

Volumetric Functional Maps

Filippo Maggioli
Pegaso University
Naples, Italy

maggioli.filippo@gmail.com

Simone Melzi
University of Milano-Bicocca
Milan, Italy

simone.melzi@unimib.it

Marco Livesu
CNR IMATI
Genoa, Italy

marco.livesu@gmail.com




Figure 1. Visual representation of our pipeline. The eigenfunctions of the LBO for volume meshes (left) are used to compute a volumetric functional map (middle). Basis alignment is exploited for several tasks: volumetric correspondences, piece-wise linear maps, and volumetric segmentation transfer (right).

Abstract

Computing volumetric correspondences between 3D shapes is a prominent tool for medical and industrial applications. In this work, we pave the way for spectral volume mapping, extending for the first time the surface-based functional maps framework. We show that the eigenfunctions of the volumetric Laplace operator define a functional space that is suitable for high-quality signal transfer. We also experiment with various techniques that edit this functional space, porting them to volume domains. We validate our method on novel volumetric datasets and on tetrahedralizations of well established surface datasets, also showcasing practical applications involving both discrete and continuous signal mapping, for segmentation transfer, mesh connectivity transfer and solid texturing. Finally, we show that the volumetric spectrum greatly improves the accuracy for classical shape matching tasks among surfaces, consistently outperforming surface-only spectral methods.

1. Introduction

The computation of volumetric correspondences plays a central role for many applications. Volume mappings enable non-invasive medical diagnosis, by warping the digital copy of a human organ into its canonical position without extracting it from the patient’s body [1]. They are also important for meshing [91], statistical shape analysis [44] and Industry 4.0, where physical simulations computed on predictive digital twins must be aligned with other digital views for optimized decision making and failure avoidance [96].

Maps between 3D objects are often represented either by exploiting one-to-one per vertex correspondences between two simplicial meshes with same connectivity, defining a so called piece-wise map (Sec. 2.1), or putting into correspondence two sets of basis functions to express the same signal in two alternative domains, as done by spectral approaches for surface meshes [87]. Both piece-wise and spectral approaches have been extensively studied in the surface setting and can nowadays be considered rather mature, as proved by numerous academic and commercial tools that incorporate them. Conversely, the equivalent task of computing volumetric correspondences largely remains a scientific open problem [91, §8.2]. It is worth noticing that the LDDMM framework has been successfully applied to match dense volumetric data (e.g., MRI/CRT scans), although diffeomorphic maps are limited to domains with grid-like structures and hardly generalize to more flexible representations like tetrahedral meshes [16, 17, 92, 112].

In recent years many attempts have been made to extend piece-wise linear approaches to volumes. This task has proved to be extremely challenging and no fully robust and efficient methods have been devised yet (Sec. 2.1). To our knowledge no attempt has been made to extend functional mapping to volumes yet (Sec. 2.2). We propose it for the first time and show that, conversely from the piece-wise linear setting, functional mapping naturally scales to volumes, providing a theoretically sound and practically useful platform for the computation of volumetric correspondences.

Using the spectrum of the volumetric Laplacian, we define a functional space that is completely agnostic of the

underlying discrete mesh, and we use it to create correspondences between solid meshes with different density and connectivity. We also show that many recent developments used to improve the functional maps framework [68, 79] can be exploited on volumes too, allowing to find an optimal balance between map quality and performance.

We validate our volumetric pipeline (Fig. 1) on a variety of datasets, including tetrahedralizations of prominent reference benchmarks widely used for shape correspondence tasks. Besides assessing the accuracy and performances of our tool on the ground truth (Sec. 5), we showcase various practical applications, including connectivity transfer of a given volume mesh to a target domain (Fig. 7), segmentation transfer from a statistical medical shape to acquired organs (Fig. 8) and solid texturing (Fig. 5).

All in all, this article opens for the first time to the computation of volumetric correspondences with spectral approaches. To this end, we believe our contributions will foster new research in this direction, positioning volume spectral mapping as a flexible and practical alternative to existing piece-wise linear approaches. To facilitate this endeavor, we release our reference code at *anonymous*.

2. Related Works

We list the main approaches to compute piece-wise maps and functional maps highlighting, for both classes of methods, existing attempts (or lack thereof) to scale to volumes.

2.1. Piece-wise Maps

Piece-wise maps define correspondences between two meshes by starting from explicit per vertex correspondences and extending the map inside mesh elements by means of barycentric interpolation. This operation is uniquely defined for simplices in any dimension [42], therefore it works for both surface (triangle) and volume (tetrahedral) meshes.

A variety of algorithms can robustly generate piece-wise surface maps to elementary domains such as triangles [33], convex polygons [35, 65, 109, 118], star-shaped polygons [64] and spheres [38, 94]. Maps between general shapes are then obtained via composition, using these elementary shapes as intermediate domains [6, 37, 46, 49, 52, 56, 63, 95, 104–107, 110, 120].

Lifting this idea to volumes is possible in principle but extremely challenging in practice. Even constructing an injective map to a convex polyhedron has proved to be surprisingly difficult [18, 61, 84]. Surface methods such as the Tutte embedding are known to fail on general volume meshes [7, 19]. To date, the only robust volume embedding algorithms are [13, 41, 83, 84], but all of them are unsuitable for real applications, due the massive use of exact numerical constructions and exponential mesh refinement [75].

A less robust but more practical alternative consists in initializing a non injective map directly with a fast ap-

proximation algorithm, and then relocate the vertices to fix flipped or vanishing elements. This problem, often referred to as mesh *untangling* [47, 91], is ill-posed in general, as valid solutions may not exist for a fixed mesh connectivity [62]. Existing volume approaches operate on convex subsets of the feasible region [48], minimize flip-preventing energies [2, 28, 36], or project onto the space of bounded distortion mappings [5]. Due to the problem’s non linearity and the absence of a feasible starting solution, none of these methods guarantees the correctness of the result. Multiple failure cases have been reported by testing the most prominent existing algorithms on available datasets [83, §6.3]. In Sec. 5.2 we show that bootstrapping this pipeline with our functional mapping dramatically increases the success rate of state of the art volume untangling algorithms.

2.2. Functional Maps

Given two discrete surfaces, shape matching aims to estimate a semantic point-wise correspondence between them. The more general (and challenging) scenario occurs when the two shapes are related by a non-rigid deformation. Many applications are related to this task, from statistical shape analysis [11] and medical imaging [72] to deformation transfer [114, 115], among many others [22, 103]. Among different approaches, functional mapping [87, 88] has attracted significant attention. This framework aims to solve the shape matching problem by estimating a correspondence between functions defined on the surfaces, rather than point-wise correspondences, which are hard to optimize for. Exploiting a basis for the space of functions, for which a standard choice arises from the extension of Fourier theory to non-Euclidean domains [54, 116], the correspondence among shapes can be compactly encoded in a matrix with dimensions equal to the basis size. Following this direction, a variety of works try to improve the framework by defining additional functional constraints [27, 85, 98], proposing alternative bases [68, 73, 74, 77, 86], or defining alternative procedures to extract point correspondences from a given functional map [32, 79, 89, 99, 100, 102]. Recently, two alternatives for scaling the functional maps approach to high-resolution meshes [70, 71] addressed the scalability issues of functional mapping. The functional maps framework additionally gives rise to a family of data-driven approaches that exploit machine learning techniques to solve the shape matching problem exploiting the functional representation as a prior [9, 14, 15, 26, 58, 108]. Despite its success, to the best of our knowledge, this framework has never been applied to volumes, leaving its potential impact fully unexplored in volume data processing.

