

Analysis and Visualization of Maps Between Shapes

Maks Ovsjanikov¹ Mirela Ben-Chen² Frederic Chazal³ Leonidas Guibas⁴

¹LIX, Ecole Polytechnique ²Technion–Israel Institute of Technology ³Geometrica, INRIA ⁴Stanford University

Abstract

In this paper we propose a method for analyzing and visualizing individual maps between shapes, or collections of such maps. Our method is based on isolating and highlighting areas where the maps induce significant distortion of a given measure in a multi-scale way. Unlike the majority of prior work which focuses on discovering maps in the context of shape matching, our main focus is on evaluating, analyzing and visualizing a given map, and the distortion(s) it introduces, in an efficient and intuitive way. We are motivated primarily by the fact that most existing metrics for map evaluation are quadratic and expensive to compute in practice, and that current map visualization techniques are suitable primarily for global map understanding, and typically do not highlight areas where the map fails to meet certain quality criteria in a multi-scale way. We propose to address these challenges in a unified way by considering the functional representation of a map, and performing spectral analysis on this representation. In particular, we propose a simple multi-scale method for map evaluation and visualization, which provides detailed multi-scale information about the distortion induced by a map, which can be used alongside existing global visualization techniques.

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

1. Introduction

Maps or correspondences are ubiquitous in geometry processing. The presence of a map between a pair of spaces enables a wide variety of applications in domains as diverse as shape parametrization, animation, reconstruction and exploration among many others. Perhaps the most crucial property of a correspondence is that it allows the transfer of information such as labels or parametrization from one space to another. It is therefore not surprising that computing and analyzing maps lies at the essence of a wide range of geometry processing applications.

At the same time *manipulating and understanding* maps between shapes, as objects in their own right, is often challenging. This is not only because the number of possible correspondences between a pair of discrete shapes is exponential but also since, even when given a correspondence or a map between a pair of shapes, it is not easy to represent it in a simple and intuitive form. Indeed, even *visualizing* a map between a pair of shapes in a way that would highlight its major properties is far from trivial.

For example, consider the two cat shapes and a map between them shown in Figure 2. This map has a low distortion in terms of geodesic distance preservation, therefore we




Figure 1: Given a point-to-point map between a pair of shapes (left), our method allows to identify and visualize major problematic areas in a multi-scale fashion (middle and right), where the corresponding points share the same color.

would expect it to be a “good” map. When visualizing this map, a common approach would be to either draw lines between corresponding points (a), or use the coordinate (x, y, z) functions on the source shape for representing corresponding points (b). Although both of these methods provide a good global understanding of the map, indicating that it roughly puts in correspondence the appropriate parts, neither of these methods easily identify problematic regions. An alternative approach would be to display some local distortion measure such as area distortion (c), or a smoothed area distortion (d). However, this also does not help to identify a major problem




Figure 2: (a) map rendered as lines joining corresponding points, (b) XYZ coordinates rendered as RGB on shape 2, and transferred to shape 1 through a given map. Both methods make it difficult to identify problematic areas. Visualizing local area distortion (c), or its smoothed version and showing the map at points attaining maximum values (d) is noisy and uninformative. (e) Our method computes functions which identify problematic regions, here the neck region is mapped to the nose.

in this map, visualized by our method in (e): the cat neck is mapped to the nose. A similar problem can be seen for a “good” map in Figure 1, where we identify that areas on the cat’s chest are mapped to the legs.

In general, constructing a *multi-scale* method for map evaluation, visualization and exploration is difficult when using the classic representation of a map as a correspondence between points. In this case one option would be to use a quadratic quality criterion (e.g. geodesic distance preservation between pairs of points). However, such criteria are usually *global*, in the sense that local shape changes, can induce large distortions of distances for pairs of points everywhere on the shape. Furthermore, quadratic quality criteria of a shape S are expensive to evaluate and often difficult to visualize, since they are defined on the product $S \times S$.

Another option is to use a *local* criterion, such as the affine distortion introduced locally by the map at every point on the shape. Unfortunately, such distortion measures can be noisy (see Figure 2(c)), and it is hard to extract from them a clear understanding of the areas on the shape where the map fails to meet certain quality criteria, possibly ordered by importance. Thus, the main goal of our work is to construct a *multi-scale* representation for a correspondence and to apply it for efficient map evaluation, visualization and exploration.

1.1. Overview

We assume that we are given two discrete shapes M and N and a map $T : M \rightarrow N$. We also assume that both shapes are endowed with measures, μ_M and μ_N , that we expect to be preserved for an optimal map T . Namely, for any measurable set $B \subseteq N$, $\mu_N(B) = \mu_M(T^{-1}(B))$, where $T^{-1}(B) = \{x, \text{ s. t. } T(x) \in B\}$. For example, if we expect the map to be area preserving, then the measures would be the volume measures on the surfaces.

The output of our visualization procedure is a set of *multi-scale distortion indicators*, which are a collection of smooth functions $w_i : N \rightarrow \mathbb{R}$ that represent areas where T distorts the measure the most. Color-coding the functions w_i on N and $w_i \circ T$ on M provides visualizations of the regions where T is least measure preserving. The visualizations are ordered

by importance, so that w_1 visualizes the most problematic region, w_2 the next most problematic and so on (see Figure 1 middle and right). In addition, we provide scalars s_i which quantify the distortion at the various scales. Note, that we do not place any assumptions on the geometry or topology of w_i , and instead provide a scale parameter k , which controls how *smooth* the functions w_i are. Large values of k allow for highly varying functions which reveal highly localized distortion, whereas medium and small k force the indicator functions to be more smooth, which is especially useful for visualization since it allows to identify global problematic regions (see Figure 4). As we focus on visualizing the problematic regions, our visualization lacks a global view of the map. Therefore, our method is complementary to standard visualization techniques.

We base our approach on the framework of *functional maps* introduced in [OBCS*12]. In that framework, instead of treating a correspondence as a relationship between points on two shapes, it is represented as a mapping between function spaces over these shapes. The key property of functional maps used in [OBCS*12] in the context of shape matching is the fact that unlike point-to-point maps, functional maps are always linear. In this paper, we will demonstrate that the spectral decomposition of a *given* functional map provides a unifying tool for analysis and visualization of maps between shapes. We will also show that our map summary provides an optimal approximation of a given map under certain conditions.

