



On the Stability of Functional Maps and Shape Difference Operators

Ruqi Huang, Frédéric Chazal, Maks Ovsjanikov

► To cite this version:

Ruqi Huang, Frédéric Chazal, Maks Ovsjanikov. On the Stability of Functional Maps and Shape Difference Operators. Computer Graphics Forum, 2017, pp.1 - 12. 10.1111/cgf.13238 . hal-01668186

HAL Id: hal-01668186

<https://inria.hal.science/hal-01668186>

Submitted on 19 Dec 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the Stability of Functional Maps and Shape Difference Operators

R. Huang¹ and F. Chazal¹ and M. Ovsjanikov²

¹DataShape, INRIA

²LIX, Ecole Polytechnique




Figure 1: (a) Given two shapes M, N and a map T between them, the functional operator V is generated as one of the shape difference operators introduced in [?]. Intuitively, the real-valued function f_2 , which is supported on 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) Given perturbed shapes N to $\tilde{N}$, we generate highlighted functions with the multi-scale framework of [?], which takes high values (indicated by warm color) in the significantly distorted region. In our paper we prove two types of consistency: horizontally, as the scale k increases, the highlighted functions remain stable, i.e., the regions where high function value takes place are consistent; vertically, at each scale, the highlighted functions are stable with respect to the changes of the input shapes.

Abstract

In this paper, we provide stability guarantees for two frameworks that are based on the notion of functional maps – the shape difference operators introduced in [?] and the framework of [?] which is used to analyze and visualize the deformations between shapes induced by a functional map. We consider two types of perturbations in our analysis: one is on the input shapes and the other is on the change in scale. In theory, we formulate and justify the robustness that has been observed in practical implementations of those frameworks. Inspired by our theoretical results, we propose a pipeline for constructing shape difference operators on point clouds and show numerically that the results are robust and informative. In particular, we show that both the shape difference operators and the derived areas of highest distortion are stable with respect to changes in shape representation and change of scale. Remarkably, this is in contrast with the well-known instability of the eigenfunctions of the Laplace-Beltrami operator computed on point clouds compared to those obtained on triangle meshes.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: —Shape Analysis.

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 [?,?], interpolation [?,?], or visualization [?]. However, even when a map between shapes is given, encoding and visualizing the

differences between them is still challenging. Approaches based on the point-to-point correspondences usually suffer from 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. The framework of functional maps, which is introduced in [?], alleviates these issues to some extent by considering more general linear mappings between functions, which can be encoded in a multi-

scale fashion with functional bases. As demonstrated in [?], functional maps provide a compact, informative representation, which can naturally incorporate tools from spectral analysis.

Based on the notion of functional maps, several approaches have been proposed to analyze pairs or collections of shapes along with maps between them. In this paper we consider two of them, which are intimately related to each other. One is the framework of *shape difference operators* introduced in [?], which encodes the differences between a pair of shapes as linear operators acting on the functions on one of the shapes (see Figure 1(a) for an illustration of one of the operators). And the other is proposed in [?], which generates a collection of multi-scale distortion functions indicating the areas on one of the shapes which undergo deformations. The latter framework can be integrated into the former in the sense that its output, which is a set of highlighted functions, correspond, in essence, to eigenfunctions of shape difference operators.

Though the theoretical formulations of both frameworks are well-established, the associated *stability analyses* remain absent. In practice, however, we observe the robustness of the outcomes of these frameworks. For example, as shown in Figure 1(b), two types of consistency are evidenced: horizontally, as the highlighted functions are consistent with respect to the change in scale; vertically, at fixed scales, the highlighted functions are stable with respect to the changes of the input shapes. In this paper, we initiate a rigorous theoretical analysis of these stability properties. In particular, our contributions are three-fold:

- We provide the first rigorous formulations and theoretical guarantees of stability properties of the shape difference operators.
- We propose a new multi-scale scheme for extracting information from the shape difference operators, which comes with rigorous stability guarantees.
- Inspired by our theoretical results, we design a practical pipeline for computing the shape difference operators on shapes represented by point clouds, and we show numerically that this pipeline is relevant and robust, even when individual spectral quantities such as eigenfunctions of the Laplace-Beltrami operator might not be.

1.1. Overview

We assume that we are given a pair M and N of connected, compact, smooth shapes without boundary. Given a map $T : M \rightarrow N$, the authors of [?] introduce a pair of linear operators acting on real-valued functions on N , each of which captures one type of differences or distortion between the two shapes induced by T . We first study the stability of these operators with respect to perturbations on metrics and measures on M and on N (Section 4).

Then we consider the multi-scale framework based on shape difference operators. For one of the shape difference operators – V as illustrated in Figure 1(a), the authors of [?] propose a functional for evaluating the deviation from a function on N to its image under V and search for a function that maximizes the functional as a distortion indicator. Then they introduce a multi-scale framework by restricting the search to a subspace spanned by the first k eigenfunctions of the Laplace-Beltrami operator (LBO) on N . Figure 2 shows typical outputs of this framework: a collection of multi-scale




Figure 2: While being more and more localized with increasing k , the functions from $k = 20$ to 100 consistently highlight the hip of the horse, whereas the ones from $k = 120$ to 200 highlight the root of its front right leg. The corresponding quantitative measurements of distortion are marked above of each shape.

highlighted functions on the shape N and a sequence of the corresponding maxima of the energy functional shown above the highlight functions with respect to different scales ranging from $k = 20$ to 200. In this example, we observe consistency in the output at different scales, which are similar to the observations from Figure 1(b). Therefore, in the second part of our analyses (Section 5), we provide a rigorous stability analysis with respect to the change in scale. One challenge, however, is that the scale in the original framework is controlled by an integer k , and as we will demonstrate in Section 5.1, the discrete nature of scale is not suitable for stability analysis. Indeed, as we show below, the result might *not* be stable with respect to changes of k . To overcome this issue, we introduce a new multi-scale framework whose scale is controlled by a continuous parameter $C \in \mathbb{R}^+$, and discuss the connection between the two multi-scale frameworks in Section 5.4. Within this continuous multi-scale framework, we provide rigorous theoretical guarantees of the stability with respect to C .