3. Background and Notation

We discretize volumes as tetrahedral meshes $\mathcal{M} = (V_{\mathcal{M}}, T_{\mathcal{M}})$, where: $V_{\mathcal{M}} \subset \mathbb{R}^3$ is a set of vertices in 3D

space; $T_{\mathcal{M}} \subset V_{\mathcal{M}}^4$ is a set of tetrahedra connecting the vertices in $V_{\mathcal{M}}$.

We refer to the external surface of $\mathcal{M}$ as the triangular mesh $\partial\mathcal{M} = (V_{\partial\mathcal{M}}, F_{\partial\mathcal{M}})$, where $V_{\partial\mathcal{M}} \subset V_{\mathcal{M}}$ is the set of surface vertices of $\mathcal{M}$; $F_{\partial\mathcal{M}} \subset V_{\partial\mathcal{M}}^3$ is the set of surface triangles defined by the tetrahedra in $T_{\mathcal{M}}$.

The inset figure depicts a section of a tetrahedral mesh, with the original external surface shown in transparency.

A scalar function $f : V_{\mathcal{M}} \rightarrow \mathbb{R}$ over a volume mesh $\mathcal{M}$ is discretized as a vector $\mathbf{f} \in \mathbb{R}^{|V_{\mathcal{M}}|}$, and its restriction $f_{\partial} : V_{\partial\mathcal{M}} \rightarrow \mathbb{R}$ to the external surface $\partial\mathcal{M}$ is discretized as a vector $\mathbf{f}_{\partial} \in \mathbb{R}^{|V_{\partial\mathcal{M}}|}$. Given a tetrahedral (or triangular) mesh $\mathcal{M}$, and given two vertices $v_i, v_j \in \mathcal{M}$, we refer to the geodesic distance between v_i and v_j over $\mathcal{M}$ as $d_{\mathcal{M}}(v_i, v_j)$. There exist several algorithms to estimate such a distance on discrete meshes [21]. For the sake of simplicity, in the remainder of the paper we will always approximate $d_{\mathcal{M}}(v_i, v_j)$ using Dijkstra’s algorithm [25] on the edge connectivity of $\mathcal{M}$.

Shape Matching and Correspondences. Given two surface meshes $\mathcal{M} = (V_{\mathcal{M}}, F_{\mathcal{M}})$ and $\mathcal{N} = (V_{\mathcal{N}}, F_{\mathcal{N}})$, the problem of *shape matching* is the problem of finding a *correspondence* $\pi : V_{\mathcal{M}} \rightarrow V_{\mathcal{N}}$ that maps vertices of $\mathcal{M}$ into vertices of $\mathcal{N}$ [66]. In most cases, the correspondence π has to satisfy semantic constraints or has to be an *isometry* [22, 103], that is a correspondence π preserving the geodesic distances between pairs of vertices. Namely

$$\forall v_i, v_j \in V_{\mathcal{M}}, \quad d_{\mathcal{M}}(v_i, v_j) = d_{\mathcal{N}}(\pi(v_i), \pi(v_j)). \quad (1)$$

Functional Maps. The problem of finding an isometric correspondence, in its general setting, is NP-hard [10]. Many solutions have been developed for computing an approximation. Among these, the *functional maps* approach has grown popular due to its high efficiency and the smoothness of the resulting correspondences [87].

Instead of searching directly for a correspondence π between two surfaces $\mathcal{M}$ and $\mathcal{N}$, the functional maps approach aims at finding a linear map $T_{\pi} : \mathcal{F}(\mathcal{N}) \rightarrow \mathcal{F}(\mathcal{M})$, induced by π , from the space $\mathcal{F}(\mathcal{N})$ of real-valued functions over $\mathcal{N}$ to the space $\mathcal{F}(\mathcal{M})$ of real-valued functions over $\mathcal{M}$. Given any function $f : V_{\mathcal{N}} \rightarrow \mathbb{R}$, the map T_{π} gives the corresponding function $T_{\pi}(f) : V_{\mathcal{M}} \rightarrow \mathbb{R}$ that maps $v_i \mapsto f(\pi(v_i))$, for each $v_i \in V_{\mathcal{M}}$.

Given two sets of (orthogonal) basis functions $\Phi_{\mathcal{M}}, \Phi_{\mathcal{N}}$ for the spaces $\mathcal{F}(\mathcal{M}), \mathcal{F}(\mathcal{N})$, respectively, a function $f : V_{\mathcal{N}} \rightarrow \mathbb{R}$ can be expressed as a linear combination of $\Phi_{\mathcal{N}}$

$$f = \sum_{i=0}^{\infty} \alpha_i \phi_{\mathcal{N}}^{(i)} = \sum_{i=0}^{\infty} \langle f, \phi_{\mathcal{N}}^{(i)} \rangle_{\mathcal{N}} \phi_{\mathcal{N}}^{(i)}, \quad (2)$$

where $\phi_{\mathcal{N}}^{(i)}$ is the i -th basis function in $\Phi_{\mathcal{N}}$ and $\langle \cdot, \cdot \rangle_{\mathcal{N}}$ denotes the inner product between functions over $\mathcal{N}$. The coefficients α_i denotes the projection of f onto the basis functions $\phi_{\mathcal{N}}^{(i)}$. Usually, $\Phi_{\mathcal{M}}, \Phi_{\mathcal{N}}$ are chosen to be the eigenfunctions of the Laplace-Beltrami operator on $\mathcal{M}$ and $\mathcal{N}$. Since T_{π} is linear, the mapping of f onto $\mathcal{M}$ is given by

$$T_{\pi}(f) = \sum_{i=0}^{\infty} \alpha_i T_{\pi}(\phi_{\mathcal{N}}^{(i)}) = \sum_{i,j=0}^{\infty} \alpha_i c_{i,j} \phi_{\mathcal{M}}^{(j)}, \quad (3)$$

where $c_{i,j} = \langle T_{\pi}(\phi_{\mathcal{N}}^{(i)}), \phi_{\mathcal{M}}^{(j)} \rangle_{\mathcal{M}}$. By truncating the two bases to k functions, the coefficients $c_{i,j}$ can be stored into a matrix $\mathbf{C} \in \mathbb{R}^{k \times k}$, leading to the matrix equation

$$T_{\pi}(f) \approx \Phi_{\mathcal{M}} \mathbf{C} \Phi_{\mathcal{N}}^{\dagger} \mathbf{f}, \quad (4)$$

where $\dagger$ denotes the Moore-Penrose pseudo-inverse. Notably, $T_{\pi}(\Phi_{\mathcal{N}}) \approx \Phi_{\mathcal{M}} \mathbf{C}$, thus by knowing the matrix $\mathbf{C}$ it is possible to extract the unknown correspondence π via a nearest-neighbor search, $\text{NN}(\Phi_{\mathcal{N}}, \Phi_{\mathcal{M}} \mathbf{C})$.

4. Method

The functional maps framework is not bound to the domain of 2-dimensional surfaces. Indeed, the Laplace-Beltrami Operator (LBO) can be defined on Riemannian manifolds of arbitrary dimension, and it always admits a spectral decomposition. We discretize the LBO of a tetrahedral mesh $\mathcal{M} = (V_{\mathcal{M}}, T_{\mathcal{M}})$ using the n -dimensional cotangent formula [20, 57] which defines the discrete LBO as a pair of matrices $\mathbf{S}$ and $\mathbf{W}$, denoted as *stiffness* and *mass* matrices, respectively. Its spectral decomposition can be obtained by solving the generalized eigenproblem $\mathbf{S}\Phi = \mathbf{W}\Phi\Lambda$.

