

ReMatching: Low-Resolution Representations for Scalable Shape Correspondence

Filippo Maggioli¹, Daniele Baieri², Emanuele Rodolà², and Simone Melzi¹

¹ University of Milano-Bicocca, Milan, Italy

² Sapienza – University of Rome, Rome, Italy

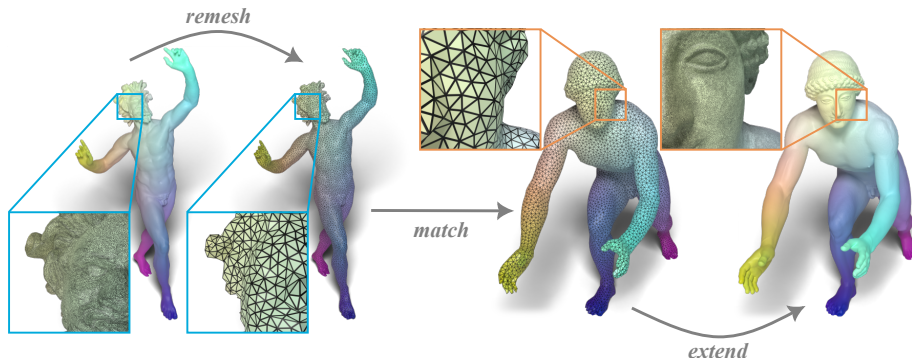


Fig. 1: Correspondence between two dense shapes computed with our method, using ZoomOut [38] on low-resolution meshes. The coordinates functions of the dancing faun statue on the left ($\sim 750k$ vertices) are transferred using a functional map to the Aphaea warrior statue on the right ($\sim 3.5M$ vertices). The colors encode the correspondence. Despite the mesh density, shown in the close-up, the computation took ~ 2 minutes.

Abstract. We introduce *ReMatching*, a novel shape correspondence solution based on the functional maps framework. Our method, by exploiting a new and appropriate *re-meshing* paradigm, can target shape-matching tasks even on meshes counting millions of vertices, where the original functional maps does not apply or requires a massive computational cost. The core of our procedure is a time-efficient remeshing algorithm which constructs a low-resolution geometry while acting conservatively on the original topology and metric. These properties allow translating the functional maps optimization problem on the resulting low-resolution representation, thus enabling efficient computation of correspondences with functional map approaches. Finally, we propose an efficient technique for extending the estimated correspondence to the original meshes. We show that our method is more efficient and effective through quantitative and qualitative comparisons, outperforming state-of-the-art pipelines in quality and computational cost.

Keywords: Shape matching · Functional maps · Geometry processing

1 Introduction and related work

The task of finding a semantically meaningful correspondence between discrete surfaces has always been a fundamental topic in the field of shape analysis. Researchers developed a wealth of solutions for this application, and new methods continue to be designed [10, 57]. Among the various approaches, the functional maps framework [47, 48] received significant attention. Rather than finding a point-wise correspondence, the functional maps framework aims to define a correspondence between functions, encoding it into a small matrix. Following this direction, a variety of works tried to improve the actual computation of the map [13, 38, 53] or to extend the framework to different types of bases [34, 39, 45].

Despite the amount of research outcomes based on the functional maps framework, both exploiting geometric properties and tools [7, 16, 17] and clever optimization techniques [46, 51, 52], the problem of computing the map is still bounded by the time complexity of sparse eigendecompositions [26]. To overcome this problem, researchers started investigating scalable solutions to the Laplace-Beltrami eigenproblem. In this regard, multi-resolution techniques [42, 43, 58, 59] proved to be very effective for scalable computation of spectral quantities for tasks such as retrieval and mesh filtering [30, 54], but their applicability to the functional map framework proved to require additional care [35]. Furthermore, recent research proposed spectral coarsening methods [32] or spectral preserving simplification techniques [28]. However, these techniques usually rely on the computation of the full-resolution spectrum, making them unsuitable for large-scale applications. Gao *et al.* [19] proposed a mixed-integer programming scalable solution to the matching problem, but they specifically target sparse correspondences only. Alternatively, Marin *et al.* [36], who provide an approach for human registration that scales across large datasets, but it is unsuitable for dense shapes.

Closely related to our work is the contribution from Magnet *et al.* [35], which, for the first time, proposes a solution for scaling the functional maps approach to high-resolution meshes. Their idea is to reduce the dimensionality of the eigenproblem, inducing a linear relationship between the functional space on the mesh and a lower-dimension functional space on a sparse sample. The extreme sparsity of the sampling, combined with its efficient computation, leads to an algorithm that can effectively deal with high-density meshes. Nevertheless, this type of reduction does not mitigate the negative influence that small components and local details at high frequency could have on the alignment of the spectra.

Finally, our pipeline relies on building robust and sparse representations of dense meshes while still keeping an approximately bijective correspondence between the high- and low-resolution shapes, and it is worth mentioning that other works have already attempted to do so. Jiang *et al.* [23] propose building a lower resolution prismatic shell preserving bijectivity with the original surface. However, their algorithm does not scale to meshes with a high triangle count, gives no control and no guarantees on the final vertex count, and requires many assump-

tions on the input shape (*e.g.*, manifoldness, orientability, no self-intersections). On the other hand, Liu *et al.* [31] propose to exploit an intrinsic error metric for decimating a mesh to a given size while keeping track of a geodesic barycentric mapping of each triangle of the resulting shape to a geodesic triangle on the input surface. Despite the robustness of the method, the simplification approach makes it unsuitable for a functional maps setting, where it is required to decimate shapes from millions of vertices to a few thousand. This limitation is further discussed in Section 4.2.

We introduce a new scalable functional map pipeline that efficiently handles meshes with high vertex density (see Figure 1), yields stable results, and is not bounded by the quality of the original triangulation. To summarize:

- we translate the matching pipeline to low-resolution representations, enabling a fast and scalable computation of functional maps between dense shapes, providing the first alternative to [35];
- we propose a new efficient and geometry-preserving remeshing algorithm specifically designed for our pipeline;
- we exploit a fast solution for extending scalar maps from the low-resolution remeshed shape to the original surface, thus efficiently obtaining suitable bases for function transfer and shape matching.

2 Background

In the discrete setting, we represent a shape as a triangular mesh $\mathcal{M} = (V, E, T)$, where (i) $V \subset \mathbb{R}^3$ is a set of vertices; (ii) $E \subset V^2$ is a set of edges between vertices; (iii) $T \subset V^3$ is a set of triangles composing the surface.