Moreover, at any fixed scale C , we prove that the new multi-scale framework is stable with respect to perturbations on the input shapes as well. Figure 3 illustrates this property: we perturb the input shapes and show the highlighted functions at the same scale $k = 50$. Note both the stability of the highlighted regions and the proximity among the maxima of the distortion energy shown above the meshes.




Figure 3: Highlighted functions at a fixed scale for different meshes. We densify the original shape (a) by adding points in the body of the horse (b) and simplify it by down-sampling the limbs (c). The corresponding distortion energy values are shown above.

As an extension, we adapt the other shape difference operator – the one capturing conformal distortion – to the multi-scale framework of [?] and prove the stability of this extension with respect to the change in scale as well (Section 5.5).

Lastly, we notice that in practice the two frameworks have so far only been constructed on shapes which are discretized as triangle meshes. In Section 6.3 we extend these constructions by designing a pipeline for computing shape difference operators on shapes represented as point clouds. As shown in Figure 4, although the eigenfunctions of the LBO generated on the mesh N and on the point cloud Y are distinct, the highlighted function generated with M, N are comparable with the one from comparing X, Y at a fixed scale. This supports the stability results we obtain in theory, and suggests a remarkable robustness of measures based on functional maps and the derived shape difference operators.

To summarize, we provide a rigorous theoretical justification for the stability of shape difference operators in the continuous setting. We also propose a new functional sub-domain construction, which we show to be more stable than the classical truncation of eigenspace, in particular leading to provably stable solutions of certain energy functionals used for highlighting distorted regions. Finally, we demonstrate the relevance of our pipeline by applying the shape difference operators to point clouds, which suggests the possibility of extending the existing frameworks to deal more general geometric objects.

1.2. Paper Organization

After discussing related works in Section 2, we introduce the preliminaries and the notations in Section 3. We then study the stability of shape difference operators in Section 4, and provide stability analysis for the framework of [?], by analyzing the perturbations of scale, in Section 5.2 (Figure 2) and of the shapes in Section 5.3 (Figure 3). We present experimental results showing the stability properties in Section 6.

2. Related Work

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

Our main focus is to perform perturbation analysis on both shape difference operators (which are linear operators, see [?] for an introduction of perturbation analysis on them) and a spectral method based on such operators. Closely related to our analysis is the framework of [?], whose authors conduct perturbation analysis on eigenspace with respect to the Laplace-Beltrami operator on shapes with missing parts. The spectral methods have long been applied in various areas: spectral clustering [?], shape analysis [?] and so on. Besides demonstrating practical usefulness of the spectral methods, providing theoretical justifications is attracting more and more research interest. 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 [?, ?, ?, ?] present theoretical guarantees on the quality (measured by some graph conductance) of the output of the respective algorithms. It is worth noting that the works above only consider the case of a single object, while in this paper, we study operators and quantities derived from pairs of shapes. From this point

of view, our work has a similar flavor to the ones by Mémoli [?, ?], who proposes metrics between shapes based on spectral invariants 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 aimed at estimating geometric quantities. For example, in [?], the authors provide a theoretical and practical analysis of stability and accuracy of normal estimation process. In [?], a sharp feature detection algorithm is presented with guarantees of stability with respect to Hausdorff noise. In the same noise model, the stability of the curvature measures is proven under certain conditions in [?]. Similar problems are also actively studied in the community of topological data analysis (TDA). The stability of persistence diagram is verified in [?], which has been instrumental in establishing a solid theoretical foundation for data analysis using topological methods. Some more recent developments in TDA also come with stability guarantees, including, e.g., the notion of distance to a measure [?].

A rich body of research has also been devoted to providing analysis for convergence properties of various discrete Laplacian operators. In [?, ?, ?] the converging behaviors of the cotangent Laplacian operators on meshes to the underlying Laplace-Beltrami operators are investigated from diverse perspectives. While in [?, ?, ?, ?], similar problems are considered in a different setting, where the discrete Laplacian operators are built on point clouds. In particular, our discretization scheme proposed in Section 6.3 is based on the result from [?], where convergence of graph Laplacian on non-uniformly sampled point clouds is proven. Lastly, we point out that unlike the frameworks of [?, ?], our scheme does not require constructing any local mesh structure.

3. Preliminaries and Notations

In this section, we introduce the fundamental notions from differential geometry involved in this work, and refer the readers to [?] for more details. Let N be a connected, compact, smooth 2-dimensional Riemannian manifold endowed with a metric g_N . The volume (or Riemannian measure) ν_N is induced by g_N . Given a positive smooth function ρ_N on N , we obtain a weighted Riemannian manifold (N, g_N, μ_N) by letting $d\mu_N = \rho_N d\nu_N$.

Remark 3.1 *In this paper, by a Riemannian manifold we mean a triple (N, g_N, ν_N) , where the volume ν_N is induced by the metric. We use the term weighted Riemannian manifold to denote (N, g_N, μ_N) , where μ_N is an arbitrary measure having a density with respect to the volume measure on N .*

The Laplace-Beltrami operator (LBO) on N , Δ_N , is semi-negative definite and self-adjoint. Since we assume that N is compact, the spectrum of Δ_N is discrete. In fact, we can order the eigenvalues of $-\Delta_N$ such that $0 = \lambda_1 < \lambda_2 \leq \dots \leq \lambda_k \leq \dots$ (only the first eigenvalue is zero as N is connected).




