{2} \rfloor - 1 \geq 0$, and one can verify that $(k + \frac{1}{2})\alpha \leq b(g) \leq (k + \frac{3}{2})\alpha$. Therefore for all $x \in C(g)$, $d(x) \in [(k + \frac{1}{2})\alpha - \varepsilon, (k + \frac{3}{2})\alpha + \varepsilon] \subset I_k$. Moreover $C(g)$ is contained in $V_k^l \subset V_k$ for some l . Indeed, if not, assume $x_1, x_2 \in C(g)$ with $x_i \in V_k^i$ for $i \in \{1, 2\}$. By the definition of V_k^i , the geodesic connecting x_1 and x_2 must pass through a point x_0 outside of V_k , which means $d_X(x_i, x_0) \geq |d(x_i) - d(x_0)| \geq \frac{\alpha}{2} - \varepsilon$. Then $d_X(x_1, x_2) \geq \alpha - 2\varepsilon$ which contradicts to the fact that $d_X(x_1, x_2) \leq d_{G'}(g, g) + \varepsilon \leq \varepsilon$. Now we construct the open set $V(g)$ as the union of the elements in the


 Figure 3.6: V with two lower neighborhoods.

open cover \mathcal{V}_0 having non-empty intersection with V_k^l , i.e.,

$$V(g) = \bigcup_{V \in \mathcal{V}_0 \text{ and } V \cap V_k^l \neq \emptyset} V.$$

In the case where $b(g) < \frac{\alpha}{2}$, we construct the open set $V(g) = V_0 \cup V_1 = d^{-1}([0, 3\alpha))$. Note in both cases, $V(g)$ is a connected open set of X . We abuse the notation and also denote $V(g)$ the subset of \mathcal{V}_0 whose union is the open set $V(g)$. What $V(g)$ represents will be clear from the context. For convenience, we call V_k^l containing $C(g)$ the center of $V(g)$. Note that it is possible that $V(g) = V(g')$ for two different vertices g, g' .

Now we obtain an intermediate open cover of X

$$\mathcal{V} = \{V(g) : g \in M \cup N\} \cup \{V \in \mathcal{V}_0 : V \notin V(g), \forall g \in M \cup N\}$$

Note as a set, \mathcal{V} does not have duplicated elements, i.e., if $V(g) = V(g')$ for $g \neq g'$, then \mathcal{V} only contains one copy of $V(g)$. We call an open set $V(g) \in \mathcal{V}$ for any $g \in M \cup N$ critical and the remaining ones regular. The following two lemmas describe the properties of the critical open sets and the regular open sets.

Lemma 3.8 *Under the hypotheses of Theorem 3.7, we have for any vertex $g \in M \cup N$,*

- i) $d(V(g)) \subset [s\alpha, (s+4)\alpha]$ for some integer $s \geq 0$, and
- ii) for any point $x \in \bigcup_{V \in \mathcal{V}_0 \setminus V(g)} V$ and any $g_x \in C(x) \subset G'$, $d_{G'}(g, g_x) \geq \frac{\alpha}{2} - 2\varepsilon$.

Proof: The claim (i) is obvious from the construction of $V(g)$. We now prove claim (ii). In the case where $b(g) < \frac{\alpha}{2}$, for any $x \in \bigcup_{V \in \mathcal{V}_0 \setminus V(g)} V$, we have $d(x) > 3\alpha$ and $b(g_x) > 3\alpha - \varepsilon$. Thus $d_{G'}(g, g_x) \geq |b(g_x) - b(g)| > 3\alpha - \varepsilon - \frac{\alpha}{2} > \frac{\alpha}{2} - 2\varepsilon$. Now consider the case where $b(g) \geq \frac{\alpha}{2}$. Assume V_k^l is the center of $V(g)$. If $d(x) \notin I_k$, then $d_X(x, y) \geq \frac{\alpha}{2} - \varepsilon$ for any point $y \in C(g)$ from the construction of $V(g)$, which implies $d_{G'}(g_x, g) \geq \frac{\alpha}{2} - 2\varepsilon$. Otherwise $d(x) \in I_k$. Then x is not in V_k^l and the geodesic from x to any point $y \in C(g)$ must pass $x_0 \notin V_k$. This implies that $d_X(x, y) > d_X(x_0, y) \geq \frac{\alpha}{2} - \varepsilon$ and $d_{G'}(g_x, g) \geq d_X(x, y) - \varepsilon \geq \frac{\alpha}{2} - 2\varepsilon$. This proves the lemma. \square

Lemma 3.9 *For any regular open set $V \in \mathcal{V}$, V is also an open set in \mathcal{V}_0 . Moreover, under the hypotheses of Theorem 3.7, it is of degree two in the nerve of $N(\mathcal{V}_0)$ with one neighboring vertex higher than V and one neighboring vertex lower than V .*

Proof: We prove the lemma by contradiction. Assume $V \in \mathcal{V}_0 \setminus \bigcup_{g \in M \cup N} V(g)$ has two neighboring vertices, say V_a, V_b , which are lower than V . Without loss of generality, assume $d(V) \subset I_j$ and $d(V_a)$ and $d(V_b)$ are subsets of I_{j-1} . Let $x_a \in V_a$ and $x_b \in V_b$ such that $(j-1)\alpha < d(x_a), d(x_b) < j\alpha$. As V_a and V_b both have non-empty intersection with V , there exist a path in $d^{-1}((j-1)\alpha, (j+2)\alpha)$. Now let $l = \inf\{s : \text{there exists a path connecting } x_a, x_b \text{ in } d^{-1}((j-1)\alpha, s] \cap (V \cup V_a \cup V_b)\}$. We have $(j+1)\alpha \leq l < (j+2)\alpha$ as V_a, V_b are disconnected.

We can choose two points $x_1, x_2 \in V \cup V_a \cup V_b$ from a path connecting x_a and x_b such that $d(x_1) = d(x_2) = l - 2\varepsilon$, and x_1, x_2 are disconnected in $d^{-1}([l - 2\varepsilon, l]) \cap (V \cup V_a \cup V_b)$ but are connected by a path, say β , in $d^{-1}([l - 2\varepsilon, l]) \cap (V \cup V_a \cup V_b)$. Obviously $d_X(x_1, x_2) \geq 2(l - (l - 2\varepsilon)) = 4\varepsilon$. Let $g_i \in C(x_i)$ for $i = 1, 2$. Then $b(g_i) \in [l - 3\varepsilon, l - \varepsilon]$ for $i = 1, 2$ and $d_{G'}(g_1, g_2) \geq d_X(x_1, x_2) - \varepsilon \geq 3\varepsilon$. Following Lemma 3.6, the path β traces out a simple path in G' denoted γ connecting g_1 and g_2 which lies in $b^{-1}(l - 4\varepsilon, l + 2\varepsilon)$. We claim γ must contain a vertex g_c in $M \cup N$. If not,

$$\begin{aligned} d_{G'}(g_1, g_2) &= |b(g_1) - b(g_2)| \\ &\leq |b(g_1) - d(x_1) + d(x_2) - b(g_2)| \\ &\leq |(b(g_1) - b(g_r)) - (d(x_1) - d(r))| + |(d(x_2) - d(r)) - (b(g_2) - b(g_r))| \\ &= |d_{G'}(g_1, g_r) - d_X(x_1, r)| + |d_{G'}(g_2, g_r) - d_X(x_2, r)| \\ &\leq 2\varepsilon \end{aligned}$$

This contradicts to the fact that $d_{G'}(g_1, g_2) \geq 3\varepsilon$. From the construction of γ , there exists a point $g \in \gamma$ so that $d_{G'}(g, g_c) \leq \varepsilon$ and there exists a point $x \in \beta \cap C(g)$. Then for any point $x_c \in C(g_c)$, we have $d_X(x, x_c) \leq 2\varepsilon$. For any vertex $x_0 \in V \cap d^{-1}(l)$, since x_0 and x are connected in $d^{-1}([l - 2\varepsilon, l]) \cap (V \cup V_a \cup V_b)$, $d_X(x_0, x) \leq 25\varepsilon$ from Theorem 3.4, and therefore $d_X(x_0, x_c) \leq 27\varepsilon$. However, since $x_0 \in V$ which is regular, $x \notin V(g_c)$. From Lemma 3.8, we have $d_X(x_0, x_c) \geq \frac{\alpha}{2} - 2\varepsilon > 27\varepsilon$. This is a contradiction. Therefore V can not have more than one neighboring vertices that are lower than V .

Using a similar argument we can also prove that V can not have more than one neighboring vertices that are higher than V . \square

We now perform a second step of merging. Two critical open set $V(g_1)$ and $V(g_2)$ in $N(\mathcal{V})$ are said to be close if there is a simple path γ in the nerve $N(\mathcal{V}_0)$ connecting the center $V_{k_1}^{l_1}$ of $V(g_1)$ and the center $V_{k_2}^{l_2}$ of $V(g_2)$ so that γ consists of at most 4 edges. If there is a regular open set along the above path, we say this regular open set connects the critical open sets $V(g_1)$ and $V(g_2)$.

We have the following properties for two close critical open sets.

Lemma 3.10 *Under the hypotheses of Theorem 3.7, we have*

- (i) *for any two vertices $g_{n_1}, g_{n_2} \in N$, $V(g_{n_1})$ and $V(g_{n_2})$ can not be close;*
- (ii) *for any $g_m \in M$, there exists at most one $g_n \in N$ such that $V(g_m)$ and $V(g_n)$ are close;*
- (iii) *if $V(g_{m_1})$ and $V(g_{m_2})$ are close for any two vertices $g_{m_1}, g_{m_2} \in M$, then there must exist a vertex $g_n \in N$ such that at least one of $V(g_{m_1})$ and $V(g_{m_2})$ is close to $V(g_n)$. Moreover, there is a path in $N(\mathcal{V}_0)$ of at most 5 edges connecting the center of $V(g_n)$ to the center of $V(g_{m_i})$ for any $i = 1, 2$.*

Proof: Let V_j^p, V_k^q are the centers of $V(g_1)$ and $V(g_2)$ respectively for $g_1, g_2 \in N \cup M$. If $V(g_1)$ and $V(g_2)$ are close, then $|k - j| \leq 4$. Assume $j \leq k$. Then for any $x_1 \in C(g_1)$ and $x_2 \in C(g_2)$, there is a path in $d^{-1}((j\alpha, (k+2)\alpha))$ connecting x_1 and x_2 . Note that $k+2-j \leq 6$. We claim that $V(g_1)$ and $V(g_2)$ are not close provided that $d_{G'}(g_1, g_2) > 12\alpha + 9\varepsilon$. Indeed, since $d_{G'}(g_1, g_2) > 12\alpha + 9\varepsilon$, the range of the function b restricted to any path connecting g_1 and g_2 in G' covers an interval of the length at least $6\alpha + 4.5\varepsilon$. This implies that the range of the function d restricted to any path in X connecting x_1 and x_2 covers an interval of the length at least $6\alpha + 0.5\varepsilon$. This means that $V(g_1)$ and $V(g_2)$ can not be close. Since $d_{G'}(g_{n_1}, g_{n_2}) \geq L > 12\alpha + 9\alpha$, $V(g_{n_1})$ and $V(g_{n_2})$ are not close. This proves (i).

Assume $V(g_m)$ is close to both $V(g_{n_1})$ and $V(g_{n_2})$ with $g_{n_1}, g_{n_2} \in N$. We have $d_{G'}(g_{n_1}, g_{n_2}) \leq d_{G'}(g_m, g_{n_1}) + d_{G'}(g_m, g_{n_2}) \leq 24\alpha + 18\varepsilon < L$, which means $g_{n_1} = g_{n_2}$. This proves (ii).

We now prove (iii). Since at most one vertex in M is added into an original edge of G' , any path in G' connecting g_{m_1} and g_{m_2} passes through at least one vertex from N . Furthermore, let γ be a geodesic in G' connecting g_{m_1} and g_{m_2} . If γ passes more than one vertices in N , $d_{G'}(g_{m_1}, g_{m_2}) \geq L > 12\alpha + 9\varepsilon$, which contradicts to the fact that $V(g_{m_1})$ and $V(g_{m_2})$ are close. Therefore γ contains exactly one vertex in N . Denote this vertex by g_n . Let $V_{k_1}^{l_1}$ and $V_{k_2}^{l_2}$ be the centers of $V(g_{m_1})$ and $V(g_{m_2})$ respectively, and δ be the simple path from $V_{k_1}^{l_1}$ to $V_{k_2}^{l_2}$ in $N(\mathcal{V}_0)$ so that δ consists of at most 4 edges, or equivalently at most five elements in \mathcal{V}_0 .

Recall $V(g_n)$ consists of a subset of \mathcal{V}_0 . We claim that δ must pass through an element in $V(g_n)$.

If the claim holds, it is easy to verify that at least one of $V(g_{m_1})$ and $V(g_{m_2})$ is close to $V(g_n)$. In addition, if we let V_k^l be the center of $V(g_n)$, then there is a path in $N(\mathcal{V}_0)$ with as most 5 edges connecting $V_{k_i}^{l_i}$ and V_k^l for any $i = 1, 2$. This proves (iii).

It remains to show the above claim. We prove by contradiction. If we let

$$V(\delta) = \{V \in \mathcal{V}_0 : V \text{ is on the path of } \delta\},$$

then $V(g_n)$ as a subset of \mathcal{V}_0 does not intersect with $V(\delta)$. We have $C(g_{m_1})$ and $C(g_{m_2})$ are contained in $V_{k_1}^{l_1}$ and $V_{k_2}^{l_2}$ respectively. For any $x_1 \in C(g_{m_1})$ and any $x_2 \in C(g_{m_2})$, there is a path β in X connecting x_1 and x_2 so that β is contained in $\bigcup_{V \in V(\delta)} V$. From Lemma 3.8, for any $x \in \beta$ and any $g_x \in C(x)$, $d_{G'}(g_x, g_n) \geq \frac{\alpha}{2} - 2\varepsilon$. From the construction in the proof of Theorem 3.4, the path β can trace out a simple path γ' in G' connecting g_{m_1} and g_{m_2} so that for any point $g \in \gamma'$, $d_{G'}(g, g_n) \geq \frac{\alpha}{2} - 3\varepsilon$. This means that γ and γ' form a loop in G' . Since β is contained in $\bigcup_{V \in V(\delta)} V$, $d(\beta)$ is contained in an interval with the length of 6α , which implies that $b(\gamma')$ is contained in an interval with the length of $6\alpha + 4\varepsilon$. Thus the length of γ' is at most $2(6\alpha + 4\varepsilon)$ as it passes through at most one vertex in N . Since γ is a geodesic in G' connecting g_{m_1} and g_{m_2} , the length of the above loop is at most $4(6\alpha + 4\varepsilon)$, which contradicts to the hypotheses of Theorem 3.7. This proves the above claim. \square

We say $g_a, g_b \in M \cup N$ are equivalent, denoted $g_a \sim_c g_b$, if there exists a finite sequence $g_a = g_1, g_2, \dots, g_k = g_b$ such that $V(g_i)$ and $V(g_{i+1})$ are close for any $i = 1, \dots, k-1$. This is an equivalence relation. From Lemma 3.10 (iii), if an equivalence class contains at least two vertices in $M \cup N$, it must contain a vertex in N . We have the following lemma

Lemma 3.11 *Under the hypotheses of Theorem 3.7, an equivalence class contains at most one vertex from N .*

Proof: If not, assume $g_{n_1} \neq g_{n_2}$ and $g_{n_1} \sim_c g_{n_2}$. Let $g_{n_1} = g_1, g_2, \dots, g_k = g_{n_2}$ be a sequence so that $V(g_i)$ and $V(g_{i+1})$ are close for any $i = 1, \dots, k-1$. WLOG, we can further assume $g_i \in M$ for $i = 2, \dots, k-1$.

We first show that $k > 5$. Assume not. let $V_{k_i}^{l_i}$ be the center of $V(g_i)$ for $i = 1, \dots, k$. From Lemma 3.10 (iii), there is a path in $N(\mathcal{V}_0)$ with at most $2 \times 5 = 10$ edges connecting $V_{k_1}^{l_1}$ to $V_{k_5}^{l_5}$. Thus for any $x_1 \in C(g_{n_1})$ and any $x_2 \in C(g_{n_2})$, there is a path β in X connecting x_1 and x_2 so that $d(\beta)$ is contained in an interval with the length at most 12α . The path β traces out a path γ in G' connecting g_{n_1} and g_{n_2} so that $b(\gamma)$ is contained in an interval with the length at most $12\alpha + 4\varepsilon$, which implies $d_{G'}(g_{n_1}, g_{n_2}) \leq 2(12\alpha + 4\varepsilon)$. This contradicts to the hypothesis concerning the lengths of the edges in G' .

Now we assume $k > 5$. Since $V(g_3)$ and $V(g_4)$ are close and $g_3, g_4 \in M$, from Lemma 3.10 (iii), there exists a $g_n \in N$ so that $V(g_n)$ is close to at least one of $V(g_3)$ and $V(g_4)$. Assume $V(g_n)$ is close to $V(g_3)$. If $V(g_n) \neq V(g_{n_1})$, we obtain a sequence of $g'_1 = g_{n_1}, g'_2 = g_2, g'_3 = g_3, g'_4 = g_n$ so that $V(g'_i)$ and $V(g'_{i+1})$ are close for any $i = 1, \dots, 3$. If $V(g_n) \neq V(g_{n_2})$, we obtain a sequence $g'_1 = g_n, g'_2 = g_3, \dots, g'_{k-1} = g_{n_2}$ so that $V(g'_i)$ and $V(g'_{i+1})$ are close for any $i = 1, \dots, k-2$. In either case, the new sequence has a length less than k . Similarly, we can obtain a shorter sequence if $V(g_n)$ is close to $V(g_4)$. One can keep shortening the sequence so that its length is no longer than 5, which however has been proven to be impossible. This proves the lemma. \square

Now we are ready to further merge the open sets in \mathcal{V} to obtain the final open cover $\tilde{\mathcal{V}}$ of X as follows. For any vertex $g_n \in N$ of G' , Let $\tilde{V}(g_n)$ be the subset of \mathcal{V} consisting of (1) $V(g_n)$, and (2) any critical open set $V(g) \subset \mathcal{V}$ with $g \sim_c g_n$, and (3) any regular open set $V \subset \mathcal{V}$ connecting two critical open sets which are equivalent to g_n . We abuse the notation and also denote $\tilde{V}(g_n)$ the open set of the union of the open sets in $\tilde{V}(g_n)$. What $\tilde{V}(g_n)$ represents will be clear from the context. Let $\tilde{\mathcal{V}}_N = \{V \in \mathcal{V} : V \in \tilde{V}(g_n) \text{ for some } g_n \in N\}$. The open cover $\tilde{\mathcal{V}} = \tilde{\mathcal{V}}_1 \cup \tilde{\mathcal{V}}_2 \cup \tilde{\mathcal{V}}_3$ of X consists of three types of open sets:

- (1) $\tilde{\mathcal{V}}_1 = \{\tilde{V}(g_n) : g_n \in N\}$;
- (2) $\tilde{\mathcal{V}}_2 = \{V(g) : g \in M \text{ and } V(g) \notin \tilde{\mathcal{V}}_N\}$
- (3) $\tilde{\mathcal{V}}_3 = \{V \in \mathcal{V} : V \text{ is regular and } V \notin \tilde{\mathcal{V}}_N\}$.

Figure 3.7 shows different types of elements in $\tilde{\mathcal{V}}$. We summarize the properties for the open cover $\tilde{\mathcal{V}}$ in the following corollary, which follows from Lemma 3.8, Lemma 3.9, Lemma 3.10, and Lemma 3.11.

Corollary 3.1 *Under the hypotheses of Theorem 3.7, the open sets in $\tilde{\mathcal{V}}$ satisfy the following properties.*

- $\tilde{V}(g_1)$ and $\tilde{V}(g_2)$ are disjoint for two different $g_1, g_2 \in N$.
- For any two open sets $\tilde{V}_1, \tilde{V}_2 \in \tilde{\mathcal{V}}_1 \cup \tilde{\mathcal{V}}_2$, any path in the 1-skeleton of the nerve $N(\tilde{\mathcal{V}})$ connecting \tilde{V}_1, \tilde{V}_2 consists of at least two elements from $\tilde{\mathcal{V}}_3$.
- Any open set $\tilde{V} \in \tilde{\mathcal{V}}_3$ is also a regular element in \mathcal{V} and thus an element in \mathcal{V}_0 . Moreover any point $g \in C(\tilde{V}) \subset G'$ is at least $\frac{\alpha}{2} - 2\varepsilon$ away from any vertex of G' .

Proposition 3.4 *Under the hypotheses of Theorem 3.7, $N(\tilde{\mathcal{V}})$ and $N(\mathcal{V}_0)$ are homotopy equivalent.*

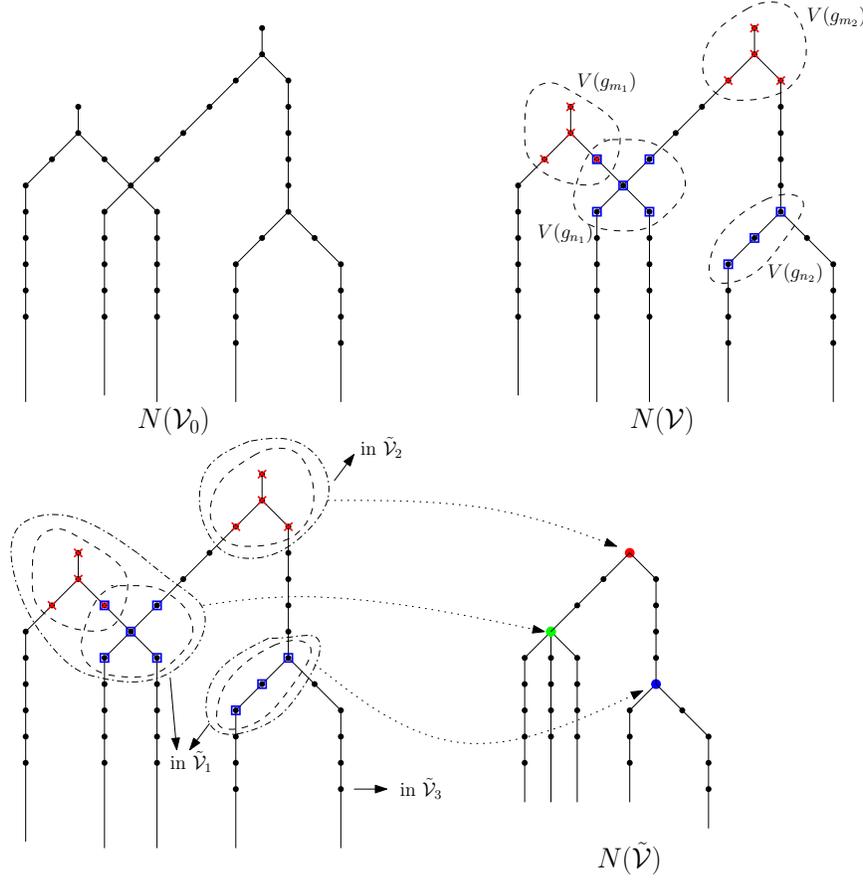


Figure 3.7: Illustration of the merging strategy: From top-left to top-right (First step): We select the unions of open sets from \mathcal{V}_0 based on the critical points w.r.t. $b(g)$ in G' and merge them respectively; From top-right to bottom-left (Second step): We further merge the unions which are mutually close. Bottom-right: the nerve of the merging result $N(\tilde{\mathcal{V}})$.

Proof:

We obtain the elements in the open covering $\tilde{\mathcal{V}}$ by merging a subset of open sets in \mathcal{V}_0 . Think of any element $\tilde{V} \in \tilde{\mathcal{V}}$ as a subset of \mathcal{V}_0 . The nerve $N(\mathcal{V}_0)$ restricted to \tilde{V} is a subgraph of $N(\mathcal{V}_0)$, whose vertex set is \tilde{V} and edge set includes the edges in $N(\mathcal{V}_0)$ with both endpoints in \tilde{V} . We call this subgraph the nerve of \tilde{V} , denoted $N(\tilde{V})$. The nerve of $N(\tilde{\mathcal{V}})$ as a topological space is the quotient space $N(\mathcal{V}_0)/\bigcup_{\tilde{V} \in \tilde{\mathcal{V}}} N(\tilde{V})$. From Proposition 0.17 in [Hatcher 2002], it is sufficient to show that $N(\tilde{V})$ is a tree for any $\tilde{V} \in \tilde{\mathcal{V}}$.

For $\tilde{V} \in \tilde{\mathcal{V}}_2 \cup \tilde{\mathcal{V}}_3$, $N(\tilde{V})$ is obviously a tree. Consider $\tilde{V} \in \tilde{\mathcal{V}}_1$. There exists a $g_n \in N$ so that $V(g_n) \subset \tilde{V}(g_n) \in \tilde{V}$. Let v_s^t be the center of $V(g_n)$. For any $g_m \in M$ and $V(g_m) \subset \tilde{V}(g_n)$, if let $v_{s'}^{t'}$ be the center of $V(g_m)$, from Lemma 3.10, $|s - s'| < 5$. Therefore, if V_l^i is the element in \tilde{V} with the smallest sub-index and V_h^j is the element in \tilde{V} with the largest sub-index, then we have $|h - l| \leq 5 + 5 + 2 + 2 = 14$. From Lemma 3.7, there is no loop in the subgraph $N(\tilde{V})$. This proves the proposition. \square

3.6.2 Construction of Open Cover for G'

In this section, we construct an open cover G' based on the open cover $\tilde{\mathcal{V}}$ of X . For an open set $V \in \mathcal{V}_0$, we construct a connected open set $U_V \subset G'$ so that $C(V) \subset U_V$ as follows. Let $l = \min\{d(V)\}$ and $u = \max\{d(V)\}$. We have $u - l \leq 2\alpha$. Let $\bar{U} = b^{-1}([l - 2\varepsilon, u + 2\varepsilon])$, and then $C(V) \subset \bar{U}$. Since $u - l + 4\varepsilon < 2\alpha + 4\varepsilon < \frac{L}{4}$, one can verify that there is no loop in \bar{U} and thus \bar{U} consists of a set of trees. We claim $C(V)$ is contained in one of the trees. Indeed, for any two $g_1, g_2 \in C(V)$, we have $l - \varepsilon < b(g_1), b(g_2) < u + \varepsilon$. Now let $x_i \in V$ so that $g_i \in C(x_i)$ for $i = 1, 2$. Let β be a path in V connecting x_1 and x_2 . Following Lemma 3.6, β can trace out a path γ in \bar{U} connecting g_1 and g_2 , which implies that $C(V)$ is contained in a tree in \bar{U} . Let U_V denote that tree. Let $\mathcal{U}_0 = \{U_V : V \in \mathcal{V}_0\}$. It is obvious that \mathcal{U}_0 is an open cover of G' . We now merge the elements in \mathcal{U}_0 to construct a new open cover according to the way in which the elements in \mathcal{V}_0 are merged to obtain $\tilde{\mathcal{V}}$. Specifically, from our construction of $\tilde{\mathcal{V}}$, any open set $\tilde{V} \in \tilde{\mathcal{V}}$ is the union of a subset of open sets in \mathcal{V}_0 . We also denote this subset using \tilde{V} . Let $U_{\tilde{V}} = \{U_V : V \in \tilde{V} \subset \mathcal{V}_0\}$. We also denote $U_{\tilde{V}}$ is the open set of the union of the open sets in $U_{\tilde{V}}$.

Consider an open set $\tilde{V} \in \tilde{\mathcal{V}}_3$. As it is also a regular open set in \mathcal{V} and thus an open set in \mathcal{V}_0 , $d(\tilde{V}) = (p\alpha, (p+2)\alpha)$ for some integer $p > 0$. From Corollary 3.1, any point in $C(\tilde{V})$ is at least $\frac{\alpha}{2} - 2\varepsilon$ away from any vertex in $M \cup N$ and any point in $U_{\tilde{V}}$ is at least $\frac{\alpha}{2} - 4\varepsilon$ away from any vertex in $M \cup N$. Thus $U_{\tilde{V}}$ is a segment in G' without any branches. We shrink $U_{\tilde{V}}$ to obtain a new open set $\tilde{U}_{\tilde{V}} = U_{\tilde{V}} \cap b^{-1}(p\alpha + 2\varepsilon, (p+2)\alpha - 2\varepsilon)$, which is also a segment in G' . For any open set $\tilde{V} \in \tilde{\mathcal{V}}_1 \cup \tilde{\mathcal{V}}_2$, let $\tilde{U}_{\tilde{V}} = U_{\tilde{V}}$. Thus we obtain

$$\tilde{\mathcal{U}} = \{\tilde{U}_{\tilde{V}} : \tilde{V} \in \tilde{\mathcal{V}}\}.$$

One can verify that $\tilde{\mathcal{U}}$ is an open cover of G' . Moreover we have the following two lemmas which relate the nerve $N(\tilde{\mathcal{V}})$ to G' .

Proposition 3.5 *Under the hypotheses of Theorem 3.7, the nerve $N(\tilde{\mathcal{V}})$ and the nerve $N(\tilde{\mathcal{U}})$ are isomorphic as graphs.*

Proof: It suffices to prove the following three claims.

- Claim (i): For any two $\tilde{V}_i, \tilde{V}_j \in \tilde{\mathcal{V}}_1 \cup \tilde{\mathcal{V}}_2$, $\tilde{U}_{\tilde{V}_i} \cap \tilde{U}_{\tilde{V}_j} = \emptyset$.

Any path in $N(\tilde{\mathcal{V}})$ connecting \tilde{V}_i and \tilde{V}_j must pass through at least two open sets in $\tilde{\mathcal{V}}_3$, which are regular open sets in \mathcal{V} . From Lemma 3.9, any regular set has two neighbors in the nerve $N(\mathcal{V})$ one lower and one higher, WLOG, assume \tilde{V}_i is higher than \tilde{V}_j . We have $\inf\{d(x) : x \in \tilde{V}_i\} \geq \alpha + \sup\{d(x) | x \in \tilde{V}_j\}$, which implies $\inf\{b(g) | g \in \tilde{U}_{\tilde{V}_i}\} \geq \alpha + \sup\{b(g) | g \in \tilde{U}_{\tilde{V}_j}\} - 2\varepsilon > \sup\{b(g) | g \in \tilde{U}_{\tilde{V}_j}\}$. Thus $\tilde{U}_{\tilde{V}_i} \cap \tilde{U}_{\tilde{V}_j} = \emptyset$.

- Claim (ii): For any two $\tilde{V}_i, \tilde{V}_j \in \tilde{\mathcal{V}}_3$, $\tilde{V}_i \cap \tilde{V}_j = \emptyset$ if and only if $\tilde{U}_{\tilde{V}_i} \cap \tilde{U}_{\tilde{V}_j} = \emptyset$.

If $\tilde{V}_i \cap \tilde{V}_j \neq \emptyset$, assume \tilde{V}_i is the only neighboring vertex in the nerve $N(\mathcal{V})$ higher than \tilde{V}_j . Let $d(\tilde{V}_j) = (p\alpha, (p+2)\alpha)$ and $d(\tilde{V}_i) = ((p+1)\alpha, (p+3)\alpha)$. Choose a point x from $\tilde{V}_i \cap \tilde{V}_j$ so that $d(x) = (p + \frac{3}{2})\alpha$. We have $C(x) \in \tilde{U}_{\tilde{V}_i} \cap \tilde{U}_{\tilde{V}_j}$, which shows $\tilde{U}_{\tilde{V}_i} \cap \tilde{U}_{\tilde{V}_j} \neq \emptyset$.

If $\tilde{V}_i \cap \tilde{V}_j = \emptyset$. Let $d(\tilde{V}_i) = (p\alpha, (p+2)\alpha)$ and $d(\tilde{V}_j) = (q\alpha, (q+2)\alpha)$. If $|p - q| \geq 2$, it is obvious that $\tilde{U}_{\tilde{V}_i} \cap \tilde{U}_{\tilde{V}_j} = \emptyset$. Now assume that $q - p \leq 1$, which forces the shortest path connecting \tilde{V}_i and

\tilde{V}_j in $N(\tilde{\mathcal{V}})$ must pass through some open set $\tilde{V} \in \tilde{\mathcal{V}}_1 \cup \tilde{\mathcal{V}}_2$. By Lemma 3.8, for any $g_i \in C(\tilde{V}_i)$ $d_{G'}(g_i, g) \geq \frac{\alpha}{2} - 2\varepsilon$ and for any $g_j \in C(\tilde{V}_j)$ $d_{G'}(g_j, g) \geq \frac{\alpha}{2} - 2\varepsilon$ for any vertex $g \in M \cup N$ such that $V(g) \in \tilde{V}$. Thus $d_{G'}(g_i, g_j) \geq \alpha - 4\varepsilon$, which implies $\tilde{U}_{\tilde{V}_i} \cap \tilde{U}_{\tilde{V}_j} = \emptyset$.

- Claim (iii): For any $\tilde{V}_i \in \tilde{\mathcal{V}}_1 \cup \tilde{\mathcal{V}}_2$ and any $\tilde{V}_j \in \tilde{\mathcal{V}}_3$, $\tilde{V}_i \cap \tilde{V}_j = \emptyset$ if and only if $\tilde{U}_{\tilde{V}_i} \cap \tilde{U}_{\tilde{V}_j} = \emptyset$.

First assume that \tilde{V}_i and \tilde{V}_j have a non-empty intersection. As $\tilde{V}_j \in \tilde{\mathcal{V}}_3$, it is a regular open set in \mathcal{V} which has one higher neighboring vertex and one lower neighboring vertex in $N(\mathcal{V}_0)$. Since \tilde{V}_j is regular, we have $d(\tilde{V}_j) = (p\alpha, (p+2)\alpha)$ for some integer $p > 0$. We know \tilde{V}_i consists of a subset of open sets in \mathcal{V}_0 and let $V \in \tilde{V}_i$ be the open set in \mathcal{V}_0 so that $V \cap \tilde{V}_j \neq \emptyset$. WLOG, assume V is the higher neighboring vertex of \tilde{V}_j and we have $d(V \cap \tilde{V}_j) \supset ((p+1)\alpha, (p+2)\alpha)$. We choose a point in $x \in \tilde{V}_j \cap V$ so that $d(x) = (p+2)\alpha - 4\varepsilon$. Since $b(C(x)) \subset ((p+2)\alpha - 5\varepsilon, (p+2)\alpha - 3\varepsilon)$, $C(x) \in \tilde{U}_{\tilde{V}_j} \cap \tilde{U}_{\tilde{V}_i}$ and thus $\tilde{U}_{\tilde{V}_i} \cap \tilde{U}_{\tilde{V}_j} \neq \emptyset$.

Second assume $\tilde{V}_i \cap \tilde{V}_j = \emptyset$. If any path in the nerve $N(\tilde{\mathcal{V}})$ connecting \tilde{V}_i and \tilde{V}_j passes through some open set in $\tilde{\mathcal{V}}_1 \cup \tilde{\mathcal{V}}_2$, then we are done based on Claim (i). Now assume there is a path β in the nerve $N(\tilde{\mathcal{V}})$ connecting \tilde{V}_i and \tilde{V}_j only passing through open sets in $\tilde{\mathcal{V}}_3$. Since any open set in $\tilde{\mathcal{V}}_3$ is a regular set in \mathcal{V} , the worst scenery is that β contains no intermediate open sets. In this worst scenery, due to the shrinking operation on $\tilde{U}_{\tilde{V}_j}$, one can verify that $\tilde{U}_{\tilde{V}_i} \cap \tilde{U}_{\tilde{V}_j} = \emptyset$.

□

Proposition 3.6 *Under the hypotheses of Theorem 3.7, $N(\tilde{\mathcal{U}})$ is homotopy equivalent to G' .*

Proof: As we have proved, $\tilde{\mathcal{U}}$ is an open covering of G' . Since any edge on the original G' has a length longer than L , one can verify that any element of $\tilde{\mathcal{U}}$ contains no loop and thus is a tree, and in particular is contractible. Furthermore, the union of any two elements of $\tilde{\mathcal{U}}$ does not contain a loop. This means that if two elements of $\tilde{\mathcal{U}}$ intersect with each other, their intersection is connected and thus contractible. Following from Nerve lemma, we have $N(\tilde{\mathcal{U}})$ is homotopy equivalent to G' . □

Proof of Theorem 3.7. From Proposition 3.4, Proposition 3.5, Proposition 3.6, it remains to show that the nerve $N(\mathcal{V}_0)$ is homotopy equivalent to the α -Reeb graph G . Indeed, we represent each node V_k^l in $N(\mathcal{V}_0)$ using a copy of the interval I_k . If $V_{k_1}^{l_1}$ and $V_{k_2}^{l_2}$ with $k_1 < k_2$ are the endpoints of an edge in $N(\mathcal{V}_0)$, then we glue the upper half of I_{k_1} to the lower half of I_{k_2} . We identify any two points which are glued together directly or indirectly. By definition, the α -Reeb graph is the quotient space of the disjoint union of these intervals. From Lemma 3.7, there are more than one node between the top node and the bottom node of any loop in $N(\mathcal{V}_0)$. Thus, we have a one-to-one correspondence between the loops in $N(\mathcal{V}_0)$ and the loops in the α -Reeb graph. This proves the theorem.

3.7 Metric Reconstruction from Discrete Sampling

In the previous sections, we propose constructions of Reeb-type graphs for a given metric space (X, d_X) which is supposed to be close to an underlying metric graph $(G', d_{G'})$ in the Gromov-Hausdorff distance. Our theoretical analysis shows that the Reeb-type graph approximation is both geometrically and topologically reliable under certain conditions.

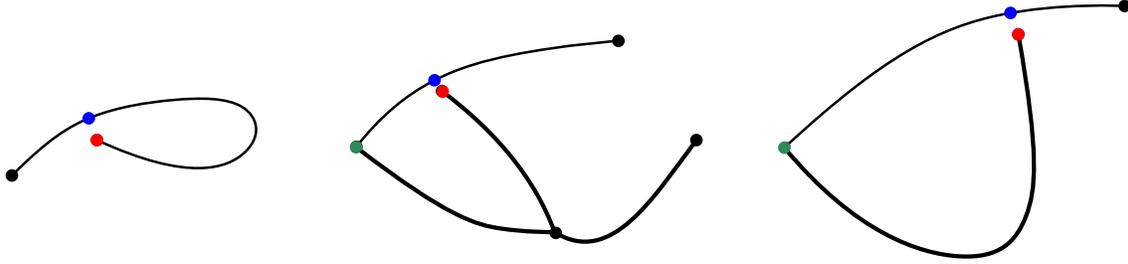


Figure 3.8: 3 types of close-self-intersection.