2.1 Functional maps

We discretize a scalar function $f : \mathcal{M} \rightarrow \mathbb{R}$ as a signal defined on the vertices V and represented as a vector $\mathbf{f} \in \mathbb{R}^{|V|}$. Hence, the Laplace-Beltrami operator $\Delta_{\mathcal{M}} : \mathcal{F}(\mathcal{M}) \rightarrow \mathcal{F}(\mathcal{M})$ is discretized as a sparse matrix $\mathbf{L}_{\mathcal{M}} \in \mathbb{R}^{|V| \times |V|}$, generally represented by means of a stiffness matrix $\mathbf{S}_{\mathcal{M}}$ and a mass matrix $\mathbf{A}_{\mathcal{M}}$ [40, 50]. The eigendecomposition $\mathbf{S}_{\mathcal{M}}\Phi_{\mathcal{M}} = \mathbf{A}_{\mathcal{M}}\Phi_{\mathcal{M}}\Lambda_{\mathcal{M}}$ of the Laplacian yields an orthonormal basis for the functional space on \mathcal{M} [29, 30] which has analogies with the Fourier basis and optimally represents smooth functions when truncated [1].

Given two shapes \mathcal{M} and \mathcal{N} , respectively with m and n vertices, a point-to-point map $\Pi : \mathcal{M} \rightarrow \mathcal{N}$ can be expressed as a binary matrix $\Pi \in \{0, 1\}^{m \times n}$ such that $\Pi(i, j) = 1$ if the vertex j -th of \mathcal{N} corresponds to the vertex i -th of \mathcal{M} , and $\Pi(i, j) = 0$ otherwise. Any point-wise correspondence Π establishes a functional map $T_{\Pi} : \mathcal{F}(\mathcal{N}) \rightarrow \mathcal{F}(\mathcal{M})$ as $T_{\Pi}(f) = f \circ \Pi$, $\forall f \in \mathcal{F}(\mathcal{N})$.

As proposed in [47], given Φ_k and Ψ_k the Laplacian bases truncated to size k respectively on \mathcal{M} and \mathcal{N} , we can project T_{Π} in these bases and compactly encode the functional map in a $k \times k$ matrix $\mathbf{C} = \Phi_k^{\top} \mathbf{A}_{\mathcal{M}} \Pi \Psi_k$, which is the linear operator that transfers the coefficients of functions from $\mathcal{F}(\mathcal{N})$ in the ones of corresponding functions of $\mathcal{F}(\mathcal{M})$ respectively computed with Φ_k and Ψ_k .

Finally, given a functional map T_{Π} , or its compact representation \mathbf{C} , we can retrieve the correspondence Π on which the map is built by the nearest neighbor search in the embedding space of the Laplacian eigenfunctions as detailed in [48].

2.2 Intrinsic Delaunay triangulations (IDT)

The IDT [3, 55] has been introduced for generalizing the Delaunay triangulation [9] to non-Euclidean metric spaces where classical algorithms [5, 60] cannot be used. IDT relies on the duality with the Voronoi diagram [15, 27], as three texels meeting at a point form a triangle whose edges cross the texels boundaries.

Definition 1. Let $(\mathcal{M}, d_{\mathcal{M}})$ be a 2-dimensional metric space (with distance function $d_{\mathcal{M}}$), and let $S = \{p_i\}_{i=1}^s \subset \mathcal{M}$ be a set of s sample points. The Voronoi decomposition (VD) of \mathcal{M} with respect to S is the collection $\{P_i\}_{i=1}^s$ such that

$$P_i = \left\{ q \in \mathcal{M} : p_i = \arg \min_{p_j \in S} (d_{\mathcal{M}}(p_j, q)) \right\}, \quad \bigcup_{i=1}^s P_i = \mathcal{M} \quad (1)$$

The points in S are called generators, and the P_i are called texels (or cells).

The intersection between two or more texels could be non-empty. In such a case, we define those texels as *adjacent*. A VD is *general* if the intersection of four or more texels is empty. If a VD is general, a *Voronoi vertex* is a point that belongs to three texels, and a *Voronoi edge* is a curve C connecting two Voronoi vertices and belonging to two texels.

It can be shown that there is a one-to-one correspondence between a Voronoi edge dividing P_i and P_j and a Delaunay edge connecting p_i to p_j [14]. This correspondence gives rise to a necessary and sufficient condition for the IDT to be a proper triangulation (i.e., it realizes a simplicial complex).

Definition 2 ([2, 15]). If S induces a general VD, the VD satisfies the closed ball property if:

- every P_i is a closed 2-ball (i.e., a topological disk without holes);
- every $P_i \cap P_j$ is either empty or a closed 1-ball (i.e., a topological segment);
- every $P_i \cap P_j \cap P_k$ is either empty or a closed 0-ball (i.e., a point).

Theorem 1 ([2, 14, 15]). If S induces a general VD, then the VD satisfies the closed ball property if and only if its dual IDT is a proper triangulation.

3 Method

We introduce a new scalable functional map pipeline handling very dense meshes and yielding stable results independently of the input triangulation’s quality.

Our method acts through three main steps. At first, we compute two low-resolution meshes that preserve the original metrics of the input shapes (Section 3.1). Then, we apply any existing pipeline, yielding a functional map between the low-resolution eigenspaces (Section 3.2). Finally, we extend the eigenfunctions to the original meshes, allowing us to use them with the functional map

estimated in the previous step to compute the desired correspondence between them (Section 3.3). Our technique is very general and works with arbitrary manifold triangulations, including meshes with degenerate geometry, non-orientable surfaces, and multiple connected components. The manifoldness requirement does not provide an obstacle to the applicability of our algorithm, as it can be efficiently enforced using existing techniques [18].

3.1 Intrinsic Delaunay remeshing

Front propagation and FPS. Naively computing the VD induced by a set of samples can quickly become very expensive. The exact geodesic algorithm introduced by Mitchell *et al.* [41] is computationally unfeasible on large meshes. Even using faster solutions, such as the heat method [8], the fast marching algorithm [25], or even the Dijkstra distance [11], searching the closest sample for each vertex is still a $\mathcal{O}(|V|s)$ operation. Instead, we take inspiration from the front propagation method proposed by Peyré *et al.* [49].

We start with a decomposition of the mesh induced by a single sample p_1 . This means that every vertex v is part of the texel P_1 and has distance $D_1[v] = d_{\mathcal{M}}(p_1, v)$ from the sample set. We then assume to have a decomposition P_1, \dots, P_{k-1} induced by samples p_1, \dots, p_{k-1} , and a vector D_{k-1} storing the distance of each vertex from the sample set. When adding a sample p_k to the VD, we can exploit the fact that $D_k = \min(D_{k-1}, d_{\mathcal{M}}(p_k, V))$. We start a front from p_k and expand it, updating the distances and assigning vertices to texel P_k until we reach all vertices such that $D_{k-1}[v] < d_{\mathcal{M}}(p_k, v)$ (see Figure 2). By propagating the front with the fast marching or Dijkstra’s algorithm, we can compute a geodesic VD with time $\mathcal{O}(|V| \log(|V|) \log(s))$.

To obtain the sample points, Peyré *et al.* [49] propose to use a geodesic farthest point sampling, achieving a geodesic uniform sampling of the surface.