Figure 4: Left: highlighted functions from the mesh setting (top) and the PCD setting (bottom) both at scale $k = 50$; Right: the 9th to the 12th eigenfunctions of the Discrete LB operator on mesh (top) and those of the Graph Laplacian on PCD (bottom).

Since N is compact and without boundary, the classic Green formula implies that for any smooth functions u, v on N .

$$\int_N u(-\Delta_N)v dV_N = \int_N \langle \nabla u, \nabla v \rangle_{g_N} dV_N \quad (1)$$

On the other hand, it is well-known that the eigenfunctions of $-\Delta_N$ form an orthonormal basis of function space $L^2(N) = \{f : \int_N f^2 dV_N < +\infty\}$, and we have the following classical result:

Proposition 3.1 Let $\{\varphi_i\}_{i \geq 1}$ be an orthonormal basis of $L^2_V(N)$ consisting of eigenfunctions of Δ_N . Then any function $u \in L^2_V(N)$ admits a decomposition $u = \sum_{i \geq 1} a_i \varphi_i$, $a_i = \int_N u \varphi_i dV_N$. Moreover:

$$\int_N u^2 dV_N = \sum_{i \geq 1} a_i^2 \quad (2)$$

If we further assume that u is differentiable, then

$$\int_N \langle \nabla u, \nabla u \rangle_{g_N} dV_N = \sum_{i \geq 1} a_i^2 \lambda_i \quad (3)$$

Here and throughout the rest of this paper we use $L^2_V(N)$ to denote the space of square integrable functions.

Functional Maps. A functional map, T_F , is simply a pull-back from the function space of N to that of M induced by the map T . Namely, given a function $w : N \rightarrow \mathbf{R}$, $T_F(w) = w \circ T$ returns a function on M . As demonstrated in [?], T_F is a linear operator across the function spaces on M and N .

Shape Difference Operators In [?], a pair of *Shape Difference Operators* was introduced, which encode the change of inner products under functional map T_F .

The area-based shape difference operator, $V : L^2(N) \rightarrow L^2(N)$, is a linear operator such that for any $f, g \in L^2(N)$,

$$\int_N fV(g) dV_N = \int_M T_F(f)T_F(g) dV_M \quad (4)$$

Rustamov and colleagues proved in [?] that such a linear operator V is well-defined for any T_F .

Note that unless T is an area-preserving map, $\int_N fg dV_N$ does not

always equal to $\int_M T_F(f)T_F(g) dV_M$, the linear operator V captures and compensates for the discrepancy.

Similarly, the so-called conformal-based shape difference operator, R , is a linear operator such that for any f, g in the Sobolev space $H_0^1(N) = \{f : \int_N f^2 + \|\nabla f\|^2 dV_N < +\infty, \int_N f dV_N = 0\}$, we have:

$$\int_N \langle \nabla f, \nabla R(g) \rangle_{g_N} dV_N = \int_M \langle \nabla T_F(f), \nabla T_F(g) \rangle_{g_M} dV_M \quad (5)$$

It follows from the Riesz representation theorem that given smooth shapes M, N and a map T , the operators V and R exist and are unique. Particularly, if M, N are 2-dimensional Riemannian manifolds without boundary, the authors of [?] show that T is locally area-preserving (resp., conformal) if and only if V (resp., R) is an identity operator.

Map Analysis In [?], an energy measuring distortions induced by a map is defined on the function space on N . Namely, for any real-valued function w on N , the authors define:

$$E(w) = \frac{\int_M T_F(w)^2 dV_M}{\int_N w^2 dV_N} \quad (6)$$

As discussed in [?], $E(w)$ should be large if $T_F(w)$ is supported on areas of M which undergo large distortion via T . Therefore, the problem of map analysis is turned into optimization of $E(w)$. Moreover, instead of optimizing $E(w)$ over all w in $L^2(N)$, a multi-scale approach is taken by adding a constraint such that w must lie in a subspace spanned by the first k eigenfunctions of $-\Delta_N$, which we denote by $S(k)$.

$$S(k) = \text{span}\{\varphi_1, \dots, \varphi_k\}. \quad (7)$$

(a,b)-closeness We now introduce our model for characterizing perturbations on the input shapes.

Definition 3.1 A Riemannian manifold $(N, \tilde{g}_N, \tilde{v}_N)$ is a -close to another one (N, g_N, v_N) if the following holds: For any $x \in N$ and any tangent vector η in $T_x N$, the tangent plane at x : $a^{-1} \leq \frac{\langle \eta, \eta \rangle_{\tilde{g}}}{\langle \eta, \eta \rangle_g} \leq a$ holds for some constant $a \geq 1$.

Definition 3.2 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, g_N, ν_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 isometric. Furthermore, the following proposition provides a quantitative relation between the perturbed and original manifolds.

Proposition 3.2 If $(N, \tilde{g}_N, \tilde{\mu}_N)$ is (a, b) -close to (N, g_N, ν_N) , then for any smooth function w on N ,

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

and

$$(ab)^{-1} d\tilde{\mu}_N \leq d\nu_M \leq abd\tilde{\mu}_N.$$

The detailed proof of this proposition and all of the other results mentioned below are provided in the supplementary material to improve readability. At the same time, we provide the outlines of the proofs of all the main theorems in the appendix.

Remark 3.2 Note that the gradient operator on a Riemannian manifold is defined directly by the metric. Thus the first inequality in this proposition is not simply a corollary of the condition of (a, b) -closeness.

Bounded-distortion Condition. Throughout our analysis in the following sections, we assume that 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 3.1 (Bounded-distortion) Let T_F be the functional map induced by $T : M \rightarrow N$, the distortions induced by T_F (or equivalently by T) are bounded:

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

$$\text{For any } w \in H_0^1(N), \int_M \langle \nabla T_F(w), \nabla T_F(w) \rangle_{g_M} d\nu_M \leq D_T \int_N \langle \nabla w, \nabla w \rangle_{g_N} d\nu_N$$

where B_T and D_T are finite positive constants.