Implementation details and a discussion on the choice of the LBO are provided in Appendix A. We highlight that most applications we address involve analyzing information at the surface (*i.e.*, the boundary of a tetrahedral mesh), compelling us to impose Neumann’s boundary conditions to prevent zero-valued eigenfunctions on the surface.

Most successful algorithms for estimating the functional map between shapes rely on the property that intrinsic structures and processes behave similarly on similar surfaces. For instance, the invariance of the LBO basis and the heat diffusion under (quasi-)isometric deformations plays a key role in stable functional map estimation [85, 86, 88]. Despite this invariance holds for continuous volumetric shapes (see Fig. 2), in a discrete setting the position and connectivity of the vertices can heavily impact the numerical stability of the Laplacian operator [53]. Appendix B provides a detailed empirical study on the invariance of the LBO in the volumetric setting, and Appendix C contains a deeper discussion on isometric deformations between volumes.




Figure 2. The first 5 non-constant volumetric LBO eigenfunctions on two humanoid non-isometric shapes

4.1. Spectral Connectivity Transfer

The success of functional mapping is also due to its ability to smooth out signals and correspondences: a truncated basis filters out high-frequency details, producing smoothly interpolated signals. This smoothness, combined with the fact that the basis is usually much smaller than the number of vertices, allows to approximate the spectral representation of noisy and partial signals. This feature has been exploited by Melzi et al. [79] to efficiently approximate a functional map using only a reduced sampling of the original surface. In our setting, we take advantage of this property to transfer connectivity between volumes.

Functional Connectivity Transfer. Let $\mathcal{M} = (V_{\mathcal{M}}, T_{\mathcal{M}})$ be a tetrahedral mesh with surface $\partial\mathcal{M} = (V_{\partial\mathcal{M}}, F_{\partial\mathcal{M}})$. Let $\partial\mathcal{N} = (V_{\partial\mathcal{N}}, F_{\partial\mathcal{N}})$ be another surface in known correspondence with $\partial\mathcal{M}$ through the (surface) map $\pi : \partial\mathcal{M} \rightarrow \partial\mathcal{N}$. There are various options to obtain a tetrahedralization $\mathcal{N}$ of the interior of $\partial\mathcal{N}$ [24, 39]. However, by no means, its connectivity will match the connectivity of $\mathcal{M}$.

Given a surface $\partial\mathcal{N}$ with induced tetrahedral mesh $\mathcal{N}$, we exploit the known correspondence $\pi : \partial\mathcal{M} \rightarrow \partial\mathcal{N}$ to induce a volumetric correspondence $\pi' : \mathcal{M} \rightarrow \mathcal{N}$. Let $\Phi_{\mathcal{M}}, \Phi_{\mathcal{N}}$ be the LBO eigenbases of the volumes $\mathcal{M}$ and $\mathcal{N}$, respectively, and let us denote with $\Phi_{\partial\mathcal{M}} = (\Phi_{\mathcal{M}})_{\partial}, \Phi_{\partial\mathcal{N}} = (\Phi_{\mathcal{N}})_{\partial}$ their restrictions to the surfaces. We stress that $\Phi_{\partial\mathcal{M}}$ is not the LBO eigenbasis of $\partial\mathcal{M}$. The functional map $\mathbf{C}$ induced by π' must be such that

$$\Phi_{\mathcal{M}}\mathbf{C} = T_{\pi'}(\Phi_{\mathcal{N}}). \quad (5)$$

We consider that the rows of $\Phi_{\partial\mathcal{M}}$ (resp. $\Phi_{\partial\mathcal{N}}$) are a subset of the rows of $\Phi_{\mathcal{M}}$ (resp. $\Phi_{\mathcal{N}}$). Recalling that π is the restriction of π' , we get

$$\Phi_{\partial\mathcal{M}}\mathbf{C} = (\Phi_{\mathcal{M}}\mathbf{C})_{\partial} = (T_{\pi'}(\Phi_{\mathcal{N}}))_{\partial} = T_{\pi}(\Phi_{\partial\mathcal{N}}). \quad (6)$$




Figure 3. Our pipeline for extrapolating interior coordinates from surface correspondences. Given the surface map π (first row), we approximate the spectral embedding of the surface coordinates of $\partial\mathcal{N}$ using the boundary restriction of the eigenfunctions of $\mathcal{M}$ (second row). Using the eigenfunctions on the entire volume, we reconstruct the interior coordinates from the spectral embedding (third row), and use these coordinates to transfer the inner connectivity (fourth row).

In general, the size of the basis (*i.e.*, the number of columns of $\Phi_{\partial\mathcal{M}}$) is significantly smaller than the number of vertices (*i.e.*, the number of rows), even if we only consider the vertices at the surface. Intuitively, this suggests that the restricted eigenfunctions $\Phi_{\partial\mathcal{M}}$ at the surface are linearly independent. Indeed, classical results from control theory support this intuition, proving that the boundary traces (*i.e.*, the restriction to the boundary) of the eigenfunctions of any elliptic operator are linearly independent in a system with Neumann boundary conditions [51, 117] (see the Appendix D for a deeper analysis of the descriptive power of the boundary traces). Therefore, $\Phi_{\partial\mathcal{M}}$ admits left inverse and the volumetric functional map can be approximated via the known surface correspondence as $\mathbf{C} \approx \Phi_{\partial\mathcal{M}}^{\dagger}T_{\pi}(\Phi_{\partial\mathcal{N}})$.

The functional map $\mathbf{C}$ can then be used with the full spectrum to transfer the coordinates of $\mathcal{N}$ onto $\mathcal{M}$. Namely

$$T_{\pi'}(x_{\mathcal{N}}) = \Phi_{\mathcal{M}}\mathbf{C}\Phi_{\mathcal{N}}^{\dagger}x_{\mathcal{N}}, \quad (7)$$

and similarly for $y_{\mathcal{N}}$ and $z_{\mathcal{N}}$. By using the values $T_{\pi'}(x_{\mathcal{N}}), T_{\pi'}(y_{\mathcal{N}}), T_{\pi'}(z_{\mathcal{N}})$ as coordinates for the vertices of $\mathcal{M}$, we get a connectivity transfer from $\mathcal{M}$ onto $\mathcal{N}$.

Spectral Extrapolation of Coordinates. A watertight surface already encodes information about the enclosed volume. Since the LBO provides smoothness guarantees in signal reconstruction, it is a good candidate for extrapolating the interior coordinates of a volume from its surface.