Definition 3. Let $(\mathcal{M}, d_{\mathcal{M}})$ be a 2-dimensional metric space, and let $s \in \mathbb{N}$ be an integer. A farthest point sampling (FPS) of size s with respect to $d_{\mathcal{M}}$ is a set $S_s \subset \mathcal{M}$ of s sample points in \mathcal{M} , incrementally built from a singleton set $S_1 = \{p_1\} \subset \mathcal{M}$ according to the following rule:

$$S_{\ell+1} = S_{\ell} \cup \left\{ \arg \max_{q \in \mathcal{M}} \min_{p \in S_{\ell}} d_{\mathcal{M}}(q, p) \right\}. \quad (2)$$

However, searching for the maximum distance in Equation (2) at every new sample brings again the algorithm to a $\mathcal{O}(|V|s)$ complexity, really becoming a

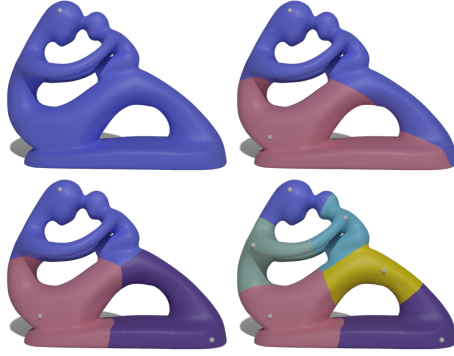


Fig. 2: The iterative front propagation from samples inducing a geodesic VD.

bottleneck for large meshes counting millions of vertices. In contrast, we propose to introduce a fixed-size binary max-heap in the front propagation algorithm, keeping track of the distances from each vertex to the sample set. During each front propagation, this heap updates the distance of each visited vertex in time $\mathcal{O}(\log(|V|))$, adding no extra complexity to the visit. Furthermore, it can gather and set to zero the farthest vertex at each iteration in just $\mathcal{O}(\log(|V|))$ time, totalling a time complexity of $\mathcal{O}(|V| \log(|V|) \log(s))$ for the entire procedure. The fast front propagation algorithm is summarized in Algorithm 1.

Algorithm 1 Geodesic FPS and its VD.

```

1: procedure VORONOI FPS( $\mathcal{M} = (V, E, T), s$ )
2:    $i \leftarrow$  random index in  $[1, |V|]$ 
3:    $S \leftarrow \{i\}$ 
4:    $D \leftarrow d_{\mathcal{M}}(V, V_i)$ 
5:    $H \leftarrow$  max-heap initialized with  $D$ 
6:    $P \leftarrow$  vector of length  $|V|$  initialized to  $i$ 
7:   for  $h \leftarrow 2$  to  $s$  do
8:      $p \leftarrow \text{FINDMAX}(H)$ 
9:      $\text{SETKEY}(H, p, 0)$ 
10:     $D_p \leftarrow 0$ 
11:     $S \leftarrow S \cup \{p\}$ 
12:    Propagate a front from  $p$ .
13:    Update  $D$ ,  $P$ , and  $H$ .
14:   end for
15: end procedure

```

Flat union property. Even if we are able to compute the VD of a farthest point sampling with high efficiency, we still cannot guarantee that the dual connectivity is a proper IDT. Ensuring the closed ball property can be computationally challenging. In particular, identifying the topology of the boundary between each pair of adjacent texels can become very costly when the number of samples is large. A possibility is to add enough samples so that the surface locally behaves like a plane. This approach, proposed by Leibon *et al.* [27] and adopted by Peyré

et al. [49], makes it easy to lose control over the number of samples, and anyway the requirements are not easy to enforce. For guaranteeing a proper IDT, while still avoiding these issues, we introduce a novel property for a VD. The proofs of our claims are provided in the supplementary materials.

Definition 4 (Flat Union Property (FUP)). *Let $(\mathcal{M}, d_{\mathcal{M}})$ be a 2-dimensional metric space. Let then $S = \{p_i\}_{i=1}^s \subset \mathcal{M}$ be a set of s sample points in \mathcal{M} , inducing a general VD $\{P_i\}_{i=1}^s$. The VD induced by S satisfies the FUP if:*

- every P_i is a closed 2-ball;
- if $P_i \cap P_j$ is not empty, then $P_i \cup P_j$ is a closed 2-ball;
- if $P_i \cap P_j \cap P_k$ is not empty, then $P_i \cup P_j \cup P_k$ is a closed 2-ball.

Theorem 2. *Let $(\mathcal{M}, d_{\mathcal{M}})$ be a 2-dimensional metric space. If a general VD $\{P_i\}_{i=1}^s$ of \mathcal{M} satisfies the FUP, it also satisfies the closed ball property.*

The FUP is much easier to verify than the closed ball property, as we only have to ensure that regions are closed 2-balls. Indeed, the following property can be exploited in this regard.

Proposition 1. *Let $\mathcal{M} = (V, E, F)$ be a manifold polygonal mesh, then \mathcal{M} is a closed 2-ball if and only if its Euler characteristic $\chi = |V| - |E| + |F|$ is 1.*

Dual mesh representation. When we compute a VD of a triangle mesh, we want to partition the vertices into disjoint sets. However, Voronoi texels are defined as regions over the surface, meaning that they should be submanifolds of \mathcal{M} . By only considering vertices, we likely end up having non-manifold geometries or non well-defined boundaries.

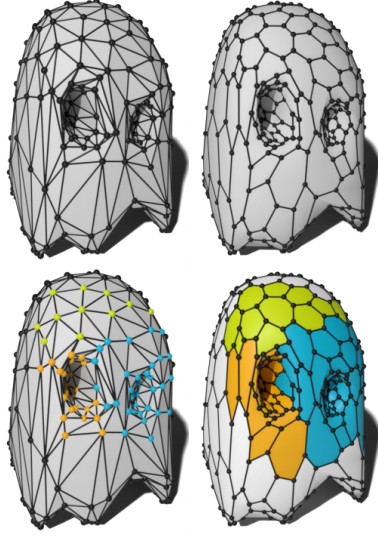


Fig. 3: Top: the primal and dual connectivities of a triangular mesh. Bottom: regions on the primal mesh are represented as sets of vertices, and the corresponding regions on the dual mesh are composed by faces.

We consider the polygonal mesh $\tilde{\mathcal{M}} = (\tilde{V}, \tilde{E}, \tilde{T})$ to be the dual mesh of $\mathcal{M} = (V, E, T)$, built by placing a vertex at each face of \mathcal{M} and forming a polygonal face for each triangle fan around a vertex of \mathcal{M} . As shown in the example from Figure 3, dual meshes allow us to define Voronoi texels by means of connected dual faces, making them actual submanifolds of the original surface. Furthermore, we can easily define boundaries between texels as paths made of dual edges.