In practice, however, we usually only have access to a finite collection of isolated discrete points \tilde{X} sampled from X and need to build a metric $d_{\tilde{X}}$ on top of it. In general, it's unlikely to recover d_X exactly with a finite sampling. Therefore it's necessary to understand the distortion between d_X and its approximation $d_{\tilde{X}}$, i.e., the Gromov-Hausdorff distance between metric spaces (X, d_X) and $(\tilde{X}, d_{\tilde{X}})$.

Particularly, in this section we consider a case in which $(X, d_X) = (G, d_G)$ is a finite metric graph embedded in a d -dimensional Euclidean space \mathbb{R}^d . And we let $\tilde{X} = \{x_1, x_2, \dots, x_n\} \subset G \subset \mathbb{R}^d$ be an ε -sampling as defined below.

Definition 3.9 An ε -sampling of the embedded metric graph (G, d_G) is a finite set of points $\tilde{X} = \{x_1, x_2, x_3, \dots, x_n\} \in G \subset \mathbb{R}^d$, such that for any $g \in G$, there exists at least one $x \in \tilde{X}$ satisfying $d_G(g, x) \leq \varepsilon$.

With a positive parameter σ , we construct a σ -neighborhood graph (see Definition 3.5), $R_\sigma = (\tilde{X}, E_\sigma)$, on top of \tilde{X} with parameter $\sigma > 0$, and we denote the metric induced by R_σ as in Definition 3.6 by $d_{\tilde{X}, \sigma}$.

On the other hand, since $\tilde{X} \subset G$, restricting the graph metric d_G to \tilde{X} give rise to another metric on \tilde{X} . The main goal of this section is to provide a uniform upper bound for $d_{\tilde{X}, \sigma}(x, y) - d_G(x, y)$. Similar analysis on smooth manifolds of dimension higher than 1 is conducted in [Bernstein 2000]. We in this section study the 1-dimensional case, which requires some different treatments.

3.7.1 Regularity Conditions.

The graph metric is intrinsic and independent of the ambient space. The sampling points, on the other hand, carry no information about the intrinsic metric. This fact arises difficulties in approximating graph metric with the sampling points. We illustrate a typical difficulty with the three singular cases of graphs embedded in \mathbb{R}^2 in Figure 3.8. In each case, the red point and the blue point are distant in the graph but close in \mathbb{R}^2 . Without prior knowledge about the graph metric, we probably end up creating a shortcut from the red points to the blue ones, therefore introduce significant distortions in metric approximation.

If the red point and the blue point are arbitrarily close to each other in the ambient space while keeping distant in the graph, then by no means can we avoid introducing the distortions. From this point of view, we propose some regularity conditions to exclude the singular situations.

To do that, we first introduce the *reach* of a manifold M embedded in \mathbb{R}^d . Given a point $x \in \mathbb{R}^d$ and a set $A \subset \mathbb{R}^d$, the distance from x to A is defined as the infimum of the distances from x to points in A , i.e., $d(x, A) = \inf_{y \in A} \|x - y\|$. The medical axis of M , denoted by $med(M)$, is defined as the following:

Definition 3.10 *The medical axis of M is the set of points which have more than one nearest neighborhoods in M , i.e.,*

$$med(M) = \{x \in \mathbb{R}^d : \text{there exist at least two points } y_1, y_2 \in M \text{ such that } d(x, M) = \|x - y_1\| = \|x - y_2\|\}$$

And finally we define the reach of M :

Definition 3.11 *The reach of M is the infimum of the distances from a point in M to the medical axis of M :*

$$reach(M) = \inf_{p \in M} d(p, med(M))$$

It's routine to assume that $reach(M)$ is lower bounded by a positive constant in the study of approximating intrinsic geometric quantities with sampling points from the ambient space, however in the case of metric graph this condition is too restrictive. For example, the reach of the second and the third graphs in Figure 3.8 are zero due to the existences of sharp corners at the green points. Thus instead of bounding the global reach of G , we bound the reaches of some specified subgraphs of G from below.

Let V_d be the vertex set of G (i.e. vertices of degree not equal to 2) and V_l be the middle points of self-loops in G . We call V_d the original vertex set of G and $V_a = V_d \cup V_l$ the augmented vertex set of G . We break the self-loops into edges from the middle points of the loop, i.e., the points in V_l . Formally, the augmented edge set E_a is defined as

$$E_a = \bigcup_i \{(v_i, v_j) : v_i, v_j \in V_a \text{ and there is no other vertex in the geodesic connecting } v_i \text{ to } v_j.\}$$

We regularize a metric graph G embedded in \mathbb{R}^d with the following conditions on the augmented vertex set V_a and the augmented edge set E_a of G . These conditions depend on three positive parameters b , τ and β , whose roles are clear in the context.

Condition 3.1 1. *For any edge $e \in E_a$, the reach of e is not smaller than τ , where τ is a positive constant.*

2. *For any two edges $e_1, e_2 \in E_a$, if $e_1 \cap e_2 = \emptyset$, then $reach(e_1 \cup e_2) \geq \tau$.*

3. *For any $v \in V_a$, let $V(v) = \{v\} \cup \{w \in V_a : (v, w) \in E_a\}$ and $E(v) = \{(v, w) : w \in V(v), w \neq v\}$ (see figure 3.9). Denote the geodesic ball of radius r centered at v by $B(v, r) = \{g \in G : d_G(g, v) \leq r\}$. For any two edges in $e_1, e_2 \in E(v)$, $reach[(e_1 \cup e_2) \setminus \bigcup_{w \in e_1 \cap e_2} B(w, r)] \geq \beta r$ holds for some positive constant $\beta \in (0, 1]$.*

4. *No edges in E_a are shorter than b , where b is a positive constant.*

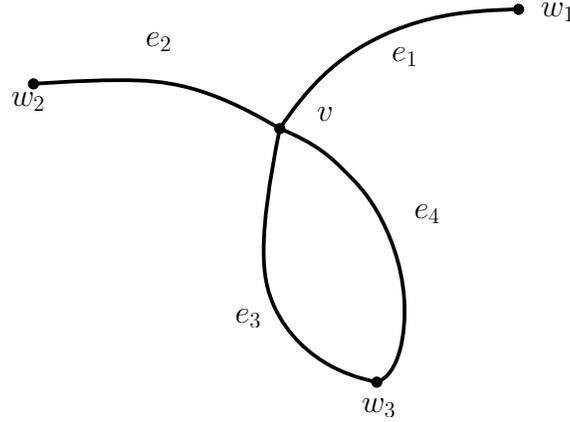


Figure 3.9: $V(v) = \{w_1, w_2, w_3\}$, $E(v) = \{e_1, e_2, e_3, e_4\}$

The first three conditions exclude the singular cases illustrated in Figure 3.8 with sufficiently large τ and β . First, thanks to condition 1, the reach of a single edge is lower bounded, meaning that if two points are distant in graph then they are distant in the ambient space. Secondly, condition 2 assures two non-intersecting edges is well-separated in Euclidean space, and therefore the case in the middle is prevented. Lastly, condition 3 takes care of the singularity between two incident edges such as the case in the right panel of Figure 3.8.

3.7.2 Upper-bound of $|d_G - d_H|$.

Instead of directly comparing $d_{\tilde{X},\sigma}$ and d_G , we further construct two more metrics on \tilde{X} as auxiliaries. We start with constructing a graph $H = (\tilde{X}, E_H)$ on \tilde{X} with parameter $\delta > 0$ as the following:

$$E_H = \{(x_i, x_j) : d_G(x_i, x_j) \leq \delta, x_i, x_j \in \tilde{X}\} \quad (3.1)$$

Then two metrics d_H, d_S are defined on top of E_H with respect to two measurements on the edges.

$$d_H(x, y) = \min_Q \sum_{i=2}^k \|x_{i-1} - x_i\|$$

$$d_S(x, y) = \min_Q \sum_{i=2}^k d_G(x_{i-1}, x_i)$$

where $Q = (x_1, x_2, \dots, x_k)$ varies over all paths along E_H connecting $x = x_1$ to $y = x_k$.

In the following, we present a uniform upper bound for $|d_G(x, y) - d_H(x, y)|$. And then in Section 3.7.3, a uniform upper bound for $|d_H(x, y) - d_{\tilde{X},\sigma}(x, y)|$ is given. Combining them we arrive at an upper bound for $|d_G(x, y) - d_{\tilde{X},\sigma}(x, y)|$.

Since $|d_H(x, y) - d_G(x, y)| \leq |d_S(x, y) - d_G(x, y)| + |d_H(x, y) - d_S(x, y)|$, we estimate $|d_H(x, y) - d_G(x, y)|$ in two steps. In step 1, we asserts that if the sampling density is high enough (or equivalently ε is small enough), then $d_S = d_G$.

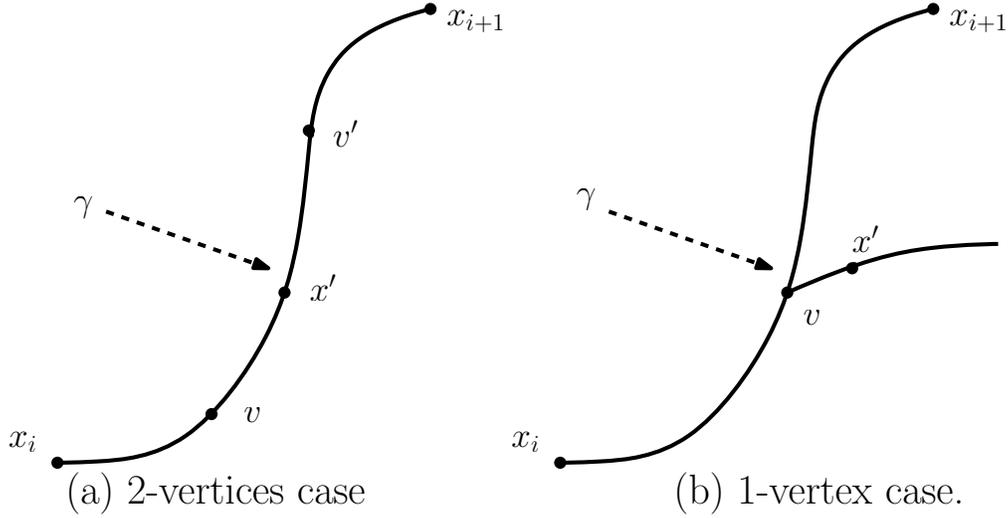


Figure 3.10: Illustration of Lemma 3.7

Proposition 3.7 Let G be a metric graph satisfying condition 3.1, and \tilde{X} be an ε -sampling of G . V_a, E_a are the augmented vertex set and augmented edge set of G . The lengths of the augmented edges in G are lower bounded by b as defined in conditions 3.1. If parameter δ is such that $b > \delta \geq 4\varepsilon$, then $d_S(x, y) = d_G(x, y)$, $\forall x, y \in \tilde{X}$.

Proof: It follows from the triangle inequality that $d_S(x, y) \geq d_G(x, y)$. Thus it remains to prove $d_G(x, y) \geq d_S(x, y)$.

Let λ be the geodesic from x to y in G , we prove that λ is identical to a connected path in H . We sort points in $\lambda \cap \tilde{X} = (x, x_2, x_3, \dots, y)$ so that $d_G(x, x) \leq d_G(x, x_2) \leq d_G(x, x_3) \leq \dots \leq d_G(x, y)$. We claim that for any two consecutive points x_i, x_{i+1} , $d_G(x_i, x_{i+1}) \leq \delta$, i.e., x_i and x_{i+1} are connected in H .

Denote by γ the geodesic between x_i, x_{i+1} . If there are two vertices, say $v, v' \in V_a$, be in the interior of γ (see Figure 3.10(a)). Then $d_G(x_i, x_{i+1}) \geq d_G(v, v') \geq b > \delta \geq 4\varepsilon$. As a consequence, there must be another sample x' between v, v' : otherwise let g be the middle point of the geodesic between v and v' , and the nearest sampling point to g is of distance larger than $\frac{b}{2} > 2\varepsilon$, violating the assumption that \tilde{X} is an ε -sampling. On the other hand, according to the construction of (x, x_2, \dots, y) , there is not another sampling point between the consecutive points x_{i-1}, x_i , and therefore there is at most one vertex in V_a in the interior of γ .

If there exists exactly one vertex of V_d (of degree larger than 2) between x_i and x_{i+1} . As mentioned above, there is not a sampling point in the interior γ . Let x' be the sampling point that is closest to v and not in the interior of γ (as depicted in Figure 3.10(b)). With the same argument before, $d_G(x_i, x') \leq 2\varepsilon$ and $d_G(x_{i+1}, x') \leq 2\varepsilon$, meaning that $d_G(x_i, x_{i+1}) \leq 4\varepsilon \leq \delta$.

If there is no vertex in V_d between x_i and x_{i+1} . Let g be the middle point of geodesic from x_i to x_{i+1} . Since \tilde{X} is an ε -sampling, $d_G(x_i, x_{i+1}) \leq d_G(g, x_i) + d_G(g, x_{i+1}) \leq 2\varepsilon < \delta$.

Noticing that (x, x_2, x_3, \dots, y) is a path in graph H connecting x to y , we have $d_S(x, y) \leq d_G(x_1, x_2) +$

$d_G(x_2, x_3) + \dots + d_G(x_{k-1}, x_k) = d_G(x_1, x_k) = d_G(x, y)$, which finishes the proof. \square

Then we proceed to bound $|d_H(x, y) - d_S(x, y)|$. A lemma by Herbert Federer [Federer 1959] as the following is repeatedly used in our analysis.

Lemma 3.12 [Federer] *Let K be a manifold embedded in \mathbb{R}^d , and $\text{reach}(K) = \tau > 0$. Let K^δ be a δ -offset of K , i.e., $K^\delta = \bigcup_{x \in K} \{y \in \mathbb{R}^d : \|x - y\| \leq \delta\}$. For any $\delta < \tau$, the map $\Pi_\delta : K^\delta \rightarrow K$ is $\frac{\tau}{\tau - \delta}$ -lipschitz. (Π_δ maps each point in K^δ to its projection in K .)*

Applying Federer's lemma to a sub-graph whose reach is bounded from below by τ , we have the following corollary.

Corollary 3.2 *Let K be a geodesically convex sub-graph of G , $\text{reach}(K) = \tau > 0$. Then for any two points x, y on K , if $\|x - y\| = l < 2\tau$, then $d_G(x, y) \leq \frac{\tau}{\tau - \frac{l}{2}} \|x - y\|$.*

Proof: We consider a parameterization of the straight line segment from x to y in \mathbb{R}^d such that $s(0) = x, s(l) = y$. The whole segment $s([0, l])$ is contained in $K^{\frac{l}{2}}$ since that for any point p in this segment, $\|p - x\| + \|p - y\| = l$. Now let a partition of $[0, l]$ be $0 = t_0 < t_1 < \dots < t_n = l$, with Federer's lemma we have $\|\Pi_{\frac{l}{2}}(s(t_{i+1})) - \Pi_{\frac{l}{2}}(s(t_i))\| \leq \frac{\tau}{\tau - \frac{l}{2}} \|t_{i+1} - t_i\|$. Adding up both side of the inequality, we obtain $d_G(x, y) \leq \frac{\tau}{\tau - \frac{l}{2}} \|x - y\|$. \square

Proposition 3.8 *Let G be a metric graph satisfying condition 3.1 and τ, b be the constants involved in the condition. Let $\text{diam}(G)$ be the diameter of graph G , i.e., $\text{diam}(G) = \max_{g_1, g_2 \in G} d_G(g_1, g_2)$. $|V_d|$ is the cardinality of the original vertex set of G . d_S and d_G are metrics induced by graph H with parameter δ as in Definition 3.1. If $\delta < \min\{\tau, b\}$, then we have*

$$0 \leq d_S(x, y) - d_H(x, y) \leq \frac{\delta}{2\tau} \text{diam}(G) + (2 + \sum_{v \in V_d} (\deg(v) - 1))\delta$$

Proof: Let $P = (x_1, \dots, x_k)$ be a path in H connecting two points $x = x_1$ and $y = x_k$ such that $d_H(x, y)$ is realized by P . We can further assume that P is a simple path in H . It follows from the definition of d_H that the path realizing the minimal length between x and y must be simple, otherwise one can remove loops without increasing the path lengths.

Let γ_i be the geodesic connecting x_{i-1} and x_i in G . Following the construction of H , the length of γ_i is not more than δ . We first claim that there must not be any point x_j with $j \geq i + 1$ in γ_i (including the end points x_{i-1}, x_i): otherwise, let x_j be on γ_i , then $d_G(x_{i-1}, x_j) \leq d_G(x_{i-1}, x_i) \leq \delta$, meaning that we can directly connect x_{i-1}, x_j to reduce the length of P . (see Figure 3.11(a) for an illustration). That contradicts the optimality of P .

A corollary of this observation is that, a point $v \in V_d$ (the vertex set of G with degree larger than 2) is contained in at most $\deg(v) - 1$ pieces of geodesic γ_i . Particularly we analyze the case of $\deg(v) = 3$ as shown in Figure 3.11(b). Let γ_i be the first geodesic containing v , and then we consider the position of x_{i+1} : it can't be in γ_i as proven before, moreover it can't be in the edge containing x_{i-1} and v , otherwise it will be

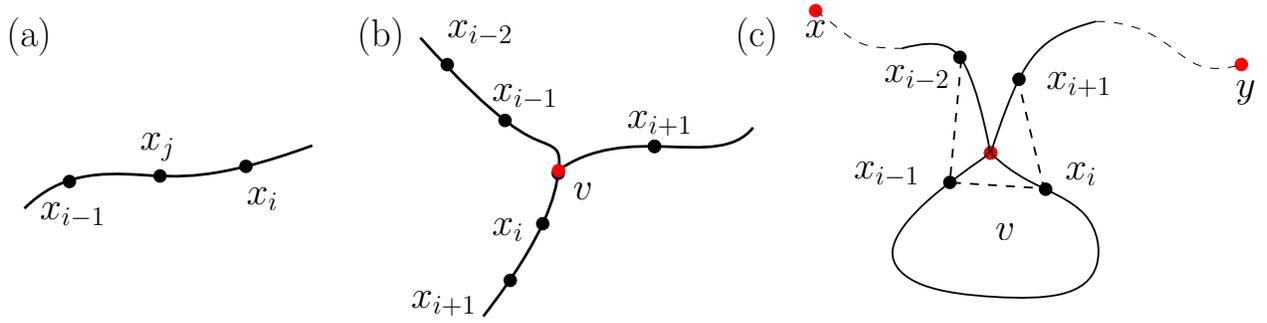


Figure 3.11: Illustrations for the proof of Proposition 3.8.

in the geodesic of some $\gamma_k, k \leq i - 1$. thus it either jumps to the third branch or stays in the same edge with x_i and $d_G(x_{i+1}, v) > d_G(x_i, v)$.

- If it's the former case, then v is contained by 2 geodesics: γ_i and γ_{i+1} . We claim that v can't be contained by $\gamma_j, j \geq i + 2$. That is because that neither of x_{j-1}, x_j can be in the *tripod* formed by x_{i-1}, x_i, x_{i+1} and v . thus γ_j must contains one of the three geodesics: γ_i, γ_{i+1} and the one connecting x_{i-1}, x_{i+1} . In either case we can show that we can connect one of x_{i-1}, x_i, x_{i+1} with one of x_j, x_{j-1} , and again that contradicts the optimality of P .
- If it's the latter case, then similarly we can prove that v can't be crossed by $\gamma_j, j \geq i + 1$.

For the case $\deg(v) > 3$, the analysis is the same (see Figure 3.11(c) for a situation that $\deg(v) = 4$ and v is crossed by 3 geodesics).

We classify all the segments $(x_{i-1}, x_i), i = 2, 3, \dots, k$ into three types. Note that since (x_{i-1}, x_i) are connected in H , $d_G(x_{i-1}, x_i) \leq \delta < b$. Thus it follows the same arguments in the proof of Proposition 3.9 that $\gamma_i, \forall i$ can not pass more than 1 vertices in the augmented vertex set V_a .

- Type A: There is no vertex in γ_i .
- Type B: There is exactly one vertex of V_d , the original vertex set of G , in γ_i .
- Type C: There is exactly one vertex of V_l , the set of middle points of self-loops in G , in γ_i .

For (x_{i-1}, x_i) in type A, we have $d_G(x_{i-1}, x_i) \leq \delta < \tau$. Following lemma 3.2, we have $d_G(x_{i-1}, x_i) \leq \frac{\tau}{\tau - \frac{\delta}{2}} \|x_{i-1} - x_i\|$, meaning that $d_H(x_{i-1}, x_i) \leq d_S(x_{i-1}, x_i) \leq \frac{\tau}{\tau - \frac{\delta}{2}} d_H(x_{i-1}, x_i)$.

For (x_{i-1}, x_i) in type B, we have $\delta \geq d_G(x_{i-1}, x_i) \geq \|x_{i-1} - x_i\| \geq 0$, thus $0 \leq d_G(x_{i-1}, x_i) - \|x_{i-1} - x_i\| \leq \delta$. Since the simple path can't pass v more than $\deg(v) - 1$ times, the number of segments of type B is then upper-bounded by $\sum_{v \in V_d} (\deg(v) - 1)$.

In the last case, since P is a simple path in H connecting x and y , P at most pass two vertices in V_l . Notice that each vertex in V_l corresponds to a self-loop of G , if there are 3 vertices in V_l on the geodesic

between x and y , then there must be at least one self-loop contains neither x nor y , meaning that P is not simple. To conclude, there are at most two segments falling into type C, for each of them, say (x_j, x_{j+1}) , $0 < d_G(x_j, x_{j+1}) - \|x_j - x_{j+1}\| \leq \delta$

Lastly, since for any path $d_H(P) \leq d_S(P)$, $d_H(x, y) \leq d_S(x, y)$, $\forall x, y \in \tilde{X}$. On the other hand, let $P = P_A \cup P_B \cup P_C$ be a decomposition of path P regarding the above classification, then $d_S(P_A) - d_H(P_A) \leq d_S(P_A) - \frac{\tau - \delta/2}{\tau} d_S(P_A) \leq \frac{\delta}{2\tau} d_S(P) \leq \frac{\delta}{2\tau} d_G(x, y) \leq \frac{\delta}{2\tau} \text{diam}(G)$. Similarly $d_S(P_B) - d_H(P_B) \leq \sum_{v \in V_d} (\text{deg}(v) - 1)\delta$, and $d_S(P_C) - d_H(P_C) \leq 2\delta$

Putting them together, we have

$$0 \leq d_S(x, y) - d_H(x, y) \leq \frac{\delta}{2\tau} \text{diam}(G) + (2 + \sum_{v \in V_d} (\text{deg}(v) - 1))\delta$$

□

Combing Proposition 3.7 and 3.8, we prove the following:

Proposition 3.9 *Let G be a metric graph embedded in \mathbb{R}^d satisfying condition 3.1, and \tilde{X} be an ε -sampling of G . Let H be a graph constructed with parameter δ as in Definition 3.1. $\text{diam}(G)$ is the diameter of G and V_d is the original vertex set of G . τ, b are parameters defined in condition 3.1. If $\min\{\tau, b\} > \delta \geq 4\varepsilon$, then*

$$|d_G(x, y) - d_H(x, y)| \leq \frac{\delta}{2\tau} \text{diam}(G) + (2 + \sum_{v \in V_d} (\text{deg}(v) - 1))\delta, \forall x, y \in \tilde{X}$$

.

3.7.3 Upper-bound of $|d_H - d_{\tilde{X}, \sigma}|$

Fixing $\delta = 4\varepsilon$, it follows from Proposition 3.9 that if ε tends to 0, then $d_H(x, y)$ converges to $d_G(x, y)$ for any $x, y \in \tilde{X}$. However in the practice, we can't construct graph H directly since the metric d_G is unknown. In the following, we prove that if G satisfies condition 3.1, then we can approximate d_G with metric of the Rips graph constructed on \tilde{X} with a proper parameter.

Before proceeding, we claim the following lemma which are used repeatedly in the proof of the upcoming theorem.

Lemma 3.13 *Let x, y be two points lying on the same edge $e \in E$ whose reach is at least $\tau > \delta$. If $\|x - y\| \leq \delta/2$, then $d_G(x, y) \leq \delta$.*

Proof: This lemma is essentially a rephrase of corollary 3.2, e is obviously a geodesically convex sub-graph of G , thus if $\|x - y\| \leq \delta/2$, $d_G(x, y) \leq \frac{\tau}{\tau - \delta/4} \|x - y\| \leq \frac{2\delta}{3} < \delta$. □

Theorem 3.8 *G is a metric graph satisfying conditions 3.1, \tilde{X} is an ε -sampling of G , δ is a positive parameter such that $\min\{b, \tau\} > \delta$, where b, τ and β below are constants defined in conditions 3.1.*

Then for any two points x, y in X , if $\|x - y\| \leq \frac{\beta}{2\beta + 2}\delta$, then $d_G(x, y) \leq \delta$.

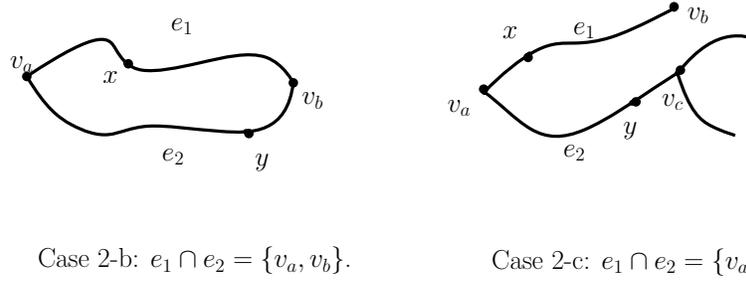


Figure 3.12: Illustration of case 2(b) and 2(c).

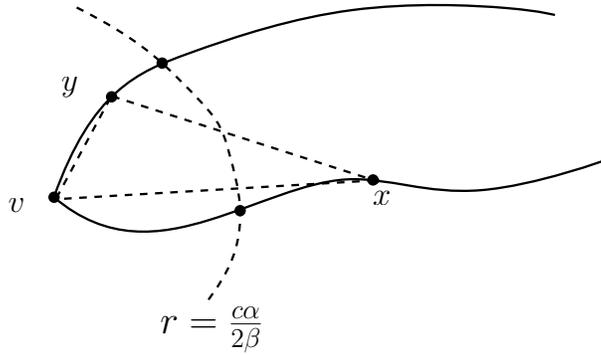


Figure 3.13: Illustration of the last scenario of case-2(b).

Proof: For the sake of clarity, in the proof we let $c = \frac{\beta}{2\beta+2} < 1/2$. We assume that x, y are on the edges $e_1, e_2 \in E_a$, where E_a is the augmented edge set of G , respectively.

The problem is analyzed on the following several cases:

1. $e_1 \cap e_2 = \emptyset$.

If $e_1 \cap e_2 = \emptyset$, we claim $\|x - y\| > c\delta$: due to condition 3.1, $reach(e_1 \cup e_2) \geq \tau > \delta$. We let $m = \inf_{x_1 \in e_1, x_2 \in e_2} \|x_1 - x_2\|$ and assume that $x'_1 \in e_1, x'_2 \in e_2$ satisfy $\|x'_1 - x'_2\| = m$ (the existences are assured by that $e_1 \cup e_2$ is compact). The interior of the Euclidean ball centered at the middle point between x'_1 and x'_2 with radius $\frac{m}{2}$ is not intersecting with $e_1 \cup e_2$, otherwise we can find another pair of points whose Euclidean distance is smaller than m , contradicting to the assumption on (x'_1, x'_2) . Therefore the center of the ball, r_c , must be on the medial axis, meaning that $\|x - y\| \geq \|x'_1 - x'_2\| = \|x'_1 - r_c\| + \|x'_2 - r_c\| \geq \tau + \tau = 2\tau > 2\delta > c\delta$.

2. $e_1 \cap e_2 \neq \emptyset$.

We first deal with the case in which e_1 and e_2 are identical, the latter two cases are illustrated in Figure 3.12.

To avoid heavy notations, we denote $e_1 \cup e_2$ by U and define $B_a = \{g \in G : d_G(g, v_a) \leq r\}$ where

$r = \frac{c\delta}{2\beta}$ (similarly define B_b and B_c with respect to v_b and v_c as illustrated in Figure 3.12).

(a) $e_1 = e_2$.

Equivalently x, y are on the same edge with positive reach larger than τ and $\|x - y\| \leq c\delta$. Directly by corollary 3.2 we know that $d_G(x, y) \leq \frac{\tau}{\tau - \|x - y\|/2} \|x - y\| \leq \frac{\tau}{\tau - c\delta/2} c\delta \leq \frac{c}{1 - c/2} \delta \leq \frac{2}{3} \delta < \delta$.

(b) $e_1 \cap e_2 = \{v_a, v_b\}$.

First, if $x, y \in U \setminus (B_a \cup B_b)$, then it follows from condition 3.1 that $\|x - y\| > 2\beta r = c\delta$.

Secondly if $x, y \in B_a \cup B_b$: either x, y are in the same geodesic ball, say B_a , then $d_G(x, y) \leq 2\frac{c\delta}{2\beta} = \frac{c\delta}{\beta} \leq \frac{\delta}{2} < \delta$; or $x \in B_a$ and $y \in B_b$, then given $d_G(v_a, v_b) \geq b > \delta$, $\|v_a - v_b\| > \delta/2$.

Thus $\|x - y\| \geq \|v_a - v_b\| - \|x - v_a\| - \|y - v_b\| > \delta/2 - \frac{c\delta}{\beta} = c\delta$.

Lastly, without loss of generality, we let $x \in U \setminus (B_a \cup B_b)$ and $y \in B_a$. If $d_G(x, v_a) > \delta$, by lemma 3.13, $\|x - v_a\| > \delta/2$. Obviously $\|v_a - y\| \leq r = \frac{c\delta}{2\beta}$, so $\|x - y\| \geq \|x - v_a\| - \|v_a - y\| > \delta/2 - \frac{c\delta}{2\beta} > c\delta$.

Now we assume $d_G(x, v_a) \leq \delta$ and $d_G(x, v_a) + d_G(v_a, y) > \delta$. Following corollary 3.2 one has $\|x - v_a\| \geq \frac{\tau - d_G(x, v_a)/2}{\tau} d_G(x, v_a) \geq \frac{\tau - \delta/2}{\tau} (\delta - d_G(v_a, y)) \geq \frac{\tau - \delta/2}{\tau} (\delta - \frac{c\delta}{2\beta})$. On the other hand $\|y - v_a\| \leq \frac{c\delta}{2\beta}$.

Then it amounts to proving

$$\frac{\tau - \delta/2}{\tau} (\delta - \frac{c\delta}{2\beta}) - \frac{c\delta}{2\beta} > c\delta.$$

Reorganizing, we have the following equivalent form:

$$\frac{\tau - \delta/2}{\tau} > (\frac{\tau - \delta/2}{2\beta\tau} + \frac{1}{2\beta} + 1)c.$$

The left-hand-side of the above inequity is not smaller than $(\delta - \delta/2)/\delta = 1/2$ (as $\tau \geq \delta$), while the right-hand-side is smaller than $(1 + 1/\beta)c = 1/2$, thus the inequality is verified, meaning that $\|x - y\| \geq \|x - v_a\| - \|y - v_a\| > c\delta$ given $d_G(x, y) > \delta$.

(c) $e_1 \cap e_2 = v_a$. First if $x, y \in U \setminus B_a$, then by the first regular condition, $\|x - y\| > 2\beta r = c\delta$.

Secondly if $x, y \in B_a$, $d_G(x, y) \leq 2r = \frac{c\delta}{\beta} < \delta/2 < \delta$.

The last scenario is $x \in U \setminus B_a$ and $y \in B_a$, which is obviously identical to the last scenario in the case of $e_1 \cap e_2 = \{v_a, v_b\}$.

□

Equipped with the above propositions, we conclude the main theorem of this section as the following:

Theorem 3.9 *If metric graph G satisfies conditions 3.1, \tilde{X} is an ε -sampling of G , if there exists a parameter σ such that $4\varepsilon \leq \sigma < \frac{\beta}{2\beta+2} \min\{\tau, b\}$. Then the Rips graph with parameter σ induces a metric on \tilde{X} , $d_{\tilde{X}, \sigma}$, as in Definition 3.5 satisfying:*

$$0 \leq d_G(x, y) - d_{\tilde{X}, \sigma}(x, y) \leq \frac{(2 + 2/\beta)\sigma}{2\tau} \text{diam}(G) + (\sum_{v \in V_d} (\text{deg}(v) - 1) + 2)(2 + 2/\beta)\sigma$$

Proof: Let $\delta = (2 + 2/\beta)\sigma$, and build a graph H with δ as in Definition 3.1. Then δ satisfies conditions in Proposition 3.9 thanks to $\delta = (2 + 2/\beta)\sigma < (2 + 2/\beta)\frac{\beta}{2\beta+2} \min\{\tau, b\} = \min\{\tau, b\}$ and $\delta > \sigma \geq 4\varepsilon$. Therefore we have $|d_G(x, y) - d_H(x, y)| \leq \frac{\delta}{2\tau} \text{diam}(G) + (\sum_{v \in V_d} (\text{deg}(v) - 1) + 2)\delta$.

We argue that R_σ is a subgraph of H . In fact, if x_i, x_j are connected in R_σ , then $\|x_i - x_j\| \leq \sigma$, according to Theorem 3.8, $d_G(x_i, x_j) \leq (2 + 2/\beta)\sigma = \delta$, therefore x_i, x_j are connected in H .

The last step is to prove that $d_{\tilde{X}, \sigma}(x, y) \geq d_H(x, y)$. Indeed if x, y are connected in G , then following the proof of Proposition 3.7 we have x, y are connected with a path $P = (x_1, x_2, \dots, x_k)$ in H where $d_G(x_i, x_{i+1}) \leq 4\varepsilon, \forall 1 \leq i \leq k - 1$. On the other hand, since $\|x_i - x_{i+1}\| \leq d_G(x_i, x_{i+1}) \leq 4\varepsilon \leq \sigma$, P is also a path in R_σ , therefore x, y are connected in both H and R_σ . Now that $R_\sigma \subset H$, any path connecting them in R_σ must connect them in H . Following the definitions of $d_{\tilde{X}, \sigma}$ and d_H , we have $d_{\tilde{X}, \sigma}(x, y) \geq d_H(x, y)$. Putting them together, we finish the proof. \square

Gromov-Hausdorff Distance between (G, d_G) and $(\tilde{X}, d_{\tilde{X}, \sigma})$. As a corollary, we derive a bound for the Gromov-Hausdorff distance between (G, d_G) and $(\tilde{X}, d_{\tilde{X}, \sigma})$. To start with, we construct a correspondence between points the two metric spaces: $C \subset G \times \tilde{X}$. For $x \in \tilde{X} \subset G$, we let $(x, x) \in C$. For each $g \in G$, we let $x(g) = \{x \in \tilde{X} \text{ s.t. } d_G(x, g) = \min_{y \in \tilde{X}} d_G(y, g)\}$ and add $(g, x(g))$ to C .

For any g_1, g_2 , since \tilde{X} is an ε -sampling of G , we have $d_G(g_i, x(g_i)) \leq \varepsilon, i = 1, 2$. It follows from Theorem 3.9 that

$$\begin{aligned} & |d_G(g_1, g_2) - d_{\tilde{X}, \sigma}(x(g_1), x(g_2))| \\ & \leq |d_G(g_1, g_2) - d_G(x(g_1), x(g_2))| + |d_G(x(g_1), x(g_2)) - d_{\tilde{X}, \sigma}(x(g_1), x(g_2))| \\ & \leq |d_G(g_1, g_2) - d_G(g_1, x(g_2))| + |d_G(g_1, x(g_2)) - d_G(g(x_1), x(g_2))| \\ & + \frac{(2 + 2/\beta)\sigma}{2\tau} \text{diam}(G) + \sum_{v \in V_d} (\text{deg}(v) - 1)(2 + 2/\beta)\sigma \\ & \leq d_G(g_2, x(g_2)) + d_G(g_1, x(g_1)) + \frac{(2 + 2/\beta)\sigma}{2\tau} \text{diam}(G) + (\sum_{v \in V_d} (\text{deg}(v) - 1) + 2)(2 + 2/\beta)\sigma \\ & \leq 2\varepsilon + \frac{(2 + 2/\beta)\sigma}{2\tau} \text{diam}(G) + (\sum_{v \in V_d} (\text{deg}(v) - 1) + 2)(2 + 2/\beta)\sigma \end{aligned}$$

After all, we obtain that if the assumptions of Theorem 3.9 are satisfied, then

$$d_{GH}(\tilde{X}, G) \leq \varepsilon + \frac{(1 + 1/\beta)\sigma}{2\tau} \text{diam}(G) + (\sum_{v \in V_d} (\text{deg}(v) - 1) + 2)(1 + 1/\beta)\sigma$$

This bound of the Gromov-Hausdorff distance between (G, d_G) and $(\tilde{X}, d_{\tilde{X}, \sigma})$ suggests two facts that impact the quality of approximation: one is the sampling density and the other is the regularity of the underlying graph G . Higher sampling density (or equivalently smaller ε) allows for a smaller σ : as the bound suggests, if ε tends to zero, we are guaranteed to recover the underlying metric d_G . On the other hand, the larger β and τ are, the smaller the bound is. Stated differently, we need denser sampling for approximating the metric of a graph with less regularity.

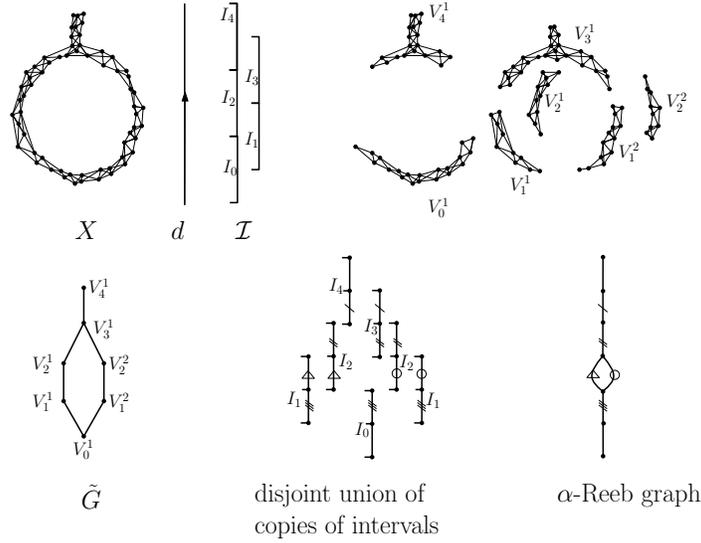


Figure 3.14: Illustration of the different steps of the algorithm for computing α -Reeb graph. In the disjoint union of copies of intervals, the subintervals marked with same labels are identified in the α -Reeb graph.