Given the ground truth correspondence $\pi : \partial\mathcal{M} \rightarrow \partial\mathcal{N}$, we can transfer the surface coordinates $x_{\partial\mathcal{N}}, y_{\partial\mathcal{N}}, z_{\partial\mathcal{N}}$ of $\partial\mathcal{N}$ onto $\partial\mathcal{M}$ as $T_\pi(x_{\partial\mathcal{N}}), T_\pi(y_{\partial\mathcal{N}}), T_\pi(z_{\partial\mathcal{N}})$. We then consider the LBO eigenbasis $\Phi_{\mathcal{M}}$ and its restriction $\Phi_{\partial\mathcal{M}}$ to the surface. Following a reasoning analogous to the discussion in the previous paragraph, we have that

$$\begin{aligned} (T_{\pi'}(x_{\mathcal{N}}))_{\partial} &\approx (\Phi_{\mathcal{M}}\Phi_{\mathcal{M}}^\dagger T_{\pi'}(x_{\mathcal{N}}))_{\partial} = \Phi_{\partial\mathcal{M}}\Phi_{\mathcal{M}}^\dagger T_{\pi'}(x_{\mathcal{N}}), \\ T_\pi(x_{\partial\mathcal{N}}) &\approx \Phi_{\partial\mathcal{M}}\Phi_{\partial\mathcal{M}}^\dagger T_\pi(x_{\partial\mathcal{N}}). \end{aligned} \quad (8)$$

Since $(T_{\pi'}(x_{\mathcal{N}}))_{\partial} = T_\pi(x_{\partial\mathcal{N}})$, then it must also be $\Phi_{\mathcal{M}}^\dagger T_{\pi'}(x_{\mathcal{N}}) \approx \Phi_{\partial\mathcal{M}}^\dagger T_\pi(x_{\partial\mathcal{N}})$. Finally, the coordinates on the interior of $\mathcal{N}$ can be extrapolated by reconstructing with the full basis:

$$T_{\pi'}(x_{\mathcal{N}}) = \Phi_{\mathcal{M}}\Phi_{\partial\mathcal{M}}^\dagger T_\pi(x_{\partial\mathcal{N}}), \quad (9)$$

and analogously for $y_{\mathcal{N}}$ and $z_{\mathcal{N}}$. A visual representation of our pipeline is shown in Fig. 3.

Concrete examples of reliable connectivity transfer using both a functional transfer of interior coordinates and their extrapolation from the spectral embedding of the surface coordinates are depicted in Fig. 7 (resp. *Transfer* and *Extrapolation*). For this specific test, we extended the LBO eigenbasis with the Coordinates Manifold Harmonics (CMH) basis [73], which also includes three additional functions encoding the extrinsic per vertex xyz coordinates and is known to be better suited to reconstruct extrinsic information. Additional experiments on the whole dataset, considering both the LBO and CMH eigenbasis are reported in Sec. 5.2.

4.2. Volume-Aware Surface Correspondence

Besides purely volumetric applications, we observe that a good functional map among volumes also induces a good functional map among their external surfaces. Intuitively, a volumetric shape is a richer structure than a surface, yielding more information both locally and globally. Therefore, we guess that functional maps among tetrahedral meshes are more accurate and informative than maps among triangular meshes. This idea is grounded from the results of Raviv et al. [97], who have shown that using volume-informed shape descriptors has positive effects in computing correspondences between surfaces. As mentioned in Sec. 4.1, the functional map that relates the LBO eigenbases of two volumetric shapes also relates their restriction to the surfaces. Consequently, the volumetric map can be used to infer a correspondence between surfaces.




Figure 4. Left: average geodesic error curves on the pairs from Su *et al.* [113] resulting from the application of FMaps [85], ZoomOut [79], and Orthoprods [68]. Right: runtime for the three methods plotted against the number of vertices in the source mesh.

Given two surface meshes $\partial\mathcal{M}, \partial\mathcal{N}$, we compute the tetrahedral meshes $\mathcal{M}, \mathcal{N}$ representing their interior volume. We then compute the functional map $\mathbf{C}$ relating their eigenbases via $\Phi_{\mathcal{M}}\mathbf{C} = T_\pi(\Phi_{\mathcal{N}})$, and we can use that same functional map for defining the surface relation $(\Phi_{\mathcal{M}})_{\partial}\mathbf{C} = (T_\pi(\Phi_{\mathcal{N}}))_{\partial}$. The surface correspondence is eventually extracted via nearest neighbor search, $\text{NN}(\Phi_{\partial\mathcal{M}}\mathbf{C}, \Phi_{\partial\mathcal{N}})$.

5. Results and Applications

We implemented our tool in MATLAB, extending the surface only implementation of the functional map from [85] and ZoomOut [79]. Our code is available at <https://github.com/filthynobleman/vol-fmaps>.

5.1. Volumetric Functional Maps

We validate our prototype and assess its ability to devise high-quality volume maps by comparing our correspondences against the ground truth. Unlike the surface setting, where a constantly growing amount of public data has become available for validation, almost no volumetric datasets are available for this specific task. We consider a set of 40 pairs of volume meshes with identical connectivity, which thus define an injective piece-wise linear map as described in Sec. 2.1. This dataset, originally proposed by [113] and then used to validate [28], covers diverse classes of objects, such as: animals, humanoids, hands, and a variety of man-made objects like glasses, lamps, laptops and scissors. For each pair of objects, we compute a volumetric functional map and extract a point-to-point correspondence as described in Sec. 4. Then, for each vertex v , we measure the geodesic error as the geodesic distance $d_{\mathcal{N}}(\pi_{\text{gt}}(v), \pi(v))$ between its mapped point $\pi(v)$ and its ground truth counterpart $\pi_{\text{gt}}(v)$. In Fig. 4 we show the cumulative curves of the geodesic error obtained with different functional map estimation algorithms, as well as their runtime in relation to the vertex count. We also provide the average geodesic error (AGE) in the legend of the figure. For comparisons, we use the well-established algorithm based on descriptor preservation (FMap) [85] and ZoomOut [79], using both the stan-




Figure 5. Coordinate transfer between two volumetric shapes using the Orthoprods basis. The normalized coordinates are treated as RGB channels for visualization (first column). For further validation, the transferred coordinates are also used for generating an error-sensitive procedural texture, which is visualized on different slices of the volume (second to fourth columns).

Table 1. Average percentage of flips using the methods described in Sec. 4.1 at varying the basis size. The average runtime (in seconds) to compute the largest basis considered is also reported.

Method	5% eigs	10% eigs	15% eigs	20% eigs	Time (s)
Transfer (LBO)	4.63%	1.50%	1.18%	0.93%	1127
Transfer (CMH)	4.75%	0.84%	0.45%	0.34%	1127
Extrapolation (LBO)	4.39%	1.45%	1.15%	0.92%	554
Extrapolation (CMH)	4.90%	0.99%	0.56%	0.42%	554

dard and the fast implementation. For the fast implementation, we use the surface points as samples to validate the claim in Sec. 4.1 that the surface encodes information about the enclosed volume. We also test our framework replacing the LBO eigenbasis with the orthogonalized eigenproducts basis (Orthoprods) proposed by Maggioli et al. [68], but instead of using the approach described by the authors for the functional map estimation, we plug the basis into the ZoomOut pipeline. As expected, the extended Orthoprods basis has higher descriptive power, but this comes with additional computational cost for aligning the larger bases. In contrast, ZoomOut achieves the best trade-off between accuracy and runtime. Overall, our results demonstrate that the volumetric functional map framework produces high-quality maps, matching the performance obtained in the surface setting. This can be further appreciated in the signal transfer example from Fig. 5, where we transfer the coordinates function between two meshes with the Orthoprods bases and use the resulting values to compute an error-sensitive procedural texture [69]. All values reported in this subsection correspond to volumetric extensions of the considered metrics. A more detailed evaluation of the maps is provided in Appendix E.




Figure 6. Left: two of the 1405 failure cases of the untangling algorithm proposed by Garanzha et al. [36], which emerged from the challenging large scale validation described in [83, §6.3]. Right: computing an initial connectivity transfer with our tool to warm start the same algorithm allows to retain a fully bijective map, thus improving the overall robustness.