This construction has the positive side effect to ensure that every VD is general. Indeed, the only place where three or more texels can meet is a dual vertex. Since texels are made of dual faces (*i.e.*, primal vertices), and the primal mesh is triangular, there is no dual vertex where more than three faces can meet.

Given the correspondence between primal and dual geometric elements, we can verify if a region is a closed 2-

ball with Proposition 1 without explicitly constructing the dual mesh. Additional details are provided in the supplementary material. After the initial FPS, we further add samples to enforce the flat union property: if two or more adjacent texels do not form a closed 2-ball, we break the connection adding samples at the Voronoi vertex or along the Voronoi edge; if a texel is not a closed 2-ball, we add a sample along its boundary to reduce its coverage.

3.2 Low-resolution matching

Let \mathcal{M} and \mathcal{N} be two triangular meshes, and let $\hat{\mathcal{M}}$ and $\hat{\mathcal{N}}$ be their low-resolution counterparts computed with the algorithm presented above (or any other alternative algorithm). Furthermore, let $\hat{\Phi}$ be a basis for the functional space over $\hat{\mathcal{M}}$, and $\hat{\Psi}$ a basis for the functional space over $\hat{\mathcal{N}}$. For example these bases can

be the truncated subset of the eigenvectors of the LBO as proposed in [47]. We assume to have some pipeline yielding a functional map \mathbf{C} , such that

$$\hat{\Psi} \mathbf{C} \approx \hat{\Pi} \hat{\Phi}, \quad (3)$$

where $\hat{\Pi} : \hat{\mathcal{M}} \rightarrow \hat{\mathcal{N}}$ is a correspondence between the vertices of $\hat{\mathcal{M}}$ and the vertices of $\hat{\mathcal{N}}$. The adopted functional maps solution can be any of the available alternatives, ranging from the original one [47] to the most recent [12]. In our experiments, we consider two of the widely adopted solutions: constrained optimization with product preservation [46] and the iterative procedure ZoomOut [38]. The first is representative of the possible optimization strategies, while the second is probably the most efficient refinement technique for functional maps estimation. Both these methods have been efficiently and effectively applied on meshes with a limited number of vertices (*i.e.*, up to 10 thousands vertices), which is exactly the setting we are in after the proposed remeshing step.

3.3 Extending the correspondence

To extend the basis from the remeshing to the original surface, we need a mapping that transports scalar fields from the low-resolution mesh $\hat{\mathcal{M}}$ to the high-resolution shape \mathcal{M} . A possible solution would be to map each triangle of $\hat{\mathcal{M}}$ to a geodesic triangle in \mathcal{M} , and then computing the geodesic barycentric coordinates of each vertex inside the triangle [31]. Unfortunately, this approach is very costly, as it requires to compute exact geodesic paths onto \mathcal{M} [56]. Instead, we approximate this mapping by projecting each vertex of \mathcal{M} onto the closest surface point of $\hat{\mathcal{M}}$. Despite this approach not being accurate in general, every triangle in $\hat{\mathcal{M}}$ corresponds to a geodesic triangle (homeomorphic to a disk) on \mathcal{M} with the same vertices. Furthermore, starting from a farthest point sampling gives evenly spaced samples, mitigating the geometric complexity of geodesic triangles.

A surface point of $\hat{\mathcal{M}}$ is either a vertex, or a weighted average of two (if on an edge) or three vertices (if on a triangle). We build a linear map $\mathbf{U}_{\mathcal{M}} \in \mathbb{R}^{m \times s}$ representing the projection of each vertex $v_i \in \mathcal{M}$ as a linear combination of at most three vertices in $\hat{\mathcal{M}}$ (*i.e.*, the vertices encoding the closest surface point).

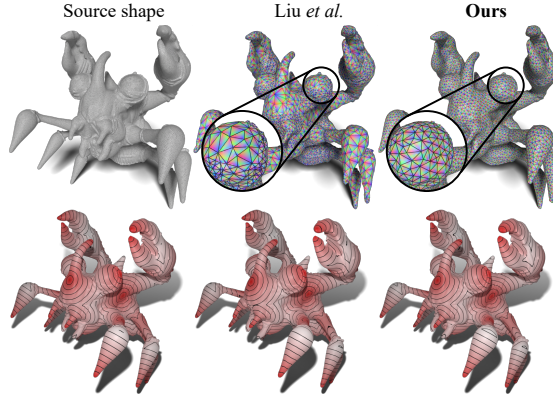


Fig. 4: Top: remeshing and mapping of the original vertices with the solution from Liu *et al.* [31] and our pipeline. Bottom: comparison of ground truth geodesic distances from multiple vertices and extended with both methods.

Figure 4 shows a comparison between our pipeline and the simplification method proposed by Liu *et al.* [31]. The top row shows the different geometries produced by the two methods in reducing the mesh size by 90% and how our mapping compares to the approximate geodesic barycentric coordinates. In the bottom row we compare the two methods in approximating a multiple source geodesic distance field, showing that they do not present appreciable differences. More discussion is provided in the supplementary materials.

Scalar functions can be evaluated at any surface point by interpolating the values at the vertices. Since every vertex of \mathcal{M} is associated to a point on the surface of $\hat{\mathcal{M}}$, we can use the linear map $\mathbf{U}_{\mathcal{M}}$ for extending scalar functions from the s vertices of $\hat{\mathcal{M}}$ to the m vertices of \mathcal{M} .

We extend the bases $\hat{\Phi}$ and $\hat{\Psi}$ to the full-resolution meshes as $\Phi = \mathbf{U}_{\mathcal{M}}\hat{\Phi}$ and $\Psi = \mathbf{U}_{\mathcal{N}}\hat{\Psi}$. Then, in a similar fashion as done in Equation (3), we search for a correspondence $\Pi : \mathcal{M} \rightarrow \mathcal{N}$ that maps vertices of \mathcal{M} to vertices of \mathcal{N} as

$$\mathbf{U}_{\mathcal{N}}\hat{\Psi}\mathbf{C} = \Psi\mathbf{C} \approx \Pi\Phi = \Pi\mathbf{U}_{\mathcal{M}}\hat{\Phi}, \quad (4)$$

where, for the sake of clarity, we explicitly write the equation in terms of the bases $\hat{\Phi}$ and $\hat{\Psi}$. Given the extended bases Φ and Ψ , and the same functional map \mathbf{C} estimated to solve Equation (3), we obtain the correspondence Π via a nearest neighbor search in the spectral space.

We notice that nearest neighbor algorithms become very slow on large meshes, effectively becoming a possible bottleneck for the pipeline. Our experiments presented in Figure 10 show that, to the scope of this paper, this solution provides satisfying performance. However, we acknowledge that establishing a relationship between (3) and (4) could lead to a more efficient expression of Π in terms of $\hat{\Pi}$, $\mathbf{U}_{\mathcal{M}}$, and $\mathbf{U}_{\mathcal{N}}$. This avenue warrants future exploration.