3.8 Algorithm

In this section, we describe an algorithm for computing the α -Reeb graph for some $\alpha > 0$. We assume the input of the algorithm includes a neighboring graph $X = (V, E)$, a function $l : E \rightarrow \mathbb{R}^+$ specifying the edge length and a parameter α . In the applications where the input is given as a set of points together with pairwise distances, i.e., a finite metric space, one can generate the neighboring graph X as a Rips graph of the input points with the parameter chosen as a fraction of α . We assume X is connected as one can apply the algorithm to each connected component otherwise.

Our algorithm, whose different steps are illustrated in Figure 3.14, can be described as follows. In the first step, we fix a node of X as the root r and then obtain the distance function $d : V \rightarrow \mathbb{R}^+$ by computing $d(v)$ as the graph distance from the node v to r . In the second step, we apply the Mapper algorithm [Singh 2007] to the nodes V with filter d to construct a graph \tilde{G} . Specifically, let $\mathcal{I} = \{(i\alpha, (i+1)\alpha), ((i+0.5)\alpha, (i+1.5)\alpha) \mid 0 \leq i \leq m\}$ so that $\bigcup_{I_k \in \mathcal{I}} I_k$ covers the range of the function d . We say an interval $I_{k_1} \in \mathcal{I}$ is lower than another interval $I_{k_2} \in \mathcal{I}$ if the midpoint of I_{k_1} is smaller than that of I_{k_2} . Now let $V_k = d^{-1}(I_k)$ and V_k^l be the l th component of V_k . Then $\{V_k^l\}_{k,l}$ is a cover of H and the graph \tilde{G} constructed by the Mapper algorithm is the 1-skeleton of the nerve of that cover. Namely, each node in \tilde{G} represents an element in $\{V_k^l\}_{k,l}$. Two nodes $V_{k_1}^{l_1}$ and $V_{k_2}^{l_2}$ are connected with an edge if $V_{k_1}^{l_1} \cap V_{k_2}^{l_2} \neq \emptyset$. In fact, when we check if $V_{k_1}^{l_1} \cap V_{k_2}^{l_2} \neq \emptyset$, we only need to check if their vertices are overlapped or not as we assume the lengths of the edges in H are fractions of α .

In the final step, we represent each node V_k^l in \tilde{G} using a copy of the interval I_k . As mentioned in the Section 3.4, α -Reeb graph is a quotient space of the disjoint union of those copies of intervals. Specifically, for an edge in \tilde{G} , let $V_{k_1}^{l_1}$ and $V_{k_2}^{l_2}$ be its endpoints. Then I_{k_1} and I_{k_2} must be partially overlapped. We

identify the overlap part of these two intervals. After identifying the overlapped intervals for all edges in \tilde{G} , the resulting quotient space is the α -Reeb graph. Algorithmically, the identification is performed as follows. We split each copy of interval I_k into two by adding a point in the middle. Now think of it as a graph with two edges and label one of them upper and the other lower. Notice that two overlapped intervals I_{k_1} and I_{k_2} can not be exactly the same. One must be lower than the other. To identify their overlapped part, we identify the upper edge of the lower interval with the lower edge of the upper interval.

The time complexity of the above algorithm is dominated by the computation of the distance function in the first step, which is $O(|E| + |V| \log |V|)$. The computation of the connected components in the second step is $O(|V| \log |V|)$ based on union-find data structure. In the final step, there are at most $O(|V|)$ number of the copies of the intervals. Based on union-find data structure, the identification can also be performed in $O(|V| \log |V|)$ time.

3.9 Experimental Results

In this section, we illustrate the performances of our algorithm in different applications.

3.9.1 Earthquake Data

The first data set was obtained from USGS Earthquake Search [[EarthquakeSearch](#)]. It consists of earthquake epicenter locations collected, between 01/01/1970 and 01/01/2010 in the rectangular area between latitudes -75 degrees and 75 degrees and longitude -170 degrees and 10 degrees with magnitude greater than 5.0.

This raw earthquake data set contains the coordinates of the epicenters of 12790 earthquakes that are mainly located around geological faults. We follow the procedure described in [[Aanjaneya 2012](#)] to remove outliers and randomly sampled 1600 landmarks. Finally, we computed a neighboring graph from these landmarks with parameter 4. The length of an edge in this graph is the Euclidean distance between its endpoints. For each connected component, we fix a root point and compute the graph distance function d to the root point as shown in Figure 3.15(a). We also set $\alpha = 4$ and apply our algorithm to the above data to obtain the α -Reeb graph. In general, the α -Reeb graph is an abstract metric graph. In this example, for the purpose of visualization, we use the coordinates of the landmarks to embed the graph into the plane as follows. Recall that for a copy of interval I_k representing the node V_k^l in \tilde{G} , we split it into two by adding a point in the middle. We embed the endpoints of the interval to the landmarks of the minimum and the maximum of the function d in V_k^l , and the point in the middle to the landmark of the median of the function d in V_k^l . Figure 3.15(b) shows the embedding of the α -Reeb graph. Note this embedding may introduce metric distortion, i.e., the Euclidean length of the edge may not reflect the length of the corresponding edge in the α -Reeb graph.

3.9.2 GPS Data

The second data set is that of 500 GPS traces tagged “Moscow” from OpenStreetMap [[Openstreetmap](#)]. Since cars move on roads, we expect the locations of cars to provide information about the metric graph structure of the Moscow road network. We first selected a metric ε -net on the raw GPS locations with $\varepsilon = 0.0001$ using furthest point sampling. Then, we computed a neighboring graph from the samples with

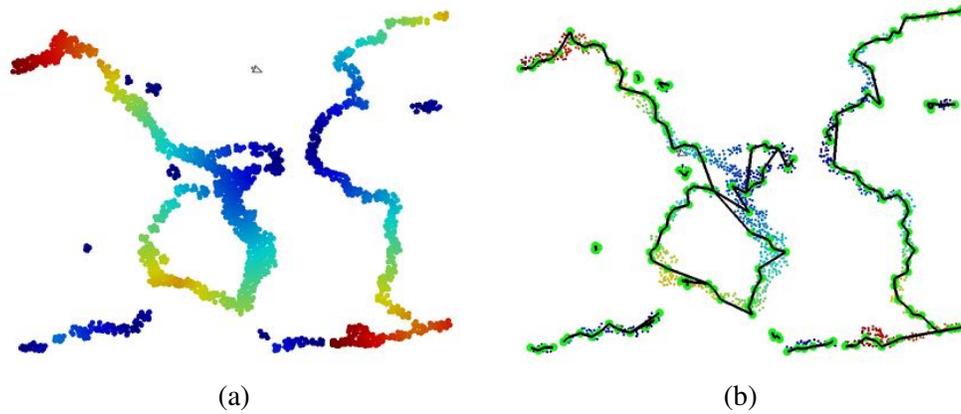


Figure 3.15: Earthquake data - (a) The distance functions d on each connected components. The value increases from cold to warm colors. (b) The reconstructed α -Reeb graph.



Figure 3.16: GPS data - (a) The distance functions d on each connected components. The value increases from cold to warm colors. (b) The reconstructed α -Reeb graph.

parameter 0.0004. Again for each connected component, we fix a root point and compute the graph distance function d to the root point as shown in Figure 3.16(a). Set α also equals 0.0004 and compute the α -Reeb graph. Again, we use the same method as above to embed the α -Reeb graph into the plane, as shown in Figure 3.16(b).

To evaluate the quality of our α -Reeb graph for each data set, we computed both original pairwise distances, and pairwise distances approximated from the constructed α -Reeb graph. For GPS traces, we randomly select 100 points as the data set is too big to compute all pairwise distances. We also evaluated the use of α -Reeb graph to speed up distance computations by showing reductions in computation time. Only

| | #OP | #OE | #N | #E | GRT | ODT | ADT | Mean | Median |
|------------|-------|--------|-------|-------|------|-------|------|-------|--------|
| GPS traces | 82541 | 313415 | 21644 | 21554 | 46.8 | 15.27 | 0.96 | 6.5% | 5.3% |
| Earthquake | 1600 | 26996 | 147 | 137 | 0.32 | 1.12 | 0.01 | 14.1% | 12.5% |

Table 3.1: #OP (#OE, #N, #E) stands for the number of original points (original edges, nodes, edges in α -Reeb graph). The graph reconstruction time (GRT) is the total time of computing distance function and reconstructing the graph. The original (ODT), respectively approximate (ADT), distance computation time shows the total time of computing these distances using the original, respectively reconstructed, graph. All times are in seconds. The last two columns show the mean and median metric distortions.

pairs of points in the same connected component are included because we obtain zero error for the pairs of vertices that are not. Statistics for the size of the reconstructed graph, error of approximate distances, and reduction in computation time are given in Table 3.1.

3.9.3 GPS Data with Crossing

The third data set we consider is also obtained from GPS traces. Roads are often split so that cars in different directions run in different lanes. In particular, this is the true for highways. In addition, when two roads cross in GPS coordinates, they may bypass through a tunnel or an evaluated bridge and thus the road network itself may not cross. Such directional information is contained in the GPS traces. We encode this directional information by stacking several consecutive GPS coordinates to form a point in a higher dimensional space. In this way, we obtain a new set of points in this higher dimension space. Then we build a neighboring graph for this new set of points based on L_2 norm and apply our algorithm to recover the road network. In particular, although the paths intersect at the cross in GPS coordinates, the road network does not and this should be detected by our algorithm.

To test the above strategy, we extract those GPS traces from the above “Moscow” dataset which pass through a highway crossing as shown in Figure 3.17(a). Since GPS records the position based on time, we resample the traces so that the distances between any two consecutive samples is the same among all traces. Then we apply the above algorithm to the resampled traces. Figure 3.17(c) and (d) show the reconstructed graph which recovers the road network of this highway crossing.

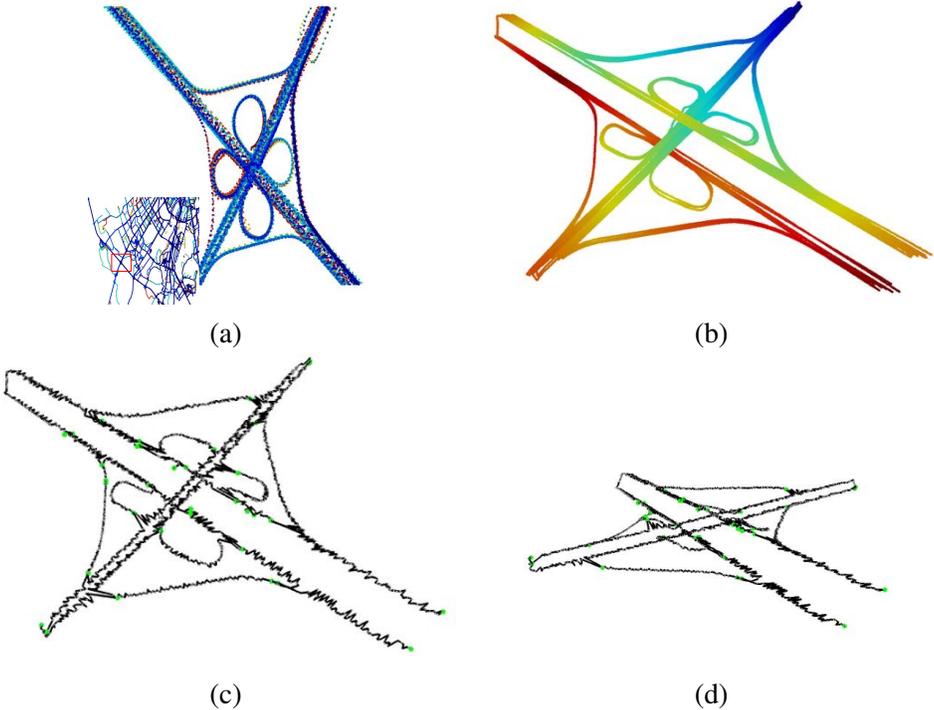


Figure 3.17: (a) GPS traces passing through a highway crossing in Moscow. (b) The distance function. (c) and (d) The reconstructed α -Reeb graph viewed from two perspectives.

On Stability of Shape Difference Operators

4.1 Introduction

Shape comparison is a fundamental problem in geometry processing. In the most general setting, this problem consists of encoding and quantifying similarities and differences across pairs or collections of shapes. This can be especially useful for shape retrieval [Tangelder 2008, Bai 2012], interpolation [Xu 2006, Von-Tycowicz 2015], or visualization [Praßni 2010]. However, even when a map between shapes is given, encoding and visualizing the differences between them are still challenging. Approaches based on the point-to-point correspondences usually suffer issues such as sensitivity to noise, difficulty of selecting an appropriate scale of analysis and inconvenient visualization. The discrete nature of point correspondences is one of the major reasons of these issues.

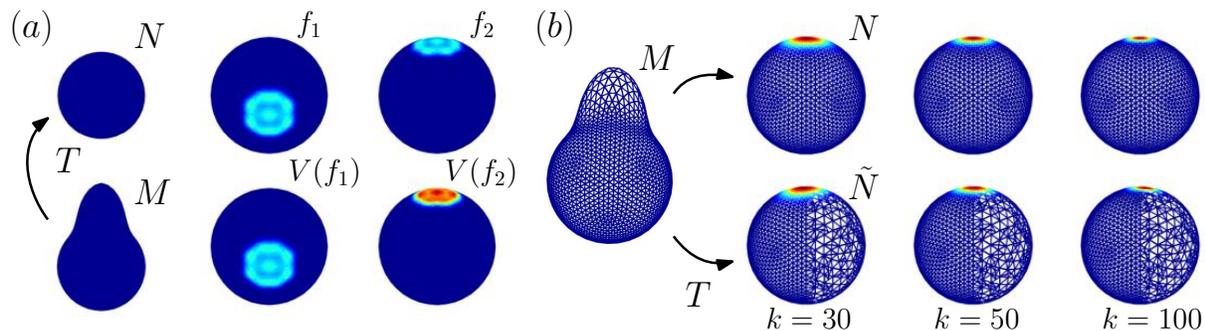


Figure 4.1: (a) Given shapes M , N and T a map between them, V is one of the shape difference operators formulated in [Rustamov 2013]. Intuitively, f_2 , which is supported in a region that undergoes deformation via T , is significantly distorted by V . Whereas f_1 , being supported in area-preserved region, remains the same after V acting on it. (b) We perturb one of the shapes, N , to \tilde{N} and generate indicators with the multi-scale framework of [Ovsjanikov 2013]. Two types of consistency are evidenced: horizontally, the scale k increases from 30 to 100, yet the indicators of each rows highlight nearly the same areas; vertically, at each scale, the indicators are stable with respect to the changes of the input shapes.

The framework of functional maps, which is introduced in [Ovsjanikov 2012], alleviates the issues to some extent by converting point-to-point correspondences into linear operators across function spaces on each of the shapes. As demonstrated in [Ovsjanikov 2012], the functional map is a compact, informative representation and suitable to incorporate with tools from spectral analysis.

In this chapter, we present theoretical analyses for two frameworks based on the notion of functional map. Given a pair of shapes M, N and a map $T : M \rightarrow N$, the framework of [Rustamov 2013] encodes differences between shapes into a pair of linear operators (so-called shape difference operators) acting on function spaces on N . The other framework, proposed in [Ovsjanikov 2013], is for map analysis and visualization. In the same setting as above, this framework generates a collection of multi-scale indicators, which are functions on N highlighting areas that deformed by T . As we will show later, the latter framework can be unified into the first one. Thus our analyses are essentially all performed on the *shape difference operators*.

The theoretical formulation of the shape difference operators is well-established, yet associated stability analyses remain absent. In practice, we observe robustness of these frameworks. For example, in Figure 4.1(b), we feed the framework of [Ovsjanikov 2013] two pairs of meshed shapes: a meshed bumpy sphere M compared to two spheres N with distinct mesh structures, and generate indicators at three different scales. Two types of consistency are evidenced: horizontally, the scale k increases from 30 to 100, yet the indicators of each rows highlight nearly the same areas; vertically at each scale, though the input meshes of N are distinct, the resulting indicators are comparable.

It is then appealing to study the stability properties of the shape difference operators. As hinted by the above numerical results, two types of stability properties are worth considering: one is with respect to perturbations of input shapes in both [Rustamov 2013] and [Ovsjanikov 2013], and the other one is regarding the changes in scale, which is peculiar to the framework of [Ovsjanikov 2013].

Though all the above frameworks assume that M and N are 3D shapes, i.e., 2-dimensional Riemannian manifolds embedded in \mathbb{R}^3 , the formulations are actually well-defined when M and N are Riemannian manifolds of an arbitrary dimension. This fact allows for potentials of these frameworks in dealing with objects beyond shapes. From this point of view, in our analyses, we assume that M and N are smooth Riemannian manifold of dimension n .

How to exploit such potentials in practical applications, however, is another interesting yet challenging problem. For example, the above frameworks have so far only constructed on triangle meshes that are only available when dealing with 3D shapes. For manifolds that are of intrinsic dimension more than 2 or embedded in \mathbb{R}^d with $d \geq 4$, the implementations are not obvious. In Chapter 5, we initiate our exploration to this problem by comparing shapes in a more primitive setting, where only discrete sampling points from the shapes are given.

4.1.1 Overview

We assume that M and N are two connect compact, smooth, n -dimensional Riemannian manifolds without boundary, endowed with metrics g_M and g_N . And let a map $T : M \rightarrow N$ be a smooth map between them.

In [Rustamov 2013], the authors introduce a pair of self-adjoint operators acting on real-value functions on N , each of which captures one type of differences between the two shapes regarding T . Particularly, for one of the shape difference operators – V as illustrated in Figure 4.1, a functional proposed in [Ovsjanikov 2013] evaluates the deviation from a function f on N to its image Vf (we will verify this connection at the beginning of Section 4.6). The maximizer of this functional is supposed to highlight the most deformed regions under map T . Instead of searching for the global maximizer, the authors maximize this functional within the space spanned by the first k eigenfunctions of the Laplace-Beltrami operator on N , and the maximizer is viewed as an indicator at scale k . By changing the scale k , a collection of indicators

and a sequence of the corresponding maxima are then obtained, allowing the users to select output(s) at one or more scales for better understandings(see Figure 4.2 for an illustration).

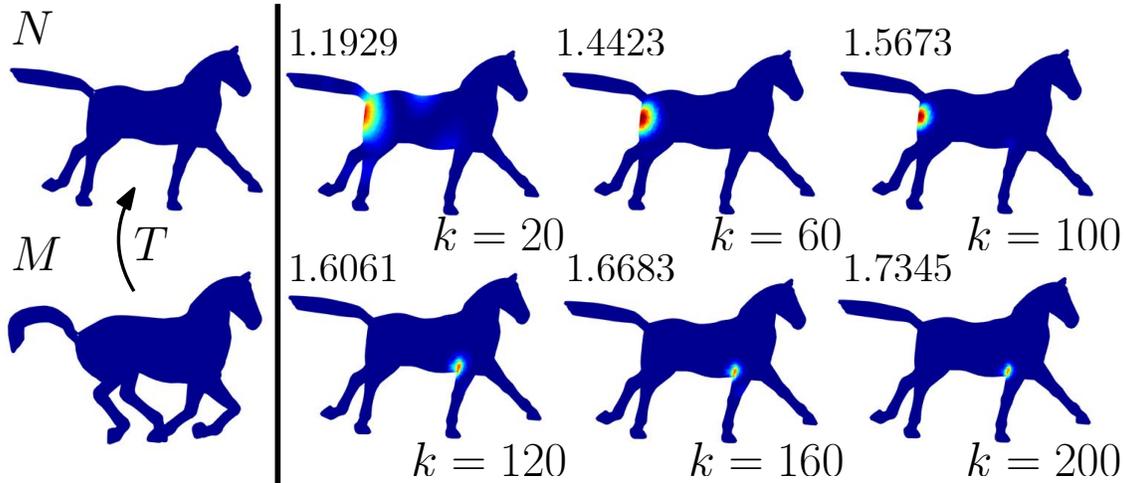


Figure 4.2: Stability across ranges of scales: the indicators from $k = 20$ to 100 consistently highlight the hip of the horse, meanwhile the ones from $k = 120$ to 200 highlight the root of its front right leg. The corresponding quantitative measurements of the distortions of the indicators are marked to the top-left of each shapes.

As mentioned before, we provide stability analyses regarding two type of perturbations.

- Perturbations with respect to the input manifolds: We start by introducing a model in Section 4.4 characterizing perturbations on the input manifolds. Then in Section 4.5, we discuss the stability of the two different shape difference operators under our perturbation model.
- Perturbations with respect to the changes in scale: As mentioned above, the multi-scale framework depends on a collection of subdomains (in [Ovsjanikov 2013], that is, the spaces spanned by the first k eigenfunctions, with $k = 1, 2, \dots$) in which we maximize a specified functional. However in Section 4.6.1, we demonstrate that the original subdomain construction indexed by k is not suitable for stability analysis and construct a new one which is closely related to the original but controlled by a continuous C . Then further in Section 4.6, we verify this stability with respect to this continuous scale C .

Especially, we extend the framework of [Ovsjanikov 2013] in Section 4.6.5 by adapting the other shape difference operator to it, which enables visualizing another type of shape differences other than the area-based one. And we as well prove that this extension enjoys a similar stability property as the original one.

Finally, in Section 4.7, we demonstrate numerical experiments that reflect the stability in practice.

4.2 Related Work

The two frameworks we analyze in this chapter are based on the notion of functional maps, which has been a key ingredient of various applications in geometry processing, to name a few, analyzing maps between shapes [Huang 2014], vector field processing [Azencot 2013, Azencot 2014] and image segmentation [Wang 2013].

Perturbation analysis has a long and rich history. In this chapter, we perform perturbation analysis on both shape difference operators (which are linear operators, see [Kato 1995] for an introduction of perturbation analysis on them) and a spectral method based on such operators. The spectral methods have long been applied in various areas: spectral clustering [von Luxburg 2006], shape analysis [Reuter 2006] and so on. Besides demonstrating practical usefulness of the spectral methods, providing theoretical justifications is attracting more and more research interests. Theoretical guarantees for spectral clustering algorithms often stem from Cheeger’s inequality, which is powerful if there exists a significant spectral gap. Assuming such a gap, several works [Kwok 2013, Lee 2012, Louis 2012, Oveis Gharan 2014, Dey 2014] present theoretical guarantees on the quality (measured by some graph conductance) of the output of the respective algorithms. It is worth noting that the above works only concern a single object, while in this paper, we study pairwise objects. From this point of view, our work has a similar flavor to the ones by Mémoli [Mémoli 2009, Mémoli 2011], who proposes metrics among shapes based on spectral invariant and discusses their robustness with respect to perturbations on the input shapes.

Beyond spectral methods, in geometric and topological data analysis, several approaches have been proposed for guaranteeing stability of the data processing and analysis techniques. In particular, stability has been theoretically proven in many works on estimating geometric quantities. In [Mitra 2004], the authors analyze both effects of practical and theoretical facts on the accuracy of normal estimation process. In [Mérigot 2011], it is assured that the sharp feature detection algorithm there is stable with respect to Hausdorff noise. In the same noise model, the stability of the curvature measures is proven under certain conditions in [Chazal 2009]. Similar problems are also actively studied in the community of topological data analysis (TDA). The stability of persistence diagram is verified in [Cohen-Steiner 2007], which lays down a solid theoretical foundation for further research in TDA. Some more recent developments in TDA come with stability assurance as well, for example the notion of distance to a measure [Chazal 2014].

A rich body of research has been devoted to provide such analysis for convergence properties of various discrete Laplacian operator. In [Wardetzky 2005, Xu 2007, Dey 2010] the converging behaviors of the cotangent Laplacian operators on meshes to the underlying Laplace-Beltrami operators are investigated from diverse perspectives. While in [Belkin 2009, Liu 2012, Hein 2007, Dey 2013b], similar problems are considered in a different setting, where the discrete Laplacian operators are built on point clouds.

4.3 Preliminaries

We first give a brief review of concepts about differential geometry and the (weighted) Laplace-Beltrami operators on manifolds. We refer the readers to [Grigoryan 2000] for a more detailed introduction (especially to the latter topic).

Then we briefly take a review over the *functional-maps-based* frameworks, which are the main focuses

in this chapter.

4.3.1 Differential Geometry

We start by defining a smooth n -dimensional manifold M . Roughly speaking, a n -dimensional manifold is a topological space that locally resembles an Euclidean space of dimension n . To give a precise definition, we introduce a chart on M :

Definition 4.1 A n -dimensional chart is a couple (U, h) where U is an open subset of M and $h : U \rightarrow \mathbb{R}^n$ is homeomorphism between U and an open subset of \mathbb{R}^n .

We call M a C -manifold of dimension n if M is a Hausdorff topological space with countable bases and each point of M belongs to a n -dimensional chart. Furthermore, we define a smooth atlas which consists of smooth, compatible charts.

Definition 4.2 A smooth atlas of M is a family of charts $\{(U_i, h_i)\}_{i \in \Gamma}$ such that

- M is covered by the open sets $U_i, i \in \Gamma$, i.e., $M \subset \bigcup_{i \in \Gamma} U_i$
- Denote by h_{ij} the transition map between chart $(U_i, h_i), (U_j, h_j)$, i.e., $h_{ij} = h_j \circ h_i^{-1} : h_i(U_i \cap U_j) \rightarrow h_j(U_i \cap U_j)$. For any pair of $i, j \in \Gamma$ satisfying $U_i \cap U_j \neq \emptyset$, the transition map h_{ij} is a smooth function (note that the domain and the image of h_{ij} are both subsets of \mathbb{R}^n).

Two smooth atlases are compatible if and only if their union is a smooth atlas. A smooth structure on M is the union of all compatible atlases. A smooth n -dimensional manifold is then a C -manifold endowed with a smooth structure.

Now we consider a point $x \in M$ and a chart (U, h) such that U contains x . This chart induces a local coordinate system (x^1, x^2, \dots, x^n) , where x^i is the i -th coordinate function of $h(y) \in \mathbb{R}^n$, where $y \in U$. The following notions from Riemannian geometry is introduced with a local coordinate system, however, most of them are independent of the selection of the local coordinate systems.

Tangent Spaces A function $f : M \rightarrow \mathbb{R}$ is smooth if for any $x \in M$, there exists a chart (U, h) such that function $f \circ h^{-1} : \mathbb{R}^n \rightarrow \mathbb{R}$ is smooth. We denote the set of smooth functions on M by $C^\infty(M)$. The tangent space at a point $x \in M$, $T_x M$, is a set of mappings $\{\eta : C^\infty(M) \rightarrow \mathbb{R}\}$ such that

- η is a linear functional on $C^\infty(M)$;
- $\forall f, g \in C^\infty(M), \eta(fg) = \eta(f)g(x) + \eta(g)f(x)$.

It is well-known that $T_x M$ is a n -dimensional linear space, and that, in chart, $\{\frac{\partial}{\partial x^i} \big|_x\}_{i=1}^n$ form a basis of $T_x M$.

Riemannian Metrics A Riemannian manifold is a manifold M endowed with a Riemannian metric, g_M , which is a family of $\{g_M(x)\}_{x \in M}$ such that at any point x , $g_M(x)$ is a symmetrical, positive definite and bilinear form on the tangent space $T_x M$, i.e., $g_M(x) : T_x M \times T_x M \rightarrow \mathbb{R}^+$.

A metric g_M at x gives rise to an inner product on $T_x M$: for $\eta, \xi \in T_x M$, $\langle \eta, \xi \rangle_{g_M(x)} = g_M(x)(\eta, \xi)$. The length of η is then defined as $\sqrt{\langle \eta, \eta \rangle_{g_M(x)}}$. Particularly, in the local coordinate system g_M admits a matrix representation, whose (i, j) -th element, $g_{i,j}$, equals $\langle \frac{\partial}{\partial x^i}, \frac{\partial}{\partial x^j} \rangle_{g_M(x)}$. Matrix $[g_{i,j}]$ is invertible, and we denote by $[g^{i,j}]$ its inverse matrix.

The matrix of a single metric g_M is depending on the choice of the local coordinate system. Interestingly, we obtain some invariance when taking another metric into consideration. Let g_M and \tilde{g}_M be two metrics on M . Provided a local coordinate system (x^1, x^2, \dots, x^n) , we denote by $[g]$ and $[\tilde{g}]$ the matrices of g_M and \tilde{g}_M in this system. Now assumed that another local coordinate system is obtained by letting $(y^1, y^2, \dots, y^n) = F(x^1, x^2, \dots, x^n)$, where $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a smooth map. The matrices with respect to y is denoted as $[g]_y$ and $[\tilde{g}]_y$. It follows from lemma 3.12 in [Grigoryan 2000] that $[g] = J(F)^T [g]_y J(F)$ and $[\tilde{g}] = J(F)^T [\tilde{g}]_y J(F)$, in which $J(F)$ is the Jacobian matrix of F . Hence we have the following equation:

$$\frac{\det[g]_y}{\det[\tilde{g}]_y} = \frac{\det(J(F)^T [g] J(F))}{\det(J(F)^T [\tilde{g}] J(F))} = \frac{\det[g]}{\det[\tilde{g}]}, \quad (4.1)$$

which suggests the ratio of determinants of the metrics is invariant to the changes of the local coordinate systems.

Gradient Given $f \in C^\infty(M)$, its gradient at x is defined as $\nabla_{g_M} f = \sum_{i,j=1}^n g^{i,j} \frac{\partial f}{\partial x^j} \frac{\partial}{\partial x^i}$. Notice that $\nabla_{g_M} f$ is spanned by $\{\frac{\partial}{\partial x^i}\}$, thus it is also a tangent vector in $T_x M$. For a lighter notation, we write $\nabla_{g_M} f$ as ∇f and by $\langle \nabla f, \nabla g \rangle_{g_M}$ we always mean $\langle \nabla_{g_M} f, \nabla_{g_M} g \rangle_{g_M}$, i.e., the inner product of the function gradients are defined upon the metric.

Riemannian Measure A Riemannian metric g_M induces a volume (also known as a Riemannian measure), ν_M , on the family of all measurable subsets of M . The following theorem assures the existence and uniqueness of ν_M . We refer the reader to [Grigoryan 2000] for a detailed proof (see Theorem 3.11 there).

Theorem 4.1 For any Riemannian manifold (M, g_M) , there exists a unique measure ν_M on $\Lambda(M)$ (the family of all measurable sets in M) such that, in any chart U ,

$$d\nu_M = \sqrt{\det[g_{i,j}]} d\lambda \quad (4.2)$$

where $[g_{i,j}]$ is the matrix of g_M and λ is the Lebesgue measure in U .

The following proposition characterizes how a perturbation on the metric on a manifold affects the associated volume and gradient norm.

Proposition 4.1 *Let g_M and \tilde{g}_M be two metrics on a n -dimensional manifold M , and $\nu_M, \tilde{\nu}_M$ be the volumes induced by them respectively. If there exists a constant $a \geq 1$ such that for any point $x \in M$ and any $\eta \in T_x M$, the following inequality holds:*

$$a^{-1} \leq \frac{\langle \eta, \eta \rangle_{g_M}}{\langle \eta, \eta \rangle_{\tilde{g}_M}} \leq a \quad (4.3)$$

, then

$$a^{-n/2} d\tilde{\nu}_M \leq d\nu_M \leq a^{n/2} d\tilde{\nu}_M \quad (4.4)$$

and for any $f \in C^\infty(M)$,

$$a^{-1} \leq \frac{\langle \nabla f, \nabla f \rangle_{g_M}}{\langle \nabla f, \nabla f \rangle_{\tilde{g}_M}} \leq a \quad (4.5)$$

Proof: To prove inequality 4.4, we show that the $a^{-n} \leq \frac{\det[g]}{\det[\tilde{g}]} \leq a^n$, where $[g]$ and $[\tilde{g}]$ are matrices of g_M and \tilde{g}_M in a common local coordinate system.

The invariance of Equation 4.1 allows us to pick freely a local coordinate system without affecting the ratio of the determinants. We then consider $\{e_1, e_2, \dots, e_n\}$, an orthonormal basis of $T_x M$, in which the matrix representation of \tilde{g}_M is an identity matrix. Thus we have $\det[\tilde{g}] = 1$, and in the following we bound $\det[g]$. Let $[g]$ be the matrix of $g_M(x)$ in basis $\{e_1, e_2, \dots, e_n\}$. If there exists an eigenvalue of $[g]$, λ , that is larger than a , then we let u be the associated eigenvector and $\eta = \sum_{i=1}^n u^i e_i$. As a consequence, $\langle \eta, \eta \rangle_{g_M} = u^T [g] u = \lambda > a$. Since $[\tilde{g}]$ is an identity matrix, we have $\langle \eta, \eta \rangle_{\tilde{g}_M} = 1$ and the ratio $\frac{\langle \eta, \eta \rangle_{g_M}}{\langle \eta, \eta \rangle_{\tilde{g}_M}}$ exceeds a , which contradicts the assumption 4.3. On the other hand, since $g_M(x)$ is positive definite, thus all its eigenvalues are positive and not larger than a . Thus the determinant of $[g]$, which equals the product of all the eigenvalues, is not larger than a^n . As a consequence of Equation 4.2, we have $d\nu_M \leq a^{n/2} d\tilde{\nu}_M$. The other side of inequality 4.4 can be proven in the same way.

To prove inequality 4.5, we notice that the definition of ∇f depends on the matrix of metric, thus we need to do some transformation. We consider $df = [g_{i,j}] \nabla f$, which is independent of g . Then we have $\langle \nabla f, \nabla f \rangle_{g_M} = ([g_{i,j}] \nabla f, \nabla f) = ([g^{i,j}] [g_{i,j}] \nabla f, [g_{i,j}] \nabla f) = ([g^{i,j}] df, df) = \langle df, df \rangle_{g_M^{-1}}$ and correspondingly $\langle \nabla f, \nabla f \rangle_{\tilde{g}_M} = \langle df, df \rangle_{\tilde{g}_M^{-1}}$.

We use the same basis (e_1, e_2, \dots, e_n) of $T_x M$ as above. For \tilde{g}_M , in this basis $[\tilde{g}_{i,j}]$ and its inverse $[\tilde{g}^{i,j}]$ are identity matrices. As we've proven that all eigenvalues of $[g_{i,j}]$ do not exceed a , we have the smallest eigenvalue of $[g^{i,j}]$ is at least a^{-1} . Therefore we derive that $\langle \nabla f, \nabla f \rangle_{g_M} = \langle df, df \rangle_{g_M^{-1}} \leq a \langle df, df \rangle_{\tilde{g}_M^{-1}} = a \langle \nabla f, \nabla f \rangle_{\tilde{g}_M}$. The other side of inequality 4.5 is proven symmetrically. \square

Weighted Riemannian Manifold For a Riemannian manifold (M, g_M) , besides the volume (or the Riemannian measure) induced by g_M , we can define other measures. For example, a measure μ_M on M defined by $d\mu_M = \rho d\nu_M$, where ρ is a smooth positive-value function on M . Such a measure is a weighted version of the volume ν_M , and it gives rise to a weighted Riemannian manifold.

Definition 4.3 *A weighted Riemannian manifold is a triple (M, g_M, μ_M) , where g_M is Riemannian metric inducing a volume ν_M and there exists a positive smooth function ρ such that $d\mu_M = \rho d\nu_M$.*

Particularly, if ρ is a constant function: $\rho(x) = 1, \forall x \in M$, then μ_M is the volume induced by g_M .

Function spaces In this part we introduce some function spaces on (weighted) Riemannian manifolds. We denote the set of all real-value functions on M by $\mathcal{F}(M)$. The set of smooth functions on M , $C^\infty(M)$, is obviously a subset of $\mathcal{F}(M)$.

Let μ_M be a measure on (M, g_M) (not necessarily a volume), by $L^2_\mu(M)$ we denote the function space in which all functions have finite square norm with respect to μ_M , i.e.,

$$L^2_\mu(M) = \{f \in \mathcal{F}(M) : \int_M f^2 d\mu_M < \infty\}$$

A special subspace of $L^2_\mu(M)$ with smoothness constraints is $H^1_{0,\mu}(M)$, which is defined on top of $H^1_\mu(M)$:

$$H^1_\mu(M) = \{f : \int_M f^2 + \langle \nabla f, \nabla f \rangle_{g_M} d\mu_M < \infty\}$$

$H^1_{0,\mu}(M)$ is the closure of all infinitely differentiable functions in H^1_μ that is compactly supported by M . An alternative understanding of $H^1_{0,\mu}(M)$ is the following:

$$H^1_{0,\mu}(M) = \{f : f \in H^1_\mu(M) \text{ and } f|_{\partial M} = 0\}.$$

Inner Products After introducing $L^2_\mu(M)$ and $H^1_{0,\mu}(M)$, we define two inner products on them respectively, which are crucial in formulating the shape difference operators (Section 4.3.5).

The first inner product takes in two functions of $L^2_\mu(M)$ and returns a real-number:

$$\langle f, g \rangle_{L^2_\mu} = \int_M fg d\mu_M, f, g \in L^2_\mu(M) \quad (4.6)$$

And we define the second inner product on function space $H^1_{0,\mu}(M)$:

$$\langle f, g \rangle_{H^1_{0,\mu}} = \int_M \langle \nabla f, \nabla g \rangle_{g_M} d\mu_M, f, g \in H^1_{0,\mu}(M) \quad (4.7)$$

4.3.2 The Laplace-Beltrami Operator

We denote the Laplace-Beltrami operator on manifold (M, g_M) by Δ_M .

