

Map-Based Exploration of Intrinsic Shape Differences and Variability

Raif M. Rustamov¹ Maks Ovsjanikov² Omri Azencot³ Mirela Ben-Chen³ Frédéric Chazal⁴ Leonidas Guibas¹

¹Stanford University ²LIX, École Polytechnique ³Technion ⁴Geometrica, INRIA




Figure 1: The notion of shape difference defined in this paper provides a way to compare deformations between shape pairs. This allows us to recognize similar expressions of shape *A* (top row) to those of shape *B* (bottom row), without correspondences between *A* and *B* and without any prior learning process.

Abstract

We develop a novel formulation for the notion of shape differences, aimed at providing detailed information about the location and nature of the differences or distortions between the two shapes being compared. Our difference operator, derived from a shape map, is much more informative than just a scalar global shape similarity score, rendering it useful in a variety of applications where more refined shape comparisons are necessary. The approach is intrinsic and is based on a linear algebraic framework, allowing the use of many common linear algebra tools (e.g. SVD, PCA) for studying a matrix representation of the operator. Remarkably, the formulation allows us not only to localize shape differences on the shapes involved, but also to compare shape differences across pairs of shapes, and to analyze the variability in entire shape collections based on the differences between the shapes. Moreover, while we use a map or correspondence to define each shape difference, consistent correspondences between the shapes are not necessary for comparing shape differences, although they can be exploited if available. We give a number of applications of shape differences, including parameterizing the intrinsic variability in a shape collection, exploring shape collections using local variability at different scales, performing shape analogies, and aligning shape collections.

CR Categories: I.3.5 [Computing Methodologies]: Computer Graphics—Computational Geometry and Object Modeling

Keywords: Shape comparison, shape matching, shape variability, data-driven methods

Links:  DL  PDF

ACM Reference Format

Rustamov, R., Ovsjanikov, M., Azencot, O., Ben-Chen, M., Chazal, F., Guibas, L. 2013. Map-Based Exploration of Intrinsic Shape Differences and Variability. *ACM Trans. Graph.* 32, 4, Article 72 (July 2013), 12 pages. DOI = 10.1145/2461912.2461959 <http://doi.acm.org/10.1145/2461912.2461959>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

2013 Copyright held by the Owner/Author. Publication rights licensed to ACM.
0730-0301/13/07-ART72 \$15.00.
DOI: <http://dx.doi.org/10.1145/2461912.2461959>

1 Introduction and Rationale

Comparing shapes is a fundamental operation in shape analysis and geometry processing, with many applications to computer graphics, including interactive shape design, shape search, and the organization of shape collections. Most approaches to comparing shapes reduce the comparison to a single number, a shape similarity score or distance. These distances can be computed either by establishing correspondences between the shapes (and therefore being able to compare the geometry at a finer scale) or by computing certain global shape descriptors and then estimating a distance in descriptor space.

In many settings, however, we may desire a more detailed understanding of how two shapes differ that goes beyond a single similarity score. Shapes can be complex objects and the very plethora of shape distances that have been proposed is testimony to the fact that no single scalar metric is able to satisfy all applications. For example, we may be interested in *where* two shapes are different and in *how* they are different. Such finer comparisons have long been important in other fields, such as industrial metrology to assess the quality of manufacturing processes, or in computational anatomy, to separate normal organ variability from disease forms for diagnostic purposes. In computer graphics, as shape collections are getting larger and larger with more objects in each category, these finer and more detailed shape comparisons are becoming important – and difficult to handle by coarse traditional techniques.

When computing maps or correspondences between shapes (including shape parametrization) the minimization of measures of shape distortion has long been used as a key optimization criterion. Yet once the map is computed, the distortion information is not stored, analyzed, or compared to that of other maps. In this paper we reverse this process. Starting from a map between two shapes, we propose a novel notion of shape differences as seen by this map, one that provides detailed information about how the shapes differ. Thus our work leverages the recent flurry of activity in algorithms for mapping shapes.

The main contribution of this paper is to give a rigorous mathematical formulation of the concept of a shape difference under a map and show how such shape differences can be computed, analyzed, and compared – thus making shape differences concrete, tangible objects that can be manipulated just like the shapes them-

selves can. Our approach is based on the following insight: in classical Riemannian geometry, local distortions induced by a map are expressed in terms of changes in the metric — which essentially is equivalent to tracking the changes in *inner products of tangent vectors* before and after these vectors are transported by the map from the source to the target shape. In contrast, in this paper, we track the changes in *inner products of real-valued functions* induced by transporting these functions from the source shape to the target shape via a functional map, as in [Ovsjanikov et al. 2012]. Our main observation is that all these changes can be captured by certain linear operators (matrices), which we call shape differences; remarkably, a single such operator works simultaneously for all pairs of functions.

Our approach has several key advantages. First, we exploit the recent functional map formulation of [Ovsjanikov et al. 2012], so our notion of a map can be quite general and incorporate mapping ambiguities due to symmetry, slippage, etc. Second, the approach is intrinsic and is not affected by the embedding of the shape in 3D. Third, when we have point-to-point correspondences, our shape difference can be directly related to classical local notions of geometric distortion, such as area or conformal distortion. Additionally, under a few assumptions and allowing for certain equivalences, the original map can be recovered from the shape difference. Fourth, we define shape difference via a linear operator formulation and discretize it into a matrix or vector form, giving us access to a wealth of linear algebra tools.

Our explicit representation of shape differences facilitates a number of challenging shape analysis tasks. For example, given two pairs of shapes, A, B and C, D , shape differences allow us to quantify how much the *change* from A to B is similar to the change from C to D , regardless of how similar A is to C . We can do these kinds of “shape analogies” only because we can compute the “difference among the differences” of the four shapes. For example, consider the face shapes in Figure 1. While the two rows of faces shown differ significantly, the relative changes between the undeformed and deformed version of each are similar, which is captured by our informative descriptors.

One of the key aspects of our shape differences is that they allow both localizing and parameterizing the variability between a single pair or of a collection of shapes. Thus, we can provide not only a canonical descriptor for a difference between a pair of shapes but also use it to analyze and visualize the source of the variability, making the interpretation of results easier and more concrete. In addition, we can now analyze the structure of shape collections based on relating the differences between the shapes and not the shapes themselves. We show several examples of the power of this approach in the paper. In particular, unlike almost all existing work, we can look at the variability of related shapes in a collection without necessarily having explicit point-to-point or landmark correspondences between the shapes (though we can use these when we have them).

After discussing related work (Section 2), we show how to formally define shape differences (Section 3) and compare them in shape collections (Section 4). We then discuss the discretization and computational aspects of shape differences (Section 5), as well as their key properties (Section 6). Finally we proceed to give a number of applications of this notion, including parameterizing the intrinsic variability in a shape collection, exploring shape collections using local variability at different scales, performing shape analogies, and aligning shape collections (Section 7).

2 Related Work

Shape differences and variability have been of interest in several scientific communities over many decades. The modern approach, generally termed *Statistical Shape Analysis*, exploits the notion of *Shape Space* introduced by D.G. Kendall [Dryden and Mardia 1998], where a standard set of key points or *landmarks* is selected on each shape and a shape is represented by its vector of landmarks after normalizing for rotation, translation and scale. Multiple shapes are analyzed jointly by first aligning their landmark vectors and then using principal components analysis (PCA) to extract the main modes of shape variation. Such learned shape variability models can also be used in segmenting shapes out of image or volume data (see, e.g. [Cootes et al. 2001]), following the active contour paradigm [Kass et al. 1988] from computer vision.

The medical research community, and especially brain anatomists, have explored many variations along this general theme, trying to compensate for the fact that exact landmarks may be hard to locate either algorithmically or manually in noisy medical images (2D or 3D). Many other shape features and shape descriptors, local and global, have been tried, including area, volume, spherical harmonics, medial axes or skeletons, etc. (see [Gerig et al. 2001; Golland et al. 2005] among many others) — see also the 2D image analysis survey by F.L. Bookstein [Bookstein 1996]. In these works, shape variability is effectively modeled by descriptor variability.

Another important issue is that not all variability carries the same significance. For example, in a population of 3D models of humans, some models may be the same human in different poses. If our goal is to understand the variability of human shapes, we must then factor out the variability due to pose variations among the subjects. Various approaches have been tried towards this end, including PCA in a Riemannian symmetric space [Fletcher et al. 2004], multivariate tensor-based morphometry using holomorphic forms [Wang et al. 2010], tensor ICA [Vasilescu and Terzopoulos 2007], the use of Laplace coordinates for points [Wuhrer et al. 2012], and others [Nain et al. 2007]. Only recently has an approach been proposed in these communities for comparing shapes intrinsically [Lai et al. 2010] using a spectral L^2 distance, but the approach suffers from the usual sign ambiguities (or more generally rotations within an eigenspace) of the eigenfunctions in spectral embeddings.