2. Related Work

Map evaluation and analysis have received a significant amount of attention in geometry processing, especially in the context of *non-rigid shape matching* where the goal is to recover the *best* map according to some quality criterion (see e.g. [BBK06, HAWG08, LF09, KCATCO*10, OMMG10, KLF11, TBW*11, SY11] among many others). Perhaps the most common quality criterion for a map between a pair of shapes is the preservation of pairwise quantities such as geodesic distances [BBK06, HAWG08, TBW*11, SY11] or spectral quantities such as the heat kernel [MHK*08, SH10, OMMG10]. Generally, such measures of quality are both ex-

pensive to compute and non-trivial to visualize. In particular, if the metric on a shape is distorted in a small location on the surface, the distances between many pairs of points can be affected. This makes localizing the error and thus providing an intuitive understanding of a map challenging.

Another way of evaluating the distortion of a map, used mostly in shape deformation and parametrization applications (see e.g. [HG00, SMW06, LZX*08, BCWG09], among others), is to consider the local affine distortion introduced by the map at every point on the shape, e.g. angular or area distortions. On the one hand these errors are efficient to compute, and easily visualized. However, although they can be aggregated into global measures of map quality, such distortion measures are too noisy to be used directly for identifying problematic regions (see e.g. Figure 2e and 2d).

The related problem of *visualizing* maps in an intuitive and multi-scale fashion has received relatively little attention, despite its importance in practice. Most visualization methods are presented in an ad-hoc manner as part of larger works whose aim is to compute the map. Common methods include line based, color based (as in Figure 2a and 2b respectively), used e.g. in [KLF11, SY11, LF09], as well as visualization of sparse correspondences using colored spheres. In the field of optical flow, where the task is to visualize dense correspondences between images, there has been some very recent work [SLKW12] which is based on color map selection. In this context, the authors assume that the correspondence is given on a regular 2D grid, and therefore the suggested solution is not easily applicable to maps between general shapes. A good map visualization method would be valuable not only as an evaluation tool, but also as a design tool for researchers investigating new algorithms for computing shape maps, especially for identifying and visualizing problematic areas. Hence, map visualization is an important tool for map “debugging”.

3. Functional Maps Refresher

We will base our approach on the *functional map representation* of a correspondence, first suggested in [OBCS*12]. In essence, functional maps represent a correspondence between two surfaces as a correspondence between *functions* defined on them, which can be encoded using a small (typically order of 100×100) matrix. Due to their algebraic structure, functional maps are more appropriate for multi-scale analysis. Here, we briefly repeat the relevant properties of this representation for completeness.

A map $T : M \rightarrow N$, which takes points on M to points on N , induces a map T_F between the function spaces of N and M , which takes functions on N to functions on M . This map is given by the pull-back $T_F(g) = g \circ T$, where $g : N \rightarrow \mathbb{R}$.

It is easy to see that although T can be arbitrary and may not even be a bijection, T_F is *linear* across function spaces and as such can be represented as a (possibly infinite) matrix




Figure 3: (a) Two copies of the uniformly sampled sphere with the identity map modified so that some points map to the north pole. Note that this map has no distortion locally (at sample points). Moreover, any method that assumes imposing a geometry of the distorted regions will fail. (b) The problematic regions highlighted using our method, showing all points (left) that are mapped to the north pole (right).

C , such that if $g = \sum_i a_i \phi_i^N$ and $T_F(g) = \sum_j b_j \phi_j^M$, where ϕ_i^M and ϕ_j^N form a basis functions for the functional spaces on M and N respectively, then $b_j = \sum_i c_{ij} a_i$, for some fixed c_{ij} independent of a_i and b_j . Moreover, if both ϕ_i^M and ϕ_j^N are orthonormal, then: $c_{ij} = \langle T_F(\phi_i^N), \phi_j^M \rangle$. If M and N are smooth manifolds, then $c_{ij} = \int_{x \in M} \phi_i^N(T(x)) \phi_j^M(x) d\mu_M(x)$, where $d\mu_M$ is a given measure on M .

4. Map Evaluation and Visualization

Given a pair of shapes M and N and a map $T : M \rightarrow N$, our goal is to visualize T by choosing of a real-valued function $w : N \rightarrow \mathbb{R}$, and displaying w on N and $w \circ T$ on M . Note that this generic approach is often used in practice, and common choices for w include composite functions of XYZ coordinates of the vertices or indicator functions of balls around landmark points (e.g. [KLF11, SY11, LF09]). However, these choices are typically done in an ad-hoc way, and can fail to highlight problematic areas (see Figures 1 and 2). Therefore one of our goals is to devise a principled way to select “optimal” functions w for map evaluation and visualization.

In particular, we select w so as to highlight the areas where the map induces highest distortion. Namely, we consider the following measure of distortion for a function w :

$$d_{\mu_M, \mu_N}(w) = \frac{\int_{x \in M} w(T(x))^2 d\mu_M}{\int_{y \in N} w(y)^2 d\mu_N}. \quad (1)$$

Intuitively, $d_{\mu_M, \mu_N}(w)$ is close to 1 if w is an indicator function of areas where T is measure preserving, and is large if T maps areas with small measure on M to areas with large measure on N . Finally, we select the optimal w as follows:

$$w_k^* = \arg \max_{w: N \rightarrow \mathbb{R}} d_{\mu_M, \mu_N}(w), \text{ s.t. } w \in \text{span}(\Phi_{1..k}^N), \quad (2)$$

where $\Phi_{1..k}^N$ are the first k eigenfunctions of the Laplace-Beltrami (LB) operator of N , and k is a parameter which we




Figure 4: The region where the map distorts the area measure the most, for varying scales k , as given by w_k^* . Note how the regions become more and more localized as k grows, and the fact that the regions do not necessarily have disk topology.

interpret as a measure of “scale.” Below we will show that w_k^* is well-defined and can be computed through the singular value decomposition of the functional representation of a given map. Furthermore, we argue that the resulting function w_k^* is a) shape-aware, b) multi-scale, c) easily computable and d) can be used to approximate a given map T in an optimal way under certain conditions.