Definition 4.4 Given a smooth Riemannian manifold (M, g_M) of dimension n , $\Delta_M : C^\infty(M) \rightarrow C^\infty(M)$ is an operator on $C^\infty(M)$, which in a local coordinate system (x^1, x^2, \dots, x^n) is defined as:

$$\Delta_M = \sum_{i,j=1}^n \frac{1}{\sqrt{\det g}} \frac{\partial}{\partial x^i} (\sqrt{\det g} g^{i,j} \frac{\partial}{\partial x^j})$$

We call (φ, λ) an eigensolution if it satisfies the following eigen equation. Especially $\varphi \in C^\infty(M)$ is called an eigenfunction matching an eigenvalue $\lambda \in \mathbb{R}$.

$$\Delta_M \varphi + \lambda \varphi = 0 \quad (4.8)$$

It is well-known that the Laplace-Beltrami operator is self-adjoint and negative semidefinite. Thus in any eigensolution, λ is guaranteed to be a non-negative real number. As we assume M to be a compact manifold, the set of eigenvalues is discrete and infinite. Thus we can sort the eigenvalues with an ascending order so that $0 \leq \lambda_1 \leq \lambda_2 \leq \lambda_3 \leq \dots$ and denote the associated eigenfunctions by $\varphi_1, \varphi_2, \varphi_3, \dots$.

The Green's first identity states:

$$\int_M \langle \nabla u, \nabla v \rangle_{g_M} d\nu_M + \int_M u \Delta_M v d\nu_M = \int_{\partial M} u \nabla v \cdot \mathbf{n} ds_M \quad (4.9)$$

where $u, v \in C^\infty(M)$, $\partial(M)$ is the boundary of M , and \mathbf{n} is the unit normal vector field to ∂M defined at each point on ∂M . Since we assume M to be boundaryless, the right hand side of Equation 4.9 is zero, resulting in the following:

For any smooth functions u, v on manifold M which is compact and without boundary.

$$\int_M u(-\Delta_M)v d\nu_M = \int_M \langle \nabla u, \nabla v \rangle_{g_M} d\nu_M \quad (4.10)$$

Considering two eigenfunctions, φ_i, φ_j associated to distinct eigenvalues λ_i, λ_j , it follows from Equation 4.10 that

$$\begin{aligned} \int_M \varphi_i(\lambda_j \varphi_j) d\nu_M &= \int_M \varphi_i(-\Delta_M)\varphi_j d\nu_M = \int_M \langle \nabla \varphi_i, \nabla \varphi_j \rangle_{g_M} d\nu_M \\ &= \int_M \varphi_j(-\Delta_M)\varphi_i d\nu_M = \int_M \varphi_j(\lambda_i \varphi_i) d\nu_M. \end{aligned}$$

Therefore we have $\int_M \varphi_i \varphi_j d\nu_M = 0$ for $\lambda_i \neq \lambda_j$.

For each eigenvalue λ_i , the functions satisfying $\Delta_M \varphi + \lambda_i \varphi = 0$ form a linear function space. If λ_i is of multiplicity 1, then this space is of dimension 1, which is spanned by $\{\varphi_i\}$. If λ_i is of multiplicity more than 1, the corresponding function space might be of dimension $k > 1$. In this case, with Gram Schmidt method we can always find $\{\varphi_i^l\}_{l=1, \dots, k}$ forms a basis of this k -dimension space and $\int_M \varphi_i^{k_1} \varphi_i^{k_2} d\nu_M = 0, 1 \leq k_1 \neq k_2 \leq k$.

In fact, the eigenfunctions the Laplace-Beltrami operator form a basis of $L^2_\nu(M)$. According to above arguments, we always assume that the basis is orthogonal. Moreover, by simply rescaling φ_i so that $\int_M \varphi_i^2 d\nu_M = 1$, we obtain an orthonormal basis of $L^2_\nu(M)$.

Proposition 4.2 *Let $\{\varphi_i\}_{i \geq 1}$ be an orthonormal basis of $L^2_\nu(M)$ consisting of eigenfunctions to $\Delta_M \varphi + \lambda \varphi = 0$. For any function $u \in L^2_\nu(M)$, it admits a eigendecomposition $u = \sum_{i \geq 1} a_i \varphi_i, a_i = \int_M u \varphi_i d\nu_M$. And we have*

$$\int_M u^2 d\nu_M = \sum_{i \geq 1} a_i^2 \quad (4.11)$$

If we further assume that u is differentiable, then

$$\int_M \langle \nabla u, \nabla u \rangle_{g_M} d\nu_M = \sum_{i \geq 1} a_i^2 \lambda_i \quad (4.12)$$

Proof: Assuming that $u = \sum_{i \geq 1} a_i \varphi_i$, we have $\int_M u \varphi_j d\nu_M = \int_M \sum_{i \geq 1} a_i \varphi_j \varphi_i d\nu_M$. It follows the orthonormality of the basis that $\int_M u \varphi_j d\nu_M = a_j$.

Similarly, $\int_M u^2 d\nu_M = \int_M (\sum_i a_i \varphi_i)^2 d\nu_M = \int_M a_i^2 \sum_{i \geq 1} \varphi_i^2 d\nu_M = \sum_{i \geq 1} a_i^2$ as $\int_M \varphi_i \varphi_j d\nu_M = 0$ and $\int_M \varphi_i^2 d\nu_M = 1$.

The Equation 4.12 follows from Equation 4.10: $\int_M \langle \nabla u, \nabla u \rangle_{g_M} d\nu_M = \int_M u(-\Delta_M)u d\nu_M = \int_M (\sum_{i \geq 1} a_i \varphi_i) (\sum_{i \geq 1} a_i \lambda_i \varphi_i) d\nu_M = \sum_{i \geq 1} a_i^2 \lambda_i \square$

Weighted Laplace-Beltrami Operator We have defined a weighted Riemannian manifold (M, g_M, μ_M) (see Definition 4.3), where $\mu_M = \rho \nu_M$. A weighted Laplace-Beltrami operator is defined as:

$$\Delta_{M,\mu} = \sum_{i,j=1}^n \frac{1}{\rho} \frac{1}{\sqrt{\det g}} \frac{\partial}{\partial x^i} (\rho \sqrt{\det g} g^{i,j} \frac{\partial}{\partial x^j}) \quad (4.13)$$

Notice that if $\rho = 1$, then $\Delta_{M,\mu}$ boils down to Δ_M .

Equation 4.10 is still valid with respect to the weighted Laplace-Beltrami operator, but with a change of measures.

$$\int_M u(-\Delta_{M,\mu})v d\mu_M = \int_M \langle \nabla u, \nabla v \rangle_{g_M} d\mu_M \quad (4.14)$$

4.3.3 Functional Maps

The framework of *functional maps* is introduced in [Ovsjanikov 2012], which converts point correspondences between two manifolds to function correspondences between function spaces on the manifolds.

Let M and N be two smooth manifolds and T be a map from M to N . A functional map, T_F , is a pull-back induced by T . Namely, given a real-value function $w \in \mathcal{F}(N)$, we define $T_F(w) = w \circ T \in \mathcal{F}(M)$. Therefore T_F is a map from $\mathcal{F}(N)$ to $\mathcal{F}(M)$. Below we review two important properties of T_F : informativeness and linearity.

First we claim that if T is bijective, then it can be fully recovered by T_F . In fact, for any point $a \in N$, let $\delta_a(x)$ be an indicator function on N such that $\delta_a(x)$ equals 1 if $x = a$ and $\delta_a(x) = 0$ otherwise. By construction, $g = T_F(\delta_a)$ satisfies that $g(y)$ equals 1 if $T(y) = a$ and 0 otherwise. Since that T is bijective, there exist a unique b such that $T(b) = a$, thus g is also an indicator function but on M . After all, for each $a \in N$, we search for a point $b \in M$ such that $T_F(\delta_a)(b)$ equals 1. Such a point b is unique and satisfies $T(b) = a$, meaning that we can completely recover T with T_F .

Secondly, T_F is a *linear* map across the two function spaces, i.e., for any $\alpha, \beta \in \mathbb{R}$ and $f, g \in \mathcal{F}(N)$, we have $T_F(\alpha f + \beta g) = \alpha T_F(f) + \beta T_F(g)$. This linearity property indicates that T_F admits a (potentially infinite) matrix representation. Suppose that $\{\varphi_i^N\}, \{\psi_j^M\}$ form a basis of $\mathcal{F}(N)$ and of $\mathcal{F}(M)$ respectively. Let $w = \sum_{i=1}^N a_i \varphi_i^N$ and accordingly $T_F(w) = \sum_{j=1}^M b_j \psi_j^M$, then there exists a unique matrix C_T such that

$C_T \mathbf{a} = \mathbf{b}$ holds for any w , where $\mathbf{a} = (a_1, a_2, \dots, a_n, \dots)$ and $\mathbf{b} = (b_1, b_2, \dots, b_n, \dots)$. We then call C_T the matrix representation of T_F with respect to basis $\{\varphi_i^N\}$ and $\{\psi_j^M\}$.

In practice, we usually consider $\{\varphi_i^N\}$, $\{\psi_j^M\}$ as eigenbasis of the Laplace-Beltrami operator on N and M respectively. And then truncation on each basis is taken, say, we only take the first k_N and k_M eigenfunctions on each shape. Then C_T is a k_M by k_N matrix instead of an infinite-dimensional one. More importantly, we can control the rank of this matrix representation to reach a trade-off between accuracy and complexity. The authors of [Ovsjanikov 2012] show that hundreds of eigenfunctions are enough to reconstruct a reasonable functional map between meshed shapes that consist of tens of thousands of vertices.

4.3.4 Map Analysis

Based on the functional maps, the authors of [Ovsjanikov 2013] present a multi-scale framework to detect and visualize the area-distortions induced by maps $T : M \rightarrow N$. Unlike the approaches which perform comparisons via point correspondences, this framework analyzes the function correspondences, which are easily converted from the point-to-point maps with the above functional map framework.

The key ingredient of this framework is a functional that measures area-distortions induced by transforming a function $w \in L^2(N)$ to $L^2(M)$ with the functional map, T_F :

$$E(w) = \frac{\int_M T_F(w)^2 d\nu_M}{\int_N w^2 d\nu_N} \quad (4.15)$$

As discussed in [Ovsjanikov 2013], if $E(w)$ is large, then w is supposed to take high absolute value within the areas where measure are less preserved when mapped from M to N via T .

Therefore, it is natural to search for functions such that $E(w)$ is large. Instead of maximizing $E(w)$ in $L^2_\nu(N)$, a multi-scale approach is taken by forcing w to reside in a subspace of $L^2_\nu(N)$ spanned by the first k eigenfunctions of Δ_N , the Laplace-Beltrami operator on N . We define $S(k)$ as the following

$$S(k) = \text{span}\{\varphi_i\}_{i=1}^k, \text{ where } \varphi_i \text{ is the } i\text{-th eigenfunction of } \Delta_N. \quad (4.16)$$

The constrained optimization problem considered in [Ovsjanikov 2013] is then:

$$\max E(w) \text{ s.t. } w \in S(k) \quad (4.17)$$

The advantage of choosing such a collection of subdomains $\{S(k)\}_{k \in \mathbb{N}^+}$ is multi-fold: (1) the span space is easy to construct and the associated constrained optimization is straightforward to solve; (2) the feasible solutions are spanned by low-frequency eigenfunctions, which are fit for the visualization purpose and stable with respect to noises of input discrete shapes. (3) the flexibility of choosing k provides a multi-scale understanding of the problem. As demonstrated in Figure 4.2, the regions highlighted by w^* become more and more localized as k , the dimension of spanned space, increases.

4.3.5 Shape Difference Operators

Another map-based approach for shape comparison is proposed in [Rustamov 2013], where a pair of *Shape Difference Operators* is introduced.

The area-based shape difference operator, $V : L^2_\nu(N) \rightarrow L^2_\nu(N)$, is defined as a linear operator such that for any f, g in square-integrable space $L^2_\nu(N) = \{f : \int_N f^2 d\nu_N < +\infty\}$,

$$\int_N fV(g)d\nu_N = \int_M T_F(f)T_F(g)d\nu_M \quad (4.18)$$

Note that unless T is an area-preserving map, equation $\int_N fg d\nu_N = \int_M T_F(f)T_F(g)d\nu_M$ does not hold for any pair of $f, g \in L^2(N)$, V captures the difference and compensates the discrepancy.

With respect to a different inner product, $\langle \cdot, \cdot \rangle_{H^1_{0,\nu}}$, the conformal-based shape difference operator, $R : H^1_{0,\nu}(N) \rightarrow H^1_{0,\nu}(N)$, is a linear operator such that for any f, g in Sobolev space $H^1_{0,\nu}(N)$,

$$\int_N \langle \nabla f, \nabla R(g) \rangle_{g_N} d\nu_N = \int_M \langle \nabla T_F(f), \nabla T_F(g) \rangle_{g_M} d\nu_M \quad (4.19)$$

It follows from the Riesz representation theorem that V, R exist and are unique. Moreover, thanks to the commutative properties of inner products 4.6 and 4.7, V and R are as well self-adjoint operators, i.e., $\int_N fV(g)d\nu_N = \int_N gV(f)d\nu_N$ and $\int_N \langle \nabla f, \nabla R(g) \rangle_{g_N} d\nu_N = \int_N \langle \nabla g, \nabla R(f) \rangle_{g_N} d\nu_N$ hold for any f, g residing in the corresponding domains.

Particularly, the authors of [Rustamov 2013] prove that T is locally area-preserving if and only if V is an identity operator. That proof (of theorem 1 therein) applies to the case where M and N are Riemannian manifolds of dimension n . Regarding R , the author of [Schumacher 2013] shows that if the dimensions of M, N are $n = 2$, then T is a *conformal map* if and only if R is an identity operator, and if $n > 2$ then T is an *isometric map* if and only if R is an identity operator. Therefore, regardless of the dimension of the input Riemannian manifolds, the shape difference operators are well-defined and both carry information about the deviation from T to an isometric map.

4.4 Perturbation Model and Bounded-distortion Condition

As mentioned in Section 4.1, one of the main results in this chapter is the stability analyses of the shape difference operators with respect to perturbations on the input manifolds. In this section, we propose our perturbation model and a bounded-distortion condition throughout our analyses.

4.4.1 Perturbation Model

Let M be a connected, compact, smooth, n -dimensional manifold without boundary endowed with a Riemannian metric g_M . And let ν_M be the volume induced by g_M . We first define (a, b) -closeness between Riemannian structures on the same smooth manifold, where a, b are positive constants not smaller than 1.

We now introduce our model for characterizing perturbations on the input shapes.

Definition 4.5 A Riemannian manifold $(N, \tilde{g}_N, \tilde{\nu}_N)$ is a -close to another one (N, g_N, ν_N) if the following holds: For any $x \in N$ and any tangent vector η in $T_x N$, the tangent plane at x : $a^{-1} \leq \frac{\langle \eta, \eta \rangle_a}{\langle \eta, \eta \rangle_{\tilde{g}}} \leq a$ holds for a constant $a \geq 1$.

Definition 4.6 A weighted Riemannian manifold (N, g_N, μ_N) is b -close to a Riemannian manifold (N, g_N, ν_N) if the following holds: μ_N is obtained by perturbing ν_N (the volume induced by g_N) with ρ_N : $d\mu_N = \rho_N d\nu_N$. And $b^{-1} \leq \rho_N \leq b$ holds for a constant $b \geq 1$.

It is clear that the a, b -closeness characterizes perturbations on the metric and on the measure, respectively. Combining them together, a weighted Riemannian manifold, $(N, \tilde{g}_N, \tilde{\mu}_N)$, is said to be (a, b) -close to a Riemannian manifold (N, g_N, ν_N) if

- $(N, \tilde{g}_N, \tilde{\mu}_N)$ is b -close to the corresponding Riemannian manifold $(N, \tilde{g}_N, \tilde{\nu}_N)$.
- $(N, \tilde{g}_N, \tilde{\nu}_N)$ is a -close to (N, g_N, ν_N) .

Intuitively, we view $(N, \tilde{g}_N, \tilde{\mu}_N)$ as a perturbed version of (N, g_N, ν_N) . It is obvious that $(1, 1)$ -closeness implies that the two are identical.

The smooth function ρ_M is usually modeled as a sampling density on manifold M . The second restriction on two Riemannian metrics might seem a bit abstract. We describe it in a special case, where M is a 2-dimensional manifold and $M_1 = (M, g_M)$, $M_2 = (M, \tilde{g}_M)$ are two Riemannian manifolds embedded in the same Euclidean space \mathbb{R}^3 . Assume that there exists $\Phi : M_1 \rightarrow M_2$ that is a diffeomorphism. Let $x \in M_1 \subset \mathbb{R}^3$ and $y = \Phi(x) \in M_2$, which are both in \mathbb{R}^3 . If g_M, \tilde{g}_M satisfy the above condition, then the tangent plane of M_1 at x is parallel to the one of M_2 at y . On the other hand, the differentiation of Φ gives rise to a push-forward between the tangent planes: a tangent vector $v \in T_x M_1$ is mapped to $D(\Phi)v \in T_y M_2$. Due to the fact that M_1, M_2 are both embedded in \mathbb{R}^3 , the inequality is re-written as:

$$a^{-1} \leq \frac{\langle v, v \rangle_{g_{\mathbb{R}^3}}}{\langle D(\Phi)v, D(\Phi)v \rangle_{g_{\mathbb{R}^3}}} \leq a,$$

where $g_{\mathbb{R}^3}$ is the Euclidean metric. In other words, denote the matrix representation of $D(\Phi)$ by A , then the eigenvalues of $A^T A$ both lie in the interval $[a^{-1}, a]$.

The following proposition characterizes the quantitative relationships between (M, g_M, ν_M) and $(M, \tilde{g}_M, \tilde{\mu}_M)$.

Proposition 4.3 If $(M, \tilde{g}_M, \tilde{\mu}_M)$ is (a, b) -close to (M, g_M, ν_M) , then for any smooth function w on M .

$$a^{-1} \leq \frac{\langle \nabla w, \nabla w \rangle_{g_M}}{\langle \nabla w, \nabla w \rangle_{\tilde{g}_M}} \leq a,$$

and

$$a^{-n/2} b^{-1} d\tilde{\mu}_M \leq d\nu_M \leq a^{n/2} b d\tilde{\mu}_M.$$

Proof: The first inequality has been proven in Proposition 4.1.

Regarding the second one, again thanks to Proposition 4.1, $a^{-n/2} d\tilde{\nu}_M \leq d\nu_M \leq a^{n/2} d\tilde{\nu}_M$. By definition of (a, b) -closeness, we have $d\nu_M \geq a^{-n/2} d\tilde{\mu}_M = a^{-n/2} \rho_M^{-1} d\tilde{\nu}_M \geq a^{-n/2} b^{-1} d\tilde{\mu}_M$. And similarly we obtain that $d\nu_M \leq a^{n/2} b d\tilde{\mu}_M$, which finishes the proof. \square

We always consider (M, g_M, ν_M) as the original input Riemannian manifold and $(M, \tilde{g}_M, \tilde{\mu}_M)$ as a perturbed one. The magnitudes of the perturbations on the metric and on the measure are controlled by

constants a and b separately. If $(M, \tilde{g}_M, \tilde{\mu}_M)$ is $(1, b)$ -close to (M, g_M, ν_M) , then $\tilde{g}_M = g_M$. On the other hand, if the former is $(a, 1)$ -close to the latter, then the perturbation is purely on metric. However in this case, according to Proposition 4.1, the measure is perturbed as well. Lastly, if the former are $(1, 1)$ -close to the latter, then they are obviously isometric.

It's worth noting that our perturbation model assumes only the basic smoothness of \tilde{g}_M and ρ_M without more restrictive constraints such as bounded higher-order derivatives, thus it allows to include a large class of perturbed manifolds.

Now let (M, g_M, ν_M) and (N, g_N, ν_N) be a pair of original Riemannian manifolds, and we perturb them to $(M, \tilde{g}_M, \tilde{\mu}_M)$ and $(N, \tilde{g}_N, \tilde{\mu}_N)$, which are (a_M, b_M) -close and (a_N, b_N) -close to the original ones respectively. We then study how the shape difference operators constructed from the perturbed pair deviate from the ones constructed from the original pair.

4.4.2 Bounded-distortion Condition

Throughout our analyses in the coming sections of this chapter, we assume the input Riemannian manifolds, (M, g_M, ν_M) and (N, g_N, ν_N) , together with the map T between them satisfy the following bounded-distortion condition.

Condition 4.1 (*Bounded-distortion*) Let T_F be the functional map induced by $T : (M, g_M, \nu_M) \rightarrow (N, g_N, \nu_N)$, the area and conformal distortion induced by T_F (or equivalently by T) are bounded:

$$\text{For any } w \in L^2_\nu(N), \int_M T_F(w)^2 d\nu_M \leq B_T \int_N w^2 d\nu_N$$

$$\text{For any } w \in H^1_{0,\nu}(N), \int_M \langle \nabla T_F(w), \nabla T_F(w) \rangle_{g_M} d\nu_M \leq D_T \int_N \langle \nabla w, \nabla w \rangle_{g_N} d\nu_N$$

where B_T and D_T are finite positive constants.

The bounded-distortion condition rules out singular maps which can be meaningless, for example, a map T that maps M to a single point of N . Though we only regularize the distortions between the original input Riemannian manifolds, the following proposition ensures the boundness of distortions between a perturbed pair.

Proposition 4.4 $(M, \tilde{g}_M, \tilde{\mu}_M), (N, \tilde{g}_N, \tilde{\mu}_N)$ are respectively (a_M, b_M) -close and (a_N, b_N) -close to (M, g_M, ν_M) and (N, g_N, ν_N) , which are smooth Riemannian manifolds of dimension n . If $(M, g_M, \nu_M), (N, g_N, \nu_N)$ and T satisfy condition 4.1, then the distortion between the perturbed manifolds are bounded as the follows:

$$\text{For any } w \in L^2_{\tilde{\mu}}(N), \int_M T_F(w)^2 d\tilde{\mu}_M \leq \tilde{B}_T \int_N w^2 d\tilde{\mu}_N$$

$$\text{For any } w \in H^1_{0,\tilde{\mu}}(N), \int_M \langle \nabla T_F(w), \nabla T_F(w) \rangle_{\tilde{g}_M} d\tilde{\mu}_M \leq \tilde{D}_T \int_N \langle \nabla w, \nabla w \rangle_{\tilde{g}_N} d\tilde{\mu}_N$$

where $\tilde{B}_T = (a_M a_N)^{n/2} b_M b_N B_T$, $\tilde{D}_T = (a_M a_N)^{1+n/2} b_M b_N D_T$ and n is the dimension of the manifolds.

Proof: According to Proposition 4.3, we have $a_M^{-n/2}b_M^{-1}d\tilde{\mu}_M \leq d\nu_M \leq a_M^{n/2}b_M d\tilde{\mu}_M$ and $a_N^{-n/2}b_N^{-1}d\tilde{\mu}_N \leq d\nu_N \leq a_N^{n/2}b_N d\tilde{\mu}_N$. Since $T_F(w)^2$ and w^2 are both non-negative on, it follows from condition 4.1 that:

$$\begin{aligned} \int_M T_F(w)^2 d\tilde{\mu}_M &\leq \int_M T_F(w)^2 a_M^{n/2} b_M d\nu_M \leq a_M^{n/2} b_M B_T \int_N w^2 d\nu_N \\ &\leq a_M^{n/2} b_M B_T \int_N w^2 a_N^{n/2} b_N d\tilde{\mu}_N = (a_M a_N)^{n/2} b_M b_N B_T \int_N w^2 d\tilde{\mu}_N \\ &= \tilde{B}_T \int_N w^2 d\tilde{\mu}_N. \end{aligned}$$

The second inequality is proven in the same way:

$$\begin{aligned} \int_M \langle \nabla T_F(w), \nabla T_F(w) \rangle_{\tilde{g}_M} d\tilde{\mu}_M &\leq \int_M a_M \langle \nabla T_F(w), \nabla T_F(w) \rangle_{g_M} d\tilde{\mu}_M \\ &\leq a_M \int_M \langle \nabla T_F(w), \nabla T_F(w) \rangle_{g_M} a_M^{n/2} b_M d\nu_M \\ &\leq a_M^{1+n/2} b_M D_T \int_N \langle \nabla w, \nabla w \rangle_{g_N} d\nu_N \\ &\leq a_M^{1+n/2} a_N^{1+n/2} b_M b_N D_T \int_N \langle \nabla w, \nabla w \rangle_{\tilde{g}_N} d\tilde{\mu}_N \\ &= \tilde{D}_T \int_N \langle \nabla w, \nabla w \rangle_{\tilde{g}_N} d\tilde{\mu}_N. \end{aligned}$$

□

As we will demonstrate soon, condition 4.1 implies that both the shape difference operators, V and R , are bounded operators, i.e., their operator norms are finite. An important fact revealed by Proposition 4.4 is that if the operators constructed from the original pair of Riemannian manifolds are bounded, then the ones constructed from a perturbed pair are still bounded as long as the perturbations are finite.

4.5 Stability of the Shape Difference Operators

As in this chapter we are concentrating on Riemannian manifolds, for the sake of simplicity, from now on we denote by M the original Riemannian manifold $(M, g_M, d\nu_M)$ and by \tilde{M} the perturbed one $(M, \tilde{g}_M, \tilde{\mu}_M)$, unless stated otherwise.

The map T between manifolds is preserved while Riemannian structures are perturbed, and it induces two area-based shape difference operators, V and \tilde{V} , such that:

$$\int_N fV(g)d\nu_N = \int_M T_F(f)T_F(g)d\nu_M, \forall f, g \in L^2_\nu(N) \quad (4.20)$$

$$\int_N f\tilde{V}(g)d\tilde{\mu}_N = \int_M T_F(f)T_F(g)d\tilde{\mu}_M, \forall f, g \in L^2_{\tilde{\mu}}(N) \quad (4.21)$$

The stability of the area-based shape difference operator is stated in the following theorem:

Theorem 4.2 *Let M, N be two smooth n -dimensional Riemannian manifolds, and T be a map from M to N . Let \tilde{M} be (a_M, b_M) -close to M and \tilde{N} be (a_N, b_N) -close to N . If a_M, a_N, b_M and b_N are finite real numbers not smaller than 1, then $L_\nu^2(N) = L_\mu^2(N)$. Moreover, if M, N and T satisfy condition 4.1, we have the following convergence:*

$$\lim_{a_M, a_N, b_M, b_N \rightarrow 1^+} \int_N (Vg - \tilde{V}g)^2 d\nu_N = 0$$

Proof: We first prove $L_\nu^2(N) = L_\mu^2(N)$ so that $\tilde{V}g$ is well-defined for $g \in L_\nu^2(N)$. According to Proposition 4.3, we have $a_N^{-n/2} b_N^{-1} d\tilde{\mu}_N \leq d\nu_N \leq a_N^{n/2} b_N d\tilde{\mu}_N$. Then for any $f \in L_\nu^2(N)$, $\int_N f^2 d\tilde{\mu}_N \leq \int_N f^2 a_N^{n/2} b_N d\nu_N = a_N^{n/2} b_N \int_N f^2 d\nu_N < \infty$, therefore $L_\nu^2(N) \subset L_\mu^2(N)$. On the other hand, one can similarly verify that $L_\mu^2(N) \subset L_\nu^2(N)$, which implies $L_\nu^2(N) = L_\mu^2(N)$.

For $f \in L_\nu^2(N) = L_\mu^2(N)$, it follows from the triangle inequality that

$$\begin{aligned} & \left| \int_N fVf d\nu_N - \int_N f\tilde{V}f d\nu_N \right| \\ & \leq \left| \int_N fVf d\nu_N - \int_N f\tilde{V}f d\tilde{\mu}_N \right| + \left| \int_N f\tilde{V}f d\tilde{\mu}_N - \int_N f\tilde{V}f d\nu_N \right| \\ & := |P_1| + |P_2| \end{aligned}$$

Then we estimate P_1 and P_2 separately. According to Proposition 4.3, measures ν_M (resp. ν_N) and $\tilde{\mu}_M$ (resp. $\tilde{\mu}_N$) satisfy

$$a_M^{-n/2} b_M^{-1} d\tilde{\mu}_M \leq d\nu_M \leq a_M^{n/2} b_M d\tilde{\mu}_M$$

and

$$a_N^{-n/2} b_N^{-1} d\tilde{\mu}_N \leq d\nu_N \leq a_N^{n/2} b_N d\tilde{\mu}_N.$$

Thus, we have

$$\begin{aligned} P_1 &= \int_M T_F(f)^2 d\nu_M - \int_M T_F(f)^2 d\tilde{\mu}_M && \text{(by definitions of } V, \tilde{V}\text{)} \\ &\geq (1 - a_M^{n/2} b_M) \int_M T_F(f)^2 d\nu_M && \text{(since } a_M^{-n/2} b_M^{-1} d\tilde{\mu}_M \leq d\nu_M\text{)} \end{aligned}$$

And similarly $P_1 \leq (1 - a_M^{-n/2} b_M^{-1}) \int_M T_F(f)^2 d\nu_M$. Noticing that $0 \leq \int_M T_F(f)^2 d\nu_M \leq B_T \int_N f^2 d\nu_N < \infty$ as $f \in L_\nu^2(N)$, we have $|P_1|$ vanishes as $a_M, b_M \rightarrow 1^+$.

Regarding P_2 , we define two complementary subsets of N with respect to f : $I^+ = \{x \in N : f(x)\tilde{V}f(x) \geq 0\}$ and $I^- = \{x \in N : f(x)\tilde{V}f(x) < 0\}$.

$$\begin{aligned} P_2 &= \int_{I^+} f\tilde{V}f(d\tilde{\mu}_N - d\nu_N) + \int_{I^-} f\tilde{V}f(d\tilde{\mu}_N - d\nu_N) \\ &\leq (1 - a_N^{-n/2} b_N^{-1}) \int_{I^+} f\tilde{V}f d\tilde{\mu}_N + (1 - a_N^{n/2} b_N) \int_{I^-} f\tilde{V}f d\tilde{\mu}_N \end{aligned}$$

And we obtain the lower bound similarly:

$$\begin{aligned} P_2 &= \int_{I^+} f \tilde{V} f (d\tilde{\mu}_N - d\nu_N) + \int_{I^-} f \tilde{V} f (d\tilde{\mu}_N - d\nu_N) \\ &\geq (1 - a_N^{n/2} b_N) \int_{I^+} f \tilde{V} f d\tilde{\mu}_N + (1 - a_N^{-n/2} b_N^{-1}) \int_{I^-} f \tilde{V} f d\tilde{\mu}_N \end{aligned}$$

If $\int_{I^+} f \tilde{V} f d\tilde{\mu}_N, \int_{I^-} f \tilde{V} f d\tilde{\mu}_N$ are both finite, then P_2 vanishes as a_N, b_N tend to 1. In fact, by the formulation of \tilde{V} , both f and $\tilde{V}f$ are functions in $L^2_{\tilde{\mu}}(N)$, thus it follows from the Cauchy-Schwarz inequality that $|\int_{I^\pm} f \tilde{V} f d\tilde{\mu}_N| \leq \sqrt{\int_{I^\pm} f^2 d\tilde{\mu}_N \int_{I^\pm} (\tilde{V}f)^2 d\tilde{\mu}_N} \leq \sqrt{\int_N f^2 d\tilde{\mu}_N \int_N (\tilde{V}f)^2 d\tilde{\mu}_N} < \infty$.

To summarize, $\int f(Vf - \tilde{V}f) d\nu_N \rightarrow 0$ for any $f \in L^2_\nu(N) = L^2_{\tilde{\mu}}(N)$. Given $f, g \in L^2_\nu(N)$, since both V and \tilde{V} are self-adjoint operators, direct computation shows that

$$4 \int_N f(Vg - \tilde{V}g) d\nu_N = \int_N (f+g)(V - \tilde{V})(f+g) d\nu_N - \int_N (f-g)(V - \tilde{V})(f-g) d\nu_N$$

As $f+g, f-g \in L^2_\nu(N)$, it follows immediately that for any pair of f, g , $\int_N f(Vg - \tilde{V}g) d\nu_N$ vanishes as a_M, b_M, a_N , and b_N converge to 1 simultaneously. Especially we let $f = Vg - \tilde{V}g \in L^2_\nu(N)$, and conclude that $\lim_{a_M, a_N, b_M, b_N \rightarrow 1^+} \int_N (Vg - \tilde{V}g)^2 d\nu_N = 0$. \square

Similar stability guarantee holds for the conformal shape difference operators as well. We start with defining the conformal shape difference operators for both pairs of manifolds.

$$\int_N \langle \nabla f, \nabla R(g) \rangle_{g_N} d\nu_N = \int_M \langle \nabla T_F(f), \nabla T_F(g) \rangle_{g_M} d\nu_M, \forall f, g \in H^1_{0,\nu}(N) \quad (4.22)$$

$$\int_N \langle \nabla f, \nabla R(g) \rangle_{\tilde{g}_N} d\tilde{\mu}_N = \int_M \langle \nabla T_F(f), \nabla T_F(g) \rangle_{\tilde{g}_M} d\tilde{\mu}_M, \forall f, g \in H^1_{0,\tilde{\mu}}(N) \quad (4.23)$$

The following theorem suggests that as a_M, b_M, a_N, b_N converge to 1 simultaneously, the norm of the gradient of $\tilde{R}f - Rf$ converges to zero.

Theorem 4.3 *Let M, N and \tilde{M}, \tilde{N} be smooth n -dimensional manifolds under the same assumptions of Theorem 4.2, then for any $f \in H^1_{0,\nu}(N)$, $\tilde{R}f$ is well-defined. Moreover, we have*

$$\lim_{a_M, b_M, a_N, b_N \rightarrow 1^+} \int_N \langle \nabla(Rf - \tilde{R}f), \nabla(Rf - \tilde{R}f) \rangle_{g_N} d\nu_N = 0$$

This theorem is proven with the same idea of proving Theorem 4.2, we refer the readers to Section 4.8 for a detailed proof.

Remark 4.1 *Our proofs for theorems 4.2 and 4.3 do not require the involved manifolds to be compact or boundaryless. The stability properties proven in this section are valid for any smooth n -dimensional Riemannian manifolds and maps satisfying condition 4.1.*

4.6 Stability of the Shape Difference Operators in a Multi-Scale Framework

In this section, we study the stability properties of the shape difference operators in the framework of [Ovsjanikov 2013], where they are employed in a multi-scale way. The main results of this section are Theorem 4.4 and Theorem 4.6, of which the proofs heavily depend on properties of the (weighted) Laplace-Beltrami operators introduced in Section 4.3.2, thus we follow the assumptions there and emphasize that in this section, (N, g_N) is a connected, compact, smooth, n -dimensional Riemannian manifold without boundary.

We start by pointing out the connection between the multi-scale framework and the shape difference operators. Recall that Definition 4.15 defines a functional measuring the area distortion induced by T for a given function $w \in L^2_\nu(N)$. Given a pair of manifolds M, N and a map $T : M \rightarrow N$, let V be the area-based shape difference operator. It follows directly from Definition 4.18 that:

$$E(w) = \frac{\int_M T_F(w)^2 d\nu_M}{\int_N w^2 d\nu_N} = \frac{\int_N wV(w) d\nu_N}{\int_N w^2 d\nu_N}$$

Since V is a positive-definite self-adjoint operator acting on $L^2_\nu(N)$, the maximum of $E(w)$ within $L^2_\nu(N)$ is simply the L^2 -norm of V . The framework of [Ovsjanikov 2013] computes the *constrained norm* of V with respect to a special collection of subdomains of $L^2_\nu(N)$ – the function spaces spanned by the first several eigenfunctions of the LB operator on N . In general, given a subdomain Ω of $L^2_\nu(N)$, the maximum of $E(w)$ constrained in Ω provides a quantitative characterization of to what extent V can distort functions in Ω . The maximizer (which we call the indicator after), w^* , is a function in Ω that is the most distorted by V .

4.6.1 An Alternative Collection of Multi-scale Subdomains

Recall that in Section 4.3.4, we review the subdomain constructed at scale k : $S(k) = \text{span}\{\varphi_1, \varphi_2, \dots, \varphi_k\}$, where φ_i is the i -th eigenfunction of the LB operator on N . Despite several advantages listed in Section 4.3.4, this subdomain construction suffers issues that are rooted in its discrete nature and is not suitable for our stability analysis.

First, we do not have a general criterion of selecting k . In general, the spectrum varies across shapes that are not isometric to each other, therefore we need to choose k according to the spectrum of the Laplace-Beltrami operator on the given input shape N .

Second, it can lead to confusing results when the truncation is not done appropriately. If a eigenvalue is of multiplicity more than 1, then as discussed in Section 4.3.2, there might be a subspace spanned by the eigenfunctions associated to this eigenvalue, which is of dimension more than 1. Though we argue that we can pick a orthogonal basis for this subspace, there does not exist a natural order for this basis functions (since the associated eigenvalues are identical) and thus we order them arbitrarily. Now if we truncate in the middle of this subspace, the resulting subdomain $S(k)$ carries the randomness of the ordering. Therefore the constrained norm with respect to the space spanned by the first several eigenfunctions is not even well-defined, In practice, we observe instability in the more subtle case of analyzing conformal differences (see Figure 4.3).

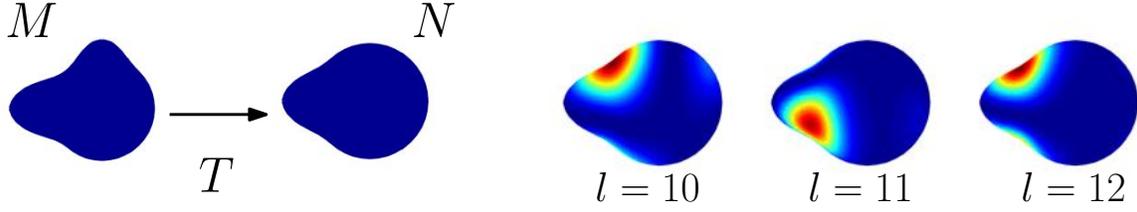


Figure 4.3: Indicators with respect to conformal shape difference operator depicted on the N at scale $l = 10, 11$ and 12 . $\lambda_9 < \lambda_{10} < \lambda_{11} \approx \lambda_{12} < \lambda_{13}$. λ_{11} is numerically close to λ_{12} (their difference is of order 10^{-5}), causing the instability in the indicator functions.

Lastly, since k is an integer, since $\lambda_{k+1} - \lambda_k$ can be large, small, or even 0, it is difficult to analyze the changes from $S(k)$ to $S(k+1)$, though this is the minimal perturbation in discrete scales already.