In particular, the following proposition suggests that this condition is satisfied in a fairly general case.

Proposition 3.3 If M, N are compact and T_F is induced by a point-wise T which is a diffeomorphism, then Condition 3.1 is satisfied.

4. Stability of the Shape Difference Operators

In this section, we first consider the stability of the shape difference operators with respect to perturbations on the metrics and the measures. For the sake of simplicity, from now on we denote by N the original Riemannian manifold $(N, g_N, d\nu_N)$ and by $\tilde{N}$ the perturbed one $(N, \tilde{g}_N, d\tilde{\mu}_N)$, unless stated otherwise.

We have defined the area-based shape difference operator V with respect to M, N and T in Eq. 4. Similarly, the perturbed pair of shapes $\tilde{M}, \tilde{N}$ together with T give rise to another shape difference operator $\tilde{V}$ acting on $L^2(\tilde{N})$, which satisfies

$$\int_N f \tilde{V}(g) d\tilde{\mu}_N = \int_M T_F(f) T_F(g) d\tilde{\mu}_M, \forall f, g \in L^2(\tilde{N}) \quad (8)$$

The stability of the area-based shape difference operator with respect to perturbations on the metrics and measures is stated in the following theorem:

Theorem 4.1 Let M, N be two smooth shapes, 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 $L^2(N) = L^2(\tilde{N})$. Moreover, if M, N and T satisfy condition 3.1, we have the following convergence guarantee

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

As mentioned above, the outline of the proof of this theorem and of the others hereinafter are provided in the appendix.

Similar stability guarantee holds for the conformal shape difference operators as well. We start with defining the conformal shape difference operator, $\tilde{R}$, for the perturbed input shapes.

$$\int_N \langle \nabla f, \nabla \tilde{R}(g) \rangle_{\tilde{g}_N} d\tilde{\mu}_N = \int_M \langle \nabla T_F(f), \nabla T_F(g) \rangle_{\tilde{g}_M} d\tilde{\mu}_M \quad (9)$$

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, which in turn means that it converges to a constant function.

Theorem 4.2 Let M, N and $\tilde{M}, \tilde{N}$ be smooth shapes under the same assumptions of Theorem 4.1, then $H_0^1(N) = H_0^1(\tilde{N})$. Moreover, we have

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

Remark 4.1 Our proofs for theorems 4.1 and 4.2 do not require the shapes involved to be compact or boundaryless. The stability properties proven in this section are valid for any pair of smooth shapes and maps satisfying Condition 3.1.

In practice, we also observe the robustness of the shape difference operators with respect to shapes with missing part(s). For instance, in Figure 5, we compare shapes that are similar to the ones in Figure 2, but with missing parts on shape M . First on the left, M_1 is obtained by cutting the head of the horse, and the highlighted function at scale $k = 50$ is plotted on N next to it, which is consistent with the one generated with full shapes. Furthermore, we removed the area highlighted on the left (the hip), and compared




Figure 5: Top row: comparing full shapes M, N without boundary. We plot the highlighted functions at scales $k = 50, 110$ respectively. Middle row: removing a large part of M does not affect the highlighted function generated from the area-based shape difference operator at the same scales. Bottom row: if we further cut the most distorted part of M_1 , the highlighted area at scale $k = 50$ shifts to the one captured at higher scale ($k = 110$) in the full shape case (See Figure 2 for comparison).

M_2 and N . Interestingly, the highlighted function at $k = 50$ detects the area that is captured in the full shape case but at higher scale (see the bottom row of Figure 2 for comparison).

5. 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 [?], where they are employed in a multi-scale way.

We start by pointing out the connection between the multi-scale framework and the shape difference operators. Recall that Eq. 6 defines a functional measuring the distortion induced by T for a given function $w \in L^2(N)$. Given a pair of manifolds M, N and a map $T : M \rightarrow N$, let V be the area-based shape difference operator formulated in Eq. 4. It follows directly from the definitions that:

$$E(w) = \frac{\int_M T_F(w)^2 dV_M}{\int_N w^2 dV_N} = \frac{\int_N wV(w) dV_N}{\int_N w^2 dV_N}.$$

Since V is a positive-definite self-adjoint operator acting on $L^2(N)$, the maximum of $E(w)$ within $L^2(N)$ is simply the L^2 -norm of V , or equivalently the largest eigenvalue of V . The framework of [?] computes the *constrained norm* of V with respect to a special collection of subdomains of $L^2(N)$: $\{S(k)\}_{k \in \mathbb{N}^+}$. In general, given a subdomain Ω of $L^2(N)$, 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 highlighted function), w^* , is a function in Ω that is the most distorted by V . An illustration of a typical output of this framework has been given in Figure 2.




Figure 6: Highlighted functions with respect to conformal-based shape difference operator depicted on shape N at scales $l = 10, 11$ and 12 . The 10th to the 13th eigenvalues are $0.5497, 0.6130, 0.6130, 0.6802$, notice the small eigen gap between λ_{11} and λ_{12} , which causes the instability in the highlighted functions.

5.1. A New Subdomain Construction

A good selection of Ω is beneficial for abstracting information from the shape difference operators.

Despite several advantages of choosing $S(k)$ demonstrated in [?], the subdomain construction suffers some issues that are rooted in its discrete nature.

First, since k must be integer, the minimal perturbation on scale is 1. In practice, we observe that the output can change a lot when k is increased by 1, i.e., the original multi-scale framework is *not* stable with respect to the changes in scale.