4.1. Shape Aware and Multi-scale Properties

Perhaps the two most important properties of our approach making it useful in practice is that our selection of optimal function w_k^* is shape-aware and multi-scale. By the former we mean that we do not place any assumptions on the geometry or the topology of the distorted region, and by latter we mean that the parameter k provides a natural way to go from very global (small k) to very local (large k) analysis, which is especially useful for map visualization.

Indeed, an alternative possibility to visualize distortion on a pair of discretized shapes would be to simply display the function $d(x) = \mu_M(x) / \mu_N(T(x))$, where $\mu_M(x)$ is the area measure at point x . Unfortunately, this naive approach does not work. First, the resulting function d can be noisy (see e.g. Figure 2c), and moreover its maxima (or even maxima of a smoothed d , Figure 2d) may not be informative. This is because local changes in T can have a catastrophic effect on d . If, for example, two nearby sample points $x_1, x_2 \in M$ are such that $\mu_M(x_1) = \mu_N(T(x_2)) = 1$ and $\mu_M(x_2) = \mu_N(T(x_1)) = \epsilon$, then $d(x_1) = \frac{\mu(x_1)}{\mu(T(x_1))} = \epsilon^{-1}$ which can be very large. However, x_1 may not correspond to a region of interest for T since by considering a slightly larger neighborhood, the distortion would be: $\frac{\mu(x_1) + \mu(x_2)}{\mu(T(x_1)) + \mu(T(x_2))} = 1$.

One may hope that by considering $d_r(x) = \mu_M(B(x, r)) / \mu_N(T(B(x, r)))$, i.e. the distortion induced by T on geodesic balls of fixed radius r , we will identify meaningful areas. To see that this may not be the case, consider the two copies of the sphere with a map between them shown in Figure 3a. The map equals to the identity everywhere except for a few points, which are mapped to the north pole. Considering $\frac{\mu_M(B(x, r))}{\mu_N(T(B(x, r)))}$ for various values of r , would place a strong assumption on the geometry and topology of the distorted area. Namely, that the errors in the map are *isotropic*, in the sense that they are equally likely to happen in all directions, depending only on the distance from the point y . In the case of Figure 3a, this will fail to

isolate the points where the map does not equal to identity. Our method, however, allows to visualize the correct areas, see Figure 3b. Anisotropy in distortion is also common in practice, as many times the symmetries of the shape have a large influence on how the errors are distributed. For example, for an elongated region such as the body of the pig in Figure 4, the errors are distributed anisotropically, due to slippage symmetries.

Thus, instead of fixing the geometry or topology of the distorted region, we assume simply that w must be *sufficiently smooth*, and by controlling the smoothness, we allow the user to control the scale at which distortion is computed and visualized. In particular, we force w to lie in the span of the first k eigenfunctions of the Laplace Beltrami operator. Since the eigenfunctions of the Laplace-Beltrami operator are ordered from being very smooth (as measured by the integral of the squared norm of the gradient) to having larger variation as k increases, the user-specified parameter k controls the smoothness of w . Thus, small values of k correspond to more “global” distortion whereas large values of k identify local changes of measure (see Figure 4).

For a discrete shape with n vertices where the measure is given as a function over the vertices, if we take $k = n$, we would get an indicator function of the point $y \in N$ which maximizes $\frac{\mu_M(T^{-1}(y))}{\mu_N(y)}$. If we take $k = 1$, we would get a constant function. Figure 4 shows a few examples of w_k^* , for a various choices of k , where the measure at each vertex is given by the area of the corresponding Voronoi region. Note that the region where the function values are high is meaningful in terms of distortion in the map. Furthermore, as k grows these regions become more localized as expected.

4.2. Computing w_k^* using Functional Maps

Given two discrete shapes M and N represented as triangle meshes and containing m and n points each, a map $T : M \rightarrow N$, and two measures μ_M and μ_N , how can we find w_k^* ? The most common way to specify a measure on discrete shapes is by assigning non-negative values to individual points. We call these functions $d\mu_M : M \rightarrow \mathbb{R}^+$ and $d\mu_N : N \rightarrow \mathbb{R}^+$ on M and N respectively. For simplicity, in this section we assume that the measures are the standard area measures, i.e. $d\mu_M(x)$ is the area of the Voronoi region of x on M . However, the results apply to any choice of measure (see Section 4.4). The following theorem ensures that computing the optimal w can




Figure 5: Visualizing the map between a cow and a giraffe from [KLF11]. We show the distortion indicators corresponding to the first three and last two singular vectors, as well as the corresponding singular values. Note how the map has problematic regions, which are easy to miss in the standard visualization (left) but are highlighted by our method.

be done simply via a singular value decomposition of the functional representation C of the map T .

Theorem 4.1 *Given a map $T : M \rightarrow N$ between two discrete shapes, the optimal function $w_k^* = \Phi_{1..k}^N \mathbf{w}$ where $\Phi_{1..k}^N$ contains the first k eigenfunctions of the LB operator on N and \mathbf{w} is the right singular vector corresponding to the largest singular value of the functional representation C of T .*

Proof First note that since d_{μ_M, μ_N} is invariant to a global scale in w , the optimization in Eq. (2) can be written as:

$$w_k^* = \arg \max_{w \in \text{span}(\Phi_{1..k}^N), \|w\|_{\mu_N}^2 = 1} \|w \circ T\|_{\mu_M}^2, \text{ where}$$

$$\|w\|_{\mu_N}^2 = \sum_{y \in N} w(y)^2 d\mu_N(y), \text{ and}$$

$$\|w \circ T\|_{\mu_M}^2 = \sum_{x \in M} w(T(x))^2 d\mu_M(x).$$

Since w is assumed to lie in the span of the first k eigenfunctions of the LB operator, we let $w = \Phi_{1..k}^N \mathbf{w}$, where $\mathbf{w} \in \mathbb{R}^k$, and remove the span constraint. This allows us to re-write, Eq. (2) as:

$$\mathbf{w}_k^* = \arg \max_{\mathbf{w} \in \mathbb{R}^k, \|\Phi_{1..k}^N \mathbf{w}\|_{\mu_M}^2 = 1} \|\Phi_{1..k}^M C \mathbf{w}\|_{\mu_M}^2, \quad (3)$$

where $\Phi_{1..k}^M$ and $\Phi_{1..k}^N$ represent the $m \times k$ and $n \times k$ matrices containing the first k eigenfunctions of the Laplace-Beltrami of M and N respectively, and C represents the functional representation of the map T . The following Lemma allows us to replace the weighted norm in Eq. (3) with the simple Euclidean norms.

