Efficient deformable shape correspondence via multiscale spectral manifold wavelets preservation

Ling Hu[*1,2], Qinsong Li[*1], Shengjun Liu[†1,3], Xinru Liu[1]
[1]Institute of Engineering Modeling and Scientific Computing, Central South University
[2]School of Mathematics and Computational Science, Hunan First Normal University
[3]State Key Laboratory of High Performance Manufacturing Complex, Central South University

Abstract

The functional map framework has proven to be extremely effective for representing dense correspondences between deformable shapes. A key step in this framework is to formulate suitable preservation constraints to encode the geometric information that must be preserved by the unknown map. For this issue, we construct novel and powerful constraints to determine the functional map, where multiscale spectral manifold wavelets are required to be preserved at each scale correspondingly. Such constraints allow us to extract significantly more information than previous methods, especially those based on descriptor preservation constraints, and strongly ensure the isometric property of the map. In addition, we also propose a remarkable efficient iterative method to alternatively update the functional maps and pointwise maps. Moreover, when we use the tight wavelet frames in iterations, the computation of the functional maps boils down to a simple filtering procedure with low-pass and various band-pass filters, which avoids time-consuming solving large systems of linear equations commonly presented in functional maps. We demonstrate on a wide variety of experiments with different datasets that our approach achieves significant improvements both in the shape correspondence quality and the computing efficiency.

1. Introduction

Finding correspondences between non-rigid shapes is a fundamental problem in computer graphics, vision and pattern recognition, with countless applications including shape comparison, texture transfer and shape interpolation, just to name a few [37].

Since many real-world deformations are approximately isometric, finding near isometric maps possesses practical significance in non-rigid shape analysis and it has been well studied in the last decades. Among extensive strategies for this task, one of the most influential techniques in recent years is the functional map framework, originally provided by [24] and extended by plenty of fellow-up works [29, 31, 28, 30]. This framework proposed to determine a functional map operator, that maps between the spaces of square-integrable functions on the respective shapes, then a high-quality pointwise correspondence can be recovered from the correspondence of special functions. The most attractive property of functional maps is finding correspondence boils down to a quite simple and efficient algebraic problem.

A key step in the functional map framework is to formulate function preservation constraints to a linear system of equations. These function preservation constraints typically encode the information (e.g., geometric or appearance) that must be globally preserved by the unknown map. In isometric shape correspondence, the commonly used constraints include those aiming at approximately aligning the given descriptor functions and commuting with the Laplace-Beltrami operators as well [24, 27, 8], or more

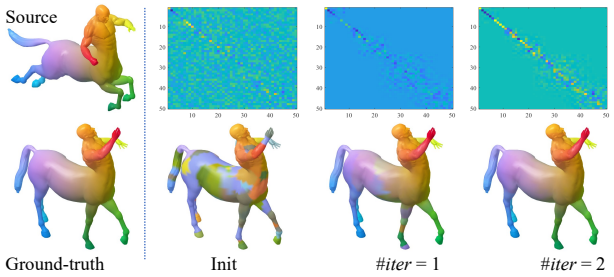

Figure 1. Demonstration of our efficient shape correspondence iterative updating. Starting with a noisy functional map between two isometric deformation shapes, we alternately optimize the functional maps and the pointwise maps. The pointwise maps are visualized by color transfer. Note that iterating twice can reach significant improvements even with a quite noisy initialization and costs 0.54 *seconds* for triangle meshes with $16K$ vertices.

[*]Joint first author.
[†]Corresponding author: shjliu.cg@csu.edu.cn (S.Liu).

powerful descriptor preservation constraints via commutativity [23] and manifold optimization [14]. However, the above descriptor based constraints still suffer severe disadvantages in correspondence quality or computing efficiency, due to the difficulties of obtaining a large set of informative and linearly independent descriptor functions and the insufficient encoding for isometric properties of deformed shapes.

Contributions. To address the above-mentioned challenges in the recent functional maps, in this paper, we concentrate on exploiting novel constraints and efficient computation ways into the pipeline of the functional maps. Our inspirations come from the powerful spectral manifold (or graph) wavelets (SMWs) [11, 15], which are constructed with a similar manner to the classical wavelets, still via signal filtering in the frequency spaces of manifolds with a set of suitable filters (functions). SMWs inherit most of the powerful properties of the classical wavelets, such as space-frequency locality, multiscale characteristics, etc, and can be computed efficiently.

The core idea of our work is that we impose SMWs preservation constraints on the functional maps, where the multiscale SMWs are required to be preserved at each scale correspondingly. This strategy resorts to the attractive properties of SMWs in non-rigid shape analysis tasks, mainly including (1) SMWs still hold the intrinsic invariance under shape isometric deformations. (2) They can efficiently and flexibly encode shape features of different frequency bands, ranging from the geometric features of the immediate neighborhoods around each point to the neighborhood with greater diffusion radii. Therefore, via a few scales of intrinsic SMWs located at each point on the shapes preserved during the mapping, the isometric properties of the maps can be sufficiently extracted and strongly ensured, while maintaining the overall linear system nature of the functional map framework. As opposed to the high-dimensional descriptors or additional regularizers in the previous works, our constraints are much simpler, more compact and efficient.

In addition, we also propose a remarkable efficient iteration way to alternatively update the computation of the functional maps and pointwise maps, where high-quality correspondences generally can be reached with only two or three iterations, as shown in Figure 1. More specifically, we show that if using suitable SMWs, more exactly, tight wavelet frames, the computation of the functional mapping matrices boils down to a simple iterative filtering procedure by various band-pass filters. This way eventually avoids the time-consuming solution of a large system of linear equations commonly presented in previous works and leads the matrices closer to the structures of isometric pointwise maps. We demonstrate on a wide range of experiments that our method helps to obtain better correspondences, especially with significantly higher computation efficiency.

2. Related works

Shape correspondence is a very well-studied area of computer graphics. Readers can be referred to the survey [34] for an in-depth view of this field. Below we review the methods most closely related to ours, mainly including spectral and iterative methods.

Point-based spectral methods. A traditional approach to correspondence problems is finding a pointwise matching between the points on two or more shapes. One prominent spectral matching strategy involves the minimization of a distortion criterion from pointwise spectral descriptor similarity [36, 3, 6, 16] or descriptors with additional integrated pairwise relations aiming for more desirable properties for the matching [38]. Some spectral methods were based on directly optimizing between spectral shape embeddings based on either adjacency or Laplacian matrices of graphs and triangle meshes [26]. Recently, these approaches have been generalized by spectral generalized multidimensional scaling [1], which explicitly formulates the minimum distortion of shape correspondence in the spectral domain. Although these approaches are mature and can often result in high-quality mappings, they still often lack flexibility and are not suited well in the presence of more general non-rigid deformations.