Second, it can lead to confusing results when k is not selected appropriately. If there is a degenerate eigenvalue, say, $\lambda_l = \lambda_{l+1} < \lambda_{l+2}$, then using the subdomain $S(l)$ can be problematic. That is because the eigenspace formed by the eigenfunctions with respect to the degenerate eigenvalue can be of more than one dimension. Truncating in this subspace introduces randomness in basis construction, therefore the space spanned by the first l eigenfunctions is not even well-defined. For example, instability in the more subtle case of analyzing conformal differences is illustrated in Figure 6.

To overcome these issues, we construct a new collection of multi-scale subdomains which evolves continuously. It follows from Proposition 3.1 that for any $w \in S(k)$, $-\int_N \langle \nabla w, \nabla w \rangle_{g_N} dV_N \leq \lambda_k \int_N w^2 dV_N$

It is then natural to consider the following multi-scale subdomains controlled by a continuous scalar-valued parameter C :

$$A(C) = \{w : \int_N \langle \nabla w, \nabla w \rangle_{g_N} dV_N \leq C \int_N w^2 dV_N\} \quad (10)$$

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*.

In particular, the following proposition indicates the relationship between the original and the new subdomain constructions.

Proposition 5.1 *If $C \geq \lambda_k$, then $S(k)$ is a proper subset of $A(C)$.*

5.2. Stability with Respect to the Changes in Scale

We first verify the stability with respect to the change in scale, which only involves the original input shapes M and N . As demonstrated in Figures 1 and 2, the results show consistency of the areas on N indicated by the highlighted functions across a range of scales.

It is then tempting to validate the stability of the maximizer of the energy. However, it is not always the case. For example, imagine that we deform the bottom of shape M in Figure 1 so that the deformations from M to N are symmetrical. In this case, at every scale, the maximum of $E(w)$ is realized by two highlighted functions w_t, w_b which highlight respectively the top and the bottom of shape N , therefore we will no longer observe consistency in the highlighted functions.

We then turn to study the stability of the *maxima* of the energy $E(w)$ with respect to the change in scale. Our stability analysis is performed on the new multi-scale framework. For a subspace $A(C)$, we define:

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

where $E(w)$ is the functional defined in Eq. 6.

Let C go through interval $[0, +\infty)$, and consider the curve $(C, \|V\|_C)$. The following theorem suggests its continuity.

Theorem 5.1 *Let us be given two connected compact smooth Riemannian manifolds M and N , and a map T between them. If M, N, T satisfy Condition 3.1, then for any positive constant $C > 0, C' = C + \varepsilon > 0$, we have:*

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

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 5.1 only makes sense when ε is close to zero. At the same time, the inequality suggests that for a perturbation of fixed magnitude $|\varepsilon|$, the larger C is, the more stable $\|V\|_C$ is.

5.3. Stability with Respect to Perturbed Inputs

On the other hand, we can also fix the scale C and add perturbations on the shapes M and N in the same way as we did in Section 4. I.e., 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. Let V and $\tilde{V}$ be the corresponding area-based shape difference operators defined in Eq. 4 and Eq. 8.

In order to define the constrained norm for $\tilde{V}$, we first construct the corresponding functional $\tilde{E}(w)$:

$$\tilde{E}(w) = \frac{\int_N w \tilde{V}(w) d\tilde{\mu}_N}{\int_N w^2 d\tilde{\mu}_N} = \frac{\int_M T_F(w)^2 d\tilde{\mu}_M}{\int_N w^2 d\tilde{\mu}_N} \quad (11)$$

The construction of the corresponding subdomain $\tilde{A}(C)$ follows the same spirit of Eq. 10:

$$\tilde{A}(C) = \{w : \int_N \langle \nabla w, \nabla w \rangle_{\tilde{g}_N} d\tilde{\mu}_N \leq C \int_N w^2 d\tilde{\mu}_N\} \quad (12)$$

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

Theorem 5.2 *Let M, N be two connected compact smooth shapes 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 3.1, then at any fixed scale C , we have:*

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

5.4. Approximating $\|V\|_C$

By investigating the behavior of the operators within the continuously evolving subdomains $A(C)$, we have a more stable and potentially richer understanding of V than that arising 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 the 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 [?], computing $\|V\|_k$ in the case where M and N are finite discrete meshed shapes is straightforward.

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

Proposition 5.2 *Let M, N and T be a pair of manifolds and a map, which satisfy Condition 3.1. IF λ_k, λ_{k+1} are two consecutive eigenvalues of the LB operator on N , then the constrained norms with respect to $A(\lambda_k)$ and $S(k)$ satisfy the following inequality:*

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

As a direct corollary, the smaller $\lambda_k / \lambda_{k+1}$ is, the better $\|V\|_k$ approximates to $\|V\|_{\lambda_k}$. It is also worth noting that this proposition indicates a general criterion of choosing a discrete scale: it is 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 clear, then the maximizer realizing $\|V\|_k$ is a nice candidate of initial guess for iterative algorithms for 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 hundred of thousands. The following proposition suggests that there is a trade-off between accuracy and complexity in this optimization.

Proposition 5.3 *For a fixed parameter C , let $\varepsilon > 0$ and λ_{l+1} be the smallest eigenvalue of the LB operator on N such that $C \leq \varepsilon \lambda_{l+1}$. Now denote $\|V\|_{C,l}$ by the optimum of the following problem:*

$$\max E(w) \text{ s.t. } w \in A(C) \cap S(l). \quad (13)$$

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

5.5. Analysis for the Conformal Shape Difference Operator

In essence, with the energy functional $E(w)$, the framework of [?] casts the problem of extracting information from the area-based shape difference operator as a series of constrained optimization problems.

Note that the framework of [?] introduces two shape difference operators which encode different types of distortion between shapes. A natural extension of the multi-scale framework of [?] is to construct parallel functionals and subdomains with respect to the conformal shape difference operators, R .