Lemma 4.2 *Let M be a discrete shape with a measure μ_M , given as a function $d\mu_M$ over the vertices of M . Then, if $\{\psi_{1..k}^M\}$ are a set of functions on M which are orthonormal with respect to $d\mu_M$:*

$$\sum_{x \in M} \psi_i^M(x) \psi_j^M(x) d\mu_M(x) = 1 \text{ if } i = j \text{ and } 0 \text{ otherwise,}$$

then for any function $w = \sum_i a_i \psi_i$:

$$\|w\|_{\mu_M}^2 = \sum_{x \in M} w^2(x) d\mu_M(x) = \sum_i a_i^2.$$

Proof Note simply that $(\sum_i a_i \psi_i(x))^2 = \sum_i a_i^2 \psi_i^2(x) + \sum_{i \neq j} a_i a_j \psi_i(x) \psi_j(x)$. Since ψ_i are assumed orthonormal, the lemma follows. \square

This Lemma can be seen as a very special case of the generic Plancherel's theorem. Most importantly for us, it holds under any choice of measure. Using Lemma 4.2, assuming that the basis functions $\Phi_{1..k}^M$ and $\Phi_{1..k}^N$ are orthonormal with respect to μ_M and μ_N respectively, we can rewrite Equation 3 as:

$$\mathbf{w}_k^* = \arg \max_{\mathbf{w} \in \mathbb{R}^k, \|\mathbf{w}\|^2 = 1} \|C \mathbf{w}\|^2. \quad (4)$$

Note that we have replaced the weighted norms $\|\cdot\|_{\mu_M}$ and $\|\cdot\|_{\mu_N}$ with the simple Euclidean norm. Now finding \mathbf{w}_k^* can be done simply using the singular value decomposition of C . In particular, if $C = U \Sigma V^T$ is the singular value decomposition of C , then it is well-known that $\mathbf{w}_k^* = V_1$, namely the right singular vector corresponding to the largest singular value. The optimal w_k^* is given by $\Phi_{1..k}^N \mathbf{w}_k^*$. \square

Figure 4 shows the functions w_k^* for a map computed by the method of Kim et al. [KLF11] and volume measure-based distortion, for several choices of k . To visualize the matching function on shape M , we use $\Phi_{1..k}^M U_1$.

4.3. Quantitative Measures

Since the largest and smallest singular values are $\sigma_{\max(\text{corresp. min})}(C) = \sup(\text{corresp. inf})_{\|w\|=1} \|Cw\|$, we can see that it is possible to evaluate the measure-preserving quality of the map T through the analysis of the singular values of its functional representation C . Indeed, the following results relate the singular values of the functional map C to the quality of the underlying map T .

Theorem 4.3 *Given two shapes, M and N , and a map $T : M \rightarrow N$, its functional representation C , which maps functions on N to functions on M is orthonormal if and only if T is measure preserving, i.e. for any measurable set $A \subset N$, $\mu_N(A) = \mu_M(T^{-1}(A))$ where μ_M and μ_N are the given measures on M and N . Moreover, if C is full-rank, i.e. $\sigma_{\min}(C) > 0$, then T must be surjective, up to sets of measure 0 on N .*

Proof See Appendix \square

Note that Theorem 5.1 in [OBCS*12] establishes one direction of Theorem 4.3 (that volume preservation implies orthonormality) which is useful for regularization when computing the map. However, the opposite direction is also true, and is useful for analyzing a map, by establishing that the

map’s quality (in terms of measure distortion) can be directly inferred from the singular values of C .

Note also that while the function corresponding to the first singular vector \mathbf{w}_k^* provides us with a visualization of the *most* distorted region, by considering the following singular vectors, we can visualize the distorted regions in order of importance. Hence, by combining the information from both the singular vectors and singular values, we have a quantitative and qualitative way of evaluating maps. Figure 5 shows the first few singular vectors of a map computed by the method in [KLF11], as well as the matching singular values, where we used $k = 30$ on M (the giraffe), and $k = 300$ on N (the cow). In addition we show the visualization coloring both surfaces using the xyz coordinates of M . While in this visualization the map seems to be of high quality, our method identifies the problematic mapping of the nose of the cow to the head and neck of the giraffe, as is seen in the first singular vector. The second and third singular vectors show the incorrect mapping on the legs, where the hoof of the cow is mapped to half the length of the giraffe’s leg. Finally, the last singular vectors (29 and 30) show how the horns of the cow are mapped to areas on the neck of the giraffe.

4.4. Changing the Measure

Note that the framework that we have described above is applicable to any pair of measures μ_M, μ_N and any map T . In most cases, we use the standard area measure. However, in many cases there are better choices for visualizing distortion. For example, for non-isometric maps some areas of stretch are expected, and we might prefer to focus on regions where extremities are not mapped to extremities.

Intuitively, by changing the measure from the volume measure in the discrete case, we wish to replace the area element on M with an arbitrary function $d\mu_M$. Or, equivalently, given two surfaces M and N we wish to compare two surfaces M_f and N_g which are *conformal* to M and N respectively, such that the conformal factor is prescribed by $f = d\mu_M$ and $g = d\mu_N$. Interestingly, for continuous 2d manifolds, the L^2 norm of the gradient of any function is invariant under conformal changes of the metric. Therefore, the use of the Laplace-Beltrami eigenfunctions is still justified, since they preserve their well-ordered smoothness properties (as measured by the norm of the gradient), meaning that restricting the function w to lie in the span of the first k eigenfunctions provides a meaningful notion of scale, allowing to concentrate from global (small k) to highly local (large k) deformation.