Functional maps. Functional map framework was originally introduced by [24] and have been extended significantly in [23, 22, 29] (see an overview in [25]) recently. These methods are based on the notion that it is often easier to obtain correspondences between functions, rather than points, Then, a good pointwise correspondence can be extracted from the functional maps [31]. Compared with the point-based methods, functional map framework has a particular advantage of being flexible and allows us to easily incorporate constraints including preservation of geometric quantities (descriptors), thus several works try to formulate more powerful constraints in functional maps to reach desirable properties, including the partial map [29, 17], more powerful descriptor constraints via commutativity [23], manifold optimization[14], continuous and orientation-preserving correspondences[28], product preservation [22], among many others. Differing from these works commonly relying on a large set of high-quality descriptor functions, we instead use the intrinsic wavelets preservation as constraints in our functional map framework, which could encode multiscale geometric features into the mapping while keeping high computation efficiency, and lead the maps closer to an isometric one.

Particularly our work is also closely related to a series of map refinement methods, originally proposed by Iterative closest point (ICP) [24]. The follow-up work BCICP[28] introduced a refinement step in both the spectral and spa-

tial domains, aiming to strongly promote the bijectivity and continuity of the point-wise maps. ZoomOut [19] and its extension consistent ZoomOut[12] recovered a higher resolution map from a lower resolution one through a simple and efficient iterative spectral up-sampling technique. Instead of progressively treating the frequency information like ZoomOut, our method processes overall frequency band information encoded in corresponding wavelet functions in each iteration.

3. Background and notation

We model a shape as a connected smooth compact two-dimensional manifold $\mathcal{M}$ (possible with boundaries) embedded into $\mathbb{R}^3$. The space of square-integrable functions on the manifold $\mathcal{M}$ are denoted by $\mathcal{L}^2(\mathcal{M}) = \left\{f : \mathcal{M} \to \mathbb{R}, \int_{\mathcal{M}} f(x)^2 dx < \infty\right\}$, where dx is the area element induced by the Riemannian metric and $< f, g >_{\mathcal{M}} = \int_{\mathcal{M}} f(x)g(x)dx$ expresses the standard inner product on $\mathcal{M}$. Any function $f(x) \in \mathcal{L}^2(\mathcal{M})$ can be equipped with an Laplace–Beltrami operator (LBO) (or Laplacian), defined as $\Delta_{\mathcal{M}} f(x) = -\mathrm{div}_{\mathcal{M}}(\nabla_{\mathcal{M}} f(x))$, where $\nabla_{\mathcal{M}}$ and $\mathrm{div}_{\mathcal{M}}$ are the intrinsic gradient and divergence respectively.

3.1. Signal processing on manifolds

As the LBO is positive-semidefinite, it admits a real eigen-decomposition $\Delta_{\mathcal{M}}\phi_i(x) = \lambda_i\phi_i(x), i = 0, 1, ...,$ with non-negative eigenvalues (or spectrum) $0{=}\lambda_0 \leq \lambda_1 \leq \ldots$ and the set of the eigenfunctions $\{\phi_i(x)\}_{i\geq 0}$ forms an complete orthonormal basis for $\mathcal{L}^2(\mathcal{M})$, w.r.t the standard inner product. Especially, the eigenvalues $\{\lambda_i\}_{i\geq 0}$ and the corresponding eigenfunctions $\{\phi_i(x)\}_{i\geq 0}$ of the LBO show obvious harmonic properties that they respectively act the roles of frequency and Fourier basis for the space $\mathcal{L}^2(\mathcal{M})$, with a order from low to high frequency. Therefore, any function $f \in \mathcal{L}^2(\mathcal{M})$ can be expressed as $f(x) = \sum_{i\geq 0} \langle f, \phi_i\rangle_{\mathcal{M}} \phi_i(x)$ and the product $\langle f, \phi_i\rangle_{\mathcal{M}}$ is called (manifold) Fourier coefficient.

The harmonic properties of $\{\lambda_i\}_{i\geq 0}$ and $\{\phi_i(x)\}_{i\geq 0}$ allow us to further define manifold filtering operator from the frequency space, i.e., given a suitable filter $g(\lambda)$, the filtering result of the function f is $\hat{f}(x) = \sum_{i\geq 0} g(\lambda_i)\langle f, \phi_i\rangle_{\mathcal{M}}\phi_i(x)$.

3.2. Functional Maps

Definition. Functional maps [24] proposed a very efficient pipeline to transfer functions between shapes and can be used to realize high-quality shape correspondence.

Let $T : \mathcal{M} \to \mathcal{N}$ be a pointwise map between two manifolds $\mathcal{M}$ and $\mathcal{N}$. Then, consider a functional map (a linear operator) $T_F : \mathcal{L}^2(\mathcal{M}) \to \mathcal{L}^2(\mathcal{N})$ that maps functions from $\mathcal{M}$ to $\mathcal{N}$, here the image of T_F is defined as $T_F(f) = f \circ T^{-1}, f \in \mathcal{L}^2(\mathcal{M})$. Assuming to be given two orthonormal bases $\{\phi_i^{\mathcal{M}}\}_{i\geq 0}$ and $\{\phi_j^{\mathcal{N}}\}_{j\geq 0}$ of $\mathcal{L}^2(\mathcal{M})$ and $\mathcal{L}^2(\mathcal{N})$ respectively, the functional map T_F, admits a matrix representation $\mathbf{C} = (c_{ji})$, the image and entries of which are

$$T_F(f) = \sum_j \sum_i \langle f, \phi_i^{\mathcal{M}}\rangle_{\mathcal{M}} \underbrace{\langle T_F(\phi_i^{\mathcal{M}}), \phi_j^{\mathcal{N}}\rangle_{\mathcal{N}}}_{c_{ji}} \phi_j^{\mathcal{N}}. \tag{1}$$

Note that, owing to be orthogonal and harmonic, the eigenfunctions of the LBO are commonly used as the bases of functional maps. As suggested in [24], actually small truncated subsets of the eigenfunctions (i.e., the first k) can approximate the map T_F significantly well. This choice usually makes the matrix $\mathbf{C} \in \mathbb{R}^{k\times k}$ small and has the low-pass effect to produce smooth correspondence.

In fact, T_F in (1) amounts to a linear transformation of the Fourier coefficients of f from basis $\{\phi_i^{\mathcal{M}}\}_{i\geq 0}$ to $\{\phi_j^{\mathcal{N}}\}_{j\geq 0}$. Let $\mathbf{a}, \mathbf{b} \in \mathbb{R}^k$ be the Fourier coefficient vectors of f and its map $T_F(f)$ respectively, then, the functional mapping in (1) boils down to a linear equation $\mathbf{b} = \mathbf{C}\mathbf{a}$. Obviously after fixing the bases of the functional spaces, determining the unknown matrix $\mathbf{C}$ becomes the most important step as $\mathbf{C}$ fully encodes the original map T_F.