In the geometry processing area there has been considerable work in comparing shapes in an indirect way, in the setting of computing good maps between shapes. This is especially true 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. [Bronstein et al. 2006; Kim et al. 2011; Sahillioğlu and Yemez 2011] among a myriad of others). Perhaps the most common such criterion for a map between a pair of shapes is preservation of pairwise quantities such as geodesic distances [Bronstein et al. 2006; Sahillioğlu and Yemez 2011] or spectral quantities such as the heat kernel (e.g. [Sharma and Horaud 2010]). Generally, such measures of quality are both expensive to compute and non-trivial to analyze, making the intuitive understanding of the difference between shapes challenging. Another way of evaluating the distortion of a map, used mostly in shape deformation and parametrization applications (see e.g. [Schaefer et al. 2006; Ben-Chen et al. 2009], 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. While such local distortion measures are efficient to compute, they can often be too noisy to be used directly for identifying problematic regions. Finally, collections on human shapes were studied in [Allen et al. 2003; Angelov et al. 2005; Hasler et al. 2009]. These papers either explore pose and human shape variability separately, or they use skeleton information to facilitate pose alignment.

Our work is also related to the large volume of research on shape similarity metrics, either map-based or descriptor-based, whose survey is beyond the scope of this paper. Shape search using on such metrics has also been intensely studied, but mostly focused on discriminating shapes under large-scale variations (e.g., cars from humans). The current effort is aimed at fine variability, which has received less attention. In a related vein, the problem of how to map shape descriptor variability back onto something semantically meaningful on the original shapes was recently addressed in [Ovsjanikov et al. 2011]. We also note that the topic of fine classification/categorization has been popular in the computer vision community in the last few years (see, e.g., [Farrell et al. 2011] and the references in the papers therein).

From our point view, all of these approaches suffer from certain drawbacks. First of all, the notion of shape difference is not made explicit — at best only a shape “distance” is defined. With that it is impossible to understand precisely where the variation happens on a shape, as each shape is treated as “atom” — typically, a point in a fixed-dimensional Euclidean space. Furthermore, it is hard to compare differences between shapes — to express “differences among the shape differences,” for the same reason. Second, large amounts of information about the shapes is ignored, and this can affect the results. For example, the connectivity of the landmarks can be just as important as their absolute positions. Third, linear methods such as PCA are most often used — a notable exception being [Kilian et al. 2007] — even when it is not clear that a flat approximation to the shape space, either locally or globally is indicated. Fourth, these works perform *extrinsic* comparisons between the shapes and do not focus on their *intrinsic geometry* which is often what carries the true semantics of the shape. Unfortunately, invariance to isometric deformations is much harder to incorporate than invariance to Euclidean transformations. Finally, unlike earlier works that require vertex-to-vertex or consistent landmark correspondences, our use of the functional framework allows us to compare shapes whose meshes may be entirely different.

Recently, we exploited the machinery of functional maps to evaluate and visualize shape maps [Ovsjanikov et al. 2013]. Several of the local distortion measures we employ here were also studied in that paper, but that work has neither the matrix formulations nor the emphasis of assessing shape differences in the context of a shape collection that we introduce.

3 Shape Differences

Shape differences are linear operators (matrices in the discrete setting) that capture the disparity between shapes M and N under a given map (T or F below) between them. We define two types of shape differences, one based on the area distortion and another based on the conformal distortion, as induced by the map. In this section, we introduce the abstract definition of shape differences, applicable both in the continuous and discrete setting, and then show how they can be computed in practice in Section 5.

3.1 Background and notation

Our formulation uses the functional maps framework [Ovsjanikov et al. 2012] to represent maps between surfaces. Namely, given two surfaces M and N , a map $T : N \rightarrow M$ between them induces a map between functions $F : L^2(M) \rightarrow L^2(N)$, where $L^2(\cdot)$ is the set of square integrable real-valued functions on a surface. This functional map F takes each function $f : M \rightarrow \mathbb{R}$ and maps it to $g : N \rightarrow \mathbb{R}$ defined as $g = F(f) = f \circ T$. As pointed out in [Ovsjanikov et al. 2012], F is a *linear* transformation between function spaces and, therefore, can be represented as a matrix in the discrete setting. It is crucial to note that functional maps are

not limited to point-to-point maps, but provide a general notion of a map that can incorporate mapping ambiguities due to symmetry, slippage, etc. In the formulations below we will directly make use of a linear functional map $F : L^2(M) \rightarrow L^2(N)$, regardless of whether or not it is associated with a point-to-point map.

3.2 Formulation

Our goal in defining the shape differences between two shapes M and N , given a functional map F is to quantify some measure of distortion induced by F between $L^2(M)$ and $L^2(N)$. We compare shapes by comparing corresponding measurements made on the function spaces of the shapes. Following a Riemannian point of view, a measurement over a shape is defined by an *inner product* of functions on the shape. Recall that such an inner product $h(\cdot, \cdot)$ is a bi-linear form taking pairs of functions into real numbers. In this paper, for every surface S , we consider the following two inner products on $L^2(S)$:

Definition 1 We define the area-based inner product as $h_a^S(f, g) = \int_S f(x)g(x) d\mu(x)$.

Definition 2 We define the conformal inner product as $h_c^S(f, g) = \int_S \nabla f(x) \cdot \nabla g(x) d\mu(x)$ on the space of differentiable functions modulo constants.

These inner products are called area-based and conformal-based because of the following result (see the appendix for a proof):

Theorem 1 Given a pair of surfaces M, N and a bijection $T : N \rightarrow M$ with the functional representation F , the following holds:

1. $h_a^M(f, g) = h_a^N(F(f), F(g)), \forall f, g$ if and only if T is locally area preserving.
2. $h_c^M(f, g) = h_c^N(F(f), F(g)), \forall f, g$ if and only if T is conformal.

If the underlying map is not locally area preserving or conformal, the stated equalities will not hold. It is natural to quantify the distortions induced by the map through the failure of these equalities, perhaps by assigning a single number measuring the discrepancy. Of course the precise notion of discrepancy will depend on the functions f and g chosen. The challenge is to encode all these numbers arising out of different f and g into a single richer notion.

Our main observation is that all these discrepancies can be captured by certain matrices (linear operators), where these matrices are not simply tables of numbers, but can be meaningfully manipulated and compared as matrices. These operators effectively compensate for the distortions caused by the map F to the measurement in question and allow us to “pull back” the measurement $h_a^N(F(f), F(g))$ on N to a measurement made on M . Having all measurements on a common space is advantageous, as we want to be able to compare and compute differences between many measurements. Technically, we can accomplish this measurement transportation by using the following consequence of the classical Riesz representation theorem from functional analysis (see the appendix for a proof):

Theorem 2 Given two shapes M, N , endowed with inner products h^M and h^N respectively, and a functional map $F : L^2(M) \rightarrow L^2(N)$, there exists a unique linear operator $D_{h^M, h^N} : L^2(M) \rightarrow L^2(M)$ satisfying:

$$h^M(f, D_{h^M, h^N}(g)) = h^N(F(f), F(g)) \forall f, g.$$

We will refer to the operator D_{h^M, h^N} as the difference between h^M and h^N .




Figure 2: Given a pair of shapes M and N (left column) and a functional map F , the shape difference V is a linear operator, which for every function f on M produces another function $V(f)$ on M which intuitively encodes how much f is distorted by F (center and right). Note that f_1 is supported in an area that deforms under the map, and f_2 in an area that does not.

The linear operator D modifies g so as to exactly compensate for the distortions introduced by the map F . It is remarkable that D is a “universal compensator” — a single such operator works simultaneously for all functions f and g . Stated differently, D depends only on the given inner products on M and N , and the functional map F . It is also important to note that D is a linear self-map of the space of functions over M , (see Figure 2 for an illustration).

Now, we can apply this theorem to inner products h_a^M and h_a^N (resp. h_c^M and h_c^N) to formally quantify the difference between them. Namely, we define the *area-based shape difference* as:

$$V_{M,N,F} = D_{h_a^M, h_a^N} \quad (1)$$

and *conformal-base shape difference* as:

$$R_{M,N,F} = D_{h_c^M, h_c^N} \quad (2)$$

Since the map is usually clear from the context, we will often use the abbreviated notation $V_{M,N}$ and $R_{M,N}$.

We stress here that both $V_{M,N}$ and $R_{M,N}$ are not numbers but operators. They yield numbers, once functions specifying the measurement of interest are given. This flexibility enables a rich set of applications as we will show in Section 7. Note also that two shape differences V_{M,N_1} and V_{M,N_2} , even if $N_1 \neq N_2$, both represent linear operators with the same domain ($L^2(M)$) and range ($L^2(M)$). This allows us to *compare* shape differences even when they are defined using maps to different shapes (see Section 4), as hinted above.