To overcome these issues, we construct a new collection of multi-scale subdomains which evolves continuously. We notice that if (φ_i, λ_i) is an eigensolution to $\Delta_N \varphi + \lambda \varphi = 0$, then $-\int_N \varphi_i \Delta_N \varphi_i d\nu_N = \lambda_i \int_N \varphi_i^2 d\nu_N$. Moreover, thanks to the orthogonality of the eigenbasis $\{\varphi_i\}_{i \geq 1}$, it's straightforward to verify that for any $w \in \text{span}\{\varphi_1, \varphi_2, \dots, \varphi_k\}$, $-\int_N w \Delta_N w d\nu_N \leq \lambda_k \int_N w^2 d\nu_N$

It's then natural to consider the following multi-scale subdomains controlled by a continuous parameter C :

$$A(C) = \left\{ w : \int_N w(-\Delta_N)w d\nu_N \leq C \int_N w^2 d\nu_N \right\} \quad (4.24)$$

In fact, thanks to the Green formula 4.10, $A(C)$ can be equivalently defined as:

$$A(C) = \left\{ w : \int_N \langle \nabla w, \nabla w \rangle_{g_M} d\nu_N \leq C \int_N w^2 d\nu_N \right\}$$

From this point of view, this expression suggests that (the normalized) Dirichlet's energy of $w \in A(C)$ is upper-bounded by C . In general, a small C prohibits large variations of w over a short distance with a global control of the magnitude of the gradient of w , therefore it forces $w \in A(C)$ to be *smooth*.

Proposition 4.5 *If $C \geq \lambda_k$, then $S(k)$ is a proper subset of $A(C)$.*

Proof: It follows from the construction that if $C \geq \lambda_k$, then $S(k)$ is a subset of $A(C)$. We further argue that $S(k)$ is a proper subset: let $w_\varepsilon = (1 - \varepsilon)\varphi_1 + \varepsilon\varphi_{k+1}$, then $w \notin S(k)$. Notice that since N is connected, $\lambda_1 = 0$, thus for ε sufficiently close to 0, $w_\varepsilon \in A(\lambda_k)$. \square

4.6.2 Stability with respect to the Changes in Scale

We first investigate the stability with respect to the changes in scale, which only involves the original input manifolds M and N . As demonstrated in Figure 4.1 and 4.2, the results show consistency of the areas on N highlighted by the indicators across a range of scales. It is then attempting to validate the stability of the

indicators. However, it is not always the case. For example, imagine that we deform the bottom of shape M in Figure 4.1 so that the deformation from M to N is symmetrical. In this case, at every scale, the maximum of $E(w)$ is realized by two indicators w_t, w_b which highlight respectively the top and the bottom of shape N , therefore we will no longer observe consistency in indicators. We then turn to study the stability of the maxima with respect to the changes in scale.

Instead of the original collection of subdomains $S(k)$ indexed by the integer k , we consider the new collection of multi-scale subdomains $A(C)$ that is controlled by the continuous parameter C . As mentioned before, the maximum of $E(w)$ constrained in a subdomain of $L^2_\nu(N)$ is a constrained norm of V . For subdomain $A(C)$, we define:

$$\|V\|_C = \max E(w) \text{ s.t. } w \in A(C)$$

Let C go through interval $[0, +\infty)$, then we obtain a curve $(C, \|V\|_C)$. The following theorem justifies its continuity.

Theorem 4.4 *Given two connected compact smooth Riemannian manifolds M and N , and a map T between them. If M, N, T satisfy condition 4.1, then for any positive constant $C > 0, C' = C + \varepsilon > 0$, we have:*

$$\left| \|V\|_{C'} - \|V\|_C \right| \leq 4B_T \sqrt{|\varepsilon|/C} + 2B_T |\varepsilon|/C.$$

We first prove an auxiliary lemma.

Lemma 4.1 *Under the same assumptions of Theorem 4.4, for any $w_0, w_1 \in L^2_\nu(N)$ such that neither $\int_N w_0^2 d\nu_N$ nor $\int_N (w_0 + w_1)^2 d\nu_N$ is zero, the absolute difference between $E(w_0 + w_1)$ and $E(w_0)$ is bounded as below:*

$$|E(w_0 + w_1) - E(w_0)| \leq 4B_T \sqrt{r} + 2B_T r$$

where B_T is the constant in condition 4.1 and $r = \frac{\int_N w_1^2 d\nu_N}{\int_N w_0^2 d\nu_N}$.

Proof: We estimate the difference by two parts:

$$\begin{aligned} |E(w_0 + w_1) - E(w_0)| &\leq \left| E(w_0 + w_1) - \frac{\int_M T_F(w_0 + w_1)^2 d\nu_M}{\int_N w_0^2 d\nu_N} \right| + \left| \frac{\int_M T_F(w_0 + w_1)^2 d\nu_M}{\int_N w_0^2 d\nu_N} - E(w_0) \right| \\ &:= P_1 + P_2 \end{aligned}$$

For P_1 , direct computation shows:

$$\begin{aligned} P_1 &\leq B_T \int_N (w_0 + w_1)^2 d\nu_N \left| \frac{1}{\int_N w_0^2 d\nu_N} - \frac{1}{\int_N (w_0 + w_1)^2 d\nu_N} \right| \\ &\leq B_T \frac{2|\int_N w_0 w_1 d\nu_N| + \int_N w_1^2 d\nu_N}{\int_N w_0^2 d\nu_N} \\ &\leq 2B_T \frac{\sqrt{\int_N w_0^2 d\nu_N \int_N w_1^2 d\nu_N}}{\int_N w_0^2 d\nu_N} + B_T \frac{\int_N w_1^2 d\nu_N}{\int_N w_0^2 d\nu_N} \\ &\leq 2B_T \sqrt{r} + B_T r \end{aligned}$$

The third line follows from the Cauchy-Schwarz inequality applied on w_0 and w_1 which are both in $L^2_\nu(N)$.

P_2 is bounded in the same way:

$$\begin{aligned} P_2 &\leq \left| \frac{\int_M T_F(2w_0w_1 + w_1^2)d\nu_M}{\int_N w_0^2d\nu_N} \right| \\ &\leq \frac{2|\int_M T_F(w_0)T_F(w_1)d\nu_M| + B_T \int_N w_1^2d\nu_N}{\int_N w_0^2d\nu_N} \\ &\leq 2B_T \frac{\sqrt{\int_N w_0^2d\nu_N \int_N w_1^2d\nu_N}}{\int_N w_0^2d\nu_N} + B_T \frac{\int_N w_1^2d\nu_N}{\int_N w_0^2d\nu_N} \\ &\leq 2B_T\sqrt{r} + B_T r \end{aligned}$$

Putting them together yields $|E(w_0 + w_1) - E(w_0)| \leq 4B_T\sqrt{r} + 2B_T r$, where $r = \frac{\int_N w_1^2d\nu_N}{\int_N w_0^2d\nu_N}$ \square

We then prove Theorem 4.4:

Proof: We first consider $\varepsilon > 0$, i.e., $C' > C$. By definition, $A(C) \subset A(C')$, thus $\|V\|_{C'} - \|V\|_C \geq 0$. We then estimate the upper bound for the difference. Given $w \in A(C + \varepsilon)$, our strategy is to construct a function $\bar{w} \in A(C)$, such that $|E(w) - E(\bar{w})|$ is bounded.

If w itself lies in $A(C)$, then it's trivial to set $\bar{w} = w$. We now consider $w \in A(C + \varepsilon) \setminus A(C)$. We assume that $\int_N w(-\Delta_N)w d\nu_N = (C + \delta) \int_N w^2 d\nu_N$, where $0 < \delta \leq \varepsilon$. Since $E(w) = E(cw)$, $\forall c \neq 0$, without loss of generality, we further assume that $w = \sum_{i \geq 1} a_i \varphi_i$ and $\sum_{i \geq 1} a_i^2 = 1$, where (φ_i, λ_i) is the i -th eigensolution to $\Delta_N \varphi + \lambda \varphi = 0$. According to Proposition 4.2, the constraint on w can be written as:

$$\sum_{i \geq 2} a_i^2 \lambda_i = C + \delta.$$

Let b_1 be a real number satisfying $b_1^2 - a_1^2 = \delta/C$ and $b_1 a_1 \geq 0$. The existence of b_1 is assured by the fact that a_1 is finite (in fact $|a_1| \leq 1$). Then we set $\bar{w} = b_1 \varphi_1 + \sum_{i \geq 2} a_i \varphi_i$. Direct computation shows that $\int_N \bar{w}(-\Delta_N)\bar{w} d\nu_N = \sum_{i \geq 2} a_i^2 \lambda_i = C + \delta$ and $\int_N \bar{w}^2 d\nu_N = b_1^2 + \sum_{i \geq 2} a_i^2 = b_1^2 + 1 - a_1^2 = 1 + \frac{\delta}{C} = \frac{C+\delta}{C} = \frac{1}{C} \int_N \bar{w}(-D_N)\bar{w} d\nu_N$, meaning that $\bar{w} \in A(C)$.

Thanks to Lemma 4.1, $|E(w) - E(\bar{w})| = |E(w) - E(w + \bar{w} - w)| \leq 4B_T\sqrt{r} + 2B_T r$, where $r = \int_N (\bar{w} - w)^2 d\nu_N / \int_N w^2 d\nu_N = (b_1 - a_1)^2$. Without loss of generality, we assume that $a_1, b_1 \geq 0$, thus $b_1 = \sqrt{a_1^2 + \delta/C} > a_1$. Moreover, $b_1 - a_1 = \sqrt{a_1^2 + \delta/C} - a_1 = \frac{\delta/C}{\sqrt{a_1^2 + \delta/C} + a_1} \leq \sqrt{\delta/C}$. Therefore $r \leq \delta/C \leq \varepsilon/C$. Now assuming that w^* is the maximizer of $E(w)$ constrained in $A(C + \varepsilon)$, the above derivation shows that there exists a $\bar{w}^* \in A(C)$ such that $|E(w^*) - E(\bar{w}^*)| \leq 4B_T\sqrt{\varepsilon/C} + 2B_T\varepsilon/C$. That implies $\|V\|_{C+\varepsilon} - \|V\|_C \leq E(w^*) - E(\bar{w}^*) \leq |E(w^*) - E(\bar{w}^*)| \leq 4B_T\sqrt{\varepsilon/C} + 2B_T\varepsilon/C$.

Regarding the case $\varepsilon < 0$, i.e., $C' < C$. We simply replace C and $C + \varepsilon$ in the previous analysis with $C - \varepsilon$ and C , respectively. With identical derivations, for each $w \in A(C)$, we construct a function $\bar{w} \in A(C - \delta)$ such that $|E(w) - E(\bar{w})| = |E(w) - E(w + \bar{w} - w)| \leq 4B_T\sqrt{r} + 2B_T r$, where $r \leq |\varepsilon|/C$. Similarly, we have $\|V\|_C - \|V\|_{C+\varepsilon} \leq 4B_T\sqrt{|\varepsilon|/C} + 2B_T|\varepsilon|/C$ for $\varepsilon < 0$.

Putting them together, we finish the proof of this theorem. \square

Notice that B_T is in fact an upper-bound for the constrained norms, i.e., $\|V\|_C \leq B_T, \forall C > 0$. Thus the inequality proven in Theorem 4.4 only makes sense when ε is close to zero. On the other hand, the inequality suggests that for a perturbation of fixed magnitude $|\varepsilon|$, the larger C is, the more stable $\|V\|_C$ is.

4.6.3 Stability With Respect To Perturbed Input Manifolds

In this section, we fix the scale C and add perturbations on M and N . As in Section 4.5, we perturb M and N to \tilde{M} and \tilde{N} , which are (a_M, b_M) -close and (a_N, b_N) -close to the unperturbed ones respectively. V and \tilde{V} are the corresponding area-based shape difference operators defined in Equation 4.20 and 4.21.

In order to define the constrained norm for \tilde{V} , we first construct the corresponding functional $\tilde{E}(w)$, which is obtained by changing measures in $E(w)$.

$$\tilde{E}(w) = \frac{\int_N w \tilde{V}(w) d\tilde{\mu}_N}{\int_N w^2 d\tilde{\mu}_N} = \frac{\int_M T_F(w)^2 d\tilde{\mu}_M}{\int_N w^2 d\tilde{\mu}_N} \quad (4.25)$$

The corresponding subdomain $\tilde{A}(C)$ can be written with respect to the weighted Laplace-Beltrami operator (see Definition 4.13). We denote by $\tilde{\Delta}_N$ the Laplace-Beltrami operator on the weighted manifold $(N, \tilde{g}_N, \tilde{\mu}_N)$ and define

$$\tilde{A}(C) = \{w : \int_N w(-\tilde{\Delta}_N)w d\tilde{\mu}_N \leq C \int_N w^2 d\tilde{\mu}_N\} \quad (4.26)$$

One advantage of our construction of subdomains is that it allows to associate subdomains with respect to different manifolds. Particularly we observe the following interleaved structures between $A(C)$ and $\tilde{A}(C)$:

Lemma 4.2 *Let \tilde{N} be a n -dimensional Riemannian manifold that is (a_N, b_N) -close to N , the two collections of subdomains defined in 4.24 and 4.26 are interleaved:*

$$A(C) \subset \tilde{A}(Ca_N^{1+n}b_N^2) \subset A(Ca_N^{2+2n}b_N^4)$$

Proof: For $w \in A(C)$, it follows from the Definition 4.24 that $\int_N \langle \nabla w, \nabla w \rangle_{g_N} d\nu_N \leq C \int_N w^2 d\nu_N$. According to Proposition 4.3, we have

$$\begin{aligned} \int_N \langle \nabla w, \nabla w \rangle_{\tilde{g}_N} d\tilde{\mu}_N &\leq \int_N a_N \langle \nabla w, \nabla w \rangle_{g_N} d\tilde{\mu}_N \\ &\leq \int_N a_N \langle \nabla w, \nabla w \rangle_{g_N} a_N^{n/2} b_N d\nu_N \\ &\leq Ca_N^{1+n/2} b_N \int_N w^2 d\nu_N \\ &\leq Ca_N^{1+n/2} b_N \int_N w^2 a_N^{n/2} b_N d\tilde{\mu}_N \\ &\leq Ca_N^{1+n} b_N^2 \int_N w^2 d\tilde{\mu}_N. \end{aligned}$$

meaning that $w \in \tilde{A}(Ca_N^{1+n}b_N^2)$, thus $A(C) \subset \tilde{A}(Ca_N^{1+n}b_N)$. Similarly we can derive the other inclusion relationship. \square

On the other hand, the functionals, $E(w)$ and $\tilde{E}(w)$ are as well related to each other.

Lemma 4.3 *For any $w \in L^2_\nu(N)$ satisfying $E(w) > 0$, the ratio of $\tilde{E}(w)$ to $E(w)$ is two-sided bounded:*

$$a_M^{-n/2} a_N^{-n/2} b_M^{-1} b_N^{-1} \leq \tilde{E}(w)/E(w) \leq a_M^{n/2} a_N^{n/2} b_M b_N$$

Proof: According to Proposition 4.3, we have

$$a_M^{-n/2} b_M^{-1} d\tilde{\mu}_M \leq d\nu_M \leq a_M^{n/2} b_M d\tilde{\mu}_M$$

and

$$a_N^{-n/2} b_N^{-1} d\tilde{\mu}_N \leq d\nu_N \leq a_N^{n/2} b_N d\tilde{\mu}_N$$

And by definition:

$$\begin{aligned} \tilde{E}(w) &= \frac{\int_M T_F(w)^2 d\tilde{\mu}_M}{\int_N w^2 d\tilde{\mu}_N} \leq \frac{\int_M T_F(w)^2 a_M^{n/2} b_M d\nu_M}{\int_N w^2 a_N^{-n/2} b_N^{-1} d\nu_N} \\ &\leq a_M^{n/2} a_N^{n/2} b_M b_N E(w). \end{aligned}$$

Symmetrically, the other side is bounded:

$$\begin{aligned} \tilde{E}(w) &= \frac{\int_M T_F(w)^2 d\tilde{\mu}_M}{\int_N w^2 d\tilde{\mu}_N} \geq \frac{\int_M T_F(w)^2 a_M^{-n/2} b_M^{-1} d\nu_M}{\int_N w^2 a_N^{n/2} b_N d\nu_N} \\ &\geq a_M^{-n/2} a_N^{-n/2} b_M^{-1} b_N^{-1} E(w). \end{aligned}$$

□

Based on the above constructions of $\tilde{A}(C)$ and $\tilde{E}(w)$, the constrained norm in the perturbed case is defined as $\|\tilde{V}\|_C = \max \tilde{E}(w)$ s.t. $w \in \tilde{A}(C)$. The main result of this section is stated in the following theorem, which claims that at each fixed scale C , the constrained norm is stable with respect to perturbations on the manifolds.

Theorem 4.5 *Let M, N be two connected compact smooth n -dimensional Riemannian manifolds without boundary, and T be a map from M to N . Let \tilde{M} (resp. \tilde{N}) be a smooth manifold that is (a_M, b_M) -close (resp. (a_N, b_N) -close) to M (resp. N). V and \tilde{V} are the area-based shape difference operators constructed with M, N and \tilde{M}, \tilde{N} respectively. If M, N, T satisfy condition 4.1, then at any fixed scale C , the following convergence is valid:*

$$\lim_{a_M, b_M, a_N, b_N \rightarrow 1^+} \|\tilde{V}\|_C = \|V\|_C$$

Proof: Given a fixed scale $C > 0$, we denote $C_1 = C a_N^{-1-n} b_N^{-2}$ and $C_2 = C a_N^{1+n} b_N^2$. Let $u_1 \in A(C_1)$, $u_2 \in A(C_2)$ and $v \in \tilde{A}(C)$ be functions satisfying $E(u_1) = \|V\|_{C_1}$, $E(u_2) = \|V\|_{C_2}$ and $\tilde{E}(v) = \|\tilde{V}\|_C$.

Thanks to Lemma 4.2, we have $A(C_1) \subset \tilde{A}(C)$, thus according to lemma 4.3:

$$\frac{\|\tilde{V}\|_C}{\|V\|_{C_1}} = \frac{\tilde{E}(v)}{E(u_1)} \geq \frac{\tilde{E}(u_1)}{E(u_1)} \geq a_M^{-n/2} a_N^{-n/2} b_M^{-1} b_N^{-1}$$

. Lemma On the other hand, as $\tilde{A}(C) \subset A(C_2)$, we have:

$$\frac{\|\tilde{V}\|_C}{\|V\|_{C_2}} = \frac{\tilde{E}(v)}{E(u_2)} \leq \frac{\tilde{E}(v)}{E(v)} \leq a_M^{n/2} a_N^{n/2} b_M b_N$$

Lastly, Putting the above two inequalities together, we have

$$a_M^{-n/2} a_N^{-n/2} b_M^{-1} b_N^{-1} \|V\|_{C_1} \leq \|\tilde{V}\|_C \leq a_M^{n/2} a_N^{n/2} b_M b_N \|V\|_{C_2}$$

It follows from Theorem 4.2 that

$$\left| \|V\|_{C_2} - \|V\|_{C_1} \right| \leq 4B_T \sqrt{\frac{C_2 - C_1}{C_1}} + 2B_T \frac{C_2 - C_1}{C_1}$$

Obviously, letting a_N, b_N tend to 1, we have $\|V\|_{C_2} \rightarrow \|V\|_{C_1}$. Moreover, as a_M, b_M tend to 1 as well, according to the squeeze theorem we have

$$\lim_{a_M, b_M, a_N, b_N \rightarrow 1^+} \|\tilde{V}\|_C = \|V\|_C$$

□

Remark 4.2 M, N, T satisfying condition 4.1 guarantees that $\|V\|_{C_1}$ and $\|V\|_{C_2}$ are finite. The finiteness of $\|\tilde{V}\|_C$ is assured by Proposition 4.4.

4.6.4 Approximating $\|V\|_C$

By investigating the behavior of operator within the continuously evolving subdomains $A(C)$, we have more stable and richer understanding of V than we've got from $S(k)$. However, in practice, calculating $\|V\|_C$ is far from being obvious. Since neither $E(w)$ nor $A(C)$ is convex, there is no guarantee on achieving the global optimum with constraint $A(C)$.

For the sake of consistency, we denote by $\|V\|_k$ the maximum of $E(w)$ within subdomain $S(k)$. As discussed in [Ovsjanikov 2013], computing $\|V\|_k$ in the case where M and N are finite discrete meshed shapes is straightforward and is solved with an eigen-decomposition of a k by k matrix.

First note that the construction of $A(C)$ and $S(k)$ are closely related. The following proposition quantifies this relationship.

Proposition 4.6 Let M, N and T be a pair of manifolds and a map between them, which satisfy condition 4.1. Let λ_k, λ_{k+1} be two consecutive eigenvalues of the LB operator on N , the constrained norms with respect to $A(\lambda_k)$ and $S(k)$ satisfy the following inequality:

$$0 \leq \|V\|_{\lambda_k} - \|V\|_k \leq 4B_T \sqrt{\lambda_k / \lambda_{k+1}} + 2B_T \lambda_k / \lambda_{k+1}$$

Proof: First of all, as shown in Proposition 4.5, $S(k)$ is a proper subset of $A(\lambda_k)$, which proves the left-side inequality.

Regarding the right-side inequality, we assume that w is the maximizer realizing $\|V\|_{\lambda_k}$ and decompose it into the eigenbasis: $w = \sum_{i \geq 1} a_i \varphi_i$. Now let $\bar{w} = \sum_{i \leq k} a_i \varphi_i$, obviously $\bar{w} \in S(k)$ and therefore $E(\bar{w}) \leq \|V\|_k$.

We then estimate the difference $E(w) - E(\bar{w})$, according to Lemma 4.1, it's upper-bounded by $4B_T \sqrt{r} + 2B_T r$, where $r = \sum_{i \geq k+1} a_i^2 / \sum_{i \geq 1} a_i^2$. As $w \in A(\lambda_k)$, it follows that $\lambda_k \sum_{i \geq 1} a_i^2 \geq \sum_{i \geq 1} a_i^2 \lambda_i \geq$

$\sum_{i \geq k+1} a_i^2 + \lambda_i \geq \lambda_{k+1} \sum_{i \geq k+1} a_i^2$, thus $r \leq \lambda_k / \lambda_{k+1}$. Plugging it back, we have $\|V\|_{\lambda_k} - \|V\|_k \leq E(w) - E(\bar{w}) \leq 4B_T \sqrt{\lambda_k / \lambda_{k+1}} + 2B_T \lambda_k / \lambda_{k+1}$. \square

As a direct corollary, if $\lambda_k / \lambda_{k+1} \approx 0$, then $\|V\|_k$ is a good approximation of $\|V\|_{\lambda_k}$.

It's worth noting that this proposition indicates a general criterion of choosing k : it's preferable to choose k such that the gap between λ_k and λ_{k+1} is significant. And as we will discuss soon, this proposition suggests that if the spectral gap is significant, then the maximizer realizing $\|V\|_k$ is a nice candidate of the initial guess for iterative algorithms maximizing $E(w)$ constrained in $A(\lambda_k)$.

Secondly, a major obstacle of optimizing within $A(C)$ is that it is of infinite dimension. Even in the discrete case, the problem scale is still determined by the number of points, which can range in the tens or hundreds of thousands. The following proposition suggests that there is a trade-off between complexity and accuracy:

Proposition 4.7 *For a fixed parameter C , let $\varepsilon > 0$ and λ_{l+1} be the smallest eigenvalue of the LB operator on N such that $C \leq \varepsilon \lambda_{l+1}$. Let*

$$\|V\|_{C,l} = \max E(w) \text{ s.t. } w \in A(C) \cap S(l) \quad (4.27)$$

, then $\|V\|_C - \|V\|_{C,l}$ is of order $\sqrt{\varepsilon}$.

Proof: Let function $w \in A(C)$ satisfies $E(w) = \|V\|_C$. We assume that $w = \sum_{i \geq 1} a_i \varphi_i$, and let \bar{w} be $\sum_{i \leq l} a_i \varphi_i$. Then obviously $E(w) - \|V\|_{C,l} \leq E(w) - E(\bar{w})$, because $\bar{w} \in A(C) \cap S(l)$ as well.

Repeating the proof of Proposition 4.6, we have $\|V\|_C - E(\bar{w}) \leq 4B_T \sqrt{r} + 2B_T r$, where $r = \sum_{i \geq l+1} a_i^2 / \sum_{i \geq 1} a_i^2$. Similarly we can deduce that $r \leq C / \lambda_{l+1} \leq \varepsilon$, therefore $\|V\|_C - \|V\|_{C,l} \leq 4B_T \sqrt{\varepsilon} + 2B_T \varepsilon$. \square

Optimization of problem 4.27 is then conducted in a space of limited dimension, l , which is controlled by the users.

A brief summary: Before we proceed to analyze the stability properties of the conformal shape difference operators under this multi-scale framework, we make a brief review of the following results on the area-based shape difference operators.

- Section 4.6.2: $(C, \|V\|_C)$ is a continuous curve and always above point sets $\{(\lambda_k, \|V\|_k)\}_{k=1}^{+\infty}$. In Figure 4.4, the black curve is continuous and always above the blue points.
- Section 4.6.3: As \tilde{M} converges to M and \tilde{N} converges to N , any point $(C, \|\tilde{V}\|_C)$ on the black curve in Figure 4.4 converges to point $(C, \|V\|_C)$ on the red dashed one.
- Section 4.6.4: Discussion on the optimization problem: $\max E(w)$ s.t. $w \in A(C)$.

4.6.5 Analysis for the Conformal Shape Difference Operator

In essence, with functional $E(w)$, the framework of [Ovsjanikov 2013] casts the problem of extracting information on shape (manifold) differences as a series of constrained optimization problems. The functional evaluates how the area-based shape difference operator *distorts* functions living in $L^2_\nu(N)$.

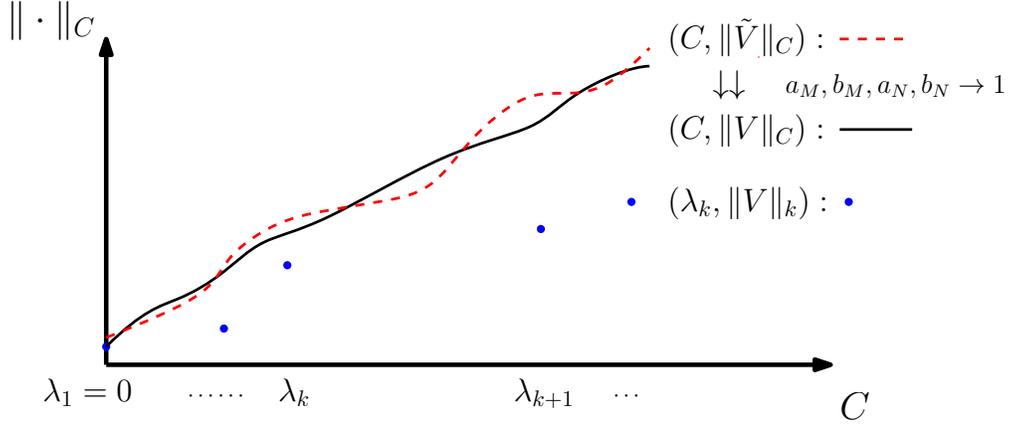


Figure 4.4: Curves representing constrained norms with respect to different subdomains construction and input manifolds, due to Proposition 4.6, if $C = \lambda_k \ll \lambda_{k+1}$, then the difference between $\|V\|_C$ and $\|V\|_k$ is well-bounded.

Note that the framework of [Rustamov 2013] introduces two shape difference operators which encode different types of differences between shapes (manifolds), a natural extension of the multi-scale framework of [Ovsjanikov 2013] is to construct parallel functionals and subdomains with respect to the conformal shape difference operators, R .

Given a pair of manifolds M, N , and a map T between them, we define a functional, F , acting on $H_{0,\nu}^1(N)$ (the domain of the conformal shape difference operator) as the following:

$$F(w) = \frac{\int_N \langle \nabla w, \nabla R(w) \rangle_{g_N} d\nu_N}{\int_N \langle \nabla w, \nabla w \rangle_{g_N} d\nu_N} = \frac{\int_M \langle \nabla T_F(w), \nabla T_F(w) \rangle_{g_M} d\nu_M}{\int_N \langle \nabla w, \nabla w \rangle_{g_N} d\nu_N}. \quad (4.28)$$

The second inequality follows from the definition of R .

The following lemma about $F(w)$ is a counterpart of Lemma 4.1.

Lemma 4.4 *Let M, N be two connected compact smooth Riemannian manifolds, and T be a map between them. If M, N, T satisfy condition 4.1, and for any $w_0, w_1 \in H_{0,\nu}^1(N)$ such that neither w_0 nor $w_0 + w_1$ is a constant function on N , then the difference between $F(w_0 + w_1)$ and $F(w_0)$ is bounded as below:*

$$|F(w_0 + w_1) - F(w_0)| \leq 4D_T \sqrt{s} + 2D_T s$$

where D_T is the constant in condition 4.1 and $s = \frac{\int_N \langle \nabla w_1, \nabla w_1 \rangle_{g_N} d\nu_N}{\int_N \langle \nabla w_0, \nabla w_0 \rangle_{g_N} d\nu_N}$.

On the other hand, modifying the multi-scale subdomain construction is necessary to suit the new functional. Let $\|R\|_C = \max F(w)$ s.t. $w \in A(C)$, and in the following we prove that $\|R\|_{C'} = \|R\|_C, \forall C' > C > 0$. In fact, given C' and a function on the boundary of $A(C')$, repeating the proof of Theorem 4.4, we find another function \bar{w} lying in $A(C)$ and $w - \bar{w}$ is a constant function. According to Lemma 4.4, we have $F(w) = F(\bar{w})$ since $\int_N \langle \nabla(w - \bar{w}), \nabla(w - \bar{w}) \rangle_{g_N} d\nu_N = 0$. Therefore, $\|R\|_C = \|R\|_{C'}$ if we use $A(C)$ as the multi-scale constraint.

To solve this issue, we construct a new subdomain, $A^{conf}(C)$, which is orthogonal to the space spanned by constant functions.

$$A^{conf}(C) = A(C) \cap \{w : \int_N w d\nu_N = 0\} \quad (4.29)$$

and let $\|R\|_C = \max F(w)$ s.t. $w \in A^{conf}(C)$.

It's worth noting that if $C < \lambda_2$, the first non-zero eigenvalue of $-\Delta_N$, then $A^{conf}(C) = \emptyset$. That follows from Proposition 4.2: let $w = \sum_{i \geq 1} a_i \varphi_i$, since $\int_N w d\nu_N = 0$, $a_1 = 0$. Therefore $\int_N w(-\Delta_N)w d\nu_N =$

$\sum_{i \geq 2} a_i^2 \lambda_i \geq \sum_{i \geq 2} a_i^2 \lambda_2 \geq \lambda_2 \sum_{i \geq 2} a_i^2 = \lambda_2 \int_N w^2 d\nu_N$. In the other words, if $C < \lambda_2$, then $A(C) \cap \{w : \int_N w d\nu_N = 0\}$ is empty. Thus C must be at least λ_2 so that $\|R\|_C$ is well-defined. In practice, it is more convenient to maximize $F(w)$ in the subdomains spanned by finite number of eigenfunctions. Following the same arguments above, we modify $S(k)$ to obtain $S^{conf}(k) = \text{span}\{\varphi_2, \varphi_3, \dots, \varphi_k\}$, where k is at least 2.

With the above formulations, we validate the stability of R with respect to the changes in scale.

Theorem 4.6 *Let M, N be two connected compact smooth Riemannian manifolds with boundary, and T be a map between them. Let λ_2 be the second eigenvalue (as well as the first non-zero eigenvalue) of $-\Delta_N$. If M, N, T satisfy condition 4.1, then for $C > \lambda_2$, $C' = C + \varepsilon > \lambda_2$ we have:*

$$\left| \|R\|_{C'} - \|R\|_C \right| \leq 4D_T \sqrt{\frac{\lambda_2 |\varepsilon|}{(C - \lambda_2)(C - |\varepsilon|)}} + 2D_T \frac{\lambda_2 |\varepsilon|}{(C - \lambda_2)(C - |\varepsilon|)}$$

Then we consider perturbations on the input manifolds. As before, we denote by \tilde{M} and \tilde{N} the perturbed version of M and N . The perturbed conformal shape difference operator, \tilde{R} , is defined in Equation 4.23. The associated functional, $\tilde{F}(w)$, is defined as the following:

$$\tilde{F}(w) = \frac{\int_M \langle \nabla T_F(w), \nabla T_F(w) \rangle_{\tilde{g}_M} d\tilde{\mu}_M}{\int_N \langle \nabla w, \nabla w \rangle_{\tilde{g}_N} d\tilde{\mu}_N}. \quad (4.30)$$

Accordingly, we define $\tilde{A}^{conf}(C) = \tilde{A}(C) \cap \{\int_N w d\tilde{\mu}_N = 0\}$ and $\|\tilde{R}\|_C = \max \tilde{F}(w)$ s.t. $w \in \tilde{A}^{conf}(C)$.

Unfortunately, the strategy of proving Theorem 4.5 does not work in the case of the conformal shape difference operators. That is because the interleaved structure described in Lemma 4.2 is not guaranteed between the new subdomains $A^{conf}(\cdot)$ and $\tilde{A}^{conf}(\cdot)$: a function satisfying $\int_N w d\nu_N = 0$ does not necessarily fulfill $\int_N w d\tilde{\mu}_N = 0$ simultaneously.

4.7 Experimental Results

In this section, we demonstrate experimental results that are related to our theoretical analyses. We conduct all the experiments on meshed shapes, i.e., discrete polygon surfaces embedded \mathbb{R}^3 . To start with, we show in the discrete case how to approach a local maximum of the constrained optimization problem $\max E(w)$ s.t. $w \in A(C)$, which is discussed in Section 4.6.4. Then the extension proposed in Section 4.6.5 is applied to detect and visualize conformal differences between a pair of shapes induced by a given map. Lastly, we demonstrate how the frameworks based on the shape difference operators react to perturbations on input shapes.

4.7.1 Approximating $\|V\|_C$

Now suppose that we are given a pair of meshed shapes, we demonstrate how to search for a *local* optimum of the constrained non-linear optimization with the *barrier function method*. Let M, N be two meshed shapes consisting of v_M and v_N vertices respectively, and let L_M (resp., L_N) and A_M (resp., A_N) be the stiffness matrix and the matrix of area elements for M (resp. N) (see [Pinkall 1993]). The functional map T_F induced by T is represented by a matrix $P \in \mathbb{R}^{v_M \times v_N}$ in the discrete setting. Let $\Phi_k \in \mathbb{R}^{v_N \times k}$ be a matrix whose columns are the first k eigenvectors solved by $L_N f = \lambda A_N f$.

Then calculating $\|V\|_C$ in this setting is equivalent to maximize the following function:

$$\max \frac{f^T P^T A_M P f}{f^T A_N f}, \text{ s.t. } \frac{f^T L_N f}{f^T A_N f} \leq C$$

Based on that a barrier function is constructed

$$G(\beta, f) = -\frac{f^T P^T A_M P f}{f^T A_N f} - \beta \log(C - \frac{f^T L_N f}{f^T A_N f})$$

As suggested in Proposition 4.6, we take the maximizer of $E(w)$ constrained in $S(k)$ as the initial guess for minimizing $G(1, f)$. After obtaining f_1 as a local minimizer, we take it as the initial guess for $G(\frac{1}{2}, f)$. The iteration continues until there is no more significant improvement or β is sufficient small. Optimizing 4.27 is done with an extra constraint that $f = \Phi_l^T a$, where $a \in \mathbb{R}^l$.

Proposition 4.7 suggests that as l tends to infinity, the optimum of problem 4.27 converges to $\|V\|_C$. We illustrate this claim by comparing a bumped sphere M with a sphere N (the same pair as in Figure 4.1), the indicators with respect to different constraints are depicted on N in Figure 4.5.

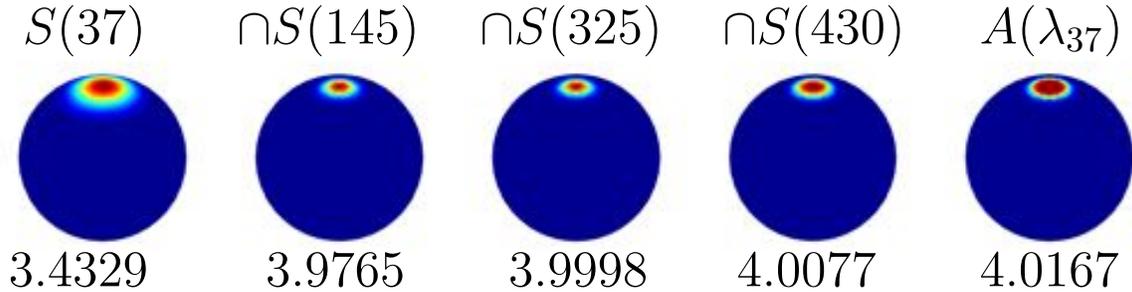


Figure 4.5: We compare a bumped sphere to a sphere (the same as in Figure 4.1) and plot on the sphere the (local) maximizers different constraints (note that only the left-most one is a global maximizer). The middle three are locally optimized by taking the constraints as intersection of $A(\lambda_{37})$ and $S(l)$, where $l = 145, 325, 430$. The eigenvalues involved are: $\lambda_{37} = 2.3857, \lambda_{145} = 8.3682, \lambda_{325} = 16.6667$, and $\lambda_{430} = 20.0167$.

As mentioned in Section 4.6.1, both subdomains $S(k)$ and $A(C)$ are designed to control the Dirichlet's Energy of feasible solutions. The difference between them is that in the former case the energy is controlled by *truncating* high frequency components while in the latter case high frequency components are allowed but

with implicit bounds on their weights. Figure 4.6 demonstrates this difference intuitively. In this experiment, we compute the local maxima/maximizers of $\|V\|_C$ with different scales C range from 0.5 to 2.