Computation of the matrix $\mathbf{C}$. To determine the unkown $\mathbf{C}$, we often generally formulate many linear constraints on $\mathbf{C}$ that allows us to recover the functional map by solving a least squares system. Assuming to be given a set of q pairs of corresponding functions $\{f_1, \cdots, f_q\} \in \mathcal{L}^2(\mathcal{M})$ and $\{g_1, \cdots, g_q\} \in \mathcal{L}^2(\mathcal{N})$, possibly derived from the constraints that correspond to descriptor and segment preservation together with the operator commutativity, we store the Fourier coefficients of these constraint functions as columns of the matrices $\mathbf{A}$ and $\mathbf{B}$ (both of size $k \times q$) respectively, then incorporate them into a linear system of equations and find the best $\mathbf{C}$ in the least squares sense, i.e.

$$\mathbf{C} = \arg\min_{\mathbf{C}} \|\mathbf{C}\mathbf{A} - \mathbf{B}\|_{\mathrm{F}}^2 + \alpha \|\Delta_{\mathcal{N}}\mathbf{C} - \mathbf{C}\Delta_{\mathcal{M}}\|_{\mathrm{F}}^2 .$$

If necessary, the solution can be refined using a certain post-processing iterative refinement.

Pointwise map recovery. As stated in [24], the original pointwise map T indeed can be recovered from T_F well. More detailedly, for each point x on $\mathcal{M}$, its correspondence on $\mathcal{N}$ is treated as $T(x) = \mathrm{argmax}_y T_F(f)(y)$, where $f = \delta_x$ is the delta-function at the point x on the shape $\mathcal{M}$. Since $\langle \delta_x, \phi_i\rangle = \phi_i(x)$, in discrete settings, if let the columns of the matrices $\Phi_{\mathcal{M}}$ and $\Phi_{\mathcal{N}}$ respectively correspond to the first k eigenfunctions of the LBOs of the shape $\mathcal{M}$ and $\mathcal{N}$, then each row of them corresponds to a Fourier coefficient vector of each point. According to Plancherel's theorem that the distance between the Fourier coefficient vectors of the functions is equal to the L_2 difference between the functions themselves, we can efficiently recover

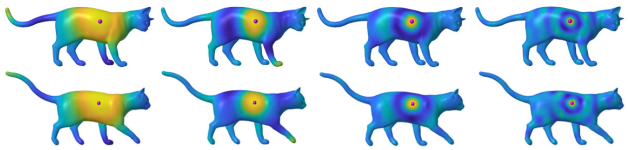

Figure 2. Isometric deformation invariance of SMWs. Several pairs of multiscale SMWs are demonstrated on two cat shapes related with isometric deformations. Each wavelet pair are respectively located at the point (the top row) and its isometric mapping image (the second row). Obviously they are isometric deformation invariant.

the pointwise map T via a nearest neighbour searching between the column vectors of $\mathbf{C}\Phi_{\mathcal{M}}^{\mathrm{T}}$ and $\Phi_{\mathcal{N}}^{\mathrm{T}}$.

4. Method

Despite plenty of observations about the functional representation $\mathbf{C}$ associated with an isometric pointwise map T have been explored [24, 33, 19] (e.g. the orthogonality of $\mathbf{C}$) recently, finding suitable and sufficient constraints on $\mathbf{C}$ to ensure the isometry of T and the computation efficiency still seems to be a difficult and open problem. For this reason, motivated by the remarkable properties of SMWs, we intend to employ them as preservation constraints to generate functional maps and also consider an efficient computation method. Most importantly, using the proposed framework, we expect to obtain a high-quality shape correspondence that outperforms the state-of-art methods.

4.1. Spectral manifold wavelets

To conquer the difficulties of translating and scaling the wavelets on irregular spaces like graphs and manifolds, Hammond et al. [11] proposed the spectral graph wavelets in analogy to the behaviors of the classical wavelets in the frequency space. where the wavelets can be simply derived from the results of filtering to special functions with various band-pass filters. Indeed, their work can easily be extended to construct wavelets on manifolds that are capable of incorporating the underlying shape geometry, just via replacing the eigen-decomposition of the graph Laplace operator with the LBO. In our framework, we turn to call them spectral manifold wavelets (SMWs).

To be specific, given a spectral filter $g(\lambda) : \mathbb{R}^+ \to \mathbb{R}^+$, a smooth and compactly supported real-valued function, the SMW at scale s and located at point y is defined as

$$\psi_{s,y}(x) = \sum_{i\geq 0} g(s\lambda_i)\phi_i^*(y)\phi_i(x). \tag{2}$$

Here, the scaling parameter s determines the support interval of the filter $g(\lambda)$, i.e., the frequency band passed. When the filter $g(s\lambda)$ journeys up from low to high frequency bands, the locality of the wavelet $\psi_{s,y}$ on the surface will increase and vice versa. In order to simultaneously capture the low-frequency features of signals, another real-valued function $h : \mathbb{R}^+ \to \mathbb{R}^+$ that satisfy $h(0) > 0$ and $h(\lambda) \to 0$ when $\lambda \to \infty$ is designed for acting as a low-pass filter. Similarly, the corresponding scaling function at point y can also be obtained via filtering, i.e.

$$\varphi_y(x) = \sum_{i\geq 0} h(\lambda_i)\phi_i^*(y)\phi_i(x).$$

In practical applications, when the scale parameter s is sampled into a finite set of discrete values, the scaling functions are able to guarantee the original signal f stably recovered from the wavelet coefficients.

Note that, SMW is actually a special kernel on manifolds and closely related to the heat kernel [36], while only differing in the filters used. As opposed to only low-pass frequency filters in the heat kernel, both low and band-pass filters are used in SMW, and thus it is equipped with better abilities to analyze shape geometric features.

4.2. Multiscale wavelet preservation

Note that, (2) shows the SMWs naturally inherit the isometric invariance of the eigen-decomposition of the LBO. Therefore, given a pair of shapes $\mathcal{M}$ and $\mathcal{N}$ related with an isometric deformation, if the pointwise map T between them is isometric, then for each point $y \in \mathcal{M}$, we have

$$T_F(\psi_{s,y}^{\mathcal{M}}) = \psi_{s,T(y)}^{\mathcal{N}}.$$

Such isometric invariance is demonstrated in Figure 2.