Matrix representation: After a choice of a basis, the linear operators defined above can be made more tangible by expressing them in terms of matrices. Indeed, when dealing with discrete shapes, the underlying function spaces are finite-dimensional vector spaces. Any inner product $h^M(\cdot, \cdot)$ can always be represented via a matrix H^M such that: $h^M(f, g) = f^\top H^M g$, where f and g are column vectors. Similarly, given the shape N with an inner product h^N , and a functional map F , we can represent $h^N(Ff, Fg) = f^\top F^\top H^N Fg$, for some matrix H^N .

When these expressions for discrete inner products are plugged into Theorem 2, we can obtain an explicit expression for the difference operator D between these inner products:

$$D = (H^M)^{-1} F^\top H^N F. \quad (3)$$

Explicit formulas for both types of shape differences under various basis choices are provided in Section 5.

Discussion: These particular shape difference formulations were chosen for a number of reasons. First, isometric and conformal maps play an important role in shape processing, and so it is important to capture these exact notions. While conformal maps are directly characterized by our framework via the requirement $R_{M,N} = I$, where I is the identity map, note that isometric maps are both area-preserving and conformal and so can be characterized by the equalities $V_{M,N} = I$ and $R_{M,N} = I$. These and several more desirable properties of our shape differences will be discussed in Section 6.

Second, inclusion of the map $F : L^2(M) \rightarrow L^2(N)$ into our formulation allows the notion of the shape difference to change and depend on the context of a particular application. For example, while a purely geometric notion of shape difference can be obtained by taking the difference induced by some type of geometrically optimal map, such a shape difference may not be optimal for studying, say, differences in human brain shapes. In this latter case, our framework allows use of the maps provided by a specialist to compute a domain-specific shape difference. This is unlike, for example, various notions of Gromov-Hausdorff distances [Bronstein et al. 2006], which are often defined with respect to some “optimal” map.

Finally, when working with shape collections, it is crucial to be able to compare the shape difference between a pair of shapes to the shape difference between another pair of shapes. We discuss this in the next section, and show that our shape difference matrices can be rigorously compared to each other in a variety of circumstances without falling into the fallacy of “comparing apples to oranges.” Key to our approach is the ability to transport or move a shape difference between two shapes to a third reference shape via connecting maps, so as to make meaningful comparisons possible.

4 Differences in Shape Collections

One of the main advantages of the shape differences defined above is that they not only encode detailed knowledge about the distortion under a given map, but also allow distortion comparisons across pairs of shapes, defining “differences between differences.” Here, we outline how such comparisons can be carried out in three different scenarios. The discussion below is valid for both kinds of shape differences; to avoid repetition we will focus on the area-based shape differences.

The first scenario arises when for shapes M , N_1 , and N_2 , one wants to compare the shape difference V_{M,N_1} to V_{M,N_2} . As mentioned earlier, both of these shape differences are linear operators with the same domain and range $L^2(M)$, and thus can be directly compared and even algebraically combined if needed.

The second scenario arises when one wants to compare the shape difference V_{M_1,N_1} to V_{M_2,N_2} , where $M_1 \neq M_2$, assuming that a linear functional map G between $L^2(M_1)$ and $L^2(M_2)$ is *known*. Note that V_{M_1,N_1} and V_{M_2,N_2} , cannot be directly compared because they are defined over different domains and ranges ($L^2(M_1)$ and $L^2(M_2)$). In order to make the comparison possible, we need to first apply a change of basis transformation to one of the matrices — this is where the cross map $G : L^2(M_1) \rightarrow L^2(M_2)$ enters the picture, to allow the transportation of the difference. By applying a matrix conjugation by G to V_{M_2,N_2} we bring it into a common basis with the other matrix, and now the matrices V_{M_1,N_1} and $G^{-1}V_{M_2,N_2}G$ can be compared and algebraically combined as needed. To avoid computing the inverse of matrix G , one can also compare matrix products GV_{M_1,N_1} and $V_{M_2,N_2}G$.

The last, third, scenario arises when one wants to compare the shape difference V_{M_1, N_1} to V_{M_2, N_2} , but a mapping between M_1 and M_2 is *not known*. If we knew the map G , we would have compared after conjugating one of the matrices by G . However, since now we do not know the map, our comparison needs to rely on the quantities that are *invariant* under matrix conjugation. It is well known that the spectrum of a matrix is such an invariant, and therefore, for comparing the shape differences we can compare the spectra of the matrices V_{M_1, N_1} and V_{M_2, N_2} .

In the third scenario, it is crucial that the shape difference matrices are represented in terms of a truncated basis (e.g. low-frequency Laplace-Beltrami eigenfunctions) spanning a subspace of smooth functions. In essence, this adds a regularization on the unknown cross map G , forcing it to be smooth. This is a benefit of the functional map representation, see [Ovsjanikov et al. 2012] for a discussion.

5 Computation

In this section we present explicit formulas for computing shape differences between two triangle mesh surfaces M and N . As it is clear from formula (Eq. 3), to compute the shape differences we need to have access to three matrices: the inner product matrices H^M, H^N and the functional map F . Since these matrices depend on the choice of a basis for the functional spaces $L^2(M)$ and $L^2(N)$, we will consider three options.

Before proceeding, let us fix our discretizations. For a surface mesh S , we discretize the area-based inner product by $h_a^S(f, g) = \sum_{x \in S} f(x)g(x)A^S(x)$ where $A^S(x)$ is the area element (Voronoi area) associated with vertex x . The Laplace-Beltrami operator is discretized as $L = (A^S)^{-1}W^S$, where A^S is the diagonal matrix of area weights and W^S is the stiffness matrix (e.g. the standard cotangent weight matrix) [Pinkall and Polthier 1993]

Option 1: Here we use the finite element “hat function” basis for both $L^2(M)$ and $L^2(N)$. In the indicator, hat function basis, the matrix associated with the area-based inner products h_a^M and h_a^N are simply the diagonal matrices of area weights at vertices, A^M and A^N . Using formula (Eq. 3) the matrix associated with the area-based shape difference is given by $V_{M,N} = (A^M)^{-1}F^T A^N F$.

To derive the conformal-based shape difference, we use Stokes’ theorem $\int \nabla f(x) \cdot \nabla g(x) d\mu(x) = - \int f(x)\Delta g(x) d\mu(x)$, where Δ is the Laplace-Beltrami operator. This is valid in the discrete case (even if there is a boundary) due to our choice of discretization. In the discrete case, this means that $h_c^M(f, g) = -f^T A^M (A^M)^{-1} W^M g$, and thus, the matrix associated with the conformal-based shape difference h_c^M is given simply by $-W$, the stiffness matrix. This implies that the conformal-based shape difference under the functional map F is given by¹: $R_{M,N} = (W^M)^{-1}F^T W^N F$.

In a special case when the surfaces M and N have identical tessellations conforming to the map T , the functional map F is simply the identity matrix. Therefore, we obtain the following formulas:

$$V_{M,N} = (A^M)^{-1}A^N \quad \text{and} \quad R_{M,N} = (W^M)^{-1}W^N.$$

These formulas shed light into the nature of our shape differences. For example, $R_{M,N}$ is seen to capture the change of the conformal

¹The conformal shape difference operator is defined on $L^2(M)$ modulo constants, and extended to the entire $L^2(M)$ by setting it to zero for constants. In the discrete setting, the same effect is achieved by using pseudo-inverses; all inverses appearing in formulas for R are pseudo-inverses.

cotangent Laplacian (without the area weights) and, as a result, of the angles of the mesh.

Option 2: Here we use the orthonormal Laplace-Beltrami bases for both $L^2(M)$ and $L^2(N)$. First note that the matrix associated with area-based inner product is simply the identity matrix because the Laplace-Beltrami basis is orthonormal. The matrix associated with the conformal-based inner product on M is the diagonal matrix $D_M = \text{diag}(-\{\lambda_i^M\})$, where λ_i^M is the i^{th} eigenvalue of the Laplacian of M ; similarly for the conformal inner product on N . Therefore, given a functional map F :

$$V_{M,N} = F^T F, \quad \text{and} \quad R_{M,N} = (D^M)^{-1}F^T D^N F. \quad (4)$$

Option 3: Here we use the orthonormal Laplace-Beltrami basis for $L^2(M)$, and hat basis for $L^2(N)$. The resulting formulas can be derived by essentially combining the derivations for the two options above to yield:

$$V_{M,N} = F^T A^N F, \quad \text{and} \quad R_{M,N} = D_M^{-1}F^T W^N F.$$

Discussion: The first option is presented here only for theoretical reasons, to show that simple and intuitive expressions exist in cases when the shapes are identically tessellated. However, this option is not practically useful: a) the obtained shape difference matrices are sensitive to noise both in the meshes and the maps, b) their sizes scale with the number of mesh vertices, and c) the computation requires the pseudo-inverse of a large matrix.