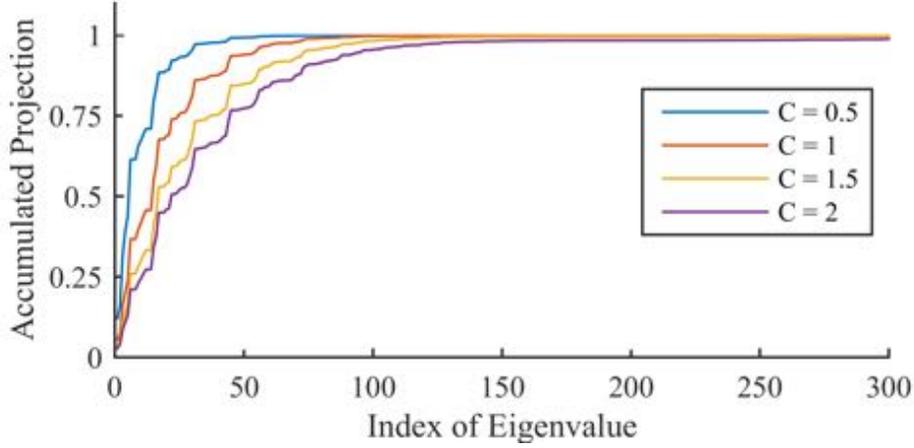


Figure 4.6: The X-axis indicates the index of eigenvalues/eigenfunctions, and the Y-axis represents the ratio

$$\sum_{i=1}^k a_i^2 / \sum_{i \geq 1} a_i^2.$$

Recall that the L^2 -norm of a function can be written as sum of square magnitudes of projections (see Proposition 4.2), i.e., $\int_N u^2 d\nu_N = \sum_{i \geq 1} a_i^2$, where $a_i = \int_N u \varphi_i d\nu_N$ and φ_i is the i -th eigenfunction of the Laplace-Beltrami operator on N . Now we compute for $k = 1 \sim 300$, the portion of L^2 -norm of the local maximizer expressed in the span space $S(k)$. The X-axis of Figure 4.6 reads the index of eigenvalues/eigenfunctions, and the Y-axis reads the ratio $\sum_{i=1}^k a_i^2 / \sum_{i \geq 1} a_i^2$. It's obvious that the four local indicators are well-expressed by the first 300 eigenfunctions (with $\lambda_{300} = 15.2029$). The blue curve indicates that the local maximizer at $C = 0.5$ is almost fully spanned by the first 50 eigenfunctions, whereas the purple curve indicates that the first 50 only represent around 75 percents of the norm of the one at $C = 2$.

4.7.2 Capturing Conformal Differences

In Section 4.6.5 we introduce a framework to analyze and visualize conformal differences, which can't be captured by the original framework in [Ovsjanikov 2013]. A simple example is demonstrated in Figure 4.7, where the map T from M to N is an area-preserving map. Therefore $E(w) = 1$ for any function w and analysis based on $E(w)$ does not provide any information about the differences between the two surfaces. On the other hand, we maximize $F(w)$ in $S^{conf}(50)$, and plot the indicator on N (see the right plot of Figure 4.7). It is obvious that this indicator highlights the regions where changes in angles take place.

In practice, we also observe stability of the multi-scale framework based on the conformal shape difference operator. We compare again the two horses in Figure 4.2, and plot the indicators that maximize $F(w)$

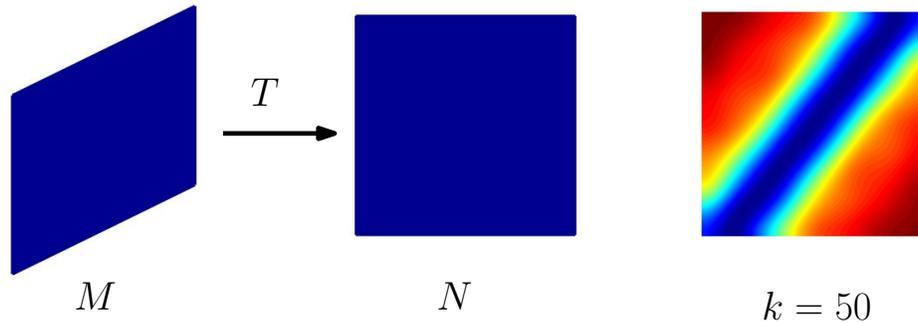


Figure 4.7: $T : M \rightarrow N$ is an area-preserving map. The conformal indicator at scale $k = 50$ captures and highlights the areas undergo conformal deformations.

within subdomains $S^{conf}(k) = \text{span}\{\varphi_2, \dots, \varphi_k\}$. The indicators at the scale k ranging from 50 to 500 are depicted in Figure 4.8. The evolution of highlighted areas is quite stable, reflecting the stability of the conformal shape difference operator with respect to the changes in scale.

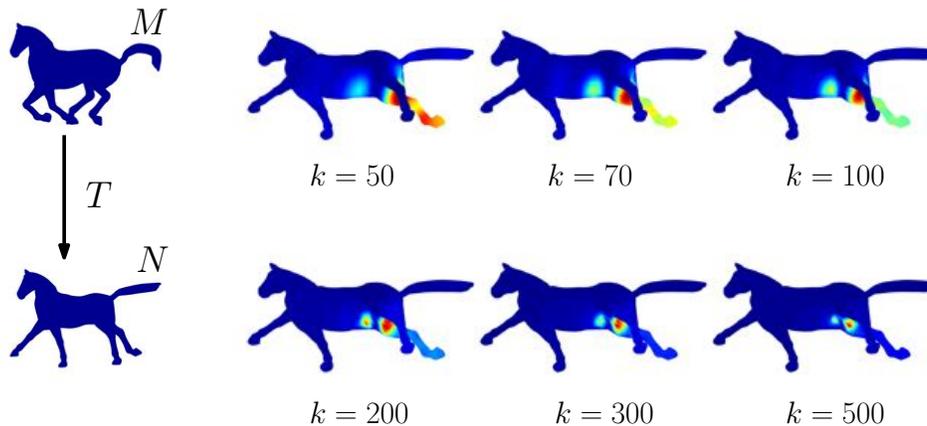


Figure 4.8: Indicators with respect to the conformal shape difference operator at k ranging from 50 to 500.

In general, the conformal differences between shapes are less intuitive to imagine and more subtle to capture. Thus this extension based on the conformal shape difference operator helps to provide a more complete picture for users to understand the distortions between shape induced by a given map.

4.7.3 Stability of the Area-based Shape Difference Operators

At the beginning of this chapter, we have shown in Figure 4.1 the robustness of the multi-scale framework in different senses. Here we consider more complicated shapes than deformed spheres.

The first example is demonstrated in Figure 4.9, we compare the two horses in Figure 4.8, but with

different meshes of horse N . At the same scale $k = 50$, the indicators with respect to different meshed N 's are similar to each other, which all consistently highlight the hip of the horse.

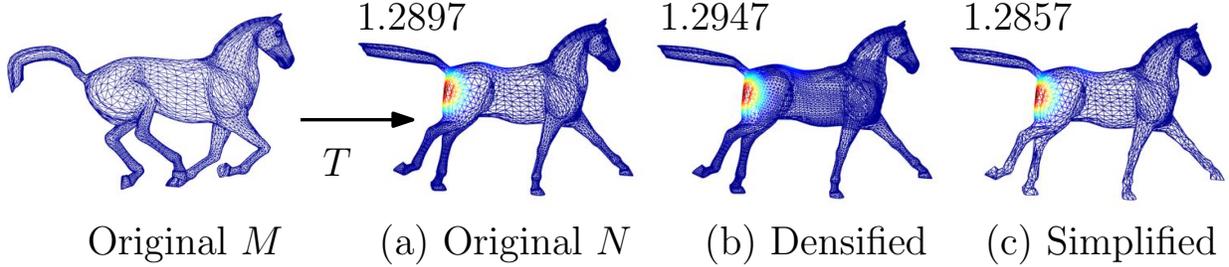


Figure 4.9: Indicators at scale $k = 50$ with different meshes. We densify the original shape (mesh-a) by adding points in the body of the horse (mesh-b) and simplify it by down-sampling the limbs (mesh-c). The corresponding distortion measurements are marked to the top-left.

Besides changing the mesh structure, we perturb the input meshes by disturbing the vertices. In this example, the vertices are especially perturbed along the normal direction so that the point-to-point correspondences are roughly preserved. We first compute the mean distance of edges of each mesh: $\bar{d}_M = 0.0144$, $\bar{d}_N = 0.0141$, and the vertex normal vectors. Given a parameter σ , we perturb a point p of mesh M to $p' = p + \sigma \bar{d}_M x_p \cdot \mathbf{n}_p$, where x_p is a one-dimensional random variable distributed normally with mean 0 and variance 1, and \mathbf{n}_p is the unit normal vector at vertex p . And we use the original mesh connectivity to connect perturbed points, since they are in a one-to-one correspondence to the unperturbed points.

We perturb *both* M and N in the same manner, and consider 4 choices of σ : 0, 0.1, 0.5, 1.0. At each level of perturbations, we generate indicators with respect to the area-based shape difference operator at 3 scales $k = 20, 50, 200$. The results are shown in Figure 4.10. We observe that when $\sigma = 0.1$, the indicators are well consistent with the ones of the first row. In fact, even when $\sigma = 0.5$, meaning that the standard deviation of the perturbations is half the mean distance, the indicators are still reasonable. At the end, we also notice that in the most noisy row, the indicator at $k = 200$ deviates from the ground-truth significantly while the first two at $k = 20, 50$ are still relevant. As we mentioned before, as k increases, the corresponding indicator is supposed to be more and more localized. The high-frequency indicators are more sensitive to the noises.

4.8 Proofs for Theorems in Section 4.6.5

Proof of Theorem 4.3 *Proof:* We first prove $H_{0,\nu}^1(N) = H_{0,\tilde{\mu}}^1(N)$ so that $\tilde{R}g$ is well-defined for $g \in H_{0,\nu}^1(N)$. Since (N, g_N, ν_N) and $(N, \tilde{g}_N, \tilde{\mu}_N)$ are on top of the same smooth topological manifold N , they share the same boundary of N . Therefore if f is zero on the boundary of the former, then it is zero on the boundary of the latter.

We now prove if $\int_N f^2 + \langle \nabla f, \nabla f \rangle_{g_N} d\nu_N < \infty$, then $\int_N f^2 + \langle \nabla f, \nabla f \rangle_{\tilde{g}_N} d\tilde{\mu}_N < \infty$ and vice

versa. First according to Proposition 4.3, we have $a_N^{-n/2}b_N^{-1}d\tilde{\mu}_N \leq d\nu_N \leq a_N^{n/2}b_Nd\tilde{\mu}_N$ and for any smooth function f , $a_N^{-1}\langle \nabla f, \nabla f \rangle_{\tilde{g}_N} \leq \langle \nabla f, \nabla f \rangle_{g_N} \leq a_N\langle \nabla f, \nabla f \rangle_{\tilde{g}_N}$. Then for any $f \in H_{0,\nu}^1(N)$, $\int_N f^2 + \langle \nabla f, \nabla f \rangle_{\tilde{g}_N} d\tilde{\mu}_N \leq \int_N f^2 a_N^{n/2} b_N d\nu_N + \int_N a_N \langle \nabla f, \nabla f \rangle_{g_N} a_N^{n/2} b_N d\nu_N \leq a_N^{n/2} b_N (\int_N f^2 d\nu_N + a_N \int_N \langle \nabla f, \nabla f \rangle_{g_N} d\nu_N) < \infty$, therefore $H_{0,\nu}^1(N) \subset H_{0,\tilde{\mu}}^1(N)$. On the other hand, one can similarly verify that $H_{0,\tilde{\mu}}^1(N) \subset H_{0,\nu}^1(N)$, which implies $H_{0,\nu}^1(N) = H_{0,\tilde{\mu}}^1(N)$.

Now considering $f \in H_{0,\nu}^1(N) = H_{0,\tilde{\mu}}^1(N)$, it follows from the triangle inequality that

$$\begin{aligned} & \left| \int_N \langle \nabla f, \nabla Rf \rangle_{g_N} d\nu_N - \int_N \langle \nabla f, \nabla \tilde{R}f \rangle_{g_N} d\nu_N \right| \\ & \leq \left| \int_N \langle \nabla f, \nabla Rf \rangle_{g_N} d\nu_N - \int_N \langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} f d\tilde{\mu}_N \right| \\ & \quad + \left| \int_N \langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} d\tilde{\mu}_N - \int_N \langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} d\nu_N \right| \\ & \quad + \left| \int_N \langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} d\nu_N - \int_N \langle \nabla f, \nabla \tilde{R}f \rangle_{g_N} d\nu_N \right| \\ & := |P_1| + |P_2| + |P_3| \end{aligned}$$

Then we estimate P_1, P_2 and P_3 separately. According to Proposition 4.3, measures ν_M (resp. ν_N) and $\tilde{\mu}_M$ (resp. $\tilde{\mu}_N$) satisfy

$$\begin{aligned} a_M^{-n/2}b_M^{-1}d\tilde{\mu}_M & \leq d\nu_M \leq a_M^{n/2}b_Md\tilde{\mu}_M \\ a_N^{-n/2}b_N^{-1}d\tilde{\mu}_N & \leq d\nu_N \leq a_N^{n/2}b_Nd\tilde{\mu}_N. \end{aligned}$$

and for any smooth functions f_M, f_N on M and N respectively,

$$\begin{aligned} a_M^{-1}\langle \nabla f_M, \nabla f_M \rangle_{\tilde{g}_M} & \leq \langle \nabla f_M, \nabla f_M \rangle_{g_M} \leq a_M\langle \nabla f_M, \nabla f_M \rangle_{\tilde{g}_M} \\ a_N^{-1}\langle \nabla f_N, \nabla f_N \rangle_{\tilde{g}_N} & \leq \langle \nabla f_N, \nabla f_N \rangle_{g_N} \leq a_N\langle \nabla f_N, \nabla f_N \rangle_{\tilde{g}_N} \end{aligned}$$

Thus, we have by the definitions of R, \tilde{R}

$$\begin{aligned} P_1 & = \int_M \langle \nabla T_F(f), \nabla T_F(f) \rangle_{g_M} d\nu_M - \int_M \langle \nabla T_F(f), \nabla T_F(f) \rangle_{\tilde{g}_M} d\tilde{\mu}_M \\ & \leq \int_M \langle \nabla T_F(f), \nabla T_F(f) \rangle_{g_M} d\nu_M - a_M^{-1} \int_M \langle \nabla T_F(f), \nabla T_F(f) \rangle_{g_M} d\tilde{\mu}_M \\ & \leq (1 - a_M^{-1-n/2}b_M^{-1}) \int_M \langle \nabla T_F(f), \nabla T_F(f) \rangle_{g_M} d\nu_M \end{aligned}$$

The lower bound of P_1 is estimated in the same way, and we have

$$P_1 \geq (1 - a_M^{1+n/2}b_M) \int_M \langle \nabla T_F(f), \nabla T_F(f) \rangle_{g_M} d\nu_M$$

Noticing that $0 \leq \int_M \langle \nabla T_F(f), \nabla T_F(f) \rangle_{g_M} d\nu_M \leq D_T \int_N \langle \nabla f, \nabla f \rangle_{g_N} d\nu_N < \infty$ as $f \in H_{0,\nu}^1(N)$, we have $|P_1|$ vanishes as $a_M, b_M \rightarrow 1^+$.

Regarding P_2 , we define two complementary subsets of N with respect to f : $I^+ = \{x \in N : \langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} \geq 0\}$ and $I^- = \{x \in N : \langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} < 0\}$.

$$\begin{aligned} P_2 &= \int_{I^+} \langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} (d\tilde{\mu}_N - d\nu_N) + \int_{I^-} \langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} (d\tilde{\mu}_N - d\nu_N) \\ &\leq (1 - a_N^{-n/2} b_N^{-1}) \int_{I^+} \langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} d\tilde{\mu}_N + (1 - a_N^{n/2} b_N) \int_{I^-} \langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} d\tilde{\mu}_N \end{aligned}$$

And we obtain the lower bound similarly:

$$\begin{aligned} P_2 &= \int_{I^+} \langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} (d\tilde{\mu}_N - d\nu_N) + \int_{I^-} \langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} (d\tilde{\mu}_N - d\nu_N) \\ &\geq (1 - a_N^{n/2} b_N) \int_{I^+} \langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} d\tilde{\mu}_N + (1 - a_N^{-n/2} b_N^{-1}) \int_{I^-} \langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} d\tilde{\mu}_N \end{aligned}$$

If the integrals of $\langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N}$ on I^+ and I^- are both finite, then P_2 vanishes as a_N, b_N tend to 1. In fact, by the formulation of \tilde{R} , both f and $\tilde{R}f$ are functions in $H_{0,\tilde{\mu}}^1(N)$, thus it follows from the Cauchy-Schwarz inequality that

$$\begin{aligned} \left| \int_{I^\pm} \langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} f d\tilde{\mu}_N \right| &\leq \sqrt{\int_{I^\pm} \langle \nabla f, \nabla f \rangle_{\tilde{g}_N} d\tilde{\mu}_N} \sqrt{\int_{I^\pm} \langle \nabla \tilde{R}f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} d\tilde{\mu}_N} \\ &\leq \sqrt{\int_N \langle \nabla f, \nabla f \rangle_{\tilde{g}_N} d\tilde{\mu}_N} \sqrt{\int_N \langle \nabla \tilde{R}f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} d\tilde{\mu}_N} \\ &< \infty \end{aligned}$$

Lastly, we argue that P_3 vanishes as a_M, a_N, b_M and b_N tend to 1. Since the metric defines a symmetrical inner products on the tangent spaces, we have $\langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} = \frac{1}{4}(\langle \nabla(f + \tilde{R}f), \nabla(f + \tilde{R}f) \rangle_{\tilde{g}_N} - \langle \nabla(f - \tilde{R}f), \nabla(f - \tilde{R}f) \rangle_{\tilde{g}_N})$ and $\langle \nabla f, \nabla \tilde{R}f \rangle_{g_N} = \frac{1}{4}(\langle \nabla(f + \tilde{R}f), \nabla(f + \tilde{R}f) \rangle_{g_N} - \langle \nabla(f - \tilde{R}f), \nabla(f - \tilde{R}f) \rangle_{g_N})$.

On the other hand, as $a_N, b_N \rightarrow 1$, the following convergence holds

$$\lim_{a_N, b_N \rightarrow 1} \langle \nabla(f \pm \tilde{R}f), \nabla(f \pm \tilde{R}f) \rangle_{\tilde{g}_N} = \langle \nabla(f \pm \tilde{R}f), \nabla(f \pm \tilde{R}f) \rangle_{g_N}$$

Therefore we have

$$\lim_{a_N, b_N \rightarrow 1} \langle \nabla f, \nabla \tilde{R}f \rangle_{\tilde{g}_N} = \langle \nabla f, \nabla \tilde{R}f \rangle_{g_N}$$

which assures that P_3 vanishes as $a_N, b_N \rightarrow 1$.

To summarize, $\int \langle \nabla f, \nabla(Rf - \tilde{R}f) \rangle_{g_N} d\nu_N \rightarrow 0$ for any $f \in H_{0,\nu}^1(N) = H_{0,\tilde{\mu}}^1(N)$. Given $f, g \in H_{0,\nu}^1(N)$, since both V and \tilde{V} are self-adjoint operators, direct computation shows that

$$\begin{aligned} 4 \int_N \langle \nabla f, \nabla(Rg - \tilde{R}g) \rangle_{g_N} d\nu_N &= \int_N \langle \nabla(f + g), \nabla(R - \tilde{R})(f + g) \rangle_{g_N} d\nu_N \\ &\quad - \int_N \langle \nabla(f - g), \nabla(R - \tilde{R})(f - g) \rangle_{g_N} d\nu_N \end{aligned}$$

As $f + g, f - g \in H_{0,\nu}^1(N)$, it follows immediately that for any pair of f, g , $\int_N \langle \nabla f, \nabla(Rg - \tilde{R}g) \rangle_{g_N} d\nu_N$ vanishes as a_M, b_M, a_N , and b_N converge to 1 simultaneously. Especially we let $f = Rg - \tilde{R}g \in H_{0,\nu}^1(N)$, and conclude that $\lim_{a_M, a_N, b_M, b_N \rightarrow 1^+} \int_N \langle \nabla(Vg - \tilde{V}g), \nabla(Vg - \tilde{V}g) \rangle_{g_N} d\nu_N = 0$. \square

Proof of Lemma 4.4 *Proof:* We estimate the difference by two parts:

$$\begin{aligned} |F(w_0 + w_1) - F(w_0)| &\leq \left| F(w_0 + w_1) - \frac{\int_M \langle \nabla T_F(w_0 + w_1), \nabla T_F(w_0 + w_1) \rangle_{g_M} d\nu_M}{\int_N \langle \nabla w_0, \nabla w_0 \rangle_{g_N} d\nu_N} \right| \\ &\quad + \left| \frac{\int_M \langle \nabla T_F(w_0 + w_1), \nabla T_F(w_0 + w_1) \rangle_{g_M} d\nu_M}{\int_N \langle \nabla w_0, \nabla w_0 \rangle_{g_N} d\nu_N} - F(w_0) \right| \\ &:= P_1 + P_2 \end{aligned}$$

For P_1 , direct computation shows:

$$\begin{aligned} P_1 &\leq D_T \int_N \langle \nabla(w_0 + w_1), \nabla(w_0 + w_1) \rangle_{g_N} d\nu_N \left| \frac{1}{\int_N \langle \nabla w_0, \nabla w_0 \rangle_{g_N} d\nu_N} - \frac{1}{\int_N \langle \nabla(w_0 + w_1), \nabla(w_0 + w_1) \rangle_{g_N} d\nu_N} \right| \\ &\leq D_T \frac{2 \left| \int_N \langle \nabla w_0, \nabla w_1 \rangle_{g_N} d\nu_N \right| + \int_N \langle \nabla w_1, \nabla w_1 \rangle_{g_N} d\nu_N}{\int_N \langle \nabla w_0, \nabla w_0 \rangle_{g_N} d\nu_N} \\ &\leq 2D_T \frac{\sqrt{\int_N \langle \nabla w_0, \nabla w_0 \rangle_{g_N} d\nu_N} \int_N \langle \nabla w_1, \nabla w_1 \rangle_{g_N} d\nu_N}{\int_N \langle \nabla w_0, \nabla w_0 \rangle_{g_N} d\nu_N} + D_T \frac{\int_N \langle \nabla w_1, \nabla w_1 \rangle_{g_N} d\nu_N}{\int_N \langle \nabla w_0, \nabla w_0 \rangle_{g_N} d\nu_N} \\ &\leq 2D_T \sqrt{s} + D_T s \end{aligned}$$

The third line follows from the Cauchy-Schwarz inequality applied on w_0 and w_1 which are both in $L_\nu^2(N)$.

P_2 is estimated in the same way, and the bound is identical:

$$\begin{aligned} P_2 &\leq \left| 2 \frac{\int_M \langle \nabla T_F(w_0), \nabla T_F(w_1) \rangle_{g_M} + \langle \nabla T_F(w_1), \nabla T_F(w_1) \rangle_{g_M} d\nu_M}{\int_N \langle \nabla w_0, \nabla w_0 \rangle_{g_N} d\nu_N} \right| \\ &\leq \frac{2 \left| \int_M \langle \nabla T_F(w_0), \nabla T_F(w_1) \rangle_{g_M} d\nu_M \right| + D_T \int_N \langle \nabla w_1, \nabla w_1 \rangle_{g_N} d\nu_N}{\int_N \langle \nabla w_0, \nabla w_0 \rangle_{g_N} d\nu_N} \\ &\leq 2D_T \frac{\sqrt{\int_N \langle \nabla w_0, \nabla w_0 \rangle_{g_N} d\nu_N} \int_N \langle \nabla w_1, \nabla w_1 \rangle_{g_N} d\nu_N}{\int_N \langle \nabla w_0, \nabla w_0 \rangle_{g_N} d\nu_N} + D_T \frac{\int_N \langle \nabla w_1, \nabla w_1 \rangle_{g_N} d\nu_N}{\int_N \langle \nabla w_0, \nabla w_0 \rangle_{g_N} d\nu_N} \\ &\leq 2D_T \sqrt{s} + D_T s \end{aligned}$$

Putting them together yields $|F(w_0 + w_1) - F(w_0)| \leq 4D_T \sqrt{s} + 2D_T s$, where $s = \frac{\int_N \langle \nabla w_1, \nabla w_1 \rangle_{g_N} d\nu_N}{\int_N \langle \nabla w_0, \nabla w_0 \rangle_{g_N} d\nu_N}$. \square

Proof of Theorem 4.6 *Proof:* The strategy of this proof is similar to the one for Theorem 4.4.

We first consider the case of $\varepsilon > 0$, i.e., $C' > C > \lambda_2$. By definition, $A^{conf}(C) \subset A^{conf}(C')$, thus $\|R\|_{C'} - \|R\|_C \geq 0$. We then estimate the upper bound for the difference. Given $w \in A^{conf}(C + \varepsilon)$, our strategy is to construct a function $\bar{w} \in A^{conf}(C)$, such that $|E(w) - E(\bar{w})|$ is uniformly bounded.

If w itself lies in $A^{conf}(C)$, then it's trivial to set $\bar{w} = w$. We now consider the case $w \in A^{conf}(C + \varepsilon) \setminus A^{conf}(C)$. Assume that $\int_N w(-\Delta_N)w d\nu_N = (C + \delta) \int_N w^2 d\nu_N$, where $0 < \delta \leq \varepsilon$. Without loss of generality, we further assume that $w = \sum_{i \geq 1} a_i \varphi_i$ and $\sum_{i \geq 1} a_i^2 = 1$, where (φ_i, λ_i) is the i -th eigensolution to $\Delta_N \varphi + \lambda \varphi = 0$. Note that $\int_N w d\nu_N = 0$, thus $a_1 = 0$, i.e., $w = \sum_{i \geq 2} a_i \varphi_i$. According to the orthogonality of eigenbasis, the constraint on w can be written as

$$\sum_{i \geq 2} a_i^2 \lambda_i = C + \delta.$$

Let b_2 be a real number satisfying $b_2^2 - a_2^2 = \delta/(C - \lambda_2)$ and $b_2 a_2 \geq 0$. The existence of b_2 is assured by the fact that a_2 is finite (in fact $|a_2| \leq 1$). Then we set $\bar{w} = b_2 \varphi_2 + \sum_{i \geq 3} a_i \varphi_i$. Direct computation shows that $\int_N \bar{w}(-D_N)\bar{w} d\nu_N = b_2^2 \lambda_2 + \sum_{i \geq 3} a_i^2 \lambda_i = C + \frac{C\delta}{C - \lambda_2}$ and $\int_N \bar{w}^2 d\nu_N = b_2^2 + \sum_{i \geq 3} a_i^2 = b_2^2 + 1 - a_2^2 = 1 + \frac{\delta}{C - \lambda_2} = \frac{1}{C} \int_N \bar{w}(-D_N)\bar{w} d\nu_N$, meaning that $\bar{w} \in A^{conf}(C)$.

Thanks to Lemma 4.4, $|F(w) - F(\bar{w})| = |F(w) - F(w + \bar{w} - w)| \leq 4D_T \sqrt{s} + 2D_T s$, where $s = \frac{\int_N \langle \nabla(\bar{w} - w), \nabla(\bar{w} - w) \rangle_{g_N} d\nu_N}{\int_N \langle \nabla w, \nabla w \rangle_{g_N} d\nu_N} = \frac{\lambda_2(b_2 - a_2)^2}{C + \delta}$. Without loss of generality, we assume that $a_2, b_2 \geq 0$, thus $b_2 = \sqrt{a_2^2 + \frac{\delta}{C - \lambda_2}} > a_2$. Moreover, $b_2 - a_2 = \sqrt{a_2^2 + \frac{\delta}{C - \lambda_2}} - a_2 = \frac{\delta/(C - \lambda_2)}{\sqrt{a_2^2 + \delta/(C - \lambda_2)} + a_2} \leq \sqrt{\delta/(C - \lambda_2)}$. Therefore $s \leq \frac{\lambda_2 \delta}{(C - \lambda_2)(C + \delta)} \leq \frac{\lambda_2 \varepsilon}{(C - \lambda_2)C}$. Now assuming that w^* is the maximizer of $F(w)$ constrained in $A^{conf}(C + \varepsilon)$, the above derivation shows that there exists a $\bar{w}^* \in A^{conf}(C)$ such that $|F(w^*) - F(\bar{w}^*)| \leq 4D_T \sqrt{s} + 2D_T s$. That implies $\|V\|_{C+\varepsilon} - \|V\|_C \leq E(w^*) - E(\bar{w}^*) \leq |E(w^*) - E(\bar{w}^*)| \leq 4D_T \sqrt{\frac{\lambda_2 \varepsilon}{C(C - \lambda_2)}} + 2D_T \frac{\lambda_2 \varepsilon}{C(C - \lambda_2)}$.

Regarding the case $\varepsilon < 0$, i.e., $C > C' > \lambda_2$. We simply replace C and $C + \varepsilon$ in the previous analysis with $C - \varepsilon$ and C . With identical derivations, for each $w \in A^{conf}(C)$, we construct a function $\bar{w} \in A^{conf}(C - \varepsilon)$ such that $|F(w) - F(\bar{w})| = |F(w) - F(w + \bar{w} - w)| \leq 4D_T \sqrt{s} + 2D_T s$, where $s \leq \frac{\lambda_2 |\varepsilon|}{(C - \varepsilon)(C - \lambda_2)}$.

Putting them together, we finish the proof of Theorem 4.6. \square

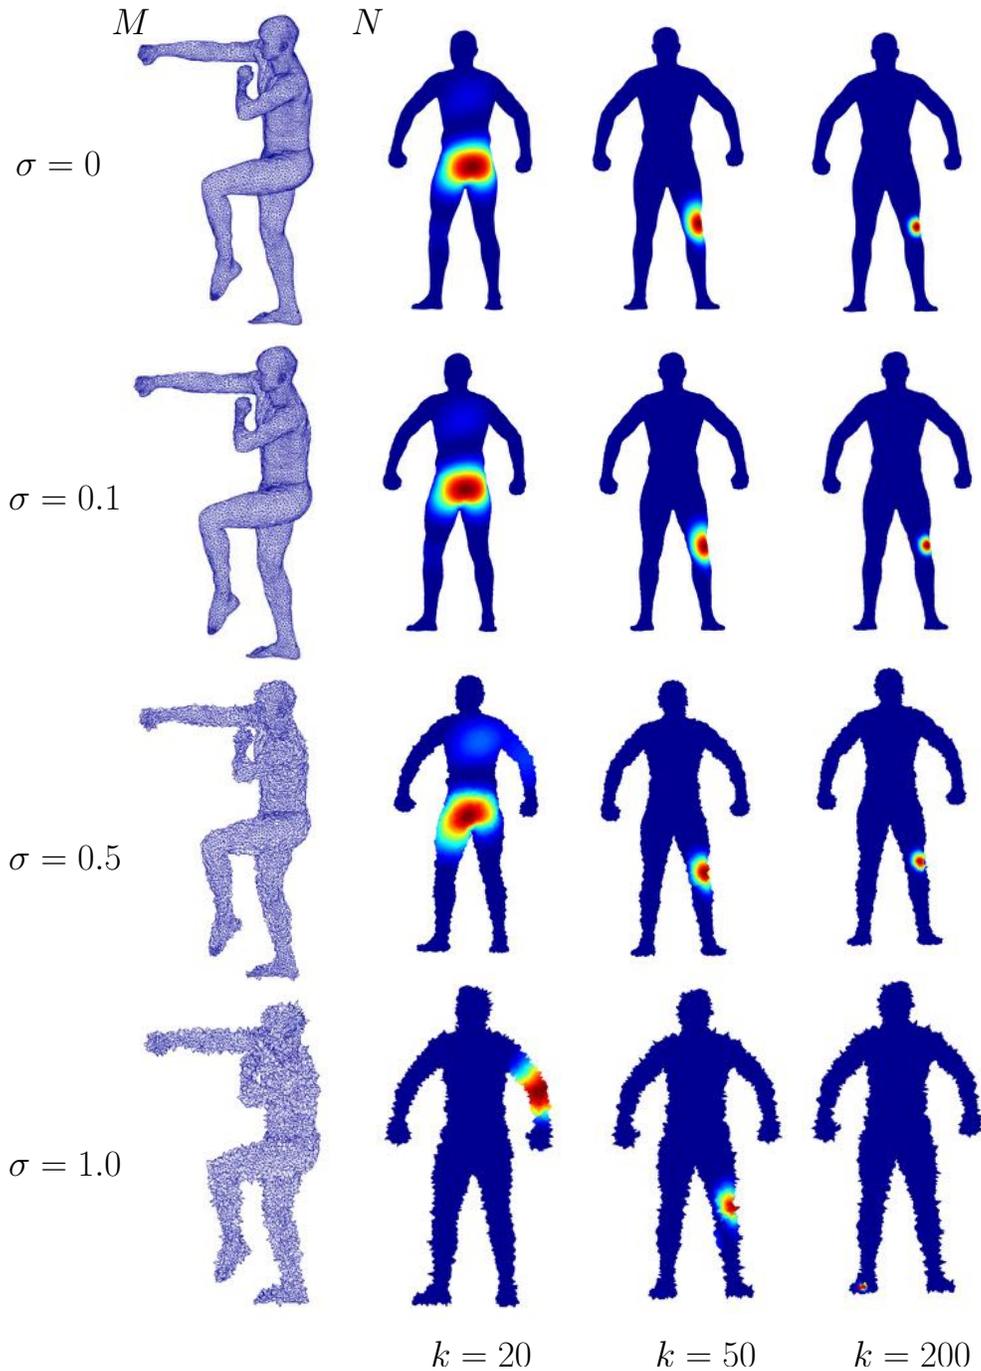


Figure 4.10: Four pairs of meshed shapes are compared, σ indicates the strength of perturbations added on each of the shapes in the same row. At each row, three indicators are plotted on mesh N , which are obtained by maximizing $E(w)$ within $S(20)$, $S(50)$ and $S(200)$ respectively. Note that the human poses in the second column stand with their backs towards us, thus the highlighted areas are the hip and the right elbow.

functional-maps-based Frameworks on Point Cloud Data

5.1 Introduction

The framework of the functional map and most of its follow-ups are proposed to deal with problems about 3D shapes, i.e., 2-dimensional Riemannian manifold embedded in \mathbb{R}^3 . Besides the fact that 3D shapes are ubiquitous in both theory and application, the other reason might be that in practice, 3D shapes can be well-approximated by polygon meshes, which allow people to estimate geometric quantities of the underlying shape from the discrete data. For example, the eigenbases with respect to the Laplace-Beltrami operators are heavily utilized in the frameworks related to the functional maps. Approximating this operator on a meshed shape is well-studied and solved by reliable and efficient methods (e.g., the cotangent scheme proposed in [Pinkall 1993]).

However, as we mentioned in Chapter 4, the frameworks of [Ovsjanikov 2012, Ovsjanikov 2013, Rustamov 2013] can be naturally adapted to a more general setting where the intrinsic dimension of input manifolds is not necessarily 2. Stated differently, these frameworks are potentially suitable for analyzing objects of higher dimensions. Constructing polygon meshes, unfortunately, is not straightforward in this more abstract setting. Moreover, only considering 2-dimensional manifolds, we would encounter the same problem when the manifolds are embedded in an Euclidean space of dimension more than 3. Even in the very original setting, generating mesh on top of a bunch of sampling points is not trivial. For instance, noises are sort of inevitable in the data acquisition, sometimes they give rise to locally non-manifold structures and thus make mesh generation a challenging task.

Due to the above limitations, in this chapter we consider implementing the above mentioned functional-maps-based frameworks on data in a more primitive form—point cloud data (PCD). There are nice works on approximating geometric quantities with PCD. The framework of [Belkin 2009] approximates the Laplace-Beltrami operator in PCD sampled from a general d -dimensional Euclidean space. However, this approach and a later one [Liu 2012] both rely on reconstructing local mesh structure. Rather than falling back to build some specific structures to associate input points, we build simply k -nearest neighborhood (k -NN) graphs on the input PCD, and for all necessary ingredients for implementing the functional-maps-based frameworks that are computed with meshes, we propose a pipeline to construct their counterparts on top of the k -NN graphs.

Among the ingredients, perhaps a reliable approximation to the Laplace-Beltrami operator is the most important. The top candidate for our pipeline is the graph Laplacian, whose convergence to the Laplace-Beltrami operator on the underlying manifold has been well studied (see [Belkin 2008, Hein 2007]). And

it has been wildly applied in practice, say, spectral clustering [von Luxburg 2006], dimensional reduction [Belkin 2002, Tenenbaum 2000a]. However, it is not so popular in the area of geometry processing, one of the main reasons is the lack of robustness: for example, in practice it is not clear how a graph Laplacian varies with respect to different k for k -NN graphs, not to mention that we are unlikely to approximate the underlying Laplace-Beltrami operator with graph Laplacian based on naive graph construction such as k -NN over finite sampling points.

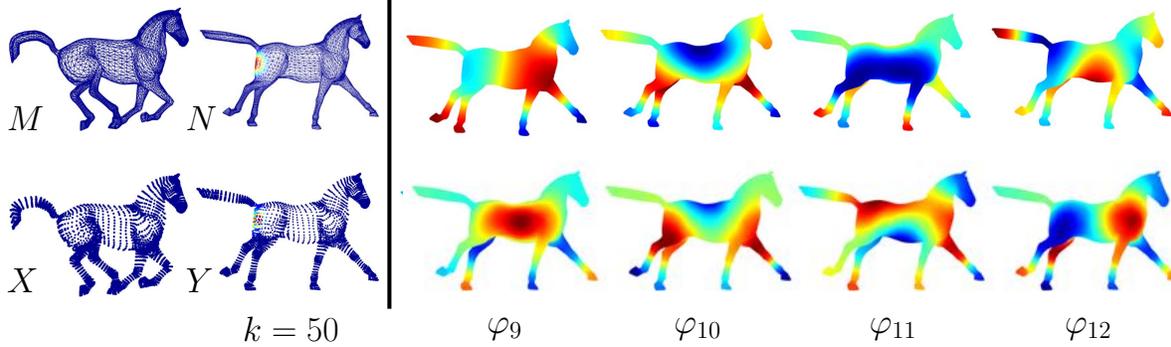


Figure 5.1: Left: We compare a pair of meshed shapes M and N and a pair of PCD X and Y , and generate indicators from the mesh setting (top) and the PCD setting (bottom) both at scale $k_s = 50$; Right: the ninth to the twelfth eigenfunction of the Discrete LB operator on mesh (top) and those of the Graph Laplacian on PCD (bottom).

On the other hand, the most important observation in this chapter is that despite its instability when applied in analyzing a *single* object, in the context of the functional-maps-based frameworks where objects are analyzed *in pair*, the results gained with the graph Laplacian are surprisingly stable. Figure 5.1 shows parts of the eigenfunctions on a horse generated from the mesh setting and from our PCD setting, where two rows are distinct from each other. However, the indicators generated from the two settings are comparable.