In practice, there is no need to compute the new surfaces $M_{d\mu_M}$ and $N_{d\mu_N}$. All we need in order to apply our machinery is to use the eigenvectors of the Laplace-Beltrami operator defined on the new surfaces. It is straightforward to see, that under a conformal change of metric of M given by the conformal factor $d\mu_M$ defined at the vertices of M , the discretization of the LB operator is given by $L_f^M = F^{-1}W^M$,

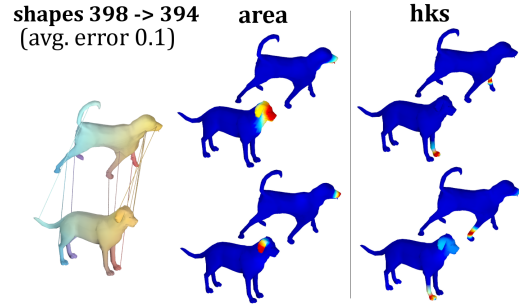


Figure 6: Visualizing two different distortion measures. We show the distortion indicators corresponding to the first two singular values, for the area measure, and a measure based on HKS. Note that HKS-based distortion is sensitive to extremities not being mapped to each other, whereas area distortion is sensitive to local scale changes.

where F is diagonal matrices representing $d\mu_M$ while W^M is the standard stiffness matrices given by the cotangent weight scheme of Meyer et al. [MDSB02]. Note that W^M is fixed and independent of $d\mu_M$.

This means that given the discrete measures $d\mu_M$ and $d\mu_N$, we can simply form the matrices: $L^M = F^{-1}W^M$ and $L^N = G^{-1}W^N$, where F and G are diagonal matrices containing $d\mu_M$ and $d\mu_N$ respectively. Then, we compute the eigenfunctions of L^M and L^N by solving the generalized eigenvalue problems: $W^M\phi^M = \lambda F\phi^M$ and similarly for N .

Note that since W is positive semi-definite and F is diagonal and positive, the resulting eigenfunctions ϕ^M, ϕ^N are orthonormal with respect to the measures $d\mu_M$ and $d\mu_N$ as remarked by Rustamov in [Rus07]. Therefore, Lemma 4.2 applies and we can solve the optimization problem in Equation 3 simply by computing the singular vectors of the functional matrix C . Note that unlike approximating functional maps with the LB operator, which, as remarked in [OBBS*12] works well for “natural” functions on manifolds, there are no restrictions in the above procedure on the measures $d\mu_M$ and $d\mu_N$, since they are used to *define* the corresponding LB operators. Although it is possible to use more accurate notions of discrete conformality (e.g. the one defined in [SSP08]), we have found this approach to be sufficient for our needs.

Figure 6 shows an example of the effects that can be identified by using different measures of distortion using our method. One example would be to take the distortion measure to be derived from the Heat Kernel Signature [SOG09]. As this function has very distinctive values for extremities, it can identify regions where extremities are not mapped to extremities, as in Figure 6 (right). The area measure, on the other hand, identifies other regions, where there are large changes in local scale, making analysis and visualization based on these measures complementary.

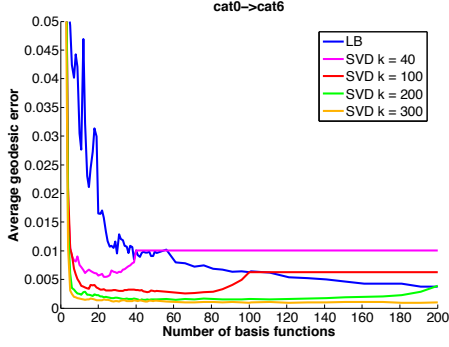


Figure 7: Reconstruction error for a point-to-point maps from their functional representation, for different values of k . Note that for $k = 300$ we can reconstruct the map using only about $r = 4$ SVD basis vectors.

4.5. Map Summary

The results above imply that the functions corresponding to the singular vectors of the functional representation of a map can be used as a compact summary of the map, which can be used for visualization, since they identify areas of maximal measure distortion. Here we also briefly note that the singular vectors provide an optimal summary of the map in an information-theoretic way – by providing an optimal low-dimensional approximation of the functional map.

Namely, suppose we are given a map $T : M \rightarrow N$, and let C be its functional representation mapping functions on N to functions on M . Now, suppose we are given a fixed budget of r functions and we would like to find an r -dimensional basis in which the action of C is well approximated. In other words, we would like to find a projection matrix V with r orthonormal columns, such that $\|Ca\| \approx \|CV_r V_r^T a\|$.

The following lemma is easy to show using the well-known properties of the singular value decomposition for low rank approximation:

Lemma 4.4 *Given a matrix $C \in \mathbb{R}^{k \times k}$ and a number $r < k$, the optimal orthogonal projection matrix $V_{opt} \in \mathbb{R}^{k \times r}$:*

$$V_{opt} = \min_{V_r} \max_{\mathbf{a}} \|Ca - CV_r V_r^T \mathbf{a}\|. \quad (5)$$

is given by first r right singular vectors of C . Moreover, in this case, for any vector \mathbf{a} , $CV_r V_r^T \mathbf{a} = U_r S_r V_r^T \mathbf{a}$ where S_r is an $r \times r$ diagonal matrix of singular values, and U_r is given by the first r left singular vectors of C .

In other words, if we are given a matrix $C \in \mathbb{R}^{k \times k}$, and we would like to approximate it by its composition with an r -dimensional projection matrix, then the optimal projection is given by the singular vectors of C .

Note that if \mathbf{a} is a vector representing a function through its coefficients in the LB basis, i.e. $f = \Phi \mathbf{a}$, $f : M \rightarrow \mathbb{R}$ then, $g = \Phi V_r V_r^T \mathbf{a}$ can be written as $g = \Psi \mathbf{b}$, where Ψ are the r basis functions given as $\Psi_i = \Phi V_i$ (the i^{th} column of V_r)

Input: Shapes M, N , a map $T : N \rightarrow M$, functions $f : M \rightarrow \mathbb{R}^+$, $g : N \rightarrow \mathbb{R}^+$, $k_M, k_N, r > 0$

Output: Functions $w_h^M, w_l^M : M \rightarrow \mathbb{R}^r$, $w_h^N, w_l^N : N \rightarrow \mathbb{R}^r$, $s_h, s_l \in \mathbb{R}^r$