5.2. Connectivity Transfer

Representing volume maps using two tetrahedral meshes that share the same connectivity but have different vertex coordinates is a common choice for many practical applications involving digital fabrication [30, 59, 122], medicine [1] and hexahedral meshing [12, 91]. As detailed in Sec. 4.1, our framework can be used to transfer the connectivity of a given volume mesh onto a target domain, either by using a functional map to transfer the vertex coordinates (§4.1 *Functional Connectivity Transfer*) or by extrapolating the connectivity from its outer surface (§4.1 *Spectral Extrapolation of Coordinates*).

If the whole spectrum is considered, the LBO basis spans the entire functional space of the mesh and the functional map is one-to-one [102]. Computing the full spectrum is computationally expensive, hence only a fraction of the eigenfunctions are typically computed. We test our prototypes on the dataset of 40 meshes released by Su et al. [113] and empirically verify that the amount of inverted elements introduced in the map gracefully decays to zero for growing numbers of eigenfunctions (see Appendix F for a detailed analysis). On average, already using 5% of the spectrum the number of flipped cells falls below 5%. It consistently decreases as basis grows, and goes below 1% if 20% of the spectrum is considered (Tab. 1), reaching full injectivity when 25% of the spectrum is used (Fig. 7). We consider both the LBO spectrum and its extension equipped with three additional functions that encode the extrinsic embedding (CMH). As shown in [73, 80] CMH is better suited to transfer discrete signals such as the surface mesh connectivity. This advantage also manifests on volumes, as reflected in the second and fourth rows of Tab. 1. For efficiency,




Figure 7. From left to right: tetrahedral mesh; surface of the target domain; ground truth connectivity transfer; connectivity transfer obtained by extrapolating the interior vertices from the surface coordinates; connectivity transfer obtained by transferring the coordinates with the functional map. Each tetrahedron t is colored according to the geometric distortion of the map, measured as $|1 - \det(J_t)|$, where J_t is the Jacobian of the affine transformation mapping t from the source to the target. All maps shown in this figure are fully injective (i.e., $\forall t, \det(J_t) > 0$).

the average running times for computing the larger bases are reported in the last column of the same table. As can be noticed the extrapolation algorithm is roughly $2\times$ faster than volumetric transfer, because in this case only the volumetric spectrum of a single mesh must be computed. Note that even using a small fraction of the spectrum, volumetric functional mapping still provides a valuable initial solution for untangling algorithms such as [28, 36], which are then required to remove the remaining inverted tetrahedra from the map. This is a standard procedure for generating injective linear maps, where mesh untangling is typically bootstrapped using the Tutte 3D embedding [41], which is known to fail on volumes [7]. Our tool provides a better initial guess than Tutte. To prove this, we considered the 1405 failures produced in the large-scale validation described in [83, §6.3]. Substituting the initial guess provided by Tutte with the one produced with our tool, we were able to obtain one-to-one maps in the 85% of their failure cases, using both LBO and CMH basis functions, dramatically increasing the success rate of this pipeline (Fig. 6).

5.3. Segmentation Transfer

Statistical shape analysis for medical data primarily benefits from accurate solutions to analyze large collections of shapes [40, 111]. The literature shows that the functional maps framework is a prominent tool for medical applications [67, 72, 76]. A standard approach to analyze collections of shapes is to determine or build a template surface which acts as a median shape to compare the entire dataset [44]. However, existing approach are limited to the analysis of surfaces, while often medical data are better represented as 3D volumes (e.g., MRI scans [101]).

Our approach allows to extend existing methodologies to deal with volumetric data. To prove this, we select




Figure 8. A segmentation of a template brain transferred to two other brains using a correspondence computed with our volumetric functional maps framework. The shapes and the segmentation are from MedShapeNet [55].

two brain shapes from the MedShapeNet dataset [55] and a population-averaged template [31]. Using TetGen [39], we compute the tetrahedral meshes representing the volumes of the brains, and segment the template using a reference segmentation [34]. We then compute a volumetric functional map from the template to each shape, extract the point-to-point correspondence and transfer the segmentation. From the results in Fig. 8, we can see that, despite the noisy segmentation in the central section, our tool can faithfully transfer the segmentation preserving its overall structure.

5.4. Shape Matching

Here we study to what extent incorporating volumetric information into the shape matching problem improves accuracy over existing surface-based functional map pipelines.

Data. We consider both the outer surfaces of our reference ground truth volumetric dataset (denoted as VOL) [113] and well established shape matching datasets, such as SHREC’19 [78], SHREC’20 [29], and TOPKIDS [50], which we all tetrahedralize using fTetWild [43]. These datasets stress shape matching algorithms in different ways, exhibiting geometric and topological noise (SHREC’19), strongly non-isometric matches (SHREC’20) and topological changes (SHREC’20 and TOPKIDS). For SHREC’20, we discard two meshes because the inside/outside system of fTetWild does not reproduce a meaningful volumetric mesh for the comparison. We refer to Appendix G for more details on the data.

Comparative Analysis. We validate the approach in Sec. 4.2 by comparing it with the standard surface-based pipeline on each dataset. To prove that our volumetric framework can be seamlessly integrated with existing extensions of functional maps, we also test it with different sets of bases, such as Coordinates Manifold Harmon-

Dataset	LBO			CMH			Orthoprods			Slowdown	Vertex Ratio
	AGE Surf.	AGE Vol.	Succ. Rate	AGE Surf.	AGE Vol.	Succ. Rate	AGE Surf.	AGE Vol.	Succ. Rate		
VOL	$1.88 \cdot 10^{-6}$	$2.92 \cdot 10^{-6}$	22.50%	$1.78 \cdot 10^{-6}$	$3.15 \cdot 10^{-6}$	20.00%	$5.82 \cdot 10^{-7}$	$3.19 \cdot 10^{-7}$	60.00%	1.92×	1.40×
SHREC'19	$1.03 \cdot 10^{-1}$	$6.64 \cdot 10^{-2}$	53.15%	$1.04 \cdot 10^{-1}$	$6.56 \cdot 10^{-2}$	54.52%	$1.11 \cdot 10^{-1}$	$7.13 \cdot 10^{-2}$	56.42%	2.03×	1.59×
TOPKIDS	$1.24 \cdot 10^{-1}$	$8.85 \cdot 10^{-2}$	60.00%	$1.23 \cdot 10^{-1}$	$8.86 \cdot 10^{-2}$	68.00%	$1.17 \cdot 10^{-1}$	$8.13 \cdot 10^{-2}$	64.00%	2.51×	2.16×
SHREC'20	$1.82 \cdot 10^{-1}$	$2.72 \cdot 10^{-1}$	14.84%	$2.70 \cdot 10^{-1}$	$1.82 \cdot 10^{-1}$	55.49%	$1.83 \cdot 10^{-1}$	$3.01 \cdot 10^{-1}$	8.24%	1.01×	0.80×

Table 2. Results of each functional maps implementation on each dataset. We report geodesic error averaged across the dataset (AGE) and percentage of single pairs where our approach obtains a better AGE (Succ. Rate). Computational overhead strongly correlates with mesh growth (right). We report: average slowdown induced by volumes and average vertex ratio between surfaces and volumes.

ics (CMH) [73] and Orthogonalized eigenproducts (Orthoprods) [68]. In all cases, we align the bases using the method based on descriptors preservation [85] and ZoomOut [79]. For a fair comparison, we use the same bases for both the surface and the volumetric approaches.