The second option has the advantages of being generally applicable and allows a smoothed approximation (by using a small (50-100) number of low-frequency eigenfunctions) when dealing with imperfect meshes and/or correspondences. The third option allows a smoothed approximation and speeds up computations in shape collections where all of the meshes are identically tessellated. Indeed, when computing shape differences from one shape to all others in such a collection, the eigenfunction basis is needed on the source mesh only. Due to basis truncation, both of the latter two options result in small shape difference matrices, making joint analysis (e.g. PCA in Section 7.1) feasible.

6 Properties

In this section we discuss a number of properties of shape differences: functoriality, informativeness, and localization and relation to point-wise measures of distortion.

Functoriality: The shape differences behave functorially under map inversions and compositions. To simplify the exposition let us assume that one uses the Laplacian basis on all shapes, and so the formulas (Eq. 4) apply. While we focus on the area based shape differences, similar arguments are valid for the conformal ones.

We start with map inversion: given the maps $F : L^2(M) \rightarrow L^2(N)$ and $F^{-1} : L^2(N) \rightarrow L^2(M)$, we want the induced shape differences satisfy $V_{N,M} = V_{M,N}^{-1}$. This cannot directly hold because the involved matrices are expressed in terms of different bases. We compute $V_{M,N} = F^T F$ and $V_{N,M} = (F^{-1})^T F^{-1}$; next, we need to apply matrix conjugation to the matrix $V_{N,M}$ to transport it into the same basis as $V_{M,N}$; this results in $F^{-1}V_{N,M}F = F^{-1}(F^{-1})^T F^{-1}F = F^{-1}(F^{-1})^T$, which indeed is the inverse of $V_{M,N}$. Thus, after transporting to a common basis, $V_{N,M} = V_{M,N}^{-1}$ holds.

As for map compositions, given the maps $F_1 : L^2(M) \rightarrow L^2(N)$, $F_2 : L^2(N) \rightarrow L^2(K)$ and the composition $F_2 F_1 : L^2(M) \rightarrow$

$L^2(K)$, we want $V_{M,K} = V_{M,N}V_{N,K}$ to hold. We compute $V_{M,N} = F_1^\top F_1$, $V_{N,K} = F_2^\top F_2$, and $V_{M,K} = (F_2 F_1)^\top F_2 F_1$. Note that only the matrix $V_{N,K}$ needs to be transported to the same basis as the other matrices; this requires conjugation by F_1 , and gives the matrix $F_1^{-1} V_{N,K} F_1 = F_1^{-1} F_2^\top F_2 F_1$. Now, the sought equality can be easily seen to hold.

The latter property can be used to speed up computations as follows. Suppose that we have a collection of shapes with functional maps $F_i : L^2(M) \rightarrow L^2(N_i)$, and we compute all of the shape differences V_{M,N_i} . Now, the shape differences between any pair of shapes N_i and N_j , after transporting to M , is given by $V_{N_i,N_j} = V_{M,N_i}^{-1} V_{M,N_j}$.

Informativeness: We have seen in Section 3.2, that $V_{M,N} = I$ and $R_{M,N} = I$, if and only if the underlying maps are area preserving and conformal respectively. Combining this with functoriality properties we have: if $V_{M,N,F} = V_{M,N,G}$ then the map $F^{-1}G$ is area preserving (resp. conformal for R). This means, in particular, that the shape difference matrices *encode* the map up to an area preserving, or conformal self-map. In other words, the shape difference matrices that we define are *fully informative* up to the given notion of distortion.

Localization and relation to existing measures: When the functional map $F : L^2(M) \rightarrow L^2(N)$ is associated with a point-to-point bijection $T : N \rightarrow M$, we can extract local distortion information from the shape difference operators. Let ρ be a compactly supported function on $\Omega \subset M$ (i.e. $\rho(x) = 0$ when $x \notin \Omega$), then $V_{M,N}\rho$ and $R_{M,N}\rho$ only depend on the restricted map $T|_{T^{-1}(\Omega)} : T^{-1}(\Omega) \rightarrow \Omega$. In other words, if the map T is modified outside the region Ω , then $V_{M,N}\rho$ and $R_{M,N}\rho$ would not change.

To prove this for, say, the conformal shape differences, note that for any $f : M \rightarrow \mathbb{R}$, the operator $R_{M,N}$ satisfies $\int_M (\nabla f)(\nabla R_{M,N}\rho) d\mu^M = \int_N \nabla(f \circ T)\nabla(\rho \circ T) d\mu^N = \int_{T^{-1}(\Omega)} \nabla(f \circ T)\nabla(\rho \circ T) d\mu^N$, where the latter equality follows from ρ being supported in Ω , and by bijectivity of T , $\rho \circ T$ being supported within $T^{-1}(\Omega)$. The last expression involves T only in an integral over the region $T^{-1}(\Omega)$, which proves the claim.

This property means that by selecting a function ρ supported within some ROI, we can use $V_{M,N}\rho$ and $R_{M,N}\rho$ as descriptors of distortion happening along this region. As a result, we can make localized comparisons between different maps such as $T_1 : N_1 \rightarrow M$ and $T_2 : N_2 \rightarrow M$; see Section 7.2 for an application.

Additionally, the area based shape differences enjoy the following two properties, see appendix for proofs. Given a region $\Omega \subset M$, and a function $\rho : M \rightarrow \mathbb{R}$ supported within this region, the support of $V_{M,N}\rho$ lies inside Ω . Finally, by letting χ be the indicator function of Ω (i.e. $\chi(x) = 1$ if $x \in \Omega$ and 0 otherwise), we can extract from $V_{M,N}$ the traditional measure of area distortion using the following formula: $\text{area}(\Omega)/\text{area}(T^{-1}(\Omega)) = h_a^M(\chi, \chi)/h_a^M(V_{M,N}\chi, \chi)$.

7 Applications

The shape differences provide a general framework with many potential applications in computer graphics, computer vision, medical imaging, structural biology, and a number of other fields that require precise comparisons between shapes. Here we explore a number of prototype applications involving collections of 3D models.

7.1 Intrinsic Shape Space

Shape differences can be used to explore variability in a collection of related shapes. For this purpose it is important to obtain a common representation that captures the landscape where the models live, to determine the “average” shape, and to visualize where variability happens directly on the shapes. We choose to represent a collection of shapes as a collection of shape differences from one of them, what we term the “base” shape. Since shape differences can be transported to different shapes, the choice of the base shape is not at all critical — it is rather like choosing an arbitrary origin when introducing a coordinate system. We will demonstrate that, by applying Principal Component Analysis (PCA) to these shape differences, we can extract the types of information outlined above.

First, we vectorize the area based and conformal shape difference matrices, apply PCA, and depict each shape’s coefficients along the two largest principal component directions. As usual, the vectors determining the PCA directions are normalized to have unit vector norm. Since we are working with matrices, this is equivalent to the unit Frobenius norm of matrices. Due to this normalization, the xy -coordinate ranges in the PCA plot are commensurable both within and across plots.

Second, we obtain the visualization on the base shape of where shape variability localizes. To this end, we convert the principal components into matrices $\{P_i\}$; the amounts of variances explained are $\{\sigma_i^2\}$. For a given function f on the base shape, let \vec{f} be its vector representation in the basis. After normalizing $\|\vec{f}\| = 1$, it is true that $\|P_i \vec{f}\|$ is small for all $i = 1, 2, \dots$, then the distortion that this function undergoes is similar between all the shapes. Note that we are not discussing the average amount of distortion, but rather the deviation of the distortion from the average. Now, we can define an aggregate amount of this deviation over all the principal components as

$$\sum_i \sigma_i^2 \|P_i \vec{f}\|^2 = \sum_i \sigma_i^2 \vec{f}^\top P_i^\top P_i \vec{f} = \vec{f}^\top M \vec{f},$$

where $M = \sum_i \sigma_i^2 P_i^\top P_i$; here the weighting by variances allows giving more importance to more prominent principal directions. To visualize what regions vary most between different shapes (again not the regions of highest distortion, but of highest variability), for every point p on the base shape, we compute the variance function as $\text{var}(p) = (\vec{f}^\top M \vec{f})^{1/2}/\|\vec{f}\|$ where \vec{f} represents the delta function centered at p . In practice we replace the delta function by the heat kernel computed at a small value of the time parameter.

Our first dataset (first column of Figure 3(a)) is comprised of 64 deformations of the unit sphere obtained by adding two protrusions to the unit sphere using normal displacement. The sizes of protrusions sample an equally spaced two-dimensional grid of values.