It is worth noting that our approach is empirical: our pipeline design is inspired by results on the convergence of the graph Laplacians in [Hein 2007]. The results characterize the limit behavior of the graph Laplacian as the size of sample tends to infinity and are proven under certain technique conditions on the underlying manifolds and the sampling density, which can hardly be guaranteed in practice.

In the experimental section, we show the relevance and robustness of applying the functional-maps-based frameworks with our pipeline directly on PCD. Our pipeline is fit for point clouds sampled from Euclidean spaces of dimension more than 3, yet in this chapter we exhibit results generated from 3D shapes, so that we can compare our results to the ones from the mesh setting.

5.2 Pipeline for PCD

Before we present our pipeline, we quickly check the *necessary* ingredients for implementing the frameworks of functional maps [Ovsjanikov 2012], shape difference operators [Rustamov 2013] and map analysis and visualization [Ovsjanikov 2013].

In the mesh setting, for an input meshed shape M consisting of n_M vertices, we need to compute W_M and A_M , which are both $n_M \times n_M$ matrices. W_M is a stiffness matrix such as the standard cotangent weighted matrix. The formulation of W_M with the cotangent scheme [Pinkall 1993] is illustrated in Figure 5.2.

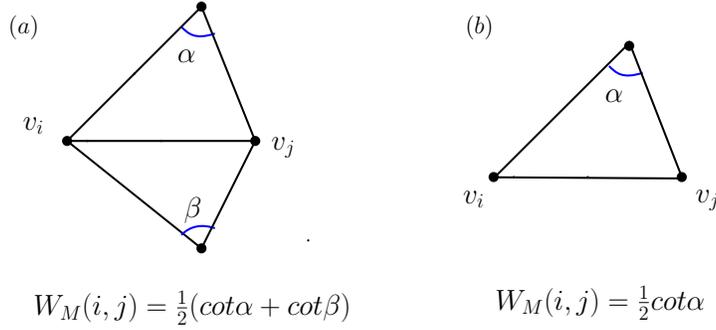


Figure 5.2: (a) $W_M(i, j)$ for an edge (v_i, v_j) that is shared by two triangles. (b) $W_M(i, j)$ for an boundary edge that belongs to a single triangle. If v_i, v_j are not connected by an edge, then $W(i, j) = 0$. Notice that in a mesh, an edge is at most shared by two triangles.

A_M is a diagonal matrix, whose (i, i) -th entry is assigned with the area element of the i -th vertex. To compute $A_M(i, i)$, we use the mixed Voronoi cell proposed in [Meyer 2003]. We collect all the triangles in the mesh containing vertex v_i . Each triangle contributes a part to the area element of v_i : if it's an acute triangle, then the contribution is one third of its area; if it is a right or an obtuse triangle, the contribution is one half of its area when v_i is opposite to the longest edge and one quarter of its area otherwise.

The discrete approximation of the Laplace-Beltrami operator is then given as $A_M^{-1}W_M$. To compute the eigenfunctions, we solve the eigenequation $W_M f = \lambda A_M f$. (We refer the readers to [Botsch 2010] for an introduction to general mesh processing.)

Now suppose we are given a point cloud $X = \{x_1, x_2, \dots, x_n\}$ consisting of n points, and X is sampled from the same smooth shape that is approximated by M . In the following we present how we compute the counterparts A_X and W_X in the PCD setting.

Building Connectivity In the mesh setting, the polygons, say, triangles, govern connections among the discrete vertices. While in the PCD case, there are two common ways to construct a graph on top of a point cloud: (1) the ε -neighborhood graph construction and (2) the k -nearest neighborhood (k -NN) graph construction. We choose the latter over the former due to two reasons: firstly we can control the sparsity of the output graphs by intuitively tuning k ; Secondly, we might face the situations where the points are non-uniformly sampled from a underlying Riemannian manifold. In such a situation, the parameter ε might be difficult to pick so that the resulting graph is connected (meaning that ε is not too small) and properly sparse (meaning that ε is not too large).

Building Weighted Graphs Fixing an integer k , we first build a k -NN graph on X , and denote the graph by $G_X = (X, E_X)$, where the edge set $E_X = \{(x_i, x_j), x_i \in N(x_j, k) \text{ or } x_j \in N(x_i, k)\}$ and $N(x_i, k)$ is the set of the nearest k neighborhoods of x_i among $X \setminus x_i$.

Then we turn this k -NN graph into a weighted graph. Typically, we use the Gaussian kernel to assign a weight to each of the connected edges in the k -NN graph above:

$$K_X(i, j) = \begin{cases} \exp(-\|x_i - x_j\|^2/2t^2) & \text{if } (x_i, x_j) \in E_X \\ 0 & \text{otherwise} \end{cases} \quad (5.1)$$

t is a parameter of K_X to be determined. We in practice let t be $\sum_{i=1}^{|X|} \sum_{x_j \in N(x_i, k)} \frac{\|x_i - x_j\|}{k|X|}$, which is depending on the scale of X .

let $D_X(i) = \frac{1}{\deg(i)} \sum_{(i,j) \in E_X} K_X(i, j)$, where $\deg(i)$ is the degree of vertex x_i in graph G_X . Then we normalize K_X with the $D_X(i)$'s:

$$\tilde{K}_X(i, j) = \frac{K_X(i, j)}{D_X(i)D_X(j)} \quad (5.2)$$

Building graph Laplacian There are various ways of defining a graph Laplacian, We use the un-normalized weighted graph Laplacian with respect to \tilde{K}_X constructed above:

$$W_X(i, j) = \begin{cases} -\tilde{K}_X(i, j) & \text{if } i \neq j \\ \sum_j \tilde{K}_X(i, j) & \text{if } i = j \end{cases} \quad (5.3)$$

Estimating Weighted Measures In the mesh setting, to each vertex p , we associate an area element, which is positively correlated to the sum of the areas of triangles around that point. In the point cloud setting, however, how to compute an area element is not clear because the surface is not regularly tessellated with polygons.

In our pipeline, to each sampling point x_i , we assign the multiplicative inverse of the sampling density estimation at x_i . There is a considerable amount of literature on density estimation from sampling, especially we chose the framework of [Biau 2011], which estimates the sampling density at point $x_i \in \mathbb{R}^d$ proportionally to:

$$\rho(x_i) \propto \left(\sum_{x_j \in N(x_i, k_n)} \|x_i - x_j\|^2 \right)^{-\frac{d}{2}}$$

The k_n in the subscript is a parameter related to n , the total number of sampling points. The convergence of the estimator to the underlying density is assured under some conditions, one of which is that $k_n \rightarrow \infty, k_n/n \rightarrow 0$ as $n \rightarrow \infty$. In practice, we simply let k_n be the same as k in our k -NN graph construction. Therefore we obtain these quantities as a by-product of constructing the k -NN graph.

Since we mostly consider 2-dimensional Riemannian manifolds embedded in \mathbb{R}^3 . A_X is then obtain by letting $A_X(i, i) = \left(\sum_{x_j \in N(x_i, k)} \|x_i - x_j\|^2 \right)^{\frac{3}{2}}$.

At the end, we view the graph Laplacian W_X as a counterpart of the stiffness matrix in the mesh setting and $A_X^{-1}W_X$ as an approximation to the Laplace-Beltrami operator, and compute the eigenfunctions by solving $W_X f = \lambda A_X f$.

5.2.1 Rationale of Our Pipeline

According to one of the main results proposed in [Hein 2007] (see ‘Main Result’ at the beginning of section 3.3 there, in the case of ‘unnormalized’ with $\lambda = 1$), the graph Laplacian constructed in Equation 5.3 is viewed as an approximation of $-\rho^{-1}\Delta$. From this point of view, the motivation of A_X above is clear: since this approximation is affected by the sampling density, we let A_X^{-1} be an approximation of ρ so that $A_X^{-1}W_X$ is an approximation of Δ .

In fact, both A_X and A_M from different settings work in a similar way. To see this, we illustrate the area elements in Figure 5.3. We assume that all the points are sampled in a flat local area. First we consider the mesh case, if the sampling is uniform, then the area elements are equal to each other. However, the area element of point v_1 (the area circled by the blue segments) is larger than that of point v_2 (area circled by the red segments). The unbalance in area distribution compensates the effect of the non-uniform sampling: obviously the sampling density around point v_1 is smaller than that around v_2 . In the point cloud setting, we plot the 7-nearest neighborhoods of x_1 and x_2 and connect them with edges in different colors. It is easy to see that the $A_X(1, 1) > A_X(2, 2)$, which mimics the relationship between the area element of v_1 and of v_2 .

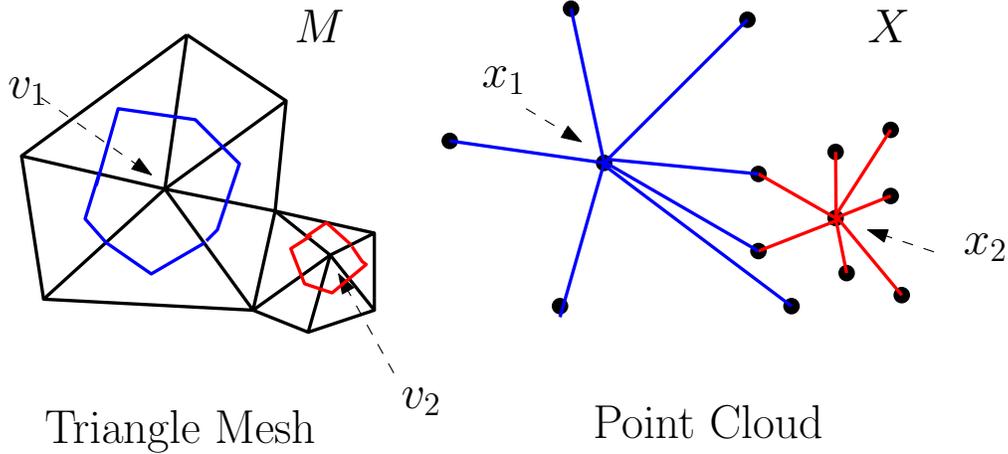


Figure 5.3: Left: the area elements in A_M compensate the effect of non-uniform sampling density around points $v_1, v_2 \in M$: larger element is assigned to the point locating in an area with lower sampling density; Right: the counterpart of A_M in our pipeline, A_X , that functions in a similar way.

On the other hand, as proven in [Bernstein 2000], under certain sampling condition, the metric induced by the graph (especially by the kNN graph) converges to the intrinsic manifold metric with high probability.

Again we emphasize that our construction is experimental without rigorous theoretical guarantees. The convergence results in [Hein 2007] is proven under several technique conditions, which in our case are not satisfied.

5.3 Experimental Results

A brief review of the three functional-maps-based frameworks has been given in Section 4.3.3, Section 4.3.4 and Section 4.3.5. In this section, we discuss the how to perform these frameworks with only point clouds sampled from interested shapes. We assume that a map T is from a source shape S_S to a target shape S_T . Let a mesh M (*resp.* N) and a point cloud X (*resp.* Y) be two representations of S_S (*resp.* S_T). The map T in the discrete case is simply a point-to-point correspondence from M to N (or from X to Y). And the corresponding functional map T_F is a pull back from the function space on N (*resp.* Y) to the one on M (*resp.* X).

In Section 5.3.1, we compute the low-rank approximations of T_F with truncated eigenbasis on both S_S and S_T and compare the results generated with M, N to those generated with X, Y .

After that, we turn to the frameworks of multi-scale map analysis and the shape difference operators. We induce T_F with the known point-to-point correspondence T and use it in the following discrete settings. Therefore, in the mesh setting, the optimization problem of 4.17 is:

$$\max E(w) = \frac{T_F(w)^T A_M T_F(w)}{w^T A_N w} \text{ s.t. } w \in \text{span}\{\phi_1^N, \phi_2^N, \dots, \phi_{k_s}^N\}$$

where ϕ_i^N is the i -th eigenvector of $A_N^{-1} W_N$. We call in this section the maximizer the area-based indicator at scale k_s .

We write the functional with respect to the conformal shape difference operator as follows (see Equation 4.28)

$$\max F(w) = \frac{T_F(w)^T W_M T_F(w)}{w^T W_N w} \text{ s.t. } w \in \text{span}\{\phi_2^N, \dots, \phi_{k_s}^N\}$$

Note that the constraint is different. We then call the maximizer the conformal indicator at scale k_s .

As described in Section 5.2, we compute A_X, W_X as the counterparts of A_M, W_M in our pipeline. Therefore by replacing them with A_X, W_X , we are able to apply this framework on PCD.

The main target of our experiments is to demonstrate the relevance and empirical robustness of implementation of the functional-maps-based frameworks on PCD with the pipeline proposed in section 5.2. Especially, we run test on several collections of shapes. The ground truths of the differences among shapes within each of the collections are visually intuitive, and we use the results from the meshed setting as ground truths.

5.3.1 Functional Maps on PCD

As shown in Figure 5.4, we are given two meshed horses M and N and a list of point-to-point correspondences T from M to N . Let X and Y be the vertex set of M and N respectively. Each horse consists of 8431 vertices, thus all the function spaces related in this experiment are vector spaces of dimension 8431. Since a functional map is a linear operator across the function spaces on different shapes, it is a 8431×8431 matrix, which is inconvenient for analysis. In the following we compute low-rank approximations of it in different ways.

We start with meshes M and N . The first 100 eigenfunctions of $W_M f = \lambda A_M f$ and $W_N f = \lambda A_N f$ are solved as the truncated bases on M and N respectively. Based on these truncated eigenbases, we approximate

and represent the functional map with a 100×100 matrix C_T . With the full information of T , T_F is faithfully represented as a permutation matrix. Thus C_T is explicitly expressed as

$$C_T(i, j) = \langle T_F(\phi_i^N), \phi_j^M \rangle_M, 1 \leq i, j \leq 100$$

where ϕ_i^N is the i -th eigenvector on N and ϕ_j^M is the j -th eigenvector on N .

On the other hand, it is also possible to compute C_T with much less information. As shown in the top-left of Figure 5.4, 6 pairs of landmarks in correspondences are selected. Following the idea in [Ovsjanikov 2012], we minimize $\|\Delta_M C_T - C_T \Delta_N\|_{Fro}$ with constraints that the 6 delta functions supported by the landmarks are preserved by C_T and denote by C_{LM} the minimizer.

Regarding the point clouds X and Y , we apply the pipeline in Section 5.2 with two k -NN graphs with $k = 15$ and $k = 40$. The eigenvectors are solutions to $W_X f = \lambda A_X f$ and $W_Y f = \lambda A_Y f$, and the following procedures are exactly the same as in the mesh setting.

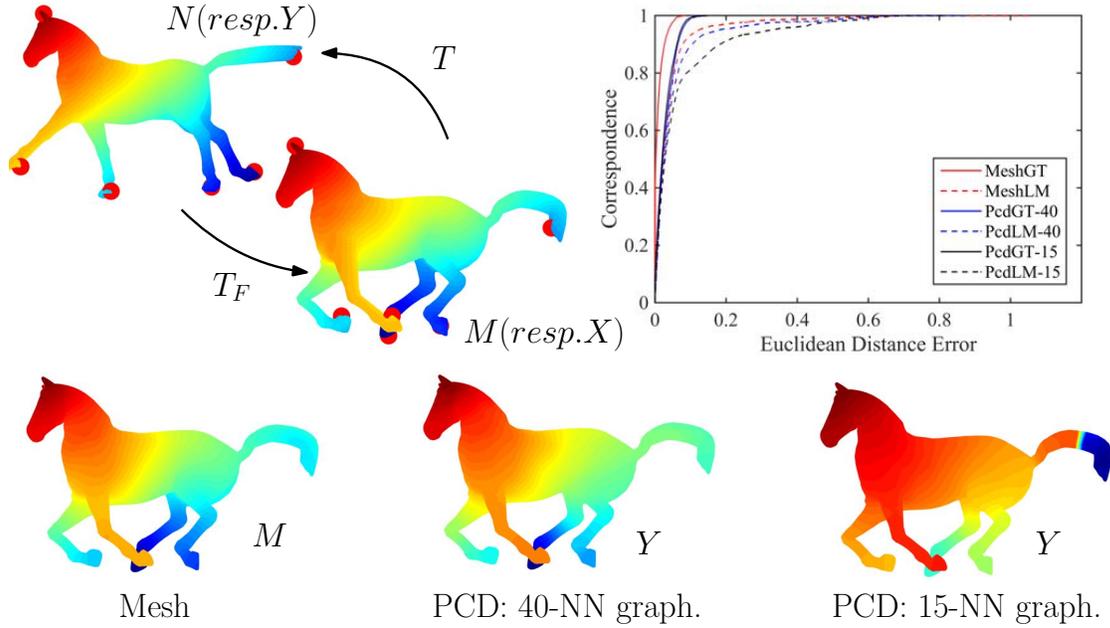


Figure 5.4: Top-left: two shapes with different representations— M , N are meshes and X , Y are point clouds; six pair of landmarks are selected in correspondences (red balls). Top-right: quantitative evaluations of functional maps approximated with different conditions and in different settings. Bottom: transforming the function f depicted in the top-left panel with functional maps approximated with landmarks in the three settings: mesh, PCD with 40-NN graph and PCD with 15-NN graph.

To evaluate the qualities of approximations with respect to different settings and methods, we employ the method from [Ovsjanikov 2012] to convert an approximated functional map, C' , to a point-to-point map T' . For each point p on M , the Euclidean distance from the ground truth $T(p)$ to $T'(p)$ (which are both on N) is computed. The chart in Figure 5.4 shows comparison of functional maps approximated in different ways. A point (e, c) on that chart reads $100 \times c\%$ of the vertices satisfying $\|T(p) - T'(p)\| \leq e$. As expected,

with the ground truth T , the approximations are significantly more accurate than the ones without. And the approximation from the mesh setting is more accurate than the other two from the PCD setting. With the ground truth, the approximations in the PCD setting with different k are indistinguishable in this chart. In the more challenging situation, we observe that the using 40-NN graph results in a better approximation that is closer to the one with mesh information.

More intuitively, in the top-left of Figure 5.4, we plot a function f and $T_F(f)$ on N and M respectively. The bottom row shows how f is transformed by functional maps computed with only landmark correspondences in different settings (for an easier comparison, we render all results with the mesh N). Again the one regarding the 40-NN graph is closer to the mesh result, and visually the function is well transformed to Y (see the middle of the bottom row of 5.4). The result regarding the 15-NN graph, however, is not satisfying.

5.3.2 Selection of k for k -NN

In our PCD pipeline, the only parameter to be determined by the users is k , the number of nearest neighborhoods for constructing k -NN graphs. Usually we choose the same k for all point clouds analyzed in pair (or in collection). The previous experiment seems to suggest large k is more preferable. In the following we compare the selections from a different point of view.

For the sake of consistency, we consider the same pair of shapes as in Section 5.3.1 and follow the notations there. Recall that we analyze and extend a framework for detecting and highlighting the differences between shapes induced by a given map in Section 4.6. In this experiment we implement this method with our pipeline and especially we build k -NN graphs on the input point clouds with different k .

As before, the results obtained with meshed shapes M and N are seen as ground truths. Regarding the point clouds X and Y , we build k -NN graphs with $k = 15, 25$ and 40 .

First of all, we compare the area-based and the conformal indicators shown in the left part of Figure 5.5. All the indicators are computed at a uniform scale $k_s = 50$. The area-based indicators along the left-most column are consistent, meaning that the results from the PCD with different parameters are relevant and consistent. Regarding the conformal case, the one of 15-NN graph fails to capture the right area, whereas the left two with larger k manage to highlight the right hind leg.

Secondly, we take a look at the corresponding eigenfunctions in the right part of Figure 5.5. As mentioned in Section 5.1, one of the issues of applying graph Laplacian in geometry processing is its instability, which is evidenced in the plot. Judging from this short window (only four eigenfunctions from the ninth to the twelfth are plotted), we notice that when the number of nearest neighborhoods changes from 15 to 25, the eigenfunctions are distinct. The last two rows are more stable with respect to a change of k , yet they are still dissimilar to the top row. Moreover, we plot eigenfunctions associated with relatively small eigenvalues, the divergences are more prominent as we pick the ones of higher frequencies.

Lastly, we make some remarks on the selection of k for k -NN. From the implementation point of view, a smaller k results in a sparser graph, which expedites the computation. However, because of the presence of noise, small k would bring instability into the construction and therefore into the result (as discussed above and shown in Section 5.3.1). In the subsequent experiments, we select $k = 40$. As will be shown, this parameter works well in analyzing other collections of shapes.

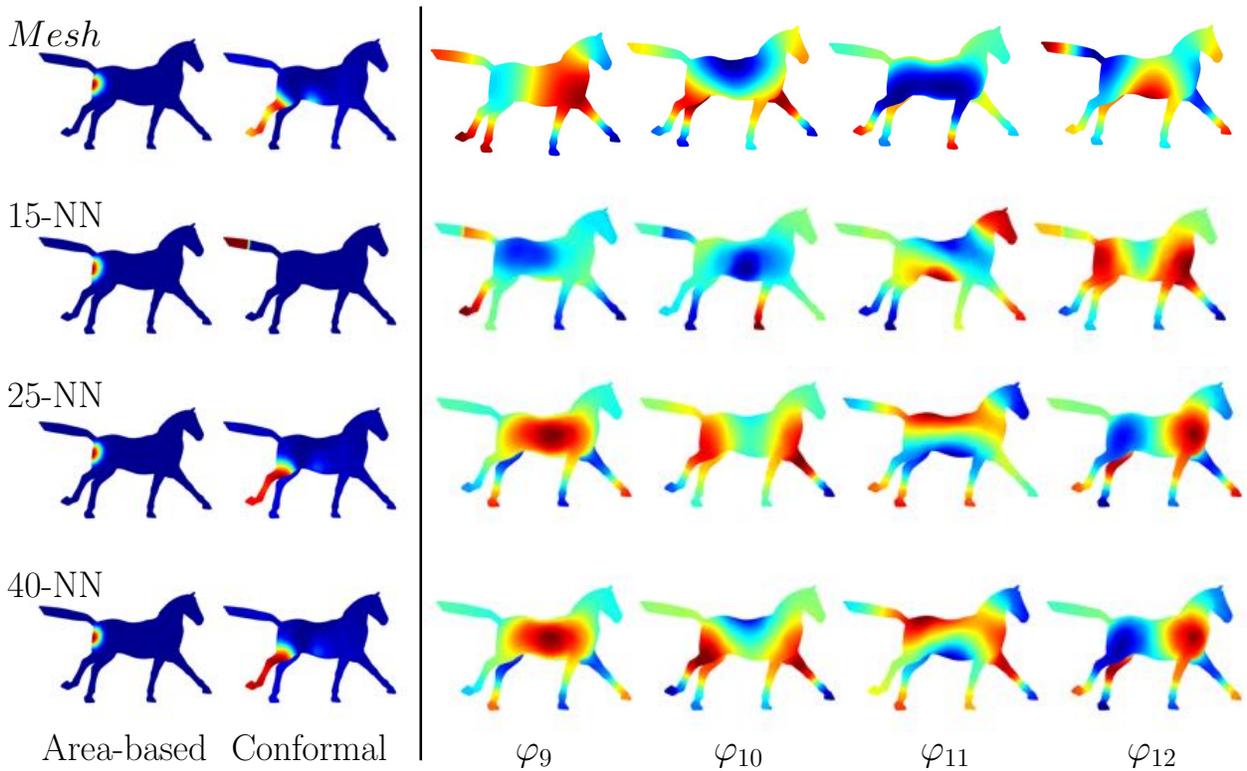


Figure 5.5: Comparison of different k for constructing k -NN graphs. The left-most two columns depict area-based and conformal indicators generated with meshes and k -NN graphs with respect to $k = 15, 25, 40$, the scales of all the indicators are fixed as $k_s = 50$. The right part shows part of the eigenfunctions on N (mesh) and on Y (PCD), which are visually distinguishable.

5.3.3 Reliability of PCD Setting in Multi-scale Framework

We’ve observed consistency between the results from the mesh setting and the PCD setting in Figure 5.1 at a single scale, now we go for multiple scales. Two near-isometric human poses in comparison are depicted in Figure 5.6(a). The top row (Figure 5.6(b), (d)) shows the indicators generated with mesh inputs: part (b) shows area differences and part (d) shows conformal differences. From left to right, the corresponding scales are respectively $k_s = 20, 60, 180$.

Again we take the vertices from the meshes as the input point clouds. In Figure 5.6(c) and (e) are the corresponding results from the PCD setting. In the case of area distortion detection, the indicators in (b) and (c) match well. Though the stability of the conformal based multi-scale framework with respect to perturbed inputs is not proven in theory (see the discussion in Section 4.6.5), the outcomes regarding the conformal differences show some consistency: the highlighted areas evolve in the same pattern along rows (d) and (e)

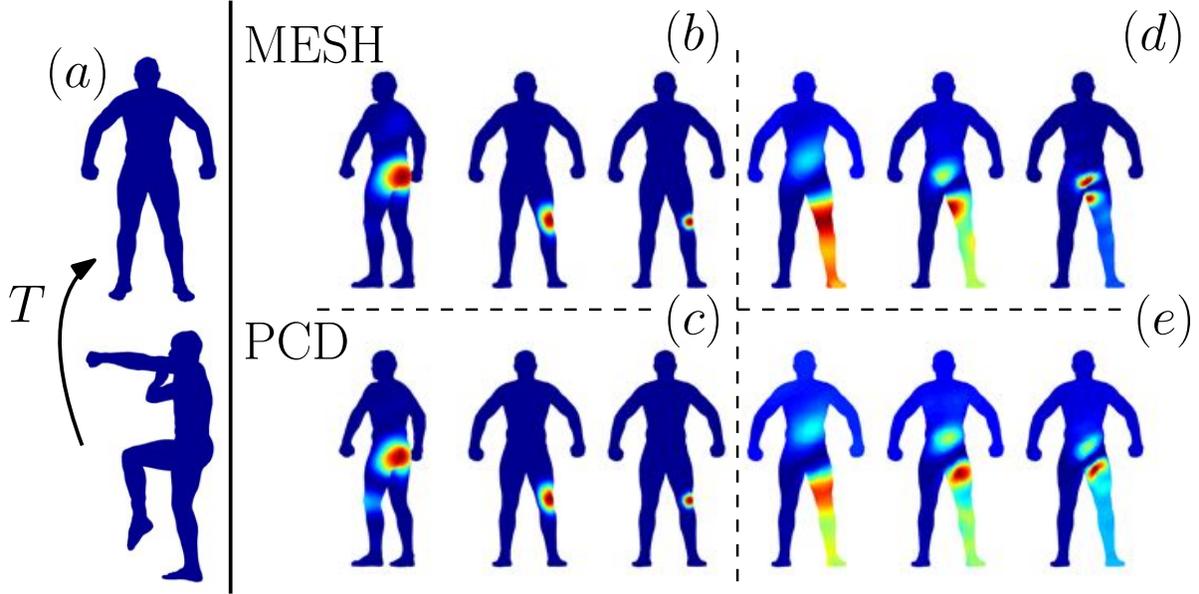


Figure 5.6: (a) shapes in comparison; Top-row on the right: indicators generated from the mesh setting—(b) are the area-based ones and (c) are the conformal ones. Bottom-row on the right: indicators generated from the PCD setting—(c) are the area-based ones and (e) are the conformal ones. In each quadrant, the scales of the indicators are $k_s = 20, 60, 180$, and increasing from left to right.

5.3.4 Analyzing Shape Collections

The above experiments shows the stability of the shape difference operators for analyzing maps between a single pair of shapes in a multi-scale way. Now we perform a higher level analysis about recovering the intrinsic structure of a collection of shapes. As we prove in Section 4.5, the shape difference operators on their own are stable with respect to perturbations on input shapes. To demonstrate that, we repeat one of the experiments in [Rustamov 2013] (see figure 3 on page 7 there), but in the PCD setting.

Given a collection of shapes with pairwise correspondences, our goal is to recover a reasonable layout in the intrinsic shape space. Two collections of shapes with different structures are considered—deformed spheres and galloping horses. In each collection we select the first shape as the base shape and construct both shape difference operators with maps from all shapes to the base shape. After approximating the area-based (resp. conformal) shape difference operators, we vectorize and embed them into \mathbb{R}^2 with the PCA algorithm.

The top row of Figure 5.7 depicts the embeddings for the deformed spheres. Both layouts uncover the grid structure of the original shape collection. The results in [Rustamov 2013] suggest that in both the area-based and the conformal cases, the variances of the first two principal components are both close to 50 percents. Regarding our result: (1) Area-based case: though the sum of percentages add up to almost 100, the grid is unbalanced and stretched along the direction of the first principle component; (2) Conformal case: balance preserved, the shapes of the first and the second rows are not well differentiated, suggesting that the operators generated in the PCD setting fail to capture the small differences.

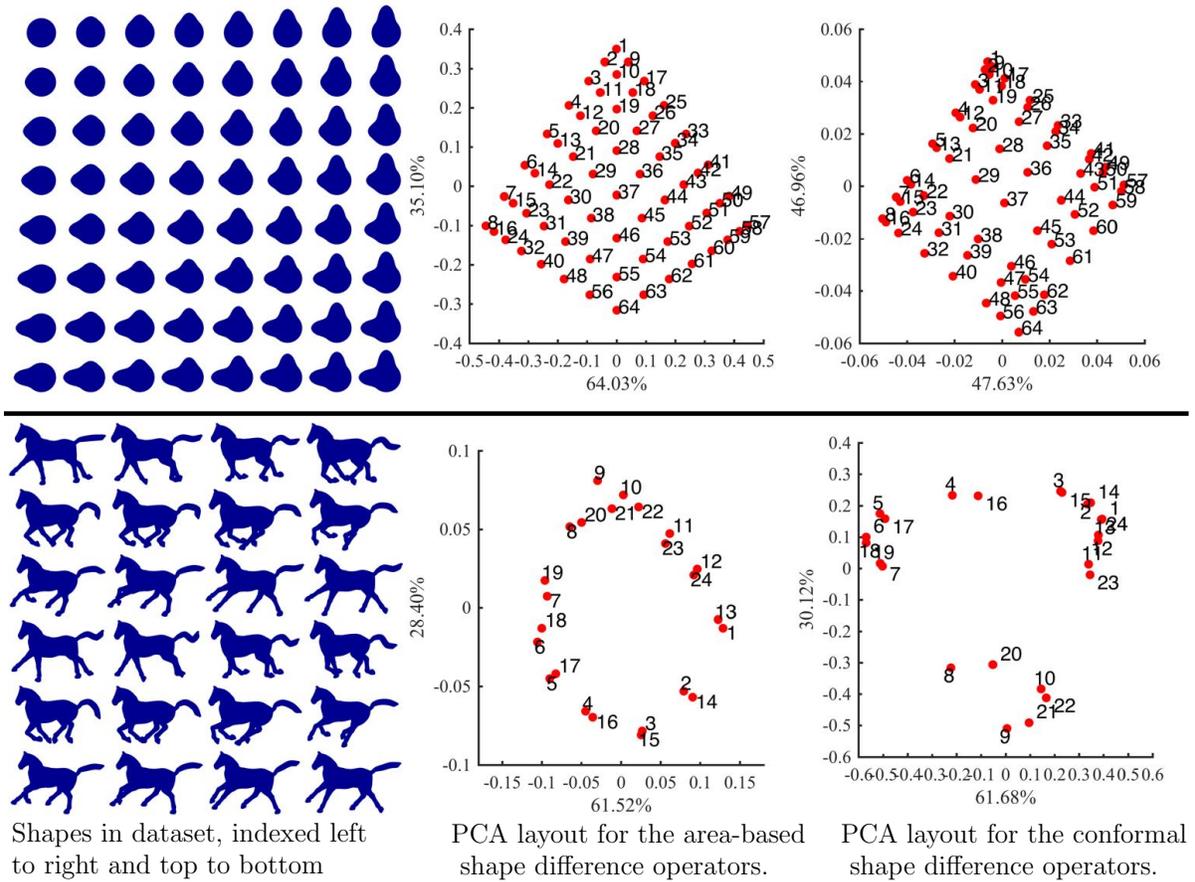


Figure 5.7: PCA plots of the two shape difference operators.

The bottom row shows the layouts for the galloping horse sequence. As demonstrated in the left panel, the sequence consists of two circles of continuous movements of the horse. Our results successfully capture the circular structure of the sequence: depicted in the layout, point i is close to point $i+12$ ($i = 1, 2, \dots, 12$). The group of $i = 1, \dots, 12$ and the group of $i = 13, \dots, 24$ each form a circle. The result also reveals the fact that there are more conformal distortions than area distortions in this data, as the range of layout in the right panel is larger than that in the middle one.

Overall, we conclude from these experiments that although the results from the PCD setting is not as accurate as those from the mesh setting, they capture most of the desired information hidden in the data. Considering that we start from a much coarser understanding of the input shapes, these results are non-trivial and remarkable.

5.3.5 Non-uniformly Sampled Data

In the previous experiments, the point clouds are nearly uniformly sampled from the underlying shapes and clean without noises or outliers. In this section, we test our method with more tricky data.

In order to test the robustness of the pipeline with respect to non-uniform sampling, a collection of synthetic data is generated as in Figure 5.8. There are two pairs of shapes: S, B and S', B' . The former pair is uniformly sampled, while in S', B' , the densified half is triply denser than the unperturbed half. Below the shapes are area-based indicators generated with 4 possible maps between spheres S, S' and bumped spheres B, B' . Obviously the four indicators all highlight the correct area, regardless of how the sampling densities are perturbed.

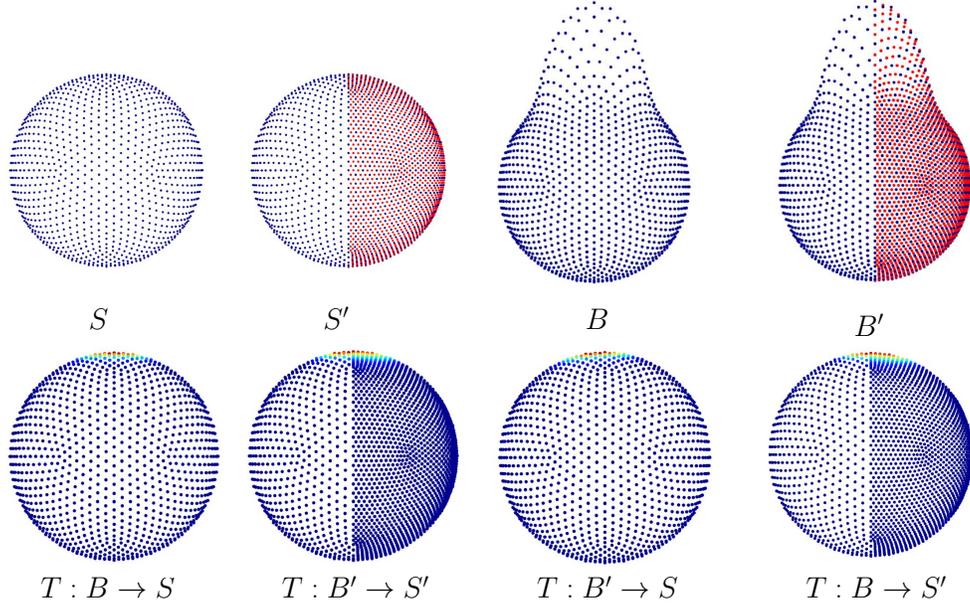


Figure 5.8: Top row: S is a uniformly sampled sphere and S' is generated by tripling sampling density on the east half of S . B' is generated by densifying B in the same fashion. Bottom row: The area-based indicators at scale $k_s = 50$, the point clouds in comparison are marked below.

Further more, we take noisy point clouds into consideration. For the sake of simplicity, we only add noises to Y , therefore we can use the same map inherited from the mesh setting (though it is not surjective any more).

We first randomly select n_p points in Y . Then for p from these n_p points, we perturb $p = (p_x, p_y, p_z) \in \mathbb{R}^3$ to $(p_x + dx, p_y + dy, p_z + dz)$ where dx, dy, dz are one-dimension random variables distributed normally with mean 0, and standard deviation d_Y (d_Y is the mean length of edges in mesh N). Repeating the displacements r times for each p , we enlarge Y to Y' with $n_p r$ more points.

In this experiment, three pairs of PCD are involved: human poses (12500 points), horses (8431 points) and cats (7207 points). In Figure 5.9, all the area-based indicators are computed at scale $k_s = 50$. The right-most two columns are indicators with respect to noisy point clouds regarding the target shape and the parameters (n_p, r) we use to generate the random noises are marked below the corresponding point clouds. Compared to the second column, where the point clouds Y are unperturbed, the indicators from noisy data are fairly robust. Note that the number of added noisy points to each point cloud is at least 3000, which is not ignorable in any case. The noisy points are clearly visible in the figure.

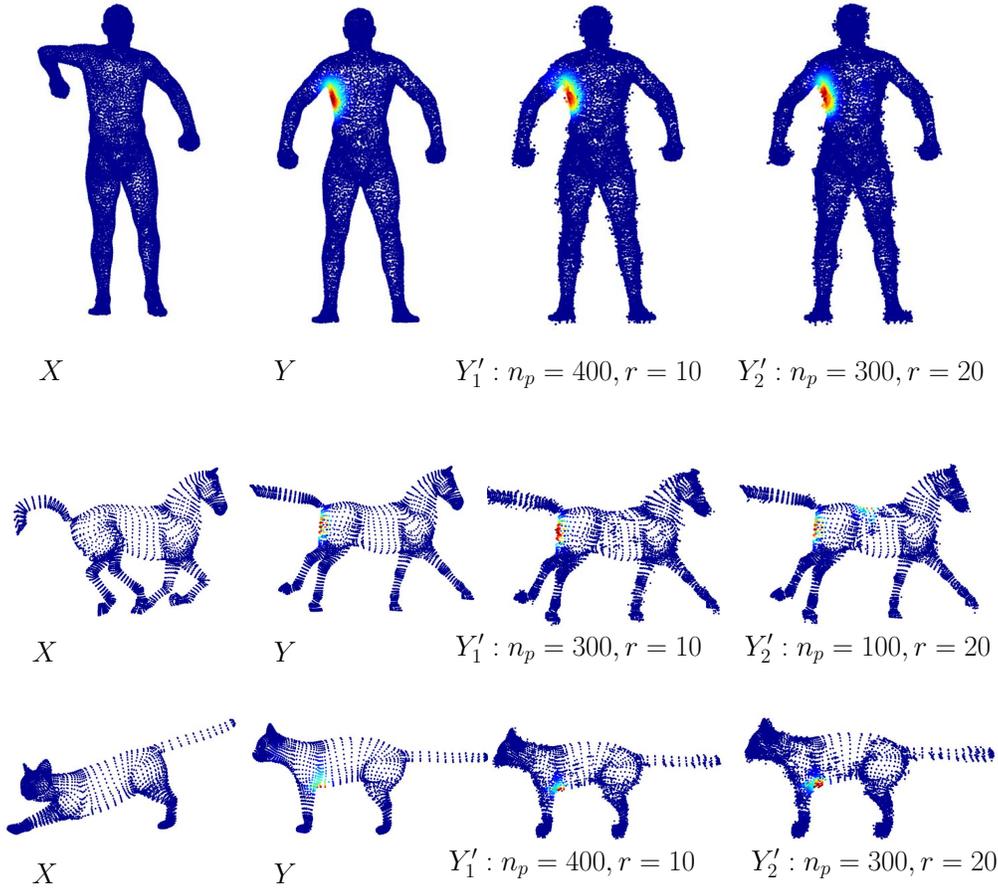


Figure 5.9: Robustness of results from the PCD setting with respect to noisy point clouds: X and Y are the original point clouds extracted from meshes. $Y'_i, i = 1, 2$ are noisy versions of Y , which are generated with the parameters marked below. The functions plotted on Y, Y'_1 and Y'_2 are area-based indicator at the same scale $k_s = 50$.