In addition, we know that the multiscale neighboring topological structure around each point x can be efficiently characterized by the corresponding wavelets, as shown in a variety of previous works [11, 16]. More specifically, the SMW with small s can capture the geometric features of the immediate neighborhoods around its located point, which correspondences to the high-frequency structural information (details) of the shape. In contrast, a larger s allows the diffusion process to spread farther in the shape, resulting in the SMW encoding the features of neighborhoods with greater radii which contains low-frequency information of the shape structure. Now considering the multiscale wavelets into the functional maps, we can further have the functional maps T_F of the isometric mapping T satisfying

$$T_F(\psi_{s_l,y}^{\mathcal{M}}) = \psi_{s_l,T(y)}^{\mathcal{N}}, l = 0, 1, \cdots, L. \tag{3}$$

Note that, for simplicity, we denote the scaling function $\varphi_y(x)$ by $\psi_{s_0,y}(x)$, whose scale parameter index of the wavelets $l = 0$. The equations in (3) mean that the structural information of each frequency band are sufficiently preserved respectively in an isometric map T.

4.3. Maps with wavelet preservation constraints

As the isometric properties of the maps can be encoded fully in the multiscale SMWs, we are inspired to employ the wavelet constraints to determine a functional map conforming with an isometric pointwise map, and then use it to realize a high-quality shape correspondence. The core idea is to ensure the isometry of the map T, we intend to impose multiscale wavelets preservation constraints to determine the unknown $\mathbf{C}$, with jointly considering various scales s. For a clearer statement, now let us consider the discrete and matrix settings.

In dicretized settings, let shapes $\mathcal{M}$ and $\mathcal{N}$ be represented by triangular meshes with m and n vertices respectively. Firstly we use the method of [20] to compute the discretized LBOs of the shape $\mathcal{M}$ and $\mathcal{N}$, which are represented as matrices $\mathbf{L}_\mathcal{M}$ and $\mathbf{L}_\mathcal{N}$ respectively, where $\mathbf{L}_\mathcal{M} = \mathbf{A}_\mathcal{M}^{-1}\mathbf{W}_\mathcal{M}$ and $\mathbf{L}_\mathcal{N} = \mathbf{A}_\mathcal{N}^{-1}\mathbf{W}_\mathcal{N}$, here $\mathbf{A}_\mathcal{M}$ is the diagonal matrix of lumped area elements and $\mathbf{W}_\mathcal{M}$ is the cotangent weight matrix. We assume the matrices $\Phi_\mathcal{M} \in \mathbb{R}^{m\times k}$ and $\Phi_\mathcal{N} \in \mathbb{R}^{n\times k}$ contain the first k eigenvectors of $\mathbf{L}_\mathcal{M}$ and $\mathbf{L}_\mathcal{N}$. Now represent the map $T : \mathcal{M} \to \mathcal{N}$ as a matrix $\mathbf{P} \in \mathbb{R}^{m\times n}$, s.t. $\mathbf{P}(i,j) = 1$ if $T(i) = j$ and 0 otherwise, where i and j are the vertex indices of $\mathcal{M}$ and $\mathcal{N}$ respectively. Further, we denote the diagonal matrix $\Lambda_\mathcal{M} = \mathrm{diag}(\lambda_0^\mathcal{M}, \lambda_1^\mathcal{M}, \cdots, \lambda_{k-1}^\mathcal{M})$ and $\Lambda_\mathcal{N} = \mathrm{diag}(\lambda_0^\mathcal{N}, \lambda_1^\mathcal{N}, \cdots, \lambda_{k-1}^\mathcal{N})$ to contain the first k eigenvalues the LBOs. With matrix settings, now we have the functional matrix $\mathbf{C}$ given by the projection of $\mathbf{P}$ onto the corresponding Fourier basis, i.e. $\mathbf{C} = \Phi_\mathcal{N}^+\mathbf{P}^\mathrm{T}\Phi_\mathcal{M}$. where $^+$ denotes the Moore-Penrose pseudo-inverse. Due to the orthogonality of LBOs' eigenvector with respect to the area-weighted inner product, we have $\Phi_\mathcal{N}^\mathrm{T}\mathbf{A}_\mathcal{N}\Phi_\mathcal{N} = \mathbf{I}$, where $\mathbf{I}$ is an identity matrix, i.e.,$\Phi_\mathcal{N}^+ = \Phi_\mathcal{N}^\mathrm{T}\mathbf{A}_\mathcal{N}$.

Given a filter $g(\lambda)$, according to (2), we can get two wavelet matrices $\Psi_s^\mathcal{M} = \Phi_\mathcal{M} g(s\Lambda_\mathcal{M})\Phi_\mathcal{M}^+ \in \mathbb{R}^{m\times m}$ and $\Psi_s^\mathcal{N} = \Phi_\mathcal{N} g(s\Lambda_\mathcal{N})\Phi_\mathcal{N}^+ \in \mathbb{R}^{n\times n}$, where each column correspondences to the SMWs at scale s and located at the corresponding vertex (index). In practice, the scale parameter s is also sampled into a set of points $\{s_l\}_{l=0}^L$.

From (2), we know the Fourier coefficient matrices of the wavelets $\Psi_s^\mathcal{M}$ and $\Psi_s^\mathcal{N}$ are $g(s\Lambda_\mathcal{M})\Phi_\mathcal{M}^+$ and $g(s\Lambda_\mathcal{N})\Phi_\mathcal{N}^+$ respectively. Thus, according to Plancherel's theorem (stated in the last part of Section 3.2), we hope to find a mapping matrix $\mathbf{P}$ and its functional mapping matrix $\mathbf{C}$ that satisfy the wavelet preservation constraints at each scale s_l. This strategy can be eventually formulated as the following optimal problem, which aims to find the best $\mathbf{C}$ and $\mathbf{P}$ in the least squares sense, i.e.

$$\begin{aligned} &\min_{\mathbf{P},\mathbf{C}} \sum_{l=0}^{L} \left\| \mathbf{C} g(s_l\Lambda_\mathcal{M})\Phi_\mathcal{M}^+ - g(s_l\Lambda_\mathcal{N})\Phi_\mathcal{N}^+\mathbf{P}^\mathrm{T} \right\|_\mathrm{F}^2, \\ &\text{s.t.} \quad \mathbf{P}\mathbf{1} = \mathbf{1}, \mathbf{P}^\mathrm{T}\mathbf{1} \le \mathbf{1}. \end{aligned} \tag{4}$$

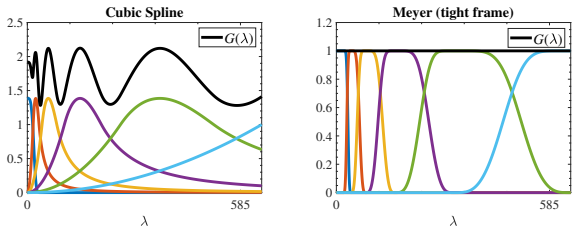

Figure 3. Filters of various types of wavelets. The colored lines in the left and right figures show the multiscale cubic spline filters (a general wavelet) and the Meyer filters (a tight wavelet frame) respectively, while the dark lines show their $G(\lambda)$ (7).