The PCA plots for area based and conformal shape differences are shown respectively in the second and third columns of Figure 3(a). These plots uncover the grid structure of the underlying shape space. The fact that the percent of variance explained by each PCA direction (shown in parenthesis for each axis) nearly add up to 100%, recovers the fact that deformations have two degrees of freedom. The percent of variance explained for both PCA directions are almost the same, meaning that deformations in both of the bumps have similar strength range. Since PCA centers data around the average, we can find the average shape for this collection by looking around the origin, which gives shapes 28, 29, 36, and 37 as the average shapes; these are exactly the four shapes in the center of the image depicting the collection (first column of Figure 3(a)). The fourth column of Figure 3(a) is the visualization of variability

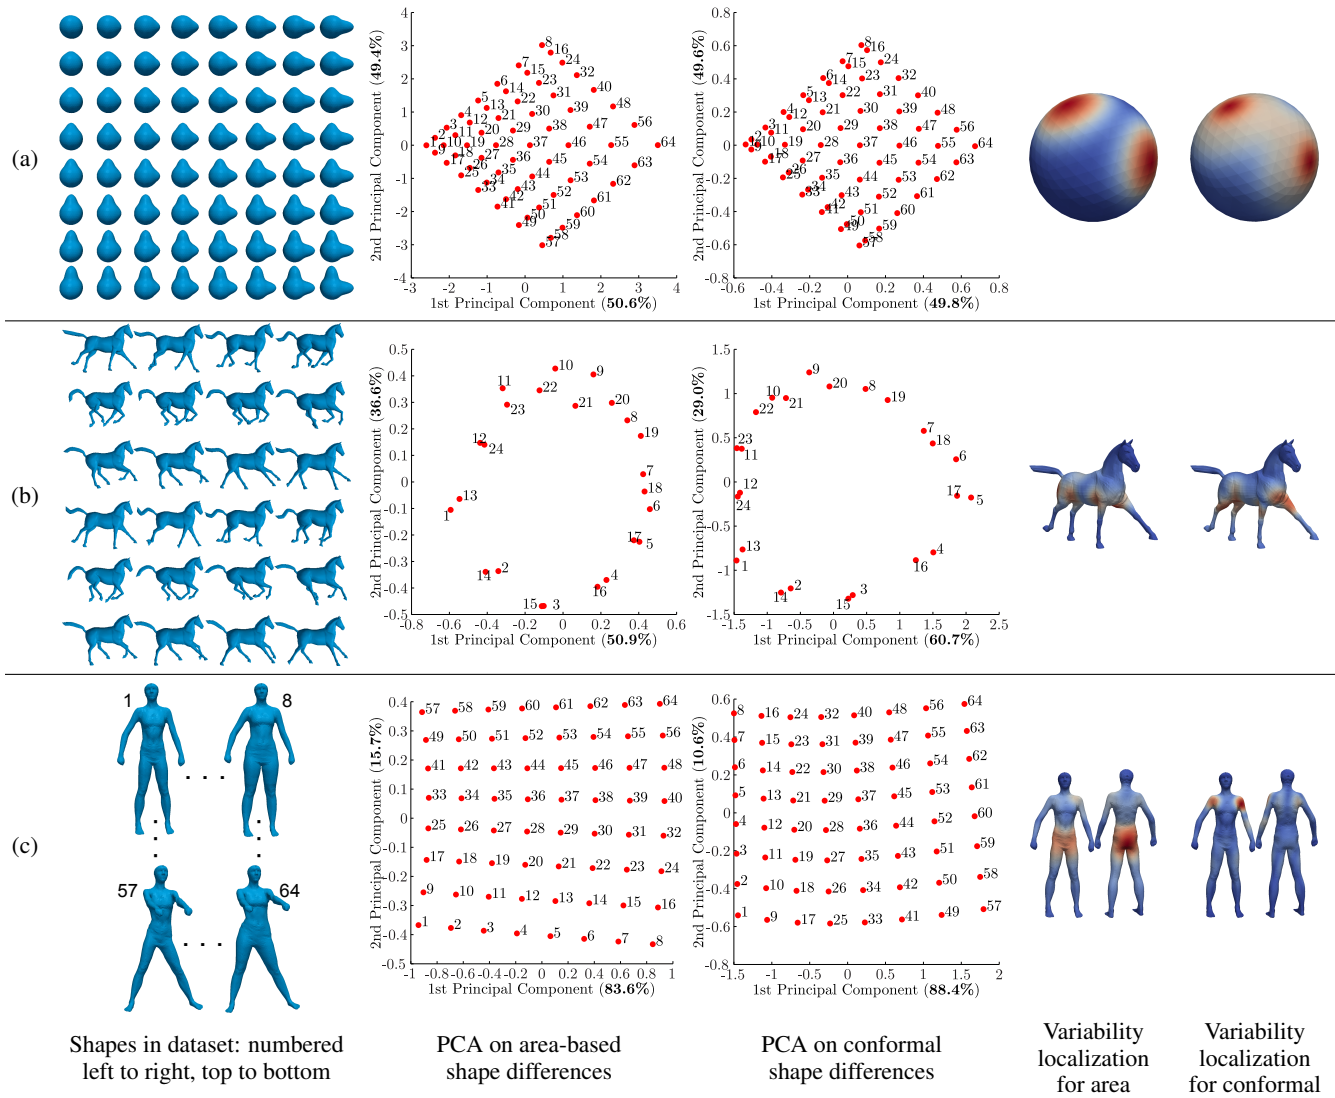


Figure 3: The proposed shape differences reveal the major variability in a collection of shapes as well as the locations (sources) of variability, useful for visualization. Each row represents a different shape collection (left), followed by PCA performed on area-based and conformal shape differences (middle) and visualization of the locations of variability color coded from red for high to blue for low variability (right).

amount (blue is small, red is large) on the base shape for both area and conformal distortions, which are both correctly identified.

Figure 3(b) depicts the results on the galloping horse sequence that was used as the frames in a video produced by Sumner and Popović [Sumner and Popović 2004]. Both of the PCA plots reveal the expected circular “topology” of this dataset, and that this circle is traversed twice. The xy -coordinate range for conformal plot is larger than that of area based PCA plot. This hints that within this collection there is more conformal than area distortion, which is likely due to different parts of body moving relative to each other, inducing higher conformal distortion than area distortion at the joints. Finally, the visualizations of variability in the last column identify the regions of variability correctly.

Figure 3(c) depicts the results on a collection of humans synthetically generated using the tools from [Hasler et al. 2009] as a black-box. This collection involves a combination of pose change and body shape variation. Namely, we sample a grid of shapes with two modes of variation: 1) hips get larger and 2) the person raises arms. Both of the PCA plots recover the two dimensional nature of the underlying shape space. From the placement of mesh id numbers in

these plots we can see that the main PCA directions for area based and conformal shape differences correspond to different modes of variation: area PCA detects hips getting larger as the main PCA direction, whereas the conformal main PCA direction corresponds to the moving arms. This can also be seen by looking at the visualizations of variability in the last column. Namely, we see that the area variability is maximal around the hips, but the conformal variability is concentrated around the shoulders.

The contrast between the area based and conformal shape differences is more dramatically demonstrated by the following experiment. We use a set of 218 shapes representing an approximately conformal deformation of the Stanford bunny into a sphere. The PCA plots for area and conformal based shape differences are overlaid in Figure 4. The conformal shape difference is by two orders of magnitude smaller than the area based shape difference. The presence of two notable PCA directions corresponds to the fact that in the beginning of the sequence the deformation concentrates at the ears of the bunny till they become rounded, and then the front of bunny is rounded.

Since functional maps provide a more general notion than point-




Figure 4: Approximately conformal deformation of a bunny into a sphere. The PCA applied to shape differences confirms the presence of large area distortion in contrast to small conformal distortion.

to-point maps, our framework is applicable in such more general settings. Here, we show an example of applying shape differences in the setting where only fuzzy correspondences between models in the collection are available. A fuzzy correspondence between two surfaces maps every point on the source surface to a probability distribution on the target. The corresponding functional map goes in the opposite direction, and is obtained by convolving a function on the target surface with the probability distributions. We generated a set of such fuzzy maps between the deformed spheres of Figure 3(a) by centering a Gaussian distribution at the corresponding point and adding noise. The first image in Figure 5 shows an example probability distribution for a single point. Despite the noise, our shape difference framework can correctly identify both the major variabilities present in the collection and their locations.

7.2 Exploring Shape Collections

Another step in understanding and using shape collections is being able to track variation of shapes on a finer level. There has been an increasing interest in this topic as evidenced by the recent papers [Ovsjanikov et al. 2011; Kim et al. 2012]. In this subsection we adopt the shape exploration approach of Kim et al.: a user paints one or several regions of interest on the shape, and then the collection is sorted according to the shape similarity within the user specified ROIs.