The results of our experiments are summarized in Tab. 2, where we show the dataset averaged geodesic error (AGE) for both frameworks and the percentage of single pairs for which our approach obtains a better AGE (Success Rate). We also report the slowdown factor induced by volumes and the ratio between the vertex count of the surface and volume meshes. Interestingly, our approach is consistently more accurate than its surface counterpart, suggesting that incorporating volumetric information into the matching process is always a rewarding choice when a volumetric discretization is available. The increased accuracy comes at the cost of additional computational, due to the introduction of the internal vertices. Inspecting the Slowdown columns in Tab. 2 we can see that there is a close relationship between the slowdown and the number of vertices introduced, suggesting that the volumetric pipeline is not significantly slower per se, but suffers from the overhead induced by the additional vertices needed to discretize the volume.

From the scale of the AGE, we can deduce that VOL is not a really challenging dataset for shape matching (watertight shapes, shared connectivity); thus, introducing volumetric information provides a real advantage only with more informative bases, such as Orthoprods. The accuracy of our approach seems related to the choice of the basis also when dealing with significantly more challenging datasets like SHREC'20. This is due to topological errors introduced by tetrahedralization. The topological changes reflect on the LBO eigenbasis and the Orthoprods basis, decreasing the accuracy. The CMH basis uses the vertex coordinates as a regularizer, compensating for the error. We provide a more detailed discussion in Appendix H.

6. Conclusions

We presented the first volumetric functional maps pipeline, showcasing real applications that exploit discrete and continuous signal transfer for solid texturing, connectivity transfer, and segmentation transfer for medical applications, also improving the accuracy of surface matching tasks. Our

practical contributions just scratch the surface of this problem and we foresee interesting avenues for future research.

Limitations. Our connectivity transfer shows great potential but is far from being perfect. Indeed, always achieving a fully injective map is technically possible, but considering a large enough portion of the spectrum may easily become computationally prohibitive. Devising smarter techniques to retain more precision for fixed spectrum size a research direction worth investigating.

The higher precision in shape matching tasks comes at the cost of an additional computational effort, because the most expensive steps in the pipeline (eigendecomposition and basis alignment) suffer from the increased vertex count. This limitation is intrinsic to our approach, but it could be mitigated by integrating scalable algorithms for computing and aligning bases specifically designed for volumes.

Future works. A theoretical study of what properties are invariant under a volume functional map must be carried out. Despite our empirical studies show that volumetric maps are more accurate than their surface counterparts, a more profound theoretical analysis would provide useful insight for further developments.

Finally, we emphasize the importance of releasing new volumetric datasets to support ongoing research on volume mapping, and in particular on spectral volume methods. As discussed in Sec. 5.4, surface datasets that are well established in the functional map community hardly extend to volumes because of the difficulty to produce tetrahedral meshes of acceptable quality, due to topological and geometric noise, (e.g., large missing parts, overlaps, multiple connected components, non-manifold configurations and self-intersections). While robust volume meshing methods such as [23, 43] can handle intersections with ease, determining semantically plausible separations between the inside and outside portions of the volume in the presence of large gaps remains an open challenge. Recently, a large scale dataset to support ongoing research on piece-wise volume mappings to convex or star-shaped domains has been released [18]. Having similar contributions to support spectral methods would greatly help practitioners and ultimately foster more research in this emerging field.

References

- [1] S Mazdak Abulnaga, Esra Abaci Turk, Mikhail Bessmeltsev, P Ellen Grant, Justin Solomon, and Polina Golland. Volumetric parameterization of the placenta to a flattened template. *IEEE transactions on medical imaging*, 41(4): 925–936, 2021. 1, 6
- [2] S Mazdak Abulnaga, Oded Stein, Polina Golland, and Justin Solomon. Symmetric volume maps: Order-invariant volumetric mesh correspondence with free boundary. *ACM Transactions on Graphics*, 42(3):1–20, 2023. 2
- [3] Yonathan Aflalo, Haim Brezis, and Ron Kimmel. On the optimality of shape and data representation in the spectral domain. *SIAM Journal on Imaging Sciences*, 8(2):1141–1160, 2015. 14, 16
- [4] Yonathan Aflalo, Haïm Brezis, Alfred Bruckstein, Ron Kimmel, and Nir Sochen. Best bases for signal spaces. *Comptes Rendus Mathématique*, 354(12):1155–1167, 2016. 14, 16
- [5] Noam Aigerman and Yaron Lipman. Injective and bounded distortion mappings in 3d. *ACM Transactions on Graphics (TOG)*, 32(4):1–14, 2013. 2
- [6] Noam Aigerman and Yaron Lipman. Hyperbolic orbifold tutte embeddings. *ACM Trans. Graph.*, 35(6):217–1, 2016. 2
- [7] Marc Alexa. Tutte embeddings of tetrahedral meshes. *Discrete & Computational Geometry*, pages 1–11, 2023. 2, 7
- [8] Marc Alexa, Philipp Herholz, Maximilian Kohlbrenner, and Olga Sorkine-Hornung. Properties of laplace operators for tetrahedral meshes. In *Computer Graphics Forum*, pages 55–68. Wiley Online Library, 2020. 14
- [9] Souhaib Attaiqi and Maks Ovsjanikov. Shape non-rigid kinematics (snk): A zero-shot method for non-rigid shape matching via unsupervised functional map regularized reconstruction, 2024. 2
- [10] Marcel Seelbach Benkner, Zorah Löhner, Vladislav Golyanik, Christof Wunderlich, Christian Theobalt, and Michael Moeller. Q-match: Iterative shape matching via quantum annealing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7586–7596, 2021. 3
- [11] Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. FAUST: Dataset and evaluation for 3D mesh registration. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, Piscataway, NJ, USA, 2014. IEEE. 2
- [12] Hendrik Brückler and Marcel Campen. Collapsing embedded cell complexes for safer hexahedral meshing. *ACM Transactions on Graphics (TOG)*, 42(6):1–24, 2023. 6
- [13] Marcel Campen, Cláudio T Silva, and Denis Zorin. Bijective maps from simplicial foliations. *ACM Transactions on Graphics (TOG)*, 35(4):1–15, 2016. 2
- [14] Dongliang Cao, Paul Roetzer, and Florian Bernard. Unsupervised learning of robust spectral shape matching. *ACM Trans. Graph.*, 42(4), 2023. 2
- [15] Dongliang Cao, Marvin Eisenberger, Nafie El Amrani, Daniel Cremers, and Florian Bernard. Spectral meets spatial: Harmonising 3d shape matching and interpolation, 2024. 2
- [16] Can Ceritoglu, Lei Wang, Lynn D Selemon, John G Csernansky, Michael I Miller, and J Tilak Ratnanather. Large deformation diffeomorphic metric mapping registration of reconstructed 3d histological section images and in vivo mr images. *Frontiers in human neuroscience*, 4:895, 2010. 1
- [17] Can Ceritoglu, Xiaoying Tang, Margaret Chow, Darian Hadjiabadi, Damish Shah, Timothy Brown, Muhammad H Burhanullah, Huong Trinh, John T Hsu, Katarina A Ament, et al. Computational analysis of lddmm for brain mapping. *Frontiers in neuroscience*, 7:151, 2013. 1
- [18] Gianmarco Cherchi and Marco Livesu. Volmap: a large scale benchmark for volume mappings to simple base domains. *Computer Graphics Forum*, 42(5), 2023. 2, 8
- [19] Kiran Chilakamarri, Nathaniel Dean, and Michael Littman. Three-dimensional tutte embedding. *Congressus Numerantium*, pages 129–140, 1995. 2
- [20] Keenan Crane. The n-dimensional cotangent formula. Online note. URL: <https://www.cs.cmu.edu/~km-crane/Projects/Other/nDCotanFormula.pdf>, 2019. 3, 14
- [21] Keenan Crane, Marco Livesu, Enrico Puppo, and Yipeng Qin. A survey of algorithms for geodesic paths and distances. *arXiv preprint arXiv:2007.10430*, 2020. 3
- [22] Bailin Deng, Yuxin Yao, Roberto M Dyke, and Juyong Zhang. A survey of non-rigid 3d registration. In *Computer Graphics Forum*, pages 559–589. Wiley Online Library, 2022. 2, 3
- [23] Lorenzo Diazzi and Marco Attene. Convex polyhedral meshing for robust solid modeling. *ACM Transactions on Graphics (TOG)*, 40(6):1–16, 2021. 8, 18
- [24] Lorenzo Diazzi, Daniele Panozzo, Amir Vaxman, and Marco Attene. Constrained delaunay tetrahedrization: A robust and practical approach. *ACM Transactions on Graphics (TOG)*, 42(6):1–15, 2023. 4, 18
- [25] Edsger W Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959. 3
- [26] Nicolas Donati, Abhishek Sharma, and Maks Ovsjanikov. Deep geometric functional maps: Robust feature learning for shape correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8592–8601, 2020. 2
- [27] Nicolas Donati, Etienne Corman, Simone Melzi, and Maks Ovsjanikov. Complex functional maps: A conformal link between tangent bundles. In *Computer Graphics Forum*, pages 317–334. Wiley Online Library, 2022. 2
- [28] Xingyi Du, Noam Aigerman, Qingnan Zhou, Shahar Z Kovalsky, Yajie Yan, Danny M Kaufman, and Tao Ju. Lifting simplices to find injectivity. *ACM Trans. Graph.*, 39(4): 120, 2020. 2, 5, 7
- [29] Roberto M Dyke, Yu-Kun Lai, Paul L Rosin, Stefano Zappalà, Seana Dykes, Daoliang Guo, Kun Li, Riccardo Marin, Simone Melzi, and Jingyu Yang. Shrec’20: Shape correspondence with non-isometric deformations. *Computers & Graphics*, 92:28–43, 2020. 7
- [30] Jimmy Etienne, Nicolas Ray, Daniele Panozzo, Samuel Hornus, Charlie CL Wang, Jonàs Martínez, Sara McMains,