We first define a functional, F , acting on $H_{0,v}^1(N)$ as the following:

$$F(w) = \frac{\int_N \langle \nabla w, \nabla R(w) \rangle_{g_N} dV_N}{\int_N \langle \nabla w, \nabla w \rangle_{g_N} dV_N} = \frac{\int_M \langle \nabla T_F(w), \nabla T_F(w) \rangle_{g_M} dV_M}{\int_N \langle \nabla w, \nabla w \rangle_{g_N} dV_N} \quad (14)$$

where w is not a constant function so that $\int_N \langle \nabla w, \nabla w \rangle_{g_N} dV_N \neq 0$.

On the other hand, modifying the multi-scale subdomain construction is necessary to suit the new functional. If we use $A(C)$ in the conformal case, then $F(w)$ is not well-defined if w is the constant function. In fact, following the same idea proving Proposition 5.1, for any $w \in A(C)$, we can find $\bar{w} \in A(C - \epsilon)$ such that $w - \bar{w}$ is a constant function, which in turn means that $F(w) = F(\bar{w})$. To obtain multi-scale results, we construct another subdomain, $A^{conf}(C)$, for the conformal case.

$$A^{conf}(C) = A(C) \cap \{w : \int_N w dV_N = 0\} \quad (15)$$

and we define $\|R\|_C = \max F(w)$ s.t. $w \in A^{conf}(C)$. It is worth noting that if $C < \lambda_2$, the second eigenvalue of $-\Delta_N$, then $A^{conf}(C)$ is empty. Thus C must be at least λ_2 so that $\|R\|_C$ is well-defined. In practice, it is easier computationally to maximize $F(w)$ in the subdomains spanned by finite number of eigenfunctions. Following the same arguments above, we modify $S(k)$ to obtain $S^{conf}(k) = \text{span}\{\phi_2, \dots, \phi_k\}$, where k must be at least 2.

After the above formulations, we validate the stability of R with respect to the change in scale.

Theorem 5.3 *Let M, N be two connected compact smooth Riemannian manifolds, and T be a map between them. Let λ_2 be the first non-zero eigenvalue of $-\Delta_N$. If M, N, T satisfy condition 3.1, then for $C > \lambda_2, C' = C + \epsilon > \lambda_2$ we have:*

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

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 Eq. 9. The associated functional, $\tilde{F}(w)$, is defined as follows:

$$\tilde{F}(w) = \frac{\int_M \langle \nabla T_F(w), \nabla T_F(w) \rangle_{\tilde{g}_M} d\tilde{\mu}_M}{\int_N \langle \nabla w, \nabla w \rangle_{\tilde{g}_N} d\tilde{\mu}_N} \quad (16)$$

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 5.2 does not work in this case. That is because the interleaved structure is not guaranteed between the new subdomains $A^{conf}(\cdot)$ and $\tilde{A}^{conf}(\cdot)$: a function satisfying $\int_N w dV_N = 0$ does not necessarily fulfill $\int_N w d\tilde{\mu}_N = 0$ simultaneously.

6. Experimental Results

In this section, we demonstrate experimental results that are related to our theoretical analyses. Notice that implementing the frameworks of [?, ?] on a pair of meshed shapes M, N requires essentially an approximation of the LBO on each of the shape. That is usually done by computing two matrices A_M, W_M , the former is a diagonal matrix whose (i, i) -th entry is the area element (see [?]) around the i -th vertex in M , and the latter is the stiffness matrix computed with the cotangent scheme (see [?]). The LBO is then approximated by $A_M^{-1} W_M$.

6.1. 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 n_M and n_N vertices respectively. The functional map T_F induced by T is represented by a matrix $P \in \mathbb{R}^{n_M \times n_N}$ in the discrete setting. Let $\Phi_k \in \mathbb{R}^{n_N \times k}$ be a matrix whose columns are the first k eigenvectors solved by $W_N f = \lambda A_N f$.

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 5.2, we take the optimizer that realizes $\|V\|_k$ 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 iterations continue until there is no more significant improvement or β is sufficiently small. Note that this method, while being easy to implement can potentially be improved with more advanced constrained optimization techniques. We leave the exploration of alternatives as an interesting direction for future work.

As mentioned in section 5.1, both subdomains $S(k)$ and $A(C)$ are designed to control the Dirichlet 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. To demonstrate this, we consider the pair of deformed spheres shown in Figure 1 and the map therein and compute the local maxima and maximizers of $\|V\|_C$ with different scales C ranging from 0 to 2.

First, we applied the algorithm described above to compute the maximum of $E(w)$ constrained in $A(C)$, and then computed the maximum constrained in $S(k), k = 1, 2, \dots, 36$, where $\lambda_{36} < 2 < \lambda_{37}$. Figure 7 shows the plots representing the maximal energy with

respect to the old and the new subdomain construction, note that since shape N (undeformed sphere) admits several sets of repeated eigenvalues, the energy plot is irregular, while the curve $(C, \|V\|_C)$ is continuous and always above the points $(\lambda_k, \|V\|_k)$. On the other hand, Figure 8 shows the portion of each of the four local maximizers expressed by the first k eigenfunctions ($k = 1 \sim 300$). It can be seen that the four local highlighted functions are well-expressed by the first 300 eigenfunctions (with $\lambda_{300} = 15.20$). 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 percent of the norm of the one at $C = 2$.




Figure 7: Plots of energy at changing scales with respect to different subdomain constructions.




Figure 8: The X-axis indicates the index of eigenvalues/eigenfunctions, and the Y-axis represents the ratio $\sum_{i=1}^k a_i^2 / \sum_{i \geq 1} a_i^2$, where the maximizer is decomposed as $\sum_{i \geq 1} a_i \phi_i$. Three eigenvalues, $\lambda_{50}, \lambda_{150}$ and λ_{250} , are labeled accordingly along the X-axis.