Our approach is based on the localization property of our shape differences, see Section 6. We pick one of the shapes in the collection as the base shape M . For a given ROI on any of the shapes, let $\vec{\rho}$ be its smoothed characteristic function expressed in terms of the function basis on the base shape. For each shape N_i in the collection, the vectors given by $V_{M,N_i}\vec{\rho}$ and $R_{M,N_i}\vec{\rho}$ carry information about the shape variability within this ROI. We concatenate these

vectors into one, and use it to interactively sort the shape collection by similarity/dissimilarity along the ROI, and for operations such as deformation magnification and interpolation. These experiments were run on the SCAPE dataset [Angelov et al. 2005] which contains 71 poses of the same subject.

Figure 6 shows two examples (separated by whitespace) of faceted exploration of human pose. In each of these examples, the first two rows show the selected ROIs and shapes that are most similar to the given shapes along these ROIs. In the third row, similarity along both of the ROIs is sought.

In addition to faceted browsing above, our approach allows the introduction of new exploration capabilities. First, the user may want to see shapes that undergo several times the magnitude of deformation in N_1 along the ROI relative to the base shape M . To this end, we multiply the localized shape differences $V_{M,N_1}\vec{\rho}$ and $R_{M,N_1}\vec{\rho}$ by the user specified amount, and sort the shapes in the collection according to the proximity to these magnified differences. Figure 7 shows two examples of this capability. In the first example, the ROI containing the knee is painted on a shape with a bent knee. Magnifying this difference with respect to the base pose, means having the knee bent even more. The shapes retrieved by our method indeed have the most severely bent knees in the dataset.

Another novel exploration capability is shape interpolation along an ROI. Given an initial (N_1) and final (N_2) shape, together with an ROI, we compute the vectors $V_{M,N_1}\vec{\rho}$ and $R_{M,N_1}\vec{\rho}$ for the initial shape, and similarly $V_{M,N_2}\vec{\rho}$ and $R_{M,N_2}\vec{\rho}$ for the final shape. Next, we produce equally spaced sample vectors between these initial and final vectors. The shapes having closest vectors to these sample vectors are retrieved. Figure 8 shows two examples of this operation.

As a qualitative comparison to the approach of [Kim et al. 2012], note that Kim et al. ROI exploration is based on rigidly aligning the ROIs of shapes during the search time, whereas our approach uses intrinsic quantities and is based on directly comparing vectors $V_{M,N_i}\vec{\rho}$ and $R_{M,N_i}\vec{\rho}$. Additionally, simply scaling/linearly interpolating these vectors leads to exaggeration and interpolation capabilities, which are not as straightforward to formulate extrinsically.

7.3 Shape Analogies

One of the higher level cognitive operations central to human thought process is constructing analogies. While very difficult to imitate, this operation has received some attention in the image processing context and has led to the concept of image analogies

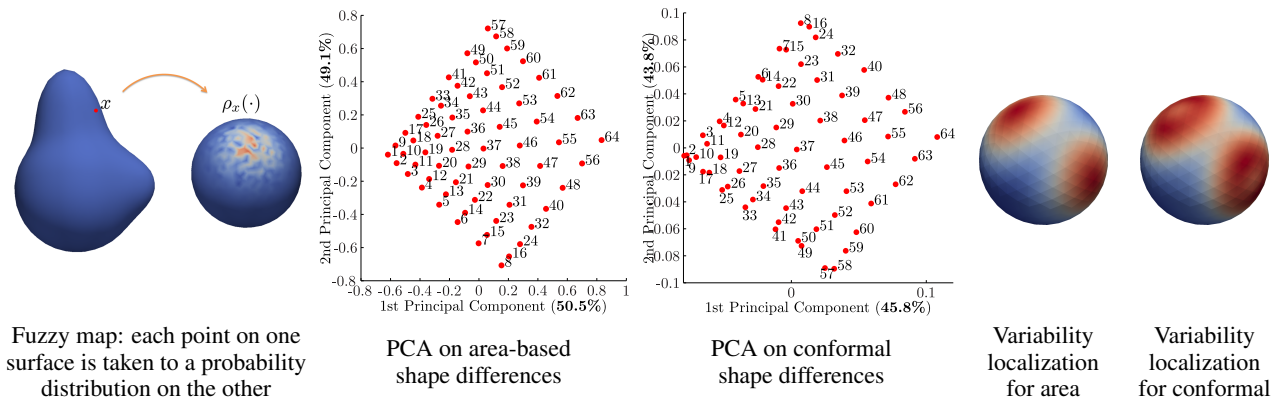


Figure 5: Given fuzzy maps between deformed spheres of Figure 3(a), the shape differences correctly identify the major variability as well as the locations of variability.




Figure 6: Faceted browsing similar to Kim et al. [2012].




Figure 7: Magnification: given a rest pose (leftmost), and an input shape (second) with an ROI, we find shapes that have 3-4 times the distortion at ROI.




Figure 8: Interpolation: given an initial pose (leftmost), and a final pose (rightmost) with an ROI, we find shapes that interpolate between the initial and final poses along the ROI.

[Hertzmann et al. 2001]. In this subsection we show that our shape differences can be used to introduce the notion of *shape analogies*.

Given a pair of 3D models A and B , and another model C , our goal is to retrieve from the collection a shape D such that D relates to C in the same way as B relates to A . In this paper we describe an approach to this problem where “in the same way” is interpreted as having the same or close shape differences. Two sets of experiments will be presented, one when a map between A and C is available, and another when such a map is not available.

We first explain how shape analogies can be obtained when a map between A and C is available. Let be G the corresponding functional map: $G : L^2(A) \rightarrow L^2(C)$. We start by computing the area $V_{A,B}$ and conformal $R_{A,B}$ shape differences between A and B . Next, for every shape X in our collection, we compute the corresponding shape differences $V_{C,X}$, $R_{C,X}$ between C and X . Among all X , we select D as

$$D = \arg \min_X \|V_{C,X}G - GV_{A,B}\|_F^2 + \|R_{C,X}G - GR_{A,B}\|_F^2,$$

here we use the Frobenius norm. Note that to compare shape differences we needed to carry out matrix conjugation for transporting the differences to a common comparison ground, which in this case can be achieved without resorting to inverses (see Section 4).

Figure 9 depicts two examples of analogies constructed using this approach on SCAPE dataset. Here, the map between A and C is known as all of the shapes in the dataset are tessellated in the same way. In the first example (left), as one goes from A to B , the hands get half raised. Therefore, we expect D to differ from C by hand being half raised as well; and indeed our approach retrieves such a pose from SCAPE. The second example involves raising the hands fully, and again our approach succeeds in finding such a pose from SCAPE. Note that the variety of retrieved poses are limited by the dataset being employed.

Figure 10 shows a similar experiment but involving multiple analogies based on the Cats and Lions dataset from [Sumner and Popović 2004]. Here the map between the base cat and base lion is known, the cats are in correspondence and so are the lions. Shape A is the base pose for the cat, and shape B is the base pose for the lion. Then, for multiple shapes C_i (poses of the cat), we find the analogies D_i (poses of the lion). In this case, we embed the V matrices after conjugation in a lower dimensional space using PCA, and compute the distances in this space. Note, that we have recovered all the correct matches between poses.

Figure 11 shows a similar experiment, using human poses from the dataset of [Hasler et al. 2009]. Here, as opposed to the SCAPE




Figure 9: Shape analogies in SCAPE: given shapes A , B , and C , we find a shape D such that the shape difference D to C is close to the shape difference B to A .

dataset, we have two humans (male and female) in different poses. Shapes A and B are the poses shown in blue in the Figure. Similarly to the cats and lions experiment, we retrieve multiple analogies, by picking a different pose C_i for the male, and extracting the matching pose D_i of the female. In this case, the dataset contains the ground-truth (which pose match) and, although matched poses are relatively different, due to human interpretation of the pose instructions given, e.g. C_1 and D_1 , we still recover the ground truth.

Our final set of experiments considers the case where a cross collection map is not available. In this case, we cannot use the conjugation method in order to bring the shape differences to the same common ground, and therefore we need to use a descriptor of the difference which is invariant to matrix conjugation. We chose the singular values of the V matrix as such a descriptor. Note, that now we have considerably less information than when a cross map is available, and therefore we need to regularize our experiment by exploiting more of the data. Here we assume that we have to “parallel” shape collections with corresponding shape variants. Instead of simply fixing A , B , C and searching for the best D such that “ D is to C like B is to A ”, we find the best permutation of the shapes which simultaneously best aligns all the shapes in one collection with their counterparts in the other collection. Namely, given n shapes, we compute the singular values of all the possible pairwise maps, and use that as a descriptor. We compare this descriptor to the singular values of all the possible pairwise maps in the other collection. To be more precise, for every one of the possible $n!$ permutations we can compute a score measuring the agreement between the descriptors. This provides a ranking of the permutations, which allows us to find all the analogies simultaneously.