4 Results

Our goal is to evaluate not only the matching accuracy of our method but also its time performance. For this task, we test it on datasets for dense correspondences, such as the SHREC19 dataset [37] and the TOSCA dataset [6]. We also introduce a novel dataset, BadTOSCA, obtained by randomly altering the vertex positions of the TOSCA meshes, to ensure that our method is stable even under circumstances where strong isometry cannot be assumed³.

Our pipeline relies on an efficient remeshing algorithm that can quickly and drastically reduce the size of a mesh by orders of magnitude while still preserving the underlying geometry and metric. In principle, any remeshing algorithm could be adopted in our pipeline, as the map for extending scalar fields discussed in Section 3.3 only requires having two input meshes. We compare our solution against well-established isotropic and anisotropic remeshing algorithms (respectively, IRM [4,21] and ARM [44]), as well as the scalable functional maps (SFM)

³ Details on the dataset generation are provided in the supplementary materials. Code for generating the dataset is available at <https://github.com/filthynobleman/bad-tosca>

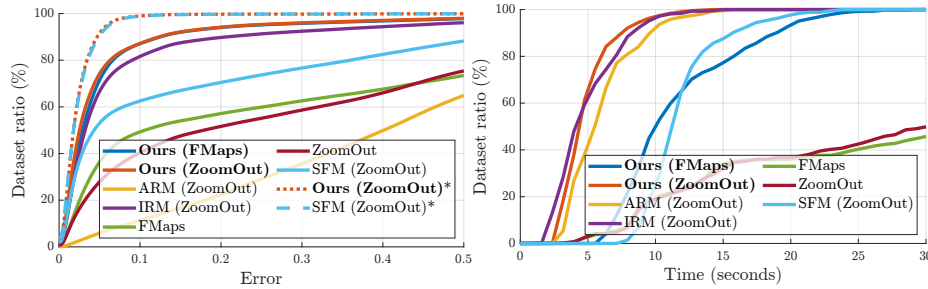


Fig. 5: Left: Accuracy curves on the SHREC19 challenge pairs for the tested methods. We also consider ideal approaches where the initial functional map for ZoomOut is given ground truth dashed lines). Right: Cumulative curves of the execution time for evaluated methods. For all methods, we also consider the time required to run the remeshing or resampling step(s).

approach proposed by Magnet *et al.* [35], which produces a subsampling of the vertices in place of a low-resolution mesh.

We implemented our remeshing algorithm in C++, using Eigen [20] and libigl [22], and the matching pipeline in MATLAB⁴.

4.1 Quality of matching

To evaluate our shape matching pipeline, we tested our method on the SHREC19 dataset [37]. We can use different shape correspondence techniques to test the entire pipeline. In particular, we compute the functional map on the low-resolution representation using both the widely adopted approach with products preservation (FMaps) [46]

and ZoomOut [38]. We compare our method with the same approaches on the full-resolution meshes, as well as against SFM. For a fair comparison, we remesh all the shapes to 3k vertices with all methods, as Magnet *et al.* [35] state that this produces the best performance with SFM. To further ease the alignment of the Laplacian spectra, we also post-process the remeshed shapes (produced with

Method	AGE ($\cdot 10^{-2}$) ↓	AUC ↑
Ours (ZoomOut)*	2.40	95.12%
SFM*	2.46	95.22%
Ours (FMaps)	6.23	88.21%
Ours (ZoomOut)	5.87	88.82%
SFM	16.84	70.19%
ARM (ZoomOut)	41.59	30.34%
IRM (ZoomOut)	8.30	84.62%
FMaps	28.31	56.15%
ZoomOut	29.36	52.02%

Table 1: Average geodesic error and area under the accuracy curve for each tested method on the SHREC19 challenge pairs. The top rows (denoted by *) show the performance of our algorithm and SFM under the assumption of having a ground truth functional map for initializing ZoomOut.

⁴ Code available at <https://github.com/filthynobleman/rematching>

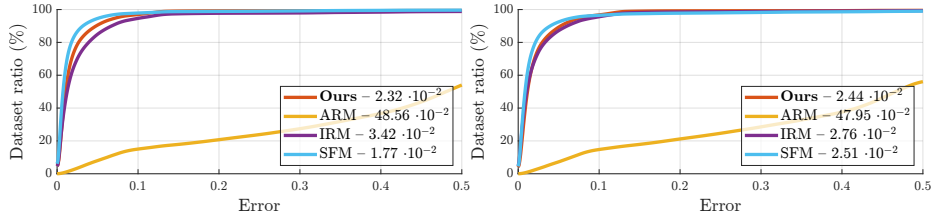


Fig. 6: Accuracy curves and average geodesic error on the TOSCA dataset (left) and the BadTOSCA dataset (right) for the tested methods.

IRM, ARM, and our algorithm) to remove small disconnected components made of few triangles before computing the extension map discussed in Section 3.3. This post-processing step cannot be performed for SFM, due to the nature of its sampling strategy. The benefits of this post-processing will be discussed in Section 4.3.

We summarize the results of our experiment in Figure 5, where we show the accuracy curves for all methods on the SHREC19 connectivity track benchmark, following the paradigm proposed by Kim *et al.* [24]. To prove the time efficiency of our pipeline, we also measure the time taken by every method on each pair of shapes in the benchmark (including the remeshing step), and we show the cumulative curves of the execution times over the entire dataset in a similar way with respect to the accuracy curves.

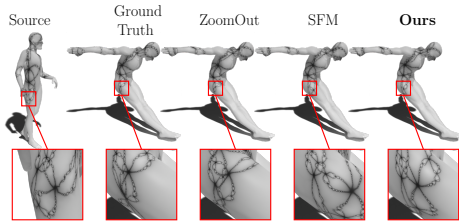


Fig. 7: Coordinate transfer between non-isometric shapes. The coordinates are used to generate highly complex patterns, whose sensitivity to the input enhances the transfer errors.

appreciated in Table 1, where we report the average geodesic error (normalized with respect to the shape diameter) and the area under the accuracy curves. In both Figure 5 and Table 1 we have also considered the idealized scenario where our method with ZoomOut as backend and SFM are initialized with a ground truth functional map. In this case, the two approaches obtain comparable results.

We also compare our method with the other approaches on the TOSCA and BadTOSCA datasets, summarizing the results in Figure 6. In the figure, we show

We see that all the low-resolution approaches provide better time and quality performance than the full-resolution methods, as the simpler geometry makes it easier and faster to align the Laplacian eigenfunctions. ARM is the only exception, since the anisotropic remeshing produces many skew triangles and a singular Laplacian matrix for all the shapes. Furthermore, our technique (with both the FMaps and ZoomOut backend) outperforms SFM and achieves better results than using the IRM remeshing strategy. This difference can be better

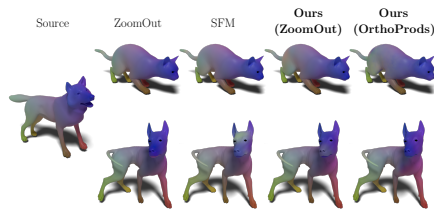


Fig. 8: Inter-class coordinate transfer from the wolf model to a cat and a dog. Our pipeline is not constrained by the standard Laplacian basis and can be used with other approaches as well.