Note that since the barrier function method is a gradient-based technique, the results depend on the initial guess and can get trapped in local maxima. Such issues will be amplified when dealing with more sophisticated input shapes, where the global maxima are not as clear as in the simple shapes demonstrated above.

In fact, the new subdomain construction enjoys better theoretical properties, while loses computational simplicity as a trade-off. However, as suggested in Proposition 5.2, the original optimization problem is closely related to the new one. From this point of view, we will use the original framework which optimizes within $S(k)$, and reduces to solving a generalized eigenvalue problem in analyzing more complicated shapes in the following experiments, which also illustrates remarkable stability.

6.2. Robustness of the Area-based Shape Difference Operator

We have observed robustness of the frameworks of [?] with respect to perturbations on the input meshes in Figures 1 and 3. In Figure 9, we generated N_s by simplifying mesh N : the former consists of 6250 vertices while the latter consists of 12499 vertices. Locally, the mesh connectivity is significantly changed, however, as shown in the top two rows of Figure 9, the highlighted functions at scales $k = 20, 50, 200$ are all consistent.

Besides changing the mesh structure, we now perturb the input meshes by introducing noise in the vertex positions. In the bottom row of Figure 9, M' and N'_s are obtained by perturbing the vertices on M and N_s along the respective normal directions. In particular, we perturb a point p to $p' = p + 1/2\bar{d}x_p\mathbf{n}_p$, where $\bar{d}$ is the mean edge length of the mesh, x_p is i.i.d $\mathcal{N}(0, 1)$ and $\mathbf{n}_p$ is the normal vector at point p . Under the additional point perturbations, the resulting highlighted functions are still consistent.




Figure 9: Top row: we compared two shapes in one-to-one correspondence, M and N , and depicted the highlighted functions at scales $k = 20, 50, 200$ on N . Middle row: we preserved M while simplifying N so that the number of vertices was reduced by half. Bottom row: the same result after perturbing the vertices of both M and N_s with Gaussian noise.

6.3. Pipeline for Point Cloud Data

Inspired by the stability of the shape difference operators and the highlighted distortion functions in theory and in the case of triangle meshes, below we aim to apply this framework to point cloud data. Approximating the LBO of a manifold with a certain Laplacian of a graph built on top of points sampled from the manifold is a problem that has been well-studied. In particular, our pipeline

takes advantage of the results in [?], where the authors show that given a point cloud X sampled from a Riemannian manifold N , the un-normalized graph Laplacian of a certain weighted graph (which we estimate with W_X below) approximates to $-\rho^{-1}\Delta_N$, where ρ is the sampling density of X . On the other hand, we use the framework in [?] to estimate the sampling density. The matrix A_X below, serves as an estimator of ρ^{-1} , therefore we use $A_X^{-1}W_X$ as an approximation of the LBO.

Our pipeline for implementing the frameworks above on shapes represented as point clouds is described in Algorithm 1, where we compute for an input point cloud X two matrices A_X, W_X and then use them as A_M, W_M in the same way as in the mesh setting. In all the experiments involving point cloud inputs, we always use $K = 40$, i.e., we compute 40-NN graphs on all point clouds.

Algorithm 1: Pipeline for Point Cloud Inputs

input : A point cloud $X = \{x_1, x_2, \dots, x_n\}$ and an integer K
output: Two matrices A_X and W_X

$A_X, K_X, W_X \leftarrow$ zero matrices of dimension $n \times n$

for $x_i \in X$ **do**

$N(x_i, K) \leftarrow$ the K nearest neighborhoods of x_i in $X \setminus x_i$

$A_X(i, i) \leftarrow (\sum_{x_j \in N(x_i, K)} \|x_i - x_j\|^2)^{3/2}$

$t \leftarrow \sum_{i=1}^n \sum_{x_j \in N(x_i, K)} \|x_i - x_j\| / Kn$

if $x_j \in N(x_i, K)$ **or** $x_i \in N(x_j, K)$ **then**

$K_X(i, j) \leftarrow \exp(-\|x_i - x_j\|^2 / 2t^2)$

for $i = [1..n]$ **do**

$d_i \leftarrow$ mean of non-zero elements in the i -th row of K_X

for $K_X(i, j) \neq 0$ **do**

$\tilde{K}_X(i, j) \leftarrow K_X(i, j) / d_i d_j$

if $i = j$ **then**

$W_X(i, j) \leftarrow \sum_j \tilde{K}_X(i, j)$

else

$W_X \leftarrow -\tilde{K}_X(i, j).$

Using these constructions, we observe that the robustness is evidenced in the results from our PCD setting as well. In Figure 10, we tested with point clouds sampled from horses (8431 points) and cats (7207 points). Point clouds X, Y come from two meshes in point-to-point correspondence, the highlighted functions on Y are shown to be consistent with the ones computed in the mesh setting on N (see the right-most column). We then resampled point clouds on shape Y while keeping the total number of points unchanged, and computed the highlighted function with respect to X and the resampled point cloud. At the end, a further step was taken to add noisy points on the resampled point clouds. Given a point cloud, we first randomly select n_p points from it. Then for a selected point p , 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 $\bar{d}$, which is the median of the distances from a point to its nearest neighbor. Repeating the

displacements r times for each selected point, we enlarge the original point cloud with $n_p r$ more points. The stability of (area-based) shape difference operator is again verified by the consistency of areas highlighted by the functions.

On the other hand, we have seen in Figure 4 that although the eigenfunctions of the graph Laplacian on the point cloud are distinct from those of the LBO on the mesh, the eigenfunctions of the shape difference operators are comparable. We further explore this by considering the pair of cats taken from the bottom row of Figure 10. In particular, we computed the eigenfunctions of the LBO on mesh N and of the graph Laplacian on Y and the perturbed resampled point cloud. The highlighted functions and part of the eigenfunctions with respect to the three representations are depicted in Figure 11. Again, changing the representation of the shape causes significant perturbations on the eigenfunctions, however, as illustrated in Figure 10, the areas indicated by the respective highlighted functions remain similar to each other.