Figure 12 shows the result of this experiment for two collections taken from the dataset in [Hasler et al. 2009]. The first collection is a group of humans in one pose, and the second is the same group of humans in a different pose. Note that, as opposed to the experiment shown in Figure 11 where we find variation between poses, here we find variation between humans, which is much more delicate. The figure shows in the top row the first collection, and in the second and third rows the best two permutations. The shapes are colored according to the ground truth, namely the same human has the same color in all rows. As is evident in the figure, in the best permutation we found the ground truth map, and aligned correctly the shape collections. The third row is also meaningful, as the second best permutation switches between the two females in the group, whose differences are more subtle compared to the other humans. Note that although the map between the two collections is available, it was not used in this experiment.

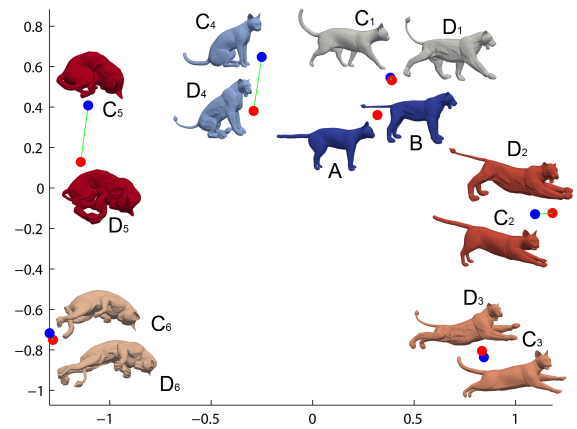


Figure 10: Multiple shape analogies between the cats and the lions. We find all the analogies — recovering the relationship between the poses of the cat and the lion.

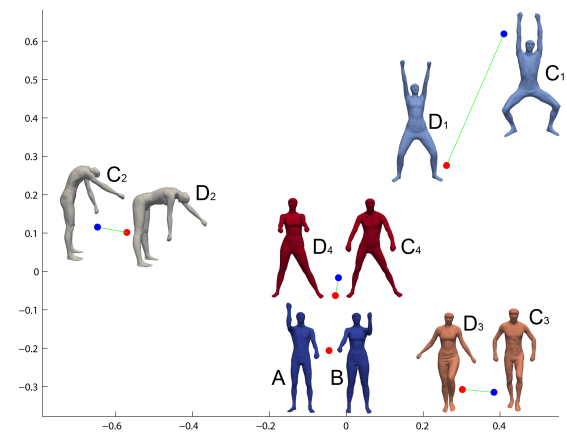


Figure 11: Multiple analogies recover the corresponding poses of the male and the female.

Figure 13 repeats this experiment for a subset of the shapes from the TOSCA dataset [Bronstein et al. 2008]. There, poses of the cat and the dog are not marked as corresponding, and in fact there are various changes in the pose between the cat and the dog — for example the tails are geometrically quite different. As in the previous figure, we show the first collection, followed by the color coded best first and second permutations. Again, we can see that the best permutation matched correctly between the poses, and the second best permutation confused between similar shapes. This demonstrates that the important information we recover is the relative change in pose between the cats within themselves and the dogs within themselves, and there is no requirement for geometric similarity between a cat and a dog.

Finally, note that Figure 1 shows another example of such space shape alignment where a cross collection map is not available.

Compute time: We report timings on a laptop with 2.4GHz Intel Core i7-2760QM processor and 4GB memory. The human dataset in Figure 3(c) contains 64 shapes, each with 6.5K vertices; we use Laplace-Beltrami bases on all of the shapes (Option 2 of Section 5). It takes about 300s to discretize the Laplacians and to compute the bases (64 eigenfunctions on the base shape, and 192 on remaining shapes). Additional 20s are required to compute functional maps and shape difference matrices (of size 64×64) between the base shape and all others. The bunny dataset (Figure 4) contains 218



Figure 12: Simultaneous analogies between all pairs in two collections (five humans in two poses), without a cross map. The color coding indicates the ground truth. The first row shows the first collection, and the second and third show the best and second best permutations respectively. Note that in the first permutation we have recovered the ground truth, and the second permutation swapped between the two females in the group, whose differences are more subtle compared to the other humans.




Figure 13: Simultaneous analogies between all pairs in two collections (five cats and five dogs from the TOSCA dataset), without a cross map. Note that our best permutation (2nd row) recovers the ground truth, and the second best permutation (3rd row) swaps between two dogs in similar poses.

shapes, each with 14K vertices. We use eigenfunctions on the base shape, and hat functions on all other shapes (Option 3 of Section 5). Discretizing Laplacians of all the meshes and computing 64 eigenfunctions on the base shape takes about 40s. Additional 10s are spent computing functional maps and shape differences (of size 64×64) between the base shape and all others. These timings are typical for all of the experiments reported in this paper.

8 Conclusion, limitation, and future work

We have shown a new formulation of the concept of shape differences and developed its mathematical properties. We also studied how the formulation can be discretized and how shape differences can be robustly and efficiently computed on meshed shapes. Shape differences can themselves be compared and used to study shape collections in new ways, allowing us to perform shape analogies or to localize or parametrize the variability in a collection.

A limitation of our method is that it computes shape differences based on externally supplied maps between shapes, and it therefore depends on the quality of these maps. Current intrinsic map technology can handle well isometric and nearly isometric shapes with some extensions to more general classes of deformations. As mapping techniques improve, we expect to be able to apply our shape

difference methods to more general classes of shapes as well.

In general, the quality of the information we get from our approach depends on the quality and density of the shape maps we start with — in other words, the quality by which we know the network defining the shape inter-relationships. What assumptions on this quality allow us to extract what type of underlying shape space structure is a topic that requires further investigation.

Spectral tools can be used on shape networks to extract functionals on the shapes that are most preserved or most variable over the network. Such functionals may aid various classification tasks, separating important from nuisance modes of variability. A lot also remains to be done on the topic of exploiting shape differences to align shape collections that only have parts with parallel structure, both in the presence and in the absence of cross maps.

Acknowledgments: We would like to thank Hao Li for the models in Figure 1, and Keenan Crane for the bunny sequence used in Figure 4. We are grateful to Max Wardetzky and Henrik Schumacher for helping with the proof of Theorem 1. The authors would like to acknowledge NSF grants CCF 1011228 and DMS 1228304, AFOSR grant FA9550-12-1-0372, a Google research award, ISF grant 699/12, ISF equipment grant, Marie Curie CIG 303511, the state of Lower-Saxony and the Volkswagen Foundation, Hannover, Germany, ANR project GIGA (ANR-09-BLAN-0331-01), the INRIA Associated team COMET, and CNRS chaire d'excellence.

References

- ALLEN, B., CURELESS, B., AND POPOVIĆ, Z. 2003. The space of human body shapes: reconstruction and parameterization from range scans. *ACM Trans. Graph.* 22, 3 (July), 587–594.
- ANGUELOV, D., SRINIVASAN, P., KOLLER, D., THRUN, S., RODGERS, J., AND DAVIS, J. 2005. Scape: shape completion and animation of people. *ACM Trans. Graph.* 24, 3, 408–416.
- BEN-CHEN, M., WEBER, O., AND GOTSMAN, C. 2009. Variational harmonic maps for space deformation. *ACM Trans. Graph.* 28, 3 (July), 34:1–34:11.
- BOOKSTEIN, F. 1996. Shape and the information in medical images: A decade of the morphometric synthesis. In *Proc. IEEE MMBIA*, 2–12.
- BREZIS, H. 2010. *Functional Analysis, Sobolev Spaces and Partial Differential Equations*. Springer-Verlag.
- BRONSTEIN, A., BRONSTEIN, M., AND KIMMEL, R. 2006. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS* 103, 5, 1168–1172.
- BRONSTEIN, A., BRONSTEIN, M., AND KIMMEL, R. 2008. *Numerical Geometry of Non-Rigid Shapes*. Springer.
- COOTES, T., TAYLOR, C., ET AL. 2001. Statistical models of appearance for medical image analysis and computer vision. In *Proc. SPIE Medical Imaging*, vol. 4322, 236–248.
- DRYDEN, I. L., AND MARDIA, K. V. 1998. *Statistical Shape Analysis*. John Wiley and Sons.
- FARRELL, R., BRANSON, S., AND WELINDER, P., 2011. First Workshop on Fine-Grained Visual Categorization (FGVC) at CVPR 2011. <http://www.fgvc.org/>.
- FLETCHER, P., LU, C., PIZER, S., AND JOSHI, S. 2004. Principal geodesic analysis for the study of nonlinear statistics of shape. *Medical Imaging, IEEE Transactions on* 23, 8, 995–1005.