Particularly, we find in the horse case that the indicator on Y'_2 is a bit off—it also highlights a part of the horse back, while the one on Y'_1 (to which 1000 more noisy points added than to Y'_2) is more consistent. We interpret this by comparing the ways noises are distributed: in Y'_1 , the noisy points are more decentralized, whereas in Y'_2 , more points are generated around each of the selected point in Y . Thus the sampling density is more distorted in Y'_2 , resulting in a less consistent indicator.

5.4 Limitation and Perspective

In this section, we show some limitations and perspectives of the pipeline proposed and empirically tested above.

5.4.1 Limitations

First, we point out some limitations of our pipeline. The first one is about scale. The way we build connectivity among points determines that the resulting graph is insensitive to local changes across objectives. On the one hand, it gives rise to robustness with respect to noisy data or even outliers (as shown in the last example in Section 5.3). On the other hand, it makes detecting differences of finer scales (i.e., larger k_s) between objectives difficult. In fact, as shown in Figure 5.7, the conformal shape difference operators from the PCD setting can not clearly differentiate the conformal changes between deformed spheres in the first and the second row depicted in the top-left panel. Plus, this limitation is the reason we fix the scale k_s relatively small (mostly $k_s = 50$).

The second limitation is that when the reach (see Definition 3.11) of the underlying shape is small, it is difficult to capture the right geometric information with the k -NN graphs. That is because the k -NN graphs are constructed with purely extrinsic information—the Euclidean distance between points. If there exist points that are distant on the shape but close in the ambient space, then the k -NN graph will connect them and introduce distortions.

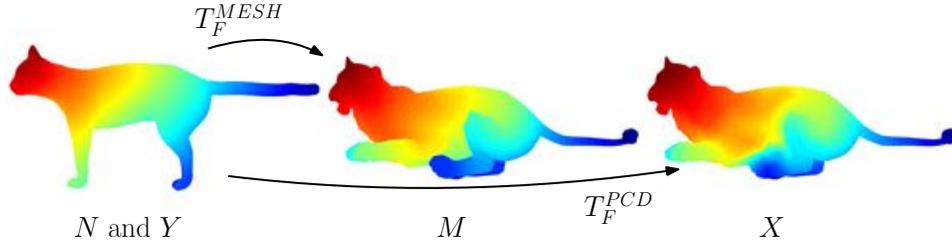


Figure 5.10: The low-rank approximation of the functional from a cat shape to a lion shape is computed with 50 eigenvectors on the cat and 150 eigenvectors on the lion. The superscripts indicate in which setting T_F is produced. The function visualized on the cat is transformed to M and X with respective functional maps. Defects are evidenced on the tip of the left hind paw in X , where the values of $T_F^{PCD} f$ are distorted as the left hind paw is close to the front part of the lion.

A typical example is illustrated in Figure 5.10. In the left-most panel of this figure, we plot a function f on N and on Y (which are the same as being rendered on the mesh). And then we repeat the procedures of the experiment in Section 5.3.1 and compute low-rank approximation of T_F with the ground truth map T from M (resp. X) to N (resp. Y). The transformed functions, $T_F^{MESH} f$ and $T_F^{PCD} f$ are plotted respectively on M and X . The left hind paw of X is supposed to be in cold color (like the left hind paw of M), however, we observe that the tip of the paw close to the abdomen is rendered with a warmer color, meaning that $T_F^{PCD} f$ is problematic in this area. Obviously, this defect is due to the huddled pose of the lion, which makes the points belonging to the abdomen and to the limbs hard to be differentiated in the k -NN graphs.

5.4.2 Beyond Shapes

We now demonstrate the potential of applying the functional-maps-based frameworks in a more abstract setting. In all above experiments, we consider point clouds sampled from closed shapes embedded in \mathbb{R}^3 . Nevertheless, our theoretical analysis applies for any pair of manifolds without constraint on its intrinsic or

embedding dimension. We now give an example showing one key advantage of our pipeline—we need no prior information about the object manifolds, and build everything from point sets.

We re-visit a classic example from [Tenenbaum 2000a], where a point cloud sampled from a Swiss-roll embedded in \mathbb{R}^3 is processed. The goal there is to *unfold* the Swiss-roll into \mathbb{R}^2 so that local distance are as preserved as possible. It is then demonstrated that by applying the *Isomap* algorithm proposed in [Tenenbaum 2000a] we obtain a significantly improved result than using other algorithm like *MDS*.

In our experiment, we obtain two embeddings in \mathbb{R}^2 , Y_{Isomap} and Y_{MDS} (see respectively the top-left and the bottom-left of Figure 5.11), from applying Isomap and MDS on the point cloud data consists of 1000 points X_{sw} (in the middle of Figure 5.11) sampled from a Swiss-roll embedded in \mathbb{R}^3 . With the ground truth of the unfolding, we can visualize and compare the two embeddings Y_{Isomap} , Y_{MDS} as being colored in the left of Figure 5.11.

Now assume that we only have point-to-point correspondences between X_{sw} and Y_{Isomap} (*resp.* Y_{MDS}). Our goal is to compare the two embeddings in \mathbb{R}^2 and figure out where the algorithms have problem in unfolding. Still we construct 40-NN graphs on the three point cloud data, and then get two pairs of shape difference operators by comparing X_{sw} to Y_{Isomap} (*resp.* Y_{MDS}).

By comparing the constrained norms of both operators with respect to the first 40 eigenvalues of the source PCD— X_{sw} , we observe a clear distinction in the conformal-based case, as shown in the right part of Figure 5.11. The constrained norm of R_{MDS} , constructed from the difference between X_{sw} and Y_{mds} is far larger than the other one (7x).

Furthermore, the corresponding conformal indicator highlights where distortions take place when mapping from X_{sw} to Y_{mds} .

The results indeed match the ground truth. First, we observe dramatic difference in conformal-based cases because if an algorithm preserves the intrinsic two-dimension structure of the Swiss-roll, then the edge connection should also be preserved. In this sense, we can conclude that the MDS algorithm is outperformed by the Isomap in this task as it connects pairs of points which are far away in the original embedding. Second, the red points marked out on the right of Figure 5.11 are actually mapped closely to their antipodal points, thus they are reasonably problematic points. The red points marked out in the isomap side, are close to the boundary of the original embedding.

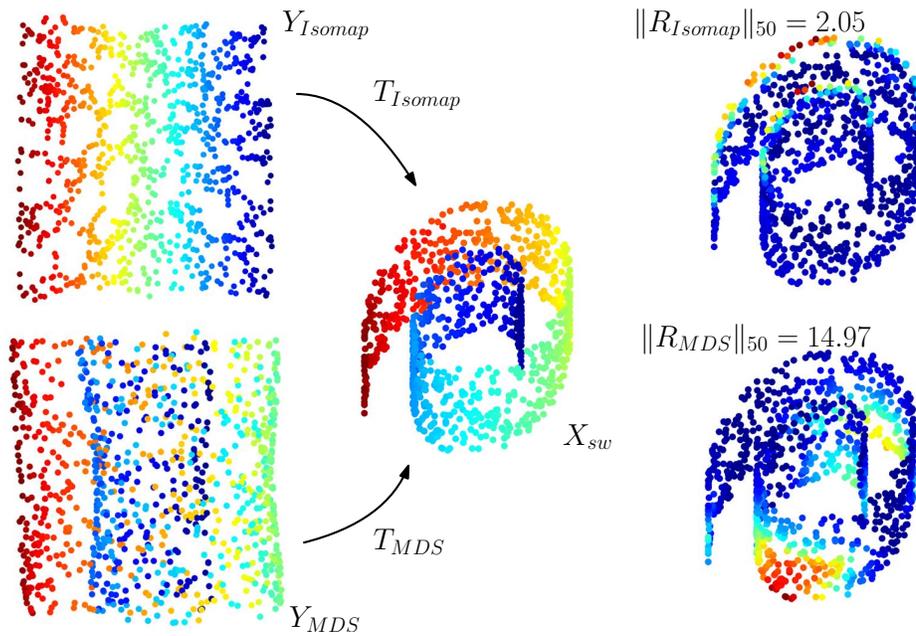


Figure 5.11: $X_{sw} \in \mathbb{R}^3$ is a point cloud sampled from a *swiss roll*. Y_{Isomap} (resp. Y_{MDS}) is the embedding of X_{sw} in \mathbb{R}^2 generated with the Isomap (resp. MDS) algorithm. The three point clouds are colored with the *ground truth*. On the right are the conformal indicator computed by comparing the two embeddings with X_{sw} respectively. The constrained norms suggest that there are more distortions introduced by the embedding regarding the MDS algorithm in this case, and the indicators manages to capture the problematic parts.

Conclusion and Perspectives

In this thesis, we study two problems on geometric data analysis—metric reconstruction for filamentary structures and shape comparison, and contribute in both theory and practice. Our investigation also raises some interesting questions and open perspective for future research. In the following, we quickly list a few of them that might be of particular interest.

Metric Reconstruction for Filamentary Structures In theory, we propose the α -Reeb graph and show that the (α -)Reeb graph of the distance function is an appropriate tool for approximating a filamentary structure, especially from the metric reconstruction point of view.

Still, there are many open problems following. For instance, with another base point, we obtain a new distance function and thus a new (α -)Reeb graph. Can we gain more information about the filamentary structure by comparing/combining (in some sense) these two graphs?

Stability Analysis on Functional-maps-based Frameworks We prove in the general case where M, N are Riemannian manifolds of dimension n that the shape difference operators and the associated multi-scale framework is stable in different senses.

Notice that the functional-maps-based frameworks take a trio (M, N, T) as input. The missing part of perturbations on the input is then about T . In fact a perturbation analysis with respect to T is of practical importance, because in the discrete setting T is combinatorial and sensitive to noises.

Performing Functional-maps-based Frameworks on PCD As shown in the experiments, our pipeline is easy to implement, and enforces less constraints on the inputs. We perform the functional-maps-based frameworks with our pipeline on some point clouds from 3D shapes and obtain results that are relevant and robust. Undoubtedly this combination extends the range of potential applications. Though there exists deficiencies due to the extrinsic nature of our approach, it is still a nice complementary to the more delicate analysis methods.

It's worth noting that this approach is purely experimental and without any theoretical guarantee. Giving theoretical guarantees for the PCD setting is both appealing and challenging.

Bibliography

- [Aanjaneya 2012] Mridul Aanjaneya, Frederic Chazal, Daniel Chen, Marc Glisse, Leonidas Guibas and Dmitriy Morozov. *Metric Graph Reconstruction From Noisy Data*. International Journal of Computational Geometry & Applications, vol. 22, no. 04, pages 305–325, 2012. (Cited on pages 11 and 44.)
- [Abraham 2007] Ittai Abraham, Mahesh Balakrishnan, Fabian Kuhn, Dahlia Malkhi, Venugopalan Ramasubramanian and Kunal Talwar. *Reconstructing approximate tree metrics*. In Proceedings of the 26th annual ACM Symposium on Principles of Distributed Computing - PODC '07, page 43, 2007. (Cited on page 11.)
- [Amenta 1998] Nina Amenta, Marshall Bern and David Eppstein. *The Crust and the β -Skeleton: Combinatorial Curve Reconstruction*. Graphical Models and Image Processing, vol. 60, no. 2, pages 125–135, 1998. (Cited on page 11.)
- [Arias-Castro 2006] Ery Arias-Castro, David L. Donoho and Xiaoming Huo. *Adaptive multiscale detection of filamentary structures in a background of uniform random points*. Annals of Statistics, vol. 34, no. 1, pages 326–349, 2006. (Cited on page 11.)
- [Azencot 2013] O. Azencot, M. Ben-Chen, F. Chazal and M. Ovsjanikov. *An Operator Approach to Tangent Vector Field Processing*. In SGP, 2013. (Cited on page 52.)
- [Azencot 2014] O. Azencot, S. Weißmann, M. Ovsjanikov, M. Wardetzky and M. Ben-Chen. *Functional Fluids on Surfaces*. In SGP, 2014. (Cited on page 52.)
- [Bai 2012] Xiang Bai, Bo Wang, Cong Yao, Wenyu Liu and Zhuowen Tu. *Co-transduction for shape retrieval*. IEEE Transactions on Image Processing, vol. 21, no. 5, pages 2747–2757, 2012. (Cited on page 49.)
- [Bauer 2014] Ulrich Bauer, Xiaoyin Ge and Yusu Wang. *Measuring Distance Between Reeb Graphs*. In Proceedings of the Thirtieth Annual Symposium on Computational Geometry, SOCG'14, pages 464:464–464:473, New York, NY, USA, 2014. ACM. (Cited on page 11.)
- [Belkin 2002] M. Belkin and P. Niyogi. *Laplacian Eigenmaps and Spectral Techniques for Embedding and Clustering*. In Advances in Neural Information Processing Systems, 2002. (Cited on page 86.)
- [Belkin 2003] Mikhail Belkin and Partha Niyogi. *Laplacian Eigenmaps for Dimensionality Reduction and Data Representation*. Neural Computation, vol. 15, no. 6, pages 1373–1396, 2003. (Cited on pages 4 and 9.)
- [Belkin 2008] Mikhail Belkin and Partha Niyogi. *Towards a theoretical foundation for Laplacian-based manifold methods*. Journal of Computer and System Sciences, vol. 74, no. 8, pages 1289 – 1308, 2008. (Cited on page 85.)

- [Belkin 2009] M. Belkin, J. Sun and Y. Wang. *Constructing Laplace operator from point clouds in R^d* . In In Proc. 20th ACM-SIAM Sympos. Discrete Algorithms, pages 1031–1040, 2009. (Cited on pages 52 and 85.)
- [Bernstein 2000] Mira Bernstein, Vin De Silva, John C. Langford and Joshua B. Tenenbaum. *Graph Approximations to Geodesics on Embedded Manifolds*, 2000. (Cited on pages 33 and 89.)
- [Biasotti 2000] Silvia Biasotti, Bianca Falcidieno and Michela Spagnuolo. *Extended Reeb Graphs for Shape Analysis and Model Compression*. Discrete Geometry for Computer Imagery, Lecture Notes in Computer Science, vol. 1953/2000, pages 185–197, 2000. (Cited on page 6.)
- [Biasotti 2008] S. Biasotti, D. Giorgi, M. Spagnuolo and B. Falcidieno. *Reeb graphs for shape analysis and applications*. Theoretical Computer Science, vol. 392, no. 1-3, pages 5–22, 2008. (Cited on page 5.)
- [Biau 2011] G. Biau, F. Chazal, D. Cohen-Steiner and C. Rodrigues L. Devroye. *A Weighted k -Nearest Neighbor Density Estimate for Geometric Inference*. Electronic Journal of Statistics, vol. 5, pages 204–237, 2011. (Cited on page 88.)
- [Botsch 2010] Mario Botsch, Leif Kobbelt, Mark Pauly, Pierre Alliez and Bruno Levy. Polygon mesh processing. AK Peters, 2010. (Cited on page 87.)
- [Burago 2001] Dmitry Burago, Sergei Wladimirowitsch Iwanow and Yuri Burago. A Course in Metric Geometry. Crm Proceedings & Lecture Notes. American Mathematical Society, 2001. (Cited on page 13.)
- [Bădoiu 2007] Mihai Bădoiu, Piotr Indyk and Anastasios Sidiropoulos. *Approximation Algorithms for Embedding General Metrics into Trees*. In Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '07, pages 512–521, Philadelphia, PA, USA, 2007. Society for Industrial and Applied Mathematics. (Cited on page 11.)
- [Carlsson 2009] Gunnar Carlsson. *Topology and data*. Bulletin of the American Mathematical Society, vol. 46, no. 2, pages 255–308, 2009. (Cited on pages 4 and 9.)
- [Carr 2003] Hamish Carr, Jack Snoeyink and Ulrike Axen. *Computing contour trees in all dimensions*. Computational Geometry: Theory and Applications, vol. 24, no. 2, pages 75–94, 2003. (Cited on page 6.)
- [Chazal 2009] F. Chazal, D. Cohen-Steiner, A. Lieutier and B. Thibert. *Stability of curvature measures*. Eurographics Symposium on Geometry Processing, vol. 28, no. 5, pages 1485–1496, 2009. (Cited on page 52.)
- [Chazal 2012] Frederic Chazal, Vin de Silva, Marc Glisse and Steve Oudot. The structure and stability of persistence modules. SpringerBriefs in Mathematics. Springer, 2012. (Cited on pages 4 and 9.)
- [Chazal 2014] Frédéric Chazal, Brittany T. Fasy, Fabrizio Lecci, Bertrand Michel, Alessandro Rinaldo and Larry Wasserman. *Robust Topological Inference: Distance To a Measure and Kernel Distance*. arXiv:1412.7197 [cs, math, stat], pages 1–32, 2014. (Cited on page 52.)

- [Chazal 2015] Frédéric Chazal, Ruqi Huang and Jian Sun. *Gromov-Hausdorff Approximation of Filamentary Structures Using Reeb-Type Graphs*. *Discrete & Computational Geometry*, vol. 53, no. 3, pages 621–649, 2015. (Cited on page 8.)
- [Chen 2010] Daniel Chen, Leonidas J Guibas, John Hershberger and Jian Sun. *Road network reconstruction for organizing paths*. *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1309–1320, 2010. (Cited on page 11.)
- [Chepoi 2008] Victor Chepoi, Feodor F. Dragan, Bertrand Estellon, Michel Habib and Yann Vaxès. *Notes on diameters, centers, and approximating trees of δ -hyperbolic geodesic spaces and graphs*. *Electronic Notes in Discrete Mathematics*, vol. 31, no. C, pages 231–234, 2008. (Cited on page 11.)
- [Choi 2010] Ena Choi, Nicholas A. Bond, Michael A. Strauss, Alison L. Coil, Marc Davis and Christopher N A Willmer. *Tracing the filamentary structure of the galaxy distribution at $z \sim 0.8$* . *Monthly Notices of the Royal Astronomical Society*, vol. 406, no. 1, pages 320–328, 2010. (Cited on page 11.)
- [Cohen-Steiner 2007] David Cohen-Steiner, Herbert Edelsbrunner and John Harer. *Stability of persistence diagrams*. *Discrete and Computational Geometry*, vol. 37, no. 1, pages 103–120, 2007. (Cited on page 52.)
- [Comaniciu 2002] Dorin Comaniciu and Peter Meer. *Mean shift: A robust approach toward feature space analysis*. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 5, pages 603–619, 2002. (Cited on page 4.)
- [Dey 2000] Tamal K Dey, Kurt Mehlhorn and Edgar A Ramos. *Curve reconstruction: Connecting dots with good reason*. *Computational Geometry*, vol. 15, no. 4, pages 229–244, 2000. (Cited on page 11.)
- [Dey 2001] Tamal K. Dey and Rephael Wenger. *Reconstructing curves with sharp corners*. *Computational Geometry: Theory and Applications*, vol. 19, no. 2-3, pages 89–99, 2001. (Cited on page 11.)
- [Dey 2010] T.K. Dey, P. Rajan and Y. Wang. *Convergence, stability, and discrete approximation of Laplace spectra*. In *In Proc. 21st ACM-SIAM Sympos. Discrete Algs.*, 2010. (Cited on page 52.)
- [Dey 2013a] Tamal Krishna Dey, Fengtao Fan and Yusu Wang. *Graph induced complex on point data*. *Proceedings of the 29th annual symposium on Symposium on computational geometry - SoCG '13*, page 107, 2013. (Cited on page 10.)
- [Dey 2013b] T.K. Dey, P. Rajan and Y. Wang. *Weighted graph laplace operator under topological noise*. In *In Proc. ACM-SIAM Symposium on Discrete Algorithms*, 2013. (Cited on page 52.)
- [Dey 2014] Tamal K. Dey, Alfred Rossi and Anastasios Sidiropoulos. *Spectral concentration, robust k -center, and simple clustering*. *CoRR*, vol. abs/1404.1008, 2014. (Cited on page 52.)
- [Dhamdhare 2006] Kedar Dhamdhare, Anupam Gupta and Harald Räcke. *Improved embeddings of graph metrics into random trees*. *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm - SODA '06*, pages 61–69, 2006. (Cited on page 11.)

- [Dieudonné 1960] Jean Dieudonné. *Foundations of modern analysis*, 1960. (Cited on page 18.)
- [EarthquakeSearch] EarthquakeSearch. <http://earthquake.usgs.gov/earthquakes/eqarchives/epic/>. (Cited on page 44.)
- [Edelsbrunner 2002] Herbert Edelsbrunner, David Letscher and Afra Zomorodian. *Topological persistence and simplification*. *Discrete and Computational Geometry*, vol. 28, no. 4, pages 511–533, 2002. (Cited on pages 4 and 9.)
- [Edelsbrunner 2010] H. Edelsbrunner and J. Harer. *Computational topology: an introduction*. American Mathematical Society, Providence, RI, 2010. (Cited on page 18.)
- [Escolano 2013] Francisco Escolano, Edwin R. Hancock and Silvia Biasotti. *Complexity fusion for indexing Reeb digraphs*. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 8047 LNCS, pages 120–127, 2013. (Cited on pages 5 and 6.)
- [Ester 1996] Martin Ester, Hans P Kriegel, Jorg Sander and Xiaowei Xu. *A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise*. In *Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining*, pages 226–231, 1996. (Cited on page 4.)
- [Fakcharoenphol 2004] Jittat Fakcharoenphol, Satish Rao and Kunal Talwar. *A tight bound on approximating arbitrary metrics by tree metrics*. In *Journal of Computer and System Sciences*, volume 69, pages 485–497, 2004. (Cited on page 11.)
- [Federer 1959] Herbert Federer. *Curvature measures*. *Transactions of the American Mathematical Society*, vol. 93, no. 3, pages 418–418, 1959. (Cited on page 37.)
- [Ge 2011] Xiaoyin Ge, I Safa, M Belkin and Yusu Wang. *Data Skeletonization via Reeb Graphs*. In *NIPS*, pages 837–845, 2011. (Cited on pages 5, 7 and 11.)
- [Genovese 2009] C. Genovese, M. Perone-Pacifico, I. Verdinelli and L. Wasserman. *On the Path Density of a Gradient Field*. *Annals of Statistics*, vol. 37, no. 6A, 2009. (Cited on page 11.)
- [Genovese 2012] Christopher R. Genovese, Marco Perone-Pacifico, Isabella Verdinelli and Larry Wasserman. *The Geometry of Nonparametric Filament Estimation*. *Journal of the American Statistical Association*, vol. 107, no. 498, pages 788–799, 2012. (Cited on page 11.)
- [Genovese 2014] Christopher R. Genovese, Marco Perone-Pacifico, Isabella Verdinelli and Larry Wasserman. *Nonparametric Ridge Estimation*. *The Annals of Statistics*, vol. 42, no. 4, pages 1–45, 2014. (Cited on page 11.)
- [Grigoryan 2000] Alexander Grigoryan. *Heat Kernel and Analysis on Manifolds*. *AMS/IP studies in advanced mathematics*. American Mathematical Soc., 2000. (Cited on pages 52 and 54.)
- [Gromov 2002] M. Gromov. *Metric structures for riemannian and non-riemannian spaces*. Birkhäuser, Modern Birkhäuser Classics, 2002. (Cited on page 12.)

- [Harvey 2010] William Harvey, Yusu Wang and Rephael Wenger. *A Randomized $O(M \log M)$ Time Algorithm for Computing Reeb Graphs of Arbitrary Simplicial Complexes*. In Proceedings of the Twenty-sixth Annual Symposium on Computational Geometry, SoCG '10, pages 267–276, New York, NY, USA, 2010. ACM. (Cited on page 10.)
- [Hatcher 2002] Allen Hatcher. *Algebraic Topology*. Group, vol. 55, no. July, page 544, 2002. (Cited on page 30.)
- [Hein 2007] Matthias Hein, Jean-Yves Audibert and Ulrike von Luxburg. *Graph Laplacians and their Convergence on Random Neighborhood Graph*. Journal of Machine Learning Research, pages 1325–1368, 2007. (Cited on pages 52, 85, 86 and 89.)
- [Hilaga 2001] Masaki Hilaga and Y Shinagawa. *Topology matching for fully automatic similarity estimation of 3D shapes*. Proceedings of the 28th annual conference on Computer graphics and interactive techniques, pages 203–212, 2001. (Cited on page 5.)
- [Huang 2014] Q. Huang, F. Wang and L. Guibas. *Functional Map Networks for Analyzing and Browsing Large Shape Collections*. ACM Trans. Graph., pages 33–44, 2014. (Cited on page 52.)
- [Kato 1995] T. Kato. *Perturbation Theory for Linear Operators*. Springer-Verlag GmbH., 1995. (Cited on page 52.)
- [Kwok 2013] Tsz Chiu Kwok, Lap Chi Lau, Yin Tat Lee, Shayan Oveis Gharan and Luca Trevisan. *Improved Cheeger’s Inequality: Analysis of Spectral Partitioning Algorithms through Higher Order Spectral Gap*. Proceedings of the 45th annual ACM symposium on Symposium on theory of computing - STOC '13, page 11, 2013. (Cited on page 52.)
- [Lafon 2004] S. Lafon. *Diffusion Maps and Geodesic Harmonics*. PhD. Thesis, Yale University, 2004. (Cited on pages 4 and 9.)
- [Lee 2012] James R Lee, Shayan Oveis Gharan and Luca Trevisan. *Multi-way spectral partitioning and higher-order cheeger inequalities*. Proceedings of the 44th symposium on Theory of Computing - STOC '12, vol. 61, no. 6, page 1117, 2012. (Cited on page 52.)
- [Liu 2012] Y. Liu, B. Prabhakaran and X. Guo. *Point-based manifold harmonics*. IEEE Transactions on visualization and computer graphics., 2012. (Cited on pages 52 and 85.)
- [Lloyd 1982] Stuart P. Lloyd. *Least Squares Quantization in PCM*. IEEE Transactions on Information Theory, vol. 28, no. 2, pages 129–137, 1982. (Cited on page 4.)
- [Louis 2012] Anand Louis, Prasad Raghavendra, Prasad Tetali and Santosh Vempala. *Many sparse cuts via higher eigenvalues*. In Proceedings of the 44th symposium on Theory of Computing - STOC '12, page 1131, 2012. (Cited on page 52.)
- [Mémoli 2009] F. Mémoli. *Spectral Gromov-Wasserstein distances for shape matching*. In 2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops 2009, pages 256–263, 2009. (Cited on page 52.)

- [Mémoli 2011] F. Mémoli. *A spectral notion of Gromov-Wasserstein distance and related methods*. Applied and Computational Harmonic Analysis, pages 363–401, 2011. (Cited on page 52.)
- [Mérigot 2011] Quentin Mérigot, Maks Ovsjanikov and Leonidas J. Guibas. *Voronoi-based curvature and feature estimation from point clouds*. IEEE Transactions on Visualization and Computer Graphics, vol. 17, no. 6, pages 743–756, 2011. (Cited on page 52.)
- [Meyer 2003] Mark Meyer, Mathieu Desbrun, Peter Schröder and Alan H. Barr. Discrete differential-geometry operators for triangulated 2-manifolds, pages 35–57. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003. (Cited on page 87.)
- [Mitra 2004] Niloy J. Mitra, An Nguyen and Leonidas Guibas. *Estimating Surface Normals in Noisy Point Cloud Data*. International Journal of Computational Geometry and Applications, vol. 14, no. 4-5, pages 261–276, 2004. (Cited on page 52.)
- [Openstreetmap] Openstreetmap. <http://www.openstreetmap.org/>. (Cited on page 44.)
- [Oveis Gharan 2014] Shayan Oveis Gharan and Luca Trevisan. *Partitioning into expanders*. Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 1256–1266, 2014. (Cited on page 52.)
- [Ovsjanikov 2012] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher and Leonidas Guibas. *Functional Maps: A Flexible Representation of Maps Between Shapes*. In ACM SIGGRAPH, 2012. (Cited on pages 6, 49, 58, 59, 85, 86 and 91.)
- [Ovsjanikov 2013] Maks Ovsjanikov, Mirela Ben-Chen, Frédéric Chazal and Leonidas Guibas. *Analysis and Visualization of Maps Between Shapes*. In Computer Graphics Forum (CGF), 2013. (Cited on pages 7, 8, 49, 50, 51, 59, 66, 72, 73, 74, 77, 85 and 86.)
- [Parsa 2013] Salman Parsa. *A Deterministic $O(m \log m)$ Time Algorithm for the Reeb Graph*. Discrete & Computational Geometry, vol. 49, no. 4, pages 864–878, 2013. (Cited on page 10.)
- [Pascucci 2003] Valerio Pascucci and Kree Cole-Mclaughlin. *Parallel computation of the topology of level sets*. Algorithmica (New York), vol. 38, no. 1, pages 249–268, 2003. (Cited on page 6.)
- [Pinkall 1993] U. Pinkall and K. Polthier. *Computing discrete minimal surfaces and their conjugates*. Experimental Mathematics, pages 15–36, 1993. (Cited on pages 76, 85 and 87.)
- [Planck 2006] Max Planck and Ulrike Von Luxburg. *A Tutorial on Spectral Clustering A Tutorial on Spectral Clustering*. Statistics and Computing, vol. 17, no. March, pages 395–416, 2006. (Cited on page 4.)
- [Pražni 2010] Jörg Stefan Pražni, Timo Ropinski, Jörg Mensmann and Klaus Hinrichs. *Shape-based transfer functions for volume visualization*. In IEEE Pacific Visualization Symposium 2010, PacificVis 2010 - Proceedings, pages 9–16, 2010. (Cited on page 49.)

- [Reeb 1946] Georges Reeb. *Sur les points singuliers d'une forme de Pfaff complètement intégrable ou d'une fonction numérique*. Comptes Rendus Acad. Sciences, vol. 222, pages 847–849, 1946. (Cited on page 5.)
- [Reuter 2006] Martin Reuter, Franz Erich Wolter and Niklas Peinecke. *Laplace-Beltrami spectra as 'Shape-DNA' of surfaces and solids*. CAD Computer Aided Design, vol. 38, no. 4, pages 342–366, 2006. (Cited on page 52.)
- [Rustamov 2013] Raif Rustamov, Maks Ovsjanikov, Omri Azencot, Mirela Ben-Chen, Frédéric Chazal and Leonidas Guibas. *Map-Based Exploration of Intrinsic Shape Differences and Variability*. In ACM SIGGRAPH, 2013. (Cited on pages 7, 49, 50, 59, 60, 74, 85, 86 and 94.)
- [Schumacher 2013] Henrik Schumacher. *Conformal maps and p-Dirichlet energies*. Tech. rep., Univ. of Gottingen, vol. 93, 2013. (Cited on page 60.)
- [Singh 2007] Gurjeet Singh, Facundo Mémoli and Gunnar Carlsson. *Topological Methods for the Analysis of High Dimensional Data Sets and 3D Object Recognition*. In Eurographics Symposium on Point-Based Graphics, pages 91–100, 2007. (Cited on pages 6, 9, 10, 11, 16 and 43.)
- [Tangelder 2008] Johan W H Tangelder and Remco C. Veltkamp. *A survey of content based 3D shape retrieval methods*. Multimedia Tools and Applications, vol. 39, no. 3, pages 441–471, 2008. (Cited on page 49.)
- [Tenenbaum 2000a] J. B. Tenenbaum, V. de Silva and J. C. Langford. *A Global Geometric Framework for Nonlinear Dimensionality Reduction*. Science., 2000. (Cited on pages 86 and 99.)
- [Tenenbaum 2000b] Joshua B Tenenbaum, Vin de Silva and John C Langford. *A Global Geometric Framework for Nonlinear Dimensionality Reduction*. Science, vol. 290, no. 5500, page 2319, 2000. (Cited on pages 4 and 9.)
- [Tupin 1998] F Tupin, H Maitre, J F Mangin, J M Nicolas and E Pechersky. *Detection of linear features in SAR images: application to road network extraction*. IEEE Transactions on Geoscience and Remote Sensing, vol. 36, no. 2, pages 434–453, mar 1998. (Cited on page 11.)
- [Van Kreveld 2006] Marc Van Kreveld, René van Oostrum, Chandrajit Bajaj, Valerio Pascucci and Dan Schikore. *Contour Trees and Small Seed Sets for Isosurface Generation*. In Topological Data Structures for Surfaces: An Introduction to Geographical Information Science, pages 71–85. 2006. (Cited on page 6.)
- [von Luxburg 2006] U. von Luxburg. *A tutorial on spectral clustering*. Technical Report 149, Max Planck Institute for Biological Cybernetics., 2006. (Cited on pages 52 and 86.)
- [Von-Tycowicz 2015] Christoph Von-Tycowicz, Christian Schulz, Hans-Peter Seidel and Klaus Hildebrandt. *Real-Time Nonlinear Shape Interpolation*. ACM Transactions on Graphics (TOG), vol. 34, no. 3, page 34, 2015. (Cited on page 49.)

-
- [Wang 2013] Fan Wang, Qixing Huang and Leonidas Guibas. *Image Co-Segmentation via Consistent Functional Maps*. In The 14th International Conference on Computer Vision (ICCV), 2013. (Cited on page 52.)
- [Wardetzky 2005] M. Wardetzky. *Covergence of the cotangent formula: An overview*. Discrete Differential Geometry., pages 89–112, 2005. (Cited on page 52.)
- [Xu 2006] Dong Xu, Hongxin Zhang, Qing Wang and Hujun Bao. *Poisson shape interpolation*. Graphical Models, vol. 68, no. 3, pages 268–281, 2006. (Cited on page 49.)
- [Xu 2007] G. Xu. *Discrete Laplace-Beltrami operators and their convergence*. Computer Aided Geometric Design, pages 398–407, 2007. (Cited on page 52.)
- [Zomorodian 2005] Afra Zomorodian and Gunnar Carlsson. *Computing persistent homology*. Discrete and Computational Geometry, vol. 33, no. 2, pages 249–274, 2005. (Cited on pages 4 and 9.)

Titre : Deux contributions à l'analyse géométrique de données : approximation de structures filamenteuses et stabilité des approches fonctionnelles pour la comparaison de formes.

Mots clés : graphe de Reeb, rapprochement métrique, cartes fonctionnelles, perturbation (mathématiques).

Résumé : En ce moment même, d'énormes quantités de données sont générées, collectées et analysées. Dans de nombreux cas, ces données sont échantillonnées sur des objets à la structure géométrique particulière. De tels objets apparaissent fréquemment dans notre vie quotidienne. Utiliser ce genre de données pour inférer la structure géométrique de tels objets est souvent ardue. Cette tâche est rendue plus difficile encore si les objets sous-jacents sont abstraits ou encore de grande dimension.

Dans cette thèse, nous nous intéressons à deux problèmes concernant l'analyse géométrique de données. Dans un premier temps, nous nous penchons sur l'inférence de la métrique de structures filamenteuses. En supposant que ces structures sont des espaces métriques proches d'un graphe métrique nous proposons une méthode, combinant les graphes de Reeb et l'algorithme Mapper, pour approximer la structure filamenteuse via un graphe de Reeb. Notre méthode peut de plus être facilement implémentée et permet de visualiser simplement le résultat.

Nous nous concentrons ensuite sur le problème de la comparaison de formes. Nous étudions un ensemble de méthodes récentes et prometteuses pour la comparaison de formes qui utilisent la notion d'applications fonctionnelles. Nos résultats théoriques montrent que ces approches sont stables et peuvent être utilisées dans un contexte plus général que la comparaison de formes comme la comparaison de variétés Riemanniennes de grande dimension.

Enfin, en nous basant sur notre analyse théorique, nous proposons une généralisation des applications fonctionnelles aux nuages de points. Bien que cette généralisation ne bénéficie par des garanties théoriques, elle permet d'étendre le champ d'application des méthodes basées sur les applications fonctionnelles.

Title: Two contributions to geometric data analysis: filamentary structures approximations, and stability properties of functional approaches for shape comparison.

Keywords: Reeb graph, metric approximation, functional maps, perturbation analysis.

Abstract: Massive amounts of data are being generated, collected and processed all the time. A considerable portion of them are sampled from objects with geometric structures. Such objects can be tangible and ubiquitous in our daily life. Inferring the geometric information from such data, however, is not always an obvious task. Moreover, it is not a rare case that the underlying objects are abstract and of high dimension, where the data inference is more challenging.

This thesis studies two problems on geometric data analysis. The first one concerns metric reconstruction for filamentary structures. We in general consider a filamentary structure as a metric space being close to an underlying metric graph, which is not necessarily embedded in some Euclidean spaces. Particularly, by combining the Reeb graph and the Mapper algorithm, we propose a variant of the Reeb graph, which not only faithfully approximates the metric of the filamentary structure but also allows for efficient implementation and convenient visualization of the result.

Then we focus on the problem of shape comparison. In this part, we study the stability properties of some recent and promising approaches for shape comparison, which are based on the notion of functional maps. Our results show that these approaches are stable in theory and potential for being used in more general setting such as comparing high-dimensional Riemannian manifolds.

Lastly, we propose a pipeline for implementing the functional-maps-based frameworks under our stability analysis on unorganised point cloud data. Though our pipeline is experimental, it undoubtedly extends the range of applications of these frameworks.