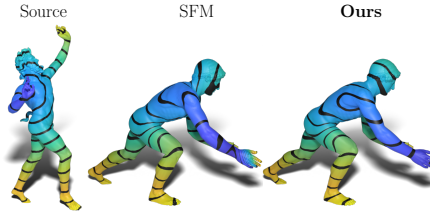


Fig. 9: Comparison of function transfer between very high-resolution models with SFM and our approach.

the accuracy curves for the tested methods on both datasets, also providing the average geodesic error normalized by the shape diameter. We see that under the favorable and unrealistic conditions offered by TOSCA, where shapes of the same class are almost perfectly isometric, SFM achieves the best results. However, when we consider the slightly altered geometry of BadTOSCA, SFM suffers a performance drop of about 50%, while our method shows more stability.

A significant benefit of the functional map approach is its independence from the difference in resolution between the source and target shape. Transferring a function with functional maps yields a much smoother and neat result than directly a point-wise correspondence. In this regard, we test our method in transferring coordinate functions between isometric and non-isometric shapes with substantial differences in resolution and triangulation. In the example from Figure 7 we generated a complex procedural fractal pattern as a function of the coordinates [33]. This pattern is very sensitive to the input, so even slight errors are highly enhanced. We see that SFM cannot precisely map the coordinates of the source mesh, shifting around and distorting the details of the pattern.

Furthermore, the example in Figure 8 shows that we can also deal with inter-class function transfer with our pipeline, outperforming SFM in a visual comparison. In the example, we also showcase that we should not necessarily rely on the standard Laplacian basis as a backend. Instead of using ZoomOut to transfer functions between the low-resolution meshes, we rely on the orthogonalized eigenproducts basis introduced by Maggioni *et al.* [34], extending it with the same technique described in Section 3.2.

Finally, to ensure that our method can scale to high-resolution meshes, we tested it using two very high-resolution 3D scans of real statues: a dancing faun counting $\sim 750k$ vertices and a warrior from the temple of Aphaea counting $\sim 3.5M$ vertices (see Figures 1 and 9). Since SFM can also deal with dense shapes, we compare it against our method. Figure 9 shows an example of the transfer of a geodesic distance function. While our method can faithfully transfer the function, SFM produces incoherent distances (*e.g.*, wrong isolines on the head and near the shoulder) and evident artefacts (*e.g.*, wrong distance on the hand).

4.2 Scalable performance

In Section 4.1 we saw that general-purpose remeshing solutions are not an optimal choice for the matching task. However, Lescoat *et al.* [28] provided a specialized mesh decimation algorithm (SMD) that reduces the complexity of the mesh while still preserving the Laplacian eigenfunctions. Furthermore, Liu *et al.* [31] proposed an elegant solution for simplifying a shape based on intrinsic error metric (IEM), which also builds a map for extending scalar functions using geodesic barycentric coordinates. We perform the evaluation on the SHREC19 dataset [37], remeshing the surfaces to 3k vertices with all the methods.

To achieve the desired scalability, the remeshing step in our pipeline must be fast enough to maintain the improvement obtained by reducing the size of the eigenproblem. In this regard, Figure 10 shows that our technique achieves better time performance than the other remeshing algorithms and the sampling strategy of SFM. In contrast, SMD and IEM are decimation algorithms that iteratively reduce the mesh size, making them orders of magnitude slower and unsuitable for large meshes. Moreover, SMD requires a pre-computation of the Laplacian eigenbasis on the original surface to guide the decimation process and preserve the spectrum, introducing additional time complexity.

4.3 Benefits of remeshing

In Section 4.1, we discussed how we remove small disconnected components after the remeshing step. The example in Figure 11 shows a comparison of mapping

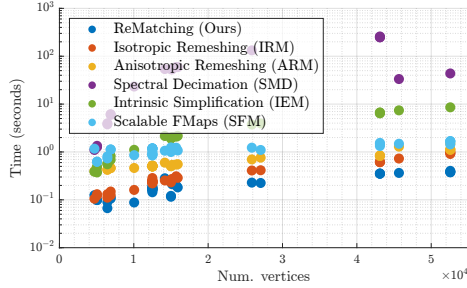


Fig. 10: Execution time (logarithmic scale on the y -axis) of different remeshing algorithms plotted against the number of vertices of the input shape (x -axis). For SFM we show the time taken for generating the sampling.

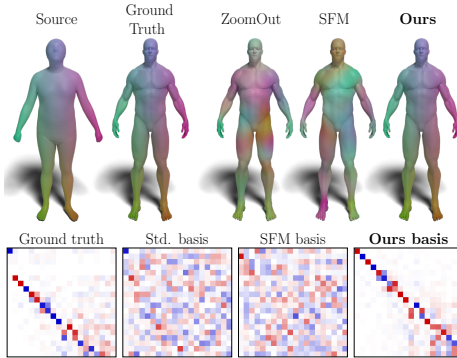


Fig. 11: Top row: coordinate transfer between a pair from SHREC19 using ZoomOut on full-resolution meshes, SFM and our approach. Bottom row: The ground truth 20×20 functional map, compared with the alignment computed with FMaps [46] from the full Laplacian eigenbasis, the SFM extended basis and our extended basis.

between a low-resolution mesh (Source), with less than 7k vertices, to a detailed mesh where eyes are represented with two disconnected components, each composed of 462 vertices out of the 27k of the entire mesh ($\sim 2\%$ of the total vertices). This is enough geometry to attract energy at lower frequencies, and when we try to align the first 20 eigenfunctions with FMaps [46], the resulting functional map turns out to be mostly noise. Initializing ZoomOut with such a map leads to meaningless correspondence, and SFM does not mitigate this issue, as it tries to preserve the original spectrum as much as possible. In contrast, by removing these components and mapping them to their closest surface points, we introduce a small error in the final correspondence, but at the same time, we ease the alignment of the initial eigenbases, resulting in more accurate and meaningful mapping. This is evident in the bottom row of Figure 11, where the ground truth functional map presents a gap in the alignment of the eigenfunctions due to the impossibility of mapping the disconnected eyes of the target shape in any point of the source shape. Using the real eigenfunctions or a close approximation yields an intense noise in the map, producing the meaningless mapping of ZoomOut and SFM. Instead, by mapping the eyes to their nearest point on the head, we can meaningfully align our extended basis.