- GERIG, G., STYNER, M., SHENTON, M., AND LIEBERMAN, J. 2001. Shape versus size: Improved understanding of the morphology of brain structures. In *Proc. MICCAI*, 24–32.
- GOLLAND, P., GRIMSON, W., SHENTON, M., AND KIKINIS, R. 2005. Detection and analysis of statistical differences in anatomical shape. *Medical Image Analysis* 9, 1, 69–86.
- HASLER, N., STOLL, C., SUNKEL, M., ROSENHANN, B., AND SEIDEL, H. 2009. A statistical model of human pose and body shape. *Computer Graphics Forum* 28, 2, 337–346.
- HERTZMANN, A., JACOBS, C. E., OLIVER, N., CURLESS, B., AND SALESIN, D. H. 2001. Image analogies. In *Proc. SIGGRAPH*, 327–340.
- KASS, M., WITKIN, A., AND TERZOPOULOS, D. 1988. Snakes: Active contour models. *Int. J. Comput. Vision* 1, 4, 321–331.
- KILIAN, M., MITRA, N. J., AND POTTMANN, H. 2007. Geometric modeling in shape space. In *Proc. SIGGRAPH*, 64:1–64:8.
- KIM, V. G., LIPMAN, Y., AND FUNKHOUSER, T. 2011. Blended intrinsic maps. *ACM Trans. Graph.* 30, 4 (July), 79:1–79:12.
- KIM, V. G., LI, W., MITRA, N. J., DIVERDI, S., AND FUNKHOUSER, T. 2012. Exploring collections of 3d models using fuzzy correspondences. *ACM Trans. Graph.* 31, 4 (July), 54:1–54:11.
- LAI, R., SHI, Y., SCHEIBEL, K., FEARS, S., WOODS, R., TOGA, A., AND CHAN, T. 2010. Metric-induced optimal embedding for intrinsic 3d shape analysis. In *CVPR*, 2871–2878.
- NAIN, D., STYNER, M., NIETHAMMER, M., LEVITT, J., SHENTON, M., GERIG, G., BOBICK, A., AND TANNENBAUM, A. 2007. Statistical shape analysis of brain structures using spherical wavelets. In *Proc. ISBI*, 209–212.
- OVSJANIKOV, M., LI, W., GUIBAS, L., AND MITRA, N. J. 2011. Exploration of continuous variability in collections of 3d shapes. *ACM Trans. Graph.* 30, 4 (July), 33:1–33:10.
- OVSJANIKOV, M., BEN-CHEN, M., SOLOMON, J., BUTSCHER, A., AND GUIBAS, L. 2012. Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.* 31, 4 (July), 30:1–30:11.
- OVSJANIKOV, M., BEN-CHEN, M., CHAZAL, F., AND GUIBAS, L. 2013. Analysis and visualization of maps between shapes. *Computer Graphics Forum*. To appear.
- PINKALL, U., AND POLTHIER, K. 1993. Computing discrete minimal surfaces and their conjugates. *Exp. Math.* 2, 1, 15–36.
- SAHILIOĞLU, Y., AND YEMEZ, Y. 2011. Coarse-to-fine combinatorial matching for dense isometric shape correspondence. *Computer Graphics Forum* 30, 5, 1461–1470.
- SCHAEFER, S., MCPHAIL, T., AND WARREN, J. 2006. Image deformation using moving least squares. *ACM Trans. Graph.* 25, 3 (July), 533–540.
- SCHUMACHER, H. 2013. Conformal maps and p -Dirichlet energies. Tech. rep., Univ. of Göttingen.
- SHARMA, A., AND HORAUD, R. 2010. Shape matching based on diffusion embedding and on mutual isometric consistency. In *Proc. NORDIA Workshop (CVPR)*, 29–36.
- SUMNER, R. W., AND POPOVIĆ, J. 2004. Deformation transfer for triangle meshes. *ACM Trans. Graph.* 23, 3 (Aug.), 399–405.
- VASILESCU, M., AND TERZOPOULOS, D. 2007. Multilinear (tensor) ICA and dimensionality reduction. *Independent Component Analysis and Signal Separation*, 818–826.
- WANG, Y., ZHANG, J., GUTMAN, B., CHAN, T., BECKER, J., AIZENSTEIN, H., LOPEZ, O., TAMBURO, R., TOGA, A., AND THOMPSON, P. 2010. Multivariate tensor-based morphometry on surfaces: Application to mapping ventricular abnormalities in hiv/aids. *Neuroimage* 49, 3, 2141–2157.
- WUHRER, S., SHU, C., AND XI, P. 2012. Posture-invariant statistical shape analysis using laplace operator. *Computers & Graphics* 36, 5, 410–416.

Appendix

Proof of Theorem 1. Note that in the discrete case this theorem simply states that (1) a bijection is area preserving if and only if the diagonal area matrices equal: $A^M = A^N$ and (2) is conformal if and only if the weight matrices equal $W^M = W^N$, which follows simply because the cotangent function is bijective on $(0, \pi)$.

In the continuous case, we are given two surfaces M and N , and a bijective map $T : N \rightarrow M$, such that $\int_{x \in M} f(x)g(x) d\mu^N(x) = \int_{x \in M} f(T(x))g(T(x)) d\mu^N(T(x)), \forall f, g : M \rightarrow \mathbb{R}$. By considering $f = g$ indicator functions on a region $\Omega \subseteq N$ it follows that $\text{Area}(\Omega) = \text{Area}(T(\Omega))$, meaning that T is locally volume-preserving. Conversely, suppose T is volume-preserving. By definition this implies that for any function $h: \int_{x \in M} h(x) d\mu^N(x) = \int_{x \in M} h(x) d\mu^N(T(x))$. Setting $h(x) = g(x)f(x)$ we get that T preserves the area-based inner product.

Suppose now that T is a conformal map between two surfaces. To simplify notation, suppose that we are simply given two Riemannian metrics G_1, G_2 on M such that $G_2 = \lambda^2 G_1$ for some smooth function λ . From local coordinate expression of the Laplace-Beltrami operator, one has $\Delta_{G_2} f = \frac{1}{\lambda^2} \Delta_{G_1} f$. Moreover, the volume form $d\text{Vol}_{G_2} = \lambda^2 d\text{Vol}_{G_1}$, and using Stokes' formula we obtain that $\int_M \langle \nabla_{G_1} f, \nabla_{G_1} g \rangle d\text{Vol}_{G_1} = - \int_M \Delta_{G_1} f g d\text{Vol}_{G_1} = - \int_M \Delta_{G_2} f g d\text{Vol}_{G_2} = \int_M \langle \nabla_{G_2} f, \nabla_{G_2} g \rangle d\text{Vol}_{G_2}$ for any pair of functions f, g . The converse — preservation of conformal-based inner product implies that the map is conformal — is proved in [Schumacher 2013].

Proof of Theorem 2. This theorem is a consequence of general Riesz representation theorem [Brezis 2010]. Namely, the presence of the linear functional map allows us to think of the inner product on shape N as simply another inner product on M . Thus, given two inner products h_1, h_2 on M our goal is to show that there exists a unique linear map D , such that: $h_1(f, g) = h_2(f, Dg), \forall f, g$. To see this note that when g is fixed, $h_1(\cdot, g)$ defines a continuous linear functional on the functional space. By the Riesz-Frechet Theorem, there exists a unique element $D(g)$ of this space such that $h_1(f, g) = h_2(f, D(g)), \forall f$. Moreover, $D(\cdot)$ is bijective and linear by linearity of $h_1(f, \cdot)$.

Proof of Two Properties of V . Let $T : N \rightarrow M$ be the underlying bijective map, and let ρ be supported in $\Omega \subset M$. For any $f \in L^2(M)$, the operator $V_{M,N}$ satisfies $\int_M f V_{M,N} \rho d\mu^M = \int_N (f \circ T)(\rho \circ T) d\mu^N$. Taking for f any function with support $\Omega' \subset \Omega^c$, gives $\int_{\Omega'} f V_{M,N} \rho d\mu^M = 0$, and so the support of $V_{M,N} \rho$ has to be contained in Ω up to a set of measure 0. This proves the first localization property for $V_{M,N}$. Now taking $f = \rho = \chi$, the characteristic function of Ω , gives $h_a^M(V_{M,N} \chi, \chi) = \int_N (\chi \circ T)(\chi \circ T) d\mu^N = \text{area}(T^{-1}(\Omega))$. As $h_a^M(\chi, \chi) = \text{area}(\Omega)$, this proves the relationship with the traditional measure of area distortion.