- Marc Alexa, Brian Wyvill, and Sylvain Lefebvre. Curvislicer: Slightly curved slicing for 3-axis printers. *ACM Transactions on Graphics (TOG)*, 38(4):1–11, 2019. 6
- [31] Alan C Evans, D Louis Collins, SR Mills, Edward D Brown, Ryan L Kelly, and Terry M Peters. 3d statistical neuroanatomical models from 305 mri volumes. In *1993 IEEE conference record nuclear science symposium and medical imaging conference*, pages 1813–1817. IEEE, 1993. 7
- [32] Danielle Ezuz and Mirela Ben-Chen. Deblurring and denoising of maps between shapes. In *Computer Graphics Forum*, pages 165–174. Wiley Online Library, 2017. 2
- [33] Ugo Finnendahl, Dimitrios Bogiokas, Pablo Robles Cervantes, and Marc Alexa. Efficient embeddings in exact arithmetic. *ACM Transactions on Graphics (TOG)*, 42(4):1–17, 2023. 2
- [34] Bruce Fischl, David H Salat, Evelina Busa, Marilyn Albert, Megan Dieterich, Christian Haselgrove, Andre Van Der Kouwe, Ron Killiany, David Kennedy, Shuna Klavenness, et al. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron*, 33(3):341–355, 2002. 7
- [35] Michael S Floater. Parametrization and smooth approximation of surface triangulations. *Computer aided geometric design*, 14(3):231–250, 1997. 2
- [36] Vladimir Garanzha, Igor Kaporin, Liudmila Kudryavtseva, François Protais, Nicolas Ray, and Dmitry Sokolov. Foldover-free maps in 50 lines of code. *ACM Transactions on Graphics (TOG)*, 40(4):1–16, 2021. 2, 6, 7
- [37] C Garner, Miao Jin, Xianfeng Gu, and Hong Qin. Topology-driven surface mappings with robust feature alignment. In *VIS 05. IEEE Visualization, 2005.*, pages 543–550. IEEE, 2005. 2
- [38] Craig Gotsman, Xianfeng Gu, and Alla Sheffer. Fundamentals of spherical parameterization for 3d meshes. In *ACM SIGGRAPH 2003 Papers*, pages 358–363. 2003. 2
- [39] Si Hang. Tetgen, a delaunay-based quality tetrahedral mesh generator. *ACM Trans. Math. Softw.*, 41(2):11, 2015. 4, 7, 18
- [40] Tobias Heimann and Hans-Peter Meinzer. Statistical shape models for 3d medical image segmentation: a review. *Medical image analysis*, 13(4):543–563, 2009. 7
- [41] Steffen Hinderink and Marcel Campen. Galaxy maps: Localized foliations for bijective volumetric mapping. *ACM Transactions on Graphics (TOG)*, 42(4):1–16, 2023. 2, 7
- [42] Kai Hormann. Barycentric interpolation. In *Approximation Theory XIV: San Antonio 2013*, pages 197–218. Springer, 2014. 2
- [43] Yixin Hu, Teseo Schneider, Bolun Wang, Denis Zorin, and Daniele Panozzo. Fast tetrahedral meshing in the wild. *ACM Transactions on Graphics (ToG)*, 39(4):117–1, 2020. 7, 8, 18
- [44] Ruqi Huang, Panos Achlioptas, Leonidas Guibas, and Maks Ovsjanikov. Limit shapes—a tool for understanding shape differences and variability in 3d model collections. In *Computer Graphics Forum*, pages 187–202. Wiley Online Library, 2019. 1, 7
- [45] Alec Jacobson. Optimized dual-volumes for tetrahedral meshes. In *Computer Graphics Forum*, page e15133. Wiley Online Library, 2024. 14
- [46] Andrei Khodakovsky, Nathan Litke, and Peter Schröder. Globally smooth parameterizations with low distortion. *ACM transactions on graphics (TOG)*, 22(3):350–357, 2003. 2
- [47] Patrick M Knupp. Hexahedral and tetrahedral mesh untangling. *Engineering with Computers*, 17:261–268, 2001. 2
- [48] Shahar Z Kovalsky, Noam Aigerman, Ronen Basri, and Yaron Lipman. Controlling singular values with semidefinite programming. *ACM Trans. Graph.*, 33(4):68–1, 2014. 2
- [49] Vladislav Kraevoy and Alla Sheffer. Cross-parameterization and compatible remeshing of 3d models. *ACM Transactions on Graphics (ToG)*, 23(3):861–869, 2004. 2
- [50] Zorah Löhner, Emanuele Rodolà, Michael M Bronstein, Daniel Cremers, Oliver Burghard, Luca Cosmo, Alexander Dieckmann, Reinhard Klein, Y Sahillioğlu, et al. Shrec’16: Matching of deformable shapes with topological noise. In *Eurographics Workshop on 3D Object Retrieval, EG 3DOR*, pages 55–60. Eurographics Association, 2016. 7
- [51] I Lasiecka and R Triggiani. Stabilization of neumann boundary feedback of parabolic equations: The case of trace in the feedback loop. *Applied Mathematics and Optimization*, 10(1):307–350, 1983. 4, 16
- [52] Aaron WF Lee, Wim Sweldens, Peter Schröder, Lawrence Cowsar, and David Dobkin. Maps: Multiresolution adaptive parameterization of surfaces. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 95–104, 1998. 2
- [53] Thibault Lescoat, Hsueh-Ti Derek Liu, Jean-Marc Thiery, Alec Jacobson, Tamy Boubekeur, and Maks Ovsjanikov. Spectral mesh simplification. In *Computer Graphics Forum*, pages 315–324. Wiley Online Library, 2020. 3, 14, 15
- [54] Bruno Lévy and Hao (Richard) Zhang. Spectral mesh processing. In *ACM SIGGRAPH 2010 Courses*, New York, NY, USA, 2010. Association for Computing Machinery. 2
- [55] Jianning Li, Zongwei Zhou, Jiancheng Yang, Antonio Pepe, Christina Gsaxner, Gijs Luijten, Chongyu Qu, Tiezheng Zhang, Xiaoxi Chen, Wenxuan Li, et al. Medshapenet—a large-scale dataset of 3d medical shapes for computer vision. *Biomedical Engineering/Biomedizinische Technik*, (0), 2024. 7
- [56] Xin Li, Yunfan Bao, Xiaohu Guo, Miao Jin, Xianfeng Gu, and Hong Qin. Globally optimal surface mapping for surfaces with arbitrary topology. *IEEE Transactions on Visualization and Computer Graphics*, 14(4):805–819, 2008. 2
- [57] Sheng-hui Liao, Ruo-feng Tong, Jin-xiang Dong, and Fudong Zhu. Gradient field based inhomogeneous volumetric mesh deformation for maxillofacial surgery simulation. *Computers & Graphics*, 33(3):424–432, 2009. 3, 14
- [58] Or Litany, Tal Remez, Emanuele Rodola, Alex Bronstein, and Michael Bronstein. Deep functional maps: Structured prediction for dense shape correspondence. In *Proceedings*