Solution. The above optimization problem (4) is not jointly convex, as both $\mathbf{P}$ and $\mathbf{C}$ are unknown and they meanwhile interact with each other. For this reason, we try to solve it approximately, via alternating the updating of $\mathbf{P}$ and $\mathbf{C}$. To be specific, firstly we fix the matrix $\mathbf{P}$ and then solve (4) with respect to $\mathbf{C}$, via solving the following problem

$$\min_{\mathbf{C}} \sum_{l=0}^{L} \left\| \mathbf{C} g(s_l\Lambda_\mathcal{M})\Phi_\mathcal{M}^+ - g(s_l\Lambda_\mathcal{N})\Phi_\mathcal{N}^+\mathbf{P}^\mathrm{T} \right\|_\mathrm{F}^2. \tag{5}$$

The above problem (5) actually has a analytical solution if we let

$$\mathbf{C} g\left(s_l\Lambda_\mathcal{M}\right)\Phi_\mathcal{M}^+ = g\left(s_l\Lambda_\mathcal{N}\right)\Phi_\mathcal{N}^+\mathbf{P}^\mathrm{T},\ l = 0, \cdots, L. \tag{6}$$

Finding a matrix $\mathbf{C}$ directly that satisfies all the conditions in (6) is still not a trivial task, due to a large system of linear equations involved. However, we find that if using suitable filters $\{g(s_l\lambda)\}_{l=0}^L$, more specifically, if the wavelets form a tight frame for the function spaces of the manifolds, the computation complexity of $\mathbf{C}$ can be significantly reduced, by avoiding solving such a large system of linear equations. Next, let us briefly introduce the contents of the tight wavelet frames on manifolds first.

As mentioned in [11], for a mesh with n vertices $\{x_i\}_{i=1}^n$, if given a set of discrete scales $\{s_l\}_{l=1}^L$ and appropriate filters $g(\lambda)$ and $h(\lambda)$, the set of wavelets $\{\psi_{s_l,x_i}\}_{l=0,i=1}^{L,\ n}$ (Note that, for simplicity, here we denote $h(\lambda)$ by $g(s_0\lambda)$ and the scaling function by wavelet ψ_{s_0,x_i}) will form a frame for the functional space of the shape, with the bounds $A = \min\limits_{\lambda\in[0,\lambda_{\max}]} G(\lambda)$ and $B = \max\limits_{\lambda\in[0,\lambda_{\max}]} G(\lambda)$, where

$$G(\lambda) = \sum_{l=0}^{L} g(s_l\lambda)^2. \tag{7}$$

Specially, if the bounds $A = B$, then such type of wavelet set is called tight frame, and further, if $A = B = 1$, it is termed Parseval frame. Figure 3 respectively shows the filters of one common wavelet and one tight wavelet frame and their $G(\lambda)$ as well.

Tight frames have plenty of attractive properties. Above all they can conserve the energy between the original and the transformed domains, thus allowing to easier and more accurate signal reconstruction [4]. In our framework, the tight wavelet frame [15] used brings more efficient computation and more accurate mapping constraints for the problem (4) than general wavelets, as we have following observation:

Remark 4.1. *If the wavelet sets* $\{\psi^{\mathcal{M}}_{s_l,x_i}\}^{L,\,m}_{l=0,i=1}$ *and* $\{\psi^{\mathcal{N}}_{s_l,x_i}\}^{L,\,n}_{l=0,i=1}$ *both derived from the filter* $g(\lambda)$ *could form tight frames for the functional spaces of the shapes* $\mathcal{M}$ *and* $\mathcal{N}$ *respectively, then the matrix* $\mathbf{C}$ *in (6) can be obtained via*

$$\mathbf{C} = \sum_{l=0}^{L} g\left(s_l\Lambda_{\mathcal{N}}\right)\Phi^{+}_{\mathcal{N}}\mathbf{P}^{\mathrm{T}}\Phi_{\mathcal{M}}g\left(s_l\Lambda_{\mathcal{M}}\right), \tag{8}$$

Proof. See Appendix A. □

Note that, compared with the general exact form of $\mathbf{C}$, i.e. $\mathbf{C} = \Phi^{+}_{\mathcal{N}}\mathbf{P}^{\mathrm{T}}\Phi_{\mathcal{M}}$, our solution to $\mathbf{C}$ in (8) can be treated as a filtering operation to this general form using various band-pass filters $g(s_l\Lambda), l = 0, \cdots, L$ and equivalent to a refinement. Note also that, this approach only relies on the products of matrices and vectors, avoiding to solve a large system of linear equations commonly presented in previous works.

We consider using the Meyer tight wavelet frame proposed by [15] (as shown in Figure 3) in our framework, as the filters of such wavelets frame get an adaptive bandwidth, which is more suitable to analyze the signals and can capture more geometric features than other frames.

After fixing $\mathbf{C}$, we try to recover the pointwise mapping matrix $\mathbf{P}$. Here, we still take the traditional nearest neighbor searching introduced in Section 3.2 to achieve this goal. That means, to find the image of the i*th* vertice of $\mathcal{M}$ on the shape $\mathcal{N}$, we can solve the optimization problem

$$T(i) = \arg\min_{j} \left\|\mathbf{C}(\Phi_{\mathcal{M}}(i))^{\mathrm{T}} - (\Phi_{\mathcal{N}}(j))^{\mathrm{T}}\right\|^2_{\mathrm{F}},$$

where $\Phi_{\mathcal{M}}(i)$ denotes the i*th* row of the eigenvector matrix $\Phi_{\mathcal{M}}$. We encode the obtained pointwise map T as a matrix $\mathbf{P}$ and update the $\mathbf{C}$ according to (8). Then iterate this procedure to continuously refine the matrices $\mathbf{C}$ and $\mathbf{P}$, until arriving desirable correspondence accuracy. Finally our pipeline is concluded in Algorithm 1, where the $maxIt$ denotes the maximum iterations.

Algorithm 1: Shape correspondence via multiscale SMWs preservation.

Input: $\Phi_{\mathcal{M}}, \Phi_{\mathcal{N}}, \{g(s\Lambda_{\mathcal{M}})\}_s, \{g(s\Lambda_{\mathcal{N}})\}_s, maxIt$
Output: $\mathbf{C}, \mathbf{P}$
Initialization: $\mathbf{P}_0$
for $k = 0$ **to** $maxIt - 1$ **do**
  $\mathbf{C}_k = \sum_s g(s\Lambda_{\mathcal{N}})\Phi^{+}_{\mathcal{N}}\mathbf{P}^{\mathrm{T}}_k\Phi_{\mathcal{M}}g(s\Lambda_{\mathcal{M}})$
  $T(i) = \min_j \left\|\mathbf{C}_k(\Phi_{\mathcal{M}}(i))^{\mathrm{T}} - (\Phi_{\mathcal{N}}(j))^{\mathrm{T}}\right\|^2_{\mathrm{F}}, \forall i$
  Encode the map T as a matrix $\mathbf{P}_{k+1}$.
end

5. Experiments and discussion

In this section, we will test our method on extensive challenging shapes from several public datasets. All experiments are tested on a PC with Intel(R) Core i7-9700K CPU at 3.60 GHz and 16.0 GB RAM, using MATLAB R2019b.