5 Conclusions

We presented a new pipeline for scalable shape matching that relies on a low-resolution representation to compute a functional map between dense shapes efficiently. A core part of this pipeline is our novel remeshing algorithm, which efficiently computes an intrinsic Delaunay triangulation of uniform surface sampling. This procedure is applied to reduce the computational cost of the functional maps framework and make it efficient also on very dense meshes. Furthermore, we use a fast and intuitive technique for extending scalar functions from low-resolution meshes to their dense counterparts, extending the computed correspondence to the original very high-resolution shapes. Our experimental evaluation proved that our procedure is effective on various types of shapes at different resolutions, outperforming other state-of-the-art solutions and solving topological issues related to the alignment of the Laplacian eigenbasis.

While using a triangulated mesh as a low-resolution representation could, in principle, be exploited with other types of shape-matching pipelines, we only tested it within the functional maps framework, using the Laplace-Beltrami eigenbasis or bases derived from it. While the current pipeline binds the method’s applicability to the setting of functional correspondences, we believe that studying its behaviour with different “backends”, as well as exploring the relationships between the low-resolution and the high-resolution correspondences, would be an interesting matter for future investigations.

Acknowledgements

This work was partially supported by the Italian Ministry of Education, Universities and Research under the grant *Dipartimenti di Eccellenza 2023-2027* of the Department of Informatics, Systems and Communication of the University of Milano-Bicocca and by the PRIN 2022 project *GEOPRIDE Geometric primitive fitting on 3D data for geometric analysis and 3D shapes*. This work is also supported by the ERC grant no.802554 *SPECGEO* and PRIN 2020 project no.2020TA3K9N *LEGO.AI*. We acknowledge the support of NVIDIA Corporation with the RTX A5000 GPUs granted through the Academic Hardware Grant Program to the University of Milano-Bicocca for the project *Learned representations for implicit binary operations on real-world 2D-3D data*.

References

1. Aflalo, Y., Brezis, H., Kimmel, R.: On the optimality of shape and data representation in the spectral domain. *SIAM Journal on Imaging Sciences* **8**(2), 1141–1160 (2015)
2. Amenta, N., Choi, S., Dey, T.K., Leekha, N.: A simple algorithm for homeomorphic surface reconstruction. In: *Proceedings of the sixteenth annual symposium on Computational geometry*. pp. 213–222 (2000)
3. Bobenko, A.I., Springborn, B.A.: A discrete laplace–beltrami operator for simplicial surfaces. *Discrete & Computational Geometry* **38**(4), 740–756 (2007)
4. Botsch, M., Pauly, M., Ross, C., Bischoff, S., Kobbelt, L.: Geometric modeling based on triangle meshes. In: *ACM SIGGRAPH 2006 Courses*, pp. 1–es (2006)
5. Bowyer, A.: Computing Dirichlet tessellations*. *The Computer Journal* **24**(2), 162–166 (01 1981)
6. Bronstein, A., Bronstein, M., Kimmel, R.: *Numerical Geometry of Non-Rigid Shapes*. Springer, New York, NY (2008)
7. Burghard, O., Dieckmann, A., Klein, R.: Embedding shapes with green’s functions for global shape matching. *Computers & Graphics* **68**, 1–10 (2017)
8. Crane, K., Weischedel, C., Wardetzky, M.: Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Transactions on Graphics (TOG)* **32**(5), 1–11 (2013)
9. Delaunay, B., et al.: Sur la sphere vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennykh Nauk* **7**(793-800), 1–2 (1934)
10. Deng, B., Yao, Y., Dyke, R.M., Zhang, J.: A survey of non-rigid 3d registration. In: *Computer Graphics Forum*. vol. 41, pp. 559–589. Wiley Online Library (2022)
11. Dijkstra, E.W.: A note on two problems in connexion with graphs. *Numerische Mathematik* **1**, 269–271 (1959)
12. Donati, N., Corman, E., Melzi, S., Ovsjanikov, M.: Complex functional maps: A conformal link between tangent bundles. In: *Computer Graphics Forum*. vol. 41, pp. 317–334. Wiley Online Library (2022)
13. Donati, N., Sharma, A., Ovsjanikov, M.: Deep geometric functional maps: Robust feature learning for shape correspondence. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8592–8601 (2020)
14. Dyer, R., Zhang, H., Möller, T.: Voronoi-delaunay duality and delaunay meshes. In: *Proceedings of the 2007 ACM symposium on Solid and physical modeling*. pp. 415–420 (2007)

15. Edelsbrunner, H., Shah, N.R.: Triangulating topological spaces. In: Proceedings of the tenth annual symposium on Computational geometry. pp. 285–292 (1994)
16. Eynard, D., Rodola, E., Glashoff, K., Bronstein, M.M.: Coupled functional maps. In: 2016 Fourth International Conference on 3D Vision (3DV). pp. 399–407. IEEE (2016)
17. Ezuz, D., Ben-Chen, M.: Deblurring and denoising of maps between shapes. In: Computer Graphics Forum. vol. 36, pp. 165–174. Wiley Online Library (2017)
18. Fei, Y., Chen, S., Su, D., Luo, J., Li, M.: A new algorithm for repairing non-manifold surfaces. In: 2013 IEEE 10th International Conference on High Performance Computing and Communications & 2013 IEEE International Conference on Embedded and Ubiquitous Computing. pp. 1704–1708. IEEE (2013)
19. Gao, M., Roetzer, P., Eisenberger, M., Löhner, Z., Moeller, M., Cremers, D., Bernard, F.: Sigma: Scale-invariant global sparse shape matching. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 645–654 (2023)
20. Guennebaud, G., Jacob, B., et al.: Eigen v3. <http://eigen.tuxfamily.org> (2010)
21. Hoppe, H., DeRose, T., Duchamp, T., McDonald, J., Stuetzle, W.: Mesh optimization. In: Proceedings of the 20th annual conference on Computer graphics and interactive techniques. pp. 19–26 (1993)
22. Jacobson, A., Panozzo, D., et al.: libigl: A simple C++ geometry processing library (2018), <https://libigl.github.io/>
23. Jiang, Z., Schneider, T., Zorin, D., Panozzo, D.: Bijective projection in a shell. *ACM Transactions on Graphics (TOG)* **39**(6), 1–18 (2020)
24. Kim, V.G., Lipman, Y., Funkhouser, T.: Blended intrinsic maps. *ACM transactions on graphics (TOG)* **30**(4), 1–12 (2011)
25. Kimmel, R., Sethian, J.A.: Computing geodesic paths on manifolds. *Proceedings of the national academy of Sciences* **95**(15), 8431–8435 (1998)
26. Kressner, D.: Numerical Methods for General and Structured Eigenvalue Problems. *Lecture Notes in Computational Science and Engineering*, Springer Berlin Heidelberg (2006)
27. Leibon, G., Letscher, D.: Delaunay triangulations and voronoi diagrams for riemannian manifolds. In: Proceedings of the sixteenth annual symposium on Computational geometry. pp. 341–349 (2000)
28. Lescoat, T., Liu, H.T.D., Thiery, J.M., Jacobson, A., Boubekur, T., Ovsjanikov, M.: Spectral mesh simplification. In: Computer Graphics Forum. vol. 39, pp. 315–324. Wiley Online Library (2020)
29. Levy, B.: Laplace-beltrami eigenfunctions towards an algorithm that "understands" geometry. In: IEEE International Conference on Shape Modeling and Applications 2006 (SMI'06). pp. 13–13 (2006)
30. Lévy, B., Zhang, H.: Spectral mesh processing. In: ACM SIGGRAPH 2010 Courses, pp. 1–312 (2010)
31. Liu, H.T.D., Gillespie, M., Chislett, B., Sharp, N., Jacobson, A., Crane, K.: Surface simplification using intrinsic error metrics. *ACM Trans. Graph.* **42**(4) (jul 2023). <https://doi.org/10.1145/3592403>, <https://doi.org/10.1145/3592403>
32. Liu, H.T.D., Jacobson, A., Ovsjanikov, M.: Spectral coarsening of geometric operators. *ACM Transactions on Graphics (TOG)* **38**(4) (jul 2019)
33. Maggioli, F., Baieri, D., Melzi, S., Rodolà, E.: Newton's fractals on surfaces via bicomplex algebra. In: ACM SIGGRAPH 2022 Posters, pp. 1–2 (2022)
34. Maggioli, F., Melzi, S., Ovsjanikov, M., Bronstein, M.M., Rodolà, E.: Orthogonalized fourier polynomials for signal approximation and transfer. In: Computer Graphics Forum. vol. 40, pp. 435–447. Wiley Online Library (2021)