$F \leftarrow \text{diag}(f); G \leftarrow \text{diag}(g);$
 $W^M \leftarrow \mathcal{L}(M); W^N \leftarrow \mathcal{L}(N);$
 $\Phi^M \leftarrow \text{eig}(F^{-1}W^M, k_M), \Phi^N \leftarrow \text{eig}(G^{-1}W^N, k_N)$
solve $\Phi_{1..k_N}^N C = \Phi_{1..k_M}^M \circ T$ for C
 $[U, S, V] \leftarrow \text{svd}(C)$
 $w_h^M \leftarrow \Phi_{1..k_M}^M V_{1..r}, w_l^M \leftarrow \Phi_{1..k_M}^M V_{k_M-r+1:k_M}$
 $w_h^N \leftarrow \Phi_{1..k_N}^N U_{1..r}, w_l^N \leftarrow \Phi_{1..k_N}^N U_{k_M-r+1:k_M}$
 $s_h \leftarrow \text{diag}(S_{1..r, 1..r}), s_l \leftarrow \text{diag}(S_{k_M-r+1:k_M, k_M-r+1:k_M})$

Algorithm 1: Visualization and evaluation of a map between shapes.

and \mathbf{b} is an r -dimensional vector given by $\mathbf{b} = V_r^T \mathbf{a}$. In other words, we can treat the approximation of C as a change of basis, where the first k Laplace-Beltrami eigenfunctions are replaced by r basis functions, each given as a linear combination of the LB eigenfunctions, with coefficients provided by columns of V_r . Note finally that C becomes a diagonal functional map when the basis functions on the two spaces are represented by $\Phi^M V_r$ and $\Phi^N U_r$.

In practice, representing a map in the modified basis has a noticeable effect in reducing the reconstruction error when capturing the point to point map through its functional representation. Figure 7 shows the average geodesic error associated with converting a point-to-point map to its functional form, using the eigenfunctions of the unweighted Laplace-Beltrami operator, and using the new basis derived from the singular vectors of C . Note that for $k = 200$, and $r = 4$ basis functions we can reconstruct the map so that the image of each point is on average within one hop neighbor of its image as given by the map (See also [OBCS*12] Figure 3 and associated text).

The figure shows another interesting phenomenon - if r is too large the reconstruction error *grows*. This happens because there are functions in the span of the first k LB eigenfunctions of M which are not in the span of the first k LB eigenfunctions of N . When r is small, the SVD procedure makes sure that these basis vectors are not used, however as $r \approx k$, there is no way to avoid these basis vectors, which introduce errors into the map.

5. Experimental Results

Our visualization and evaluation framework is described in Algorithm 1. Here $\text{diag}(f)$ constructs a diagonal matrix out of a function f , and $\text{diag}(F)$ extracts the diagonal from the matrix F , $\mathcal{L}(M)$ is the cotangent weights Laplacian of the shape M , $\text{eig}(M, k)$ returns the eigenvectors of M matching the smallest k eigenvalues, and $\text{svd}(C)$ computes the singular value decomposition $C = USV^T$.

The computed functions w_h^M (or w_l^M) are used to visualize regions in M which stretch (or shrink) the measure f under

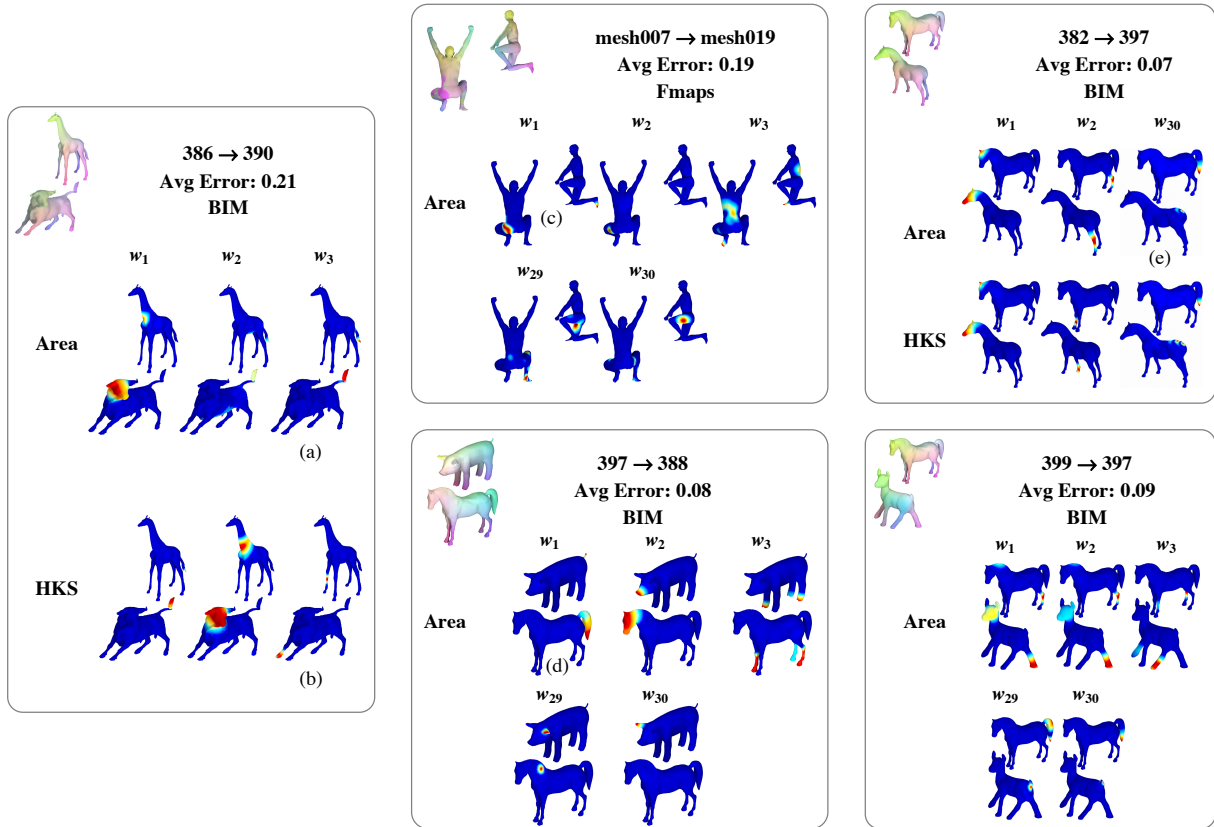


Figure 8: Visualizing the distortion of a few maps from [KLF11] and [OBCS*12]. See the text for details.

the map. Every column in $w_{h,l}^M$ and $w_{h,l}^N$ represent functions which correspond under the map, thus color coding these functions on the shapes M and N allows us to see in which regions the distortions are concentrated, and how these regions are mapped. The values $s_h \geq 1$ and $s_l \leq 1$ are a quantitative measure of distortion, where large values of $\|s_{h,l} - 1\|$ indicate large distortion.