5.1. Implement details

Datasets. We will use five public domain datasets across standard to challenging ones to test our method. FAUST [5] is composed of 10 poses of 10 human subjects with significant variability between different human subjects. SCAPE [2] contains 71 registered meshes of a particular human subject in different poses. TOSCA [7] contains a total of 80 shapes divided into 8 classes across animals to human bodies with varying resolutions. SHREC'16 Topology[18] includes 26 non-intersecting manifold kids shapes that are produced by merging self-intersecting parts of the shapes from KIDS dataset[32]. SHREC'16 Partiality benchmark [9] spans different classes and includes two sub-datasets (cuts and holes) to exemplify different kinds of partiality.

Evaluation methodology. We use the correspondence quality characteristics (CQC) curves [13] , which exhibit the percentages of matches that are at most r in geodesic error, as well as the average geodesic error of all the points on the shape $\mathcal{M}$, to measure the correspondence quality.

Parameter settings. To balance the performance and the computation efficiency well, we generally set $k = 100$, $L = 5$ and take 3 iterations in the all following experiments. Detailed discussion about the parameter choices can be found in Appendix B.

5.2. Results and comparisons

In this section, we first show the comparisons with four most similar methods where functional mapping refinements and iterations are used, including ICP [24], CPD [31], BCICP [28] and ZoomOut [19], and also compare with other elegant state-of-the-art strategies for shape correspondence like SmoothShells [10] both w.r.t matching quality and computing efficiency.

Comparison with similar pipelines. We first test the performances of all methods on three benchmark datasets FAUST, SCAPE and TOSCA. For all competitors, we take the matching results of using the nearest searching in SHOT [35] descriptor spaces as their initialization and use the settings and codes provided online by their authors. All the

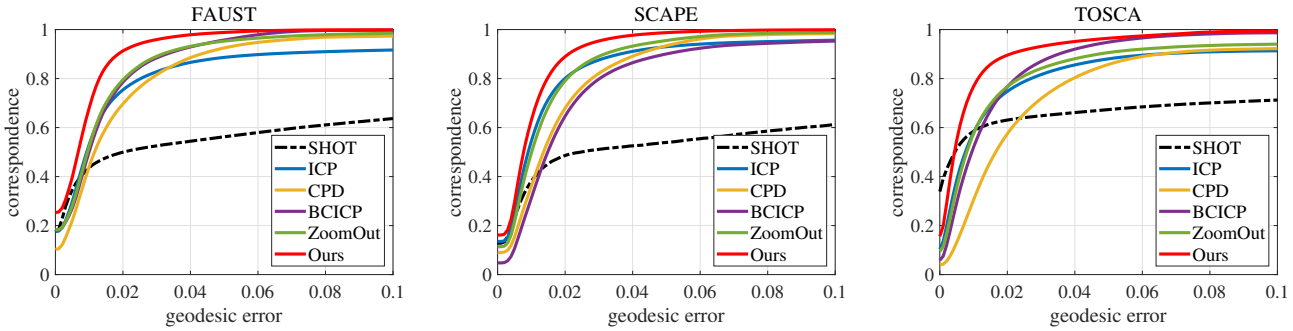

Figure 4. Matching performance comparisons on three benchmark datasets FAUST, SCAPE and TOSCA, evaluated by CQC curves.

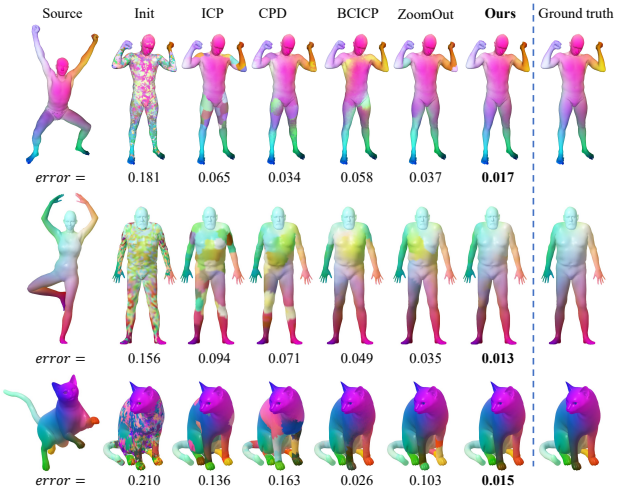

Figure 5. Matching accuracy comparisons via color transfer. The shapes are respectively from the SCAPE (top row), FAUST (second row) and TOSCA (bottom row), with average geodesic errors shown below each shape. Less error and color distortion happened in our method compared with the others.

Table 1. Quantitative comparison with representative functional map based refinement with the same initialization. Only refinement runtime is counted.

	Average error ($\times 10^{-3}$)			Average runtime (s)		
Method \Dataset	FAUST	SCAPE	TOSCA	FAUST	SCAPE	TOSCA
SHOT (Init)	137.8	157.0	148.7	—	—	—
ICP	37.9	23.9	64.0	5.8	19.2	251.6
CPD	23.5	21.9	*61.7*	31.4	123.9	*20.1*
BCICP	14.1	29.3	*18.8*	292.8	449.8	*68.36*
ZoomOut	16.8	16.5	29.1	2.01	2.32	3.88
Ours	**10.9**	**11.3**	**16.1**	**0.37**	**0.67**	**1.94**

CQC curves and the average errors on these three datasets are shown in Figure 4 and Table 1 respectively. We can find that our method improves the correspondence quality to other competitors. Besides, Figure 5 demonstrates more visual comparisons via color transfer between the shapes based on the respective correspondence. The results announce our advantages in matching accuracy once again. We also compare the matching performance on SHREC'16 Partiality benchmark and SHREC'16 Topology [18], two particularly challenging datasets with plenty of additional variations of broken, missing connectivity or topological noises respectively. As shown in Figure 6 and 7, our approach still achieves superior accuracy to other methods even under challenging circumstances. Note that, our method still outperforms the PFM [9], a sophisticated work that devotes to dealing with partial correspondence and represents the state-of-the-art in this field.