35. Magnet, R., Ovsjanikov, M.: Scalable and efficient functional map computations on dense meshes. In: *Computer Graphics Forum*. vol. 42, pp. 89–101. Wiley Online Library (2023)
36. Marin, R., Corona, E., Pons-Moll, G.: Nicp: Neural icp for 3d human registration at scale. In: *European Conference on Computer Vision*. Springer (2024)
37. Melzi, S., Marin, R., Rodolà, E., Castellani, U., Ren, J., Poulenard, A., Wonka, P., Ovsjanikov, M.: Shrec 2019: Matching humans with different connectivity. In: *Eurographics Workshop on 3D Object Retrieval*. vol. 7, p. 3. The Eurographics Association (2019)
38. Melzi, S., Ren, J., Rodolà, E., Sharma, A., Wonka, P., Ovsjanikov, M.: Zoomout: spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics (TOG)* **38**(6), 1–14 (2019)
39. Melzi, S., Rodolà, E., Castellani, U., Bronstein, M.M.: Localized manifold harmonics for spectral shape analysis. In: *Computer Graphics Forum*. vol. 37, pp. 20–34. Wiley Online Library (2018)
40. Meyer, M., Desbrun, M., Schröder, P., Barr, A.H.: Discrete Differential-Geometry Operators for Triangulated 2-Manifolds. In: *Visualization and mathematics III*, pp. 35–57. Springer, New York, NY (2003)
41. Mitchell, J.S., Mount, D.M., Papadimitriou, C.H.: The discrete geodesic problem. *SIAM Journal on Computing* **16**(4), 647–668 (1987)
42. Nasikun, A., Brandt, C., Hildebrandt, K.: Fast approximation of laplace-beltrami eigenproblems. In: *Computer Graphics Forum*. vol. 37, pp. 121–134. Wiley Online Library (2018)
43. Nasikun, A., Hildebrandt, K.: The hierarchical subspace iteration method for laplace-beltrami eigenproblems. *ACM Transactions on Graphics (TOG)* **41**(2), 1–14 (2022)
44. Nivolières, V., Lévy, B., Geuzaine, C.: Anisotropic and feature sensitive triangular remeshing using normal lifting. *Journal of Computational and Applied Mathematics* **289**, 225–240 (2015)
45. Nogneng, D., Melzi, S., Rodola, E., Castellani, U., Bronstein, M., Ovsjanikov, M.: Improved functional mappings via product preservation. In: *Computer Graphics Forum*. vol. 37, pp. 179–190. Wiley Online Library (2018)
46. Nogneng, D., Ovsjanikov, M.: Informative descriptor preservation via commutativity for shape matching. In: *Computer Graphics Forum*. vol. 36, pp. 259–267. Wiley Online Library (2017)
47. Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L.: Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (ToG)* **31**(4), 1–11 (2012)
48. Ovsjanikov, M., Corman, E., Bronstein, M., Rodolà, E., Ben-Chen, M., Guibas, L., Chazal, F., Bronstein, A.: Computing and processing correspondences with functional maps. In: *SIGGRAPH ASIA 2016 Courses*, pp. 1–60 (2016)
49. Peyré, G., Cohen, L.D.: Geodesic remeshing using front propagation. *International Journal of Computer Vision* **69**, 145–156 (2006)
50. Pinkall, U., Polthier, K.: Computing Discrete Minimal Surfaces and their Conjugates. *Experimental mathematics* **2**(1), 15–36 (1993)
51. Ren, J., Melzi, S., Ovsjanikov, M., Wonka, P., et al.: Maptree: recovering multiple solutions in the space of maps. *ACM Transactions on Graphics (TOG)* **39**(6), 264–1 (2020)
52. Ren, J., Melzi, S., Wonka, P., Ovsjanikov, M.: Discrete optimization for shape matching. In: *Computer Graphics Forum*. vol. 40, pp. 81–96. Wiley Online Library (2021)

53. Ren, J., Poulenard, A., Wonka, P., Ovsjanikov, M.: Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics (ToG)* **37**(6), 1–16 (2018)
54. Reuter, M., Wolter, F.E., Peinecke, N.: Laplace–beltrami spectra as ‘shape-dna’ of surfaces and solids. *Computer-Aided Design* **38**(4), 342–366 (2006)
55. Rivin, I.: Euclidean structures on simplicial surfaces and hyperbolic volume. *Annals of mathematics* **139**(3), 553–580 (1994)
56. Rustamov, R.M.: Barycentric coordinates on surfaces. In: *Computer Graphics Forum*. vol. 29, pp. 1507–1516. Wiley Online Library (2010)
57. Sahillioğlu, Y.: Recent advances in shape correspondence. *The Visual Computer* **36**(8), 1705–1721 (2020)
58. Shoham, M., Vaxman, A., Ben-Chen, M.: Hierarchical functional maps between subdivision surfaces. In: *Computer Graphics Forum*. vol. 38, pp. 55–73. Wiley Online Library (2019)
59. Vaxman, A., Ben-Chen, M., Gotsman, C.: A multi-resolution approach to heat kernels on discrete surfaces. In: *ACM SIGGRAPH 2010 papers*, pp. 1–10 (2010)
60. Watson, D.F.: Computing the n-dimensional Delaunay tessellation with application to Voronoi polytopes*. *The Computer Journal* **24**(2), 167–172 (01 1981)