We have experimented with a few maps from the map collection created by the methods in [KLF11] and [OBCS*12]. Figure 8 shows the visualization of a few such maps, using different distortion measures. For all our experiments we took $k_M = 30, k_N = 300$. We always pick a large k_N to minimize the loss due to reconstruction error on N , and have found that for visualization purposes $k_M = 30$ provides a good trade-off of smoothness versus accuracy. We visualize the functions w_l, w_h , and also show the average geodesic error (see [KLF11] Figure 7 for a visualization of this value). As can be seen in the figure, this standard error metric is not always informative, as there are problematic regions for which this error metric is small. We can further see that different distortion measures locate different type of errors in the map, prioritizing them differently.

Note specifically how w_3 differs when computing it using the area measure vs. when computing it using HKS (marked

(a) and (b) respectively in the figure). Using the area measure (a), we get a function localized on the tail area, which indeed scaled considerably. Using the HKS measure (b) we identify the fact that a cylindrical area on the leg of the giraffe is mapped to an extremity on the bull.

Our method is most effective when the map analyzed is of high quality, but has some problematic regions. However, it can also be useful even when the map is bad enough so that standard color coding of the coordinate function can be used to identify problems. Consider for example Figure 8 (c) where both the coordinate color coding and w_1 identify a discontinuity: the right foot mapped to the top of the left leg. However our method also identifies additional areas (e.g. in w_{29}), which are hard to see in the standard visualization.

As we are primarily interested in maps between non-rigid and often non-isometric shapes, there are bound to be distortions which although they are large, are “acceptable”, as the mapping is semantically meaningful. See for example (d): there is a very big area distortion, however the map is good in that region, since it nicely maps between the tails of the pig and the horse. Furthermore, as the maps we investigate are not necessarily bijections, interesting phenomena appear. For example, by looking at the lower singular vectors (e.g. w_{30} in (e)), we can identify regions where the map fails to be




Figure 9: Compact visualization of 50 maps from deformed humans onto an undeformed shape. These functions demonstrate the areas with maximum distortion, on the shapes where this area was distorted the most.

surjective - namely maps a large area on one shape to a zero (or very small) area on the second shape.

Computing this visualization is very efficient in practice since it only requires computing the SVD of a moderate sized matrix C . The most costly operation is computing the first k eigenfunctions of the LB operator. Since $k \approx 100 - 300$ this step does takes about 10-30 seconds on a standard PC for shapes with up to 50k points.

5.1. Shape Collections

An interesting application of our visualization technique is that it allows to visualize *collections* of maps compactly. Namely, suppose we have a set of maps from n shapes: S_1, \dots, S_n to a given shape S_0 . This gives us n functional maps C_i , which can be represented as matrices $C_i \in \mathbb{R}^{k \times k}$. Now, we can construct a matrix $C \in \mathbb{R}^{nk \times k}$, where each block of size $k \times k$ corresponds to one of the i matrices C_i . Note that given a function $f : S_0 \rightarrow \mathbb{R}$ represented as a vector $\mathbf{a} \in \mathbb{R}^k$ of basis coefficients, we have $C\mathbf{a} = \mathbf{b}$ where $\mathbf{b} \in \mathbb{R}^{nk \times 1}$ and each block of size k corresponds to $C_i\mathbf{a}$, i.e. the coefficients of $T_i \circ f$ on shape S_i . Thus: $\mathbf{a}_1 = \arg \max_{\|\mathbf{a}\|=1} \|C\mathbf{a}\|$ represents the coefficients of the function f which is *jointly most distorted* by the maps C_i .

To illustrate this phenomenon, we considered 50 deformed human shapes from the SCAPE dataset, and their

correspondences onto the undeformed shape. We represented each map as a 200×30 functional mapping and constructed a 10000×30 matrix C as described above. We computed the first singular vectors as described above, and picked the deformed shapes where the norm of the coefficient vector was the largest. Figure 9 shows the functions corresponding to the first singular vectors of C , visualized on the shapes we picked. We have repeated this experiment for the 19 shapes from the animals collection in [KLF11], shown in Figure 10. We identified some errors in the map which were consistent across the collection, e.g. the nose of the dog was mapped to the neck in a number of cases. We note that this joint map analysis is closely related to standard mesh-based inverse kinematics [SZGP05] except that our method is completely *intrinsic* and does not rely on computations in Euclidean space.

5.2. Limitations

Our method has several limitations. First, in terms of usability, our approach generates a large number of images (by varying the scale and the measures), which might create visual overload for a user. This is especially true when visualizing a collection of maps. It could be more appropriate to have an interactive means of navigating the different images, so that the user can explore various problematic regions under different measures. Second, the topology of the problematic region might have more than one connected component, in which case the color coding visualization does not provide enough information to identify which region is mapped to which. This could be easily solved by post-processing the multi-scale distortion indicators to identify non-zero regions. In addition, our method has no information on the semantics of the map, and therefore could identify regions which are highly distorted, yet semantically mapped correctly. Finally, by construction, our method does not visualize the global effect of the map, e.g. whether a symmetric flip is present, and as such is complementary to existing visualization methods.

6. Discussion

As the research on finding maps between shapes advances, it is critical that the research on evaluation methods follows closely. This work demonstrated an attempt to visualize and evaluate maps between shapes in a multi-scale way, by leveraging the machinery provided by the functional map representation and its spectral decomposition. We have also provided preliminary evidence that this decomposition is useful for visualizing collections of maps. In general, we believe that this approach may open the door for research on “map processing”: the study of maps between shapes as objects in their own right.