In addition, our method brings great improvements in computing efficiency to other methods. As with an efficient iteration and the computation times be mainly paid on the products of matrices and vectors in (8) and nearest searching for pointwise map recovery, without any time-consuming solutions of linear equations or calculations of geodesic distances, our pipeline posses a comparably efficient computation. Besides, our framework can use the approximate nearest neighbor searching FLANN library [21] instead of exact nearest neighbor searching and the subsampling approaches provided by ZoomOut [19] to accelerate the recovery of the pointwise maps. The computing times on three benchmark datasets are demonstrated in Table 1, which shows our superiority in computing efficiency.

Comparison with other strategies. We also compare with other shape correspondence strategies like SmoothShells [10], which is an elegant and state-of-the-art method that combines geometric and spectral alignment by embedding the input shapes into an extrinsic-intrinsic product space. The corresponding quantitative and visual comparisons are shown in Table 2 and Figure 8 respectively. Our approach produces comparable results with SmoothShells while providing great efficiency improvements.

Robustness. We test our robustness on shapes with different conditions, as shown in Figure 9. The CQC curves (left) show the comparison on the SCAPE remesh dataset from [28], where the shapes have different number of vertices and often significantly different triangulation (connectivity). Moreover, the results on shapes with different resolutions and correspondence initialization are also shown in the right figure. Obviously our performance is still fully guaranteed in these cases which are closer to origin SCAPE

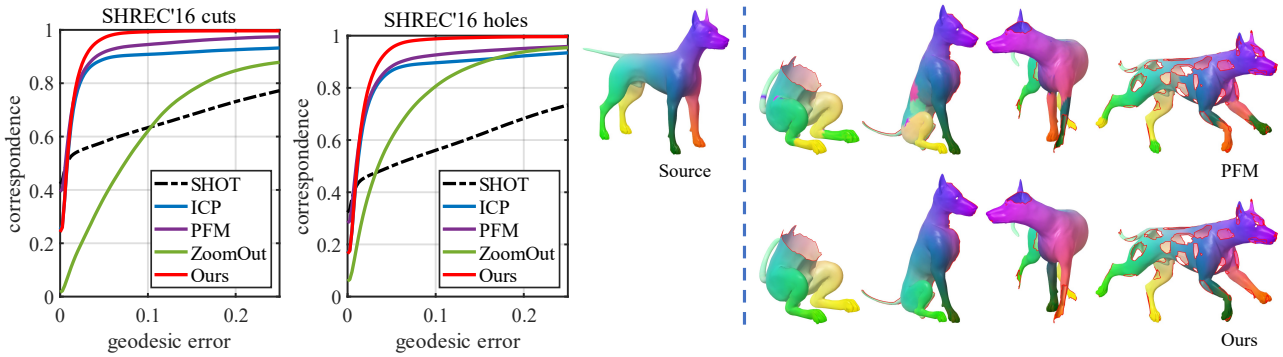

Figure 6. Matching performance comparisons on the SHREC'16 Partiality [9] benchmark. All CQC curves are shown on the left two figures, while the visual comparison with PFM [29] via color transfer are shown on the right. The color distortions disappear in our method as opposed to PFM.

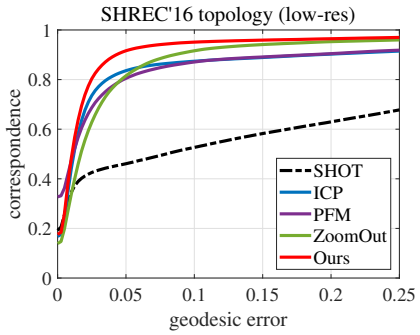

Figure 7. Matching performance comparisons on the SHREC'16 Topology.

Figure 8. Visual Comparison with SmoothShells [10]. Our method produces a more exact result and runs more faster than SmoothShells, where the runtime of ours and the Smoothshells are $10.2s$ and $350.8s$ respectively.

Table 2. Quantitative comparison with SmoothShells. The runtime contains pre-computing and refinement time.

	Average error ($\times 10^{-3}$)			Average runtime (s)		
Method \Dataset	FAUST	SCAPE	TOSCA	FAUST	SCAPE	TOSCA
SmoothShells	12.4	15.8	17.2	142.4	216.1	678.2
Ours	10.9	11.3	16.1	2.4	4.6	58.1
Improv Ours	12.1%	28.5%	6.4%	**58**×	**46**×	**10**×

dataset which has the same vertices and triangulation.

6. Conclusion

We have presented a novel approach for shape correspondence that incorporated novel constraints into the functional map framework. The core idea is that multiscale spectral manifold wavelets are required to be preserved at each scale respectively in the functional maps. Such constraints have been proven to be able to strongly ensure the isometry of the underlying pointwise maps. We also provided an efficient iteration strategy for the determinations of the functional maps and the pointwise maps, where high-quality correspondence can be achieved with only three iterations. Our method produced state-of-the-art correspondence results on extensive challenging datasets. The implementation to replicate our results will be released upon the publication.

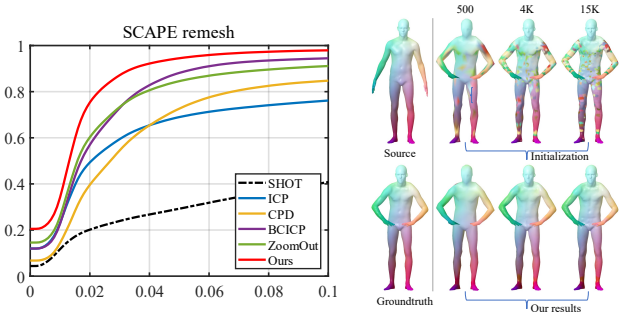

Figure 9. Robustness demonstration. The left CQC lines show the correspondence results on the SCAPE remesh dataset, while the right figure shows our results on shapes with different resolutions and correspondence initialization (represented via color transfer).

Acknowledgments

We would acknowledge the anonymous reviewers for their valuable comments. This work was supported by the Hunan Science Fund for Distinguished Young Scholars (No.2019JJ20027), the Natural Science Foundation of China (No.61572527,11802090), and the Hunan Provincial Innovation Foundation For Postgraduate (No.CX20190092).

References

[1] Yonathan Aflalo, Anastasia Dubrovina, and Ron Kimmel. Spectral generalized multi-dimensional scaling. *International Journal of Computer Vision*, 118(3):380–392, 2016.

[2] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. Scape: Shape completion and animation of people. *ACM Trans. Graph.*, 24(3):408–416, July 2005.

[3] M. Aubry, U. Schlickewei, and D. Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 1626–1633, 2011.

[4] John J Benedetto and Matthew Fickus. Finite normalized tight frames. *Advances in Computational Mathematics*, 18(2-4):357–385, 2003.

[5] F. Bogo, J. Romero, M. Loper, and M. J. Black. Faust: Dataset and evaluation for 3d mesh registration. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3794–3801, 2014.

[6] D. Boscaini, J. Masci, M. M. Bronstein, and D. Cremers. Anisotropic Diffusion Descriptors. *Computer Graphics Forum*, 35(2):431–441, 2016.