6.4. Analyzing Shape Collections

The experiment above shows the stability of the shape difference operators for analyzing maps between a single pair of shapes in a multi-scale way. As we prove in Section 4, the shape difference operators are stable with respect to perturbations on the input shapes. To demonstrate this, we repeat one of the experiments in [?] (see Figure 3 on page 7 therein), but in the point cloud setting. We compute the shape difference operators and then vectorize them so that we can apply PCA. The PCA embeddings in $\mathbb{R}^2$ are depicted in the right two columns of Figure 12.

The top row of Figure 12 depicts the embeddings for the deformed spheres. Both layouts uncover the grid structure of the original shape collection. The results in [?] suggest that in both area and conformal cases, the variances of the first two principal components are evenly close to 50 percents. In our results: (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 principal component; (2) Conformal case: balance preserved, the shapes of the first and the second rows are not well differentiated, suggesting that the operators are less sensitive to small changes.

The bottom row shows the layouts for the galloping horse sequence, which consists of two cycles of continuous movement of the horse. Our results successfully capture the circular structure of the sequence, as depicted in the layout. The plot also reveals the fact that there are more conformal distortions than area distortions in this data, as the range of layout in the third column is larger than that in the second one.

Overall, we conclude from these experiments that although the results from the PCD setting are not always as accurate as those from the mesh setting, our results capture most of the basic and significant information hidden in the data. Considering that we start from a much coarser understanding of the input shapes, these results are non-trivial and quite remarkable, especially given the well-known instability in the eigenfunctions of the LBO.




Figure 10: Robustness of results from the PCD setting: X and Y are the original point clouds extracted from meshes. In the third column, we resampled point clouds on the polygonal mesh, and then we added noisy point displacements to that in the forth column, which are both generated with $n_p = 400, r = 20$, i.e. there are 8000 more noisy points (see the text for details). All the area-based highlighted functions are computed at scale $k = 50$. On the rightmost column are the highlighted functions from the mesh setting.




Figure 11: Left: the highlighted functions on the mesh (top), the noiseless point cloud (middle) and the noisy point cloud (bottom). The 11th to the 15th eigenfunctions of the LBO on mesh (top), the graph Laplacian constructed on top of the original point cloud (middle) and of the resampled noisy point cloud (bottom).

7. Conclusion and Future Work

In this paper we present two types of stability guarantees for the shape difference operators. We also introduce a new multi-scale scheme for extracting information from the shape difference operators, which is provably stable in contrast to the original one proposed in [?]. From a practical point of view, we present a pipeline for constructing shape difference operators on point clouds, which extends the range of applications of the related frameworks.

Several follow-up problems arise along our investigation. We especially remark the optimization problem attached to our new multi-scale scheme. As the new scheme provides more stable results in theory, it is appealing to design an efficient implementation. It is as well appealing to consider more rigorous analysis of our pipeline for point cloud data.

Acknowledgement The authors would like to acknowledge the ANR TopData (ANR project TopData ANR-13-BS01-0008), the ERC Gudhi project (Geometric Understanding in Higher Dimensions), Marie-Curie CIG-334283, a CNRS chaire d'excellence,

chaire Jean Marjoullet from Ecole Polytechnique, FUI project TAND-DEM 2, a grant from the Direction Générale de l'Armement (DGA) and a Google Focused Research Award.

References

Appendix A:

Here we outline the proofs for the main theorems in this paper. We refer interested readers to the corresponding supplemental material for detailed proofs.

Theorem 4.1 and Theorem 4.2: First we can prove $L^2(N) = L^2(\tilde{N})$ with proposition 3.2. To verify the convergence, we first prove that $\int_N fVfdv_N - \int_N f\tilde{V}fdv_N$ vanishes as a_M, a_N, b_M, b_N converge to 1 simultaneously. Then due to the fact that V is self-adjoint, we prove $\int_N gVfdv_N - \int_N g\tilde{V}fdv_N$ vanishes under the same condition. Lastly, we let $g = Vf - \tilde{V}f$ and finish the proof.

The idea of proving in the case of conformal shape difference operator is analogous to the area-based one, but proving




Figure 12: PCA plot of the two shape difference operators.

$\int_N \langle \nabla f, \nabla Rf \rangle_{g_N} d\nu_N - \int_N \langle \nabla f, \nabla \tilde{R}f \rangle_{g_N} d\nu_N \rightarrow 0$ is slightly more complicated as both the measure and the inner-product are perturbed.

Theorem 5.1 and Theorem 5.3: Given parameters $C' > C$, our strategy is to find for any function $w \in A(C')$ a function $\tilde{w} \in A(C)$, such that $|E(w) - E(\tilde{w})|$ is upper-bounded by some variable with respect to $C' - C$, which vanishes as $C' \rightarrow C$. Regarding the conformal case, we apply the same idea, i.e., find for any function $w \in A^{conf}(C')$ a function $\tilde{w} \in A^{conf}(C)$, such that $|F(w) - F(\tilde{w})|$ is uniformly bounded by a variable depending on $C' - C$.

Theorem 5.2: The key observations to proving this theorem are: first, $A(C)$ and $\tilde{A}(C)$ are interleaving, i.e., for any $C > 0$ we can find a C' such that $A(C) \subset A(C')$ and vice versa; second, given a $w \in L^2(N) = L^2(\tilde{N})$, the ratio of $E(w)$ to $\tilde{E}(w)$ is two-side bounded with respect to a_M, b_M, a_N and b_N . The theorem is obvious then after verifying those observations.