By adding this additional level of abstraction, it could potentially be possible to perform new types of analysis, for example to consider map analogies (A is to B as C is to D), or to compare between objects which are not directly comparable (such as images and surfaces), by comparing the set of

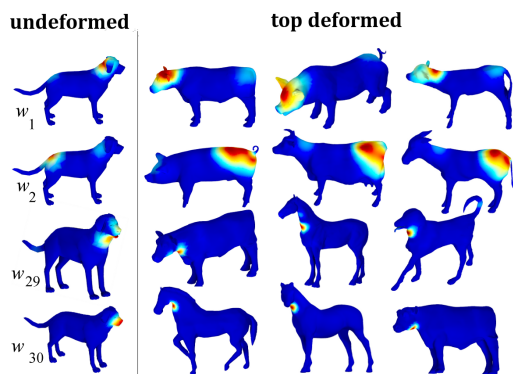


Figure 10: Compact visualization of 19 maps from different animals onto a dog shape from [KLF11].

maps that acts on these objects. Furthermore, it could potentially be possible to define “map descriptors” by using the spectral decomposition of the maps, and to perform “map retrieval” on a collection of maps. In essence, investigating maps between objects provides a unique perspective on the properties of these objects, and thus we believe that developing algorithms for map understanding is crucial for the ultimate goal of shape understanding.

Acknowledgements. The authors would like to acknowledge NSF grants NSF CCF 1161480 and DMS 1228304, AFOSR grant FA9550-12-1-0372, FP7-CIG grant 303511 and ISF grant no. 699/12., the ANR project GIGA (ANR-09-BLAN-0331-01), a Google Research award, the Max Planck Center for Visual Computing and Communication and the INRIA Associated team “COMET.”

References

- [BBK06] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS* 103, 5 (2006), 1168–1172. 2
- [BCWG09] BEN-CHEN M., WEBER O., GOTSMAN C.: Variational harmonic maps for space deformation. *ACM Trans. Graph.* 28, 3 (July 2009), 34:1–34:11. 3
- [HAWG08] HUANG Q.-X., ADAMS B., WICKE M., GUIBAS L. J.: Non-rigid registration under isometric deformations. *CGF (Proc. SGP)* 27, 5 (2008), 1449–1457. 2
- [HG00] HORMANN K., GREINER G.: MIPS: An efficient global parametrization method. In *Curve and Surface Design*, Innovations in Applied Mathematics. 2000, pp. 153–162. 3
- [KCATCO*10] KIN-CHUNG AU O., TAI C.-L., COHEN-OR D., ZHENG Y., FU H.: Electors voting for fast automatic shape correspondence. *Comp. Graph. Forum* 29, 2 (2010), 645–654. 2
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. *ACM Trans. Graph. (Proc. SIGGRAPH)* 30, 4 (2011). 2, 3, 5, 6, 8, 9, 10
- [LF09] LIPMAN Y., FUNKHOUSER T.: Möbius voting for surface correspondence. In *Proc. of SIGGRAPH* (2009), vol. 28:3, pp. 72:1–72:12. 2, 3
- [LZX*08] LIU L., ZHANG L., XU Y., GOTSMAN C., GORTLER S. J.: A local/global approach to mesh parameterization. In *Proc. SGP* (2008), pp. 1495–1504. 3
- [MDSB02] MEYER M., DESBRUN M., SCHRÖDER P., BARR A. H.: Discrete differential geometry operators for triangulated 2-manifolds. In *Proc. VisMath’02* (Berlin, Germany, 2002). 6
- [MHK*08] MATEUS D., HORAUD R. P., KNOSSOW D., CUZZOLIN F., BOYER E.: Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Proc. CVPR* (2008). 2
- [OBCS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.* 31, 4 (July 2012), 30:1–30:11. 2, 3, 5, 6, 7, 8, 10
- [OMMG10] OVSJANIKOV M., MERIGOT Q., MEMOLI F., GUIBAS L.: One point isometric matching with the heat kernel. *CGF* 29, 5 (2010), 1555–1564. 2
- [Rus07] RUSTAMOV R. M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Proceedings of SGP* (2007), pp. 225–233. 6
- [SH10] SHARMA A., HORAUD R. P.: Shape matching based on diffusion embedding and on mutual isometric consistency. In *Proc. NORDIA Workshop (CVPR)* (June 2010). 2
- [SLKW12] SELLENT A., LAUER P.-S., KONDERMANN D., WINGBERMÜHLE J.: A toolbox to visualize dense image correspondences (stereo disparities & optical flow). <http://hci.iwr.uni-heidelberg.de/Static/correspondenceVisualization/Visualization.pdf> (2012). 3
- [SMW06] SCHAEFER S., MCPHAIL T., WARREN J.: Image deformation using moving least squares. *ACM Trans. Graph.* 25, 3 (July 2006), 533–540. 3
- [SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum* 28, 5 (2009), 1383–1392. 6
- [SSP08] SPRINGBORN B., SCHRÖDER P., PINKALL U.: Conformal equivalence of triangle meshes. In *ACM SIGGRAPH ASIA 2008 courses* (2008), ACM, p. 18. 6
- [SY11] SAHILLIOĞLU Y., YEMEZ Y.: Coarse-to-fine combinatorial matching for dense isometric shape correspondence. *Computer Graphics Forum* 30, 5 (2011), 1461–1470. 2, 3
- [SZGP05] SUMNER R. W., ZWICKER M., GOTSMAN C., POPOVIĆ J.: Mesh-based inverse kinematics. In *ACM SIGGRAPH* (2005), pp. 488–495. 9
- [TBW*11] TEVS A., BERNER A., WAND M., IHRKE I., SEIDEL H.-P.: Intrinsic shape matching by planned landmark sampling. In *Proc. Eurographics* (2011), p. to appear. 2

Appendix

6.1. Proof of Theorem 4.3

We only need to prove that if C is orthonormal then T is locally volume preserving the other implications being Theorem 5.1 in [OBCS*12]. If C is orthonormal then for any function f on M we have $\langle Cf, Cf \rangle_N = \langle f, f \rangle_M$. Assume that T is not locally volume preserving, i.e. there exists a set $A \subset M$ such that $\mu_M(A) \neq \mu_N(T^{-1}(A))$ where μ_M and μ_N are the volume measures on M and N . If we denote by f_A the indicatrix function of A , $f_A(x) = 1$ if $x \in A$ and $f_A(x) = 0$ otherwise, Cf_A is the indicatrix function of $T^{-1}(A)$ and we then have $\langle f_A, f_A \rangle_M = \mu_M(A)$ is not equal to $\langle Cf_A, Cf_A \rangle_N = \mu_N(T^{-1}(A))$: a contradiction. The second part of the theorem is proven in the same way: since all the singular values are positive any subset of M with non zero measure must have a non empty pre-image.