[7] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. *Numerical geometry of non-rigid shapes*. Springer Science & Business Media, 2008.

[8] Étienne Corman, Maks Ovsjanikov, and Antonin Chambolle. Supervised descriptor learning for non-rigid shape matching. In *European Conference on Computer Vision*, pages 283–298, 2014.

[9] L. Cosmo, E. Rodolà, M. M. Bronstein, A. Torsello, D. Cremers, and Y. Sahillioglu. Partial Matching of Deformable Shapes. In *Eurographics Workshop on 3D Object Retrieval*. The Eurographics Association, 2016.

[10] Marvin Eisenberger, Zorah Lähner, and Daniel Cremers. Smooth Shells: Multi-Scale Shape Registration with Functional Maps. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 12265–12274, 2020.

[11] David K Hammond, Pierre Vandergheynst, and Remi Gribonval. Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis*, 30(2):129–150, 2011.

[12] Ruqi Huang, Jing Ren, Peter Wonka, and Maks Ovsjanikov. Consistent zoomout: Efficient spectral map synchronization. *Computer Graphics Forum*, 39(5):265–278, 2020.

[13] Vladimir G. Kim, Yaron Lipman, and Thomas Funkhouser. Blended intrinsic maps. *ACM Trans. Graph.*, 30(4), July 2011.

[14] Artiom Kovnatsky, Klaus Glashoff, and Michael M. Bronstein. Madmm: A generic algorithm for non-smooth optimization on manifolds. In *European Conference on Computer Vision*, 2016.

[15] Nora Leonardi and Dimitri Van De Ville. Tight wavelet frames on multislice graphs. *IEEE Transactions on Signal Processing*, 61(13):3357–3367, 2013.

[16] Qinsong Li, Ling Hu, Shengjun Liu, Dangfu Yang, and Xinru Liu. Anisotropic Spectral Manifold Wavelet Descriptor. *Computer Graphics Forum*, 40(1):81–96, 2021.

[17] O. Litany, E. Rodolà, A. M. Bronstein, and M. M. Bronstein. Fully Spectral Partial Shape Matching. *Computer Graphics Forum*, 36(2):247–258, 2017.

[18] Z. Lähner, E. Rodolà, M. M. Bronstein, D. Cremers, O. Burghard, L. Cosmo, A. Dieckmann, R. Klein, and Y. Sahillioglu. Matching of Deformable Shapes with Topological Noise. In *Eurographics Workshop on 3D Object Retrieval*. The Eurographics Association, 2016.

[19] Simone Melzi, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, and Maks Ovsjanikov. Zoomout: Spectral upsampling for efficient shape correspondence. *ACM Trans. Graph.*, 38(6), Nov. 2019.

[20] Mark Meyer, Mathieu Desbrun, Peter Schröder, and Alan H Barr. Discrete differential-geometry operators for triangulated 2-manifolds. In *Visualization and mathematics III*, pages 35–57. Springer, 2003.

[21] Marius Muja and David G Lowe. Scalable nearest neighbor algorithms for high dimensional data. *IEEE transactions on pattern analysis and machine intelligence*, 36(11):2227–2240, 2014.

[22] Dorian Nogneng, Simone Melzi, Emanuele Rodola, Umberto Castellani, Michael M Bronstein, and Maks Ovsjanikov. Improved functional mappings via product preservation. *Computer Graphics Forum*, 37(2):179–190, 2018.

[23] Dorian Nogneng and Maks Ovsjanikov. Informative Descriptor Preservation via Commutativity for Shape Matching. *Computer Graphics Forum*, 36(2):259–267, 2017.

[24] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: A flexible representation of maps between shapes. *ACM Trans. Graph.*, 31(4), July 2012.

[25] Maks Ovsjanikov, Etienne Corman, Michael Bronstein, Emanuele Rodolà, Mirela Ben-Chen, Leonidas Guibas, Frederic Chazal, and Alex Bronstein. Computing and processing correspondences with functional maps. In *ACM SIGGRAPH 2017 Courses*, SIGGRAPH '17, New York, NY, USA, 2017. Association for Computing Machinery.

[26] Maks Ovsjanikov, Quentin Mérigot, Facundo Mémoli, and Leonidas Guibas. One point isometric matching with the heat kernel. *Computer Graphics Forum*, 29(5):1555–1564, 2010.

[27] J. Pokrass, A. M. Bronstein, M. M. Bronstein, P. Sprechmann, and G. Sapiro. Sparse modeling of intrinsic correspondences. *Computer Graphics Forum*, 32(2):459–468, 2013.

[28] Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. Continuous and orientation-preserving correspondences via functional maps. *ACM Trans. Graph.*, 37(6), Dec. 2018.

[29] E. Rodolà, L. Cosmo, M. M. Bronstein, A. Torsello, and D. Cremers. Partial Functional Correspondence. *Computer Graphics Forum*, 36(1):222–236, 2017.

[30] E. Rodolà, Z. Lähner, A. M. Bronstein, M. M. Bronstein, and J. Solomon. Functional Maps Representation On Product Manifolds. *Computer Graphics Forum*, 38(1):678–689, 2019.

[31] E. Rodolà, M. Moeller, and D. Cremers. Regularized Pointwise Map Recovery from Functional Correspondence. *Computer Graphics Forum*, 36(8):700–711, 2017.

[32] E. Rodolà, S. Bulò, T. Windheuser, M. Vestner, and D. Cremers. Dense non-rigid shape correspondence using random forests. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 4177–4184, 2014.

[33] Raif M. Rustamov, Maks Ovsjanikov, Omri Azencot, Mirela Ben-Chen, Frédéric Chazal, and Leonidas Guibas. Map-based exploration of intrinsic shape differences and variability. *ACM Trans. Graph.*, 32(4), July 2013.

[34] Yusuf Sahillioğlu. Recent advances in shape correspondence. *The Visual Computer*, 36(8):1705–1721, 2020.

[35] Samuele Salti, Federico Tombari, and Luigi Di Stefano. Shot: Unique signatures of histograms for surface and texture description. *Computer Vision and Image Understanding*, 125:251–264, 2014.

[36] Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum*, 28(5):1383–1392, 2010.

[37] Oliver van Kaick, Hao Zhang, Ghassan Hamarneh, and Daniel Cohen-Or. A survey on shape correspondence. *Computer Graphics Forum*, 30(6):1681–1707, 2011.

[38] M. Vestner, Z. Lähner, A. Boyarski, O. Litany, R. Slossberg, T. Remez, E. Rodola, A. Bronstein, M. Bronstein, R. Kimmel, and D. Cremers. Efficient deformable shape correspondence via kernel matching. In *2017 International Conference on 3D Vision (3DV)*, pages 517–526, 2017.