- of the IEEE international conference on computer vision, pages 5659–5667, 2017. [2](#)
- [59] Tao Liu, Tianyu Zhang, Yongxue Chen, Yuming Huang, and Charlie CL Wang. Neural slicer for multi-axis 3d printing. *ACM Transactions on Graphics (TOG)*, 43(4):1–15, 2024. [6](#)
- [60] Marco Livesu. cinolib: a generic programming header only c++ library for processing polygonal and polyhedral meshes. *Transactions on Computational Science XXXIV*, pages 64–76, 2019. [18](#)
- [61] Marco Livesu. A Mesh Generation Perspective on Robust Mappings. In *Smart Tools and Applications in Graphics (STAG)*. The Eurographics Association, 2020. [2](#)
- [62] Marco Livesu. Mapping surfaces with earcut. *arXiv preprint arXiv:2012.08233*, 2020. [2](#)
- [63] Marco Livesu. Scalable mesh refinement for canonical polygonal schemas of extremely high genus shapes. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 27(1):254–260, 2021. [2](#)
- [64] Marco Livesu. Advancing front surface mapping. *Computer Graphics Forum*, 43(2), 2024. [2](#)
- [65] Marco Livesu. Stripe embedding: Efficient maps with exact numeric computation. *ACM Transactions on Graphics (TOG)*, 43(6):1–14, 2024. [2](#)
- [66] Sven Loncaric. A survey of shape analysis techniques. *Pattern recognition*, 31(8):983–1001, 1998. [3](#)
- [67] Francesca Maccarone, Giorgio Longari, Giulio Viganò, Denis Peruzzo, Filippo Maggioli, and Simone Melzi. S4A: Scalable Spectral Statistical Shape Analysis. In *Smart Tools and Applications in Graphics - Eurographics Italian Chapter Conference*. The Eurographics Association, 2024. [7](#)
- [68] Filippo Maggioli, Simone Melzi, Maksim Ovsjanikov, Michael M Bronstein, and Emanuele Rodolà. Orthogonalized fourier polynomials for signal approximation and transfer. In *Computer Graphics Forum*, pages 435–447. Wiley Online Library, 2021. [2](#), [5](#), [6](#), [8](#), [19](#)
- [69] Filippo Maggioli, Daniele Baieri, Simone Melzi, and Emanuele Rodolà. Newton’s fractals on surfaces via bi-complex algebra. In *ACM SIGGRAPH 2022 Posters*, New York, NY, USA, 2022. Association for Computing Machinery. [6](#)
- [70] Filippo Maggioli, Daniele Baieri, Emanuele Rodolà, and Simone Melzi. Rematching: Low-resolution representations for scalable shape correspondence. In *European Conference on Computer Vision*, pages 183–200. Springer, 2025. [2](#)
- [71] Robin Magnet and Maks Ovsjanikov. Scalable and efficient functional map computations on dense meshes. In *Computer Graphics Forum*, pages 89–101. Wiley Online Library, 2023. [2](#)
- [72] Robin Magnet, Kevin Bloch, Maxime Taverne, Simone Melzi, Maya Geoffroy, Roman H Khonsari, and Maks Ovsjanikov. Assessing craniofacial growth and form without landmarks: A new automatic approach based on spectral methods. *Journal of Morphology*, 284(8):e21609, 2023. [2](#), [7](#)
- [73] Riccardo Marin, Simone Melzi, Pietro Musoni, Filippo Bardon, Marco Tarini, Umberto Castellani, et al. Cmh: Coordinates manifold harmonics for functional remeshing. In *3DOR@ Eurographics*, pages 63–70, 2019. [2](#), [5](#), [6](#), [8](#)
- [74] Riccardo Marin, Marie-Julie Rakotosaona, Simone Melzi, and Maks Ovsjanikov. Correspondence learning via linearly-invariant embedding, 2020. [2](#)
- [75] Federico Meloni, Gianmarco Cherchi, Riccardo Scateni, Livesu Marco, et al. To what extent are existing volume mapping algorithms practically useful? In *ITALIAN CHAPTER CONFERENCE*, pages 1–8. The Eurographics Association, 2024. [2](#)
- [76] Simone Melzi, Alessandro Mella, Letizia Squarcina, Marcella Bellani, Cinzia Perlini, Mirella Ruggeri, Carlo Alfredo Altamura, Paolo Brambilla, and Umberto Castellani. Functional maps for brain classification on spectral domain. In *Spectral and Shape Analysis in Medical Imaging: First International Workshop, SeSAMI 2016, Held in Conjunction with MICCAI 2016, Athens, Greece, October 21, 2016, Revised Selected Papers 1*, pages 25–36. Springer, 2016. [7](#)
- [77] Simone Melzi, Emanuele Rodolà, Umberto Castellani, and Michael M Bronstein. Localized manifold harmonics for spectral shape analysis. In *Computer Graphics Forum*, pages 20–34. Wiley Online Library, 2018. [2](#)
- [78] Simone Melzi, Riccardo Marin, Emanuele Rodolà, Umberto Castellani, Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. Shrec 2019: Matching humans with different connectivity. In *Eurographics Workshop on 3D Object Retrieval*, page 3. The Eurographics Association, 2019. [7](#)
- [79] Simone Melzi, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, and Maks Ovsjanikov. Zoomout: spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics (TOG)*, 38(6):1–14, 2019. [2](#), [4](#), [5](#), [8](#)
- [80] Simone Melzi, Riccardo Marin, Pietro Musoni, Filippo Bardon, Marco Tarini, and Umberto Castellani. Intrinsic/extrinsic embedding for functional remeshing of 3d shapes. *Computers & Graphics*, 88:1–12, 2020. [6](#)
- [81] Luca Moschella, Simone Melzi, Luca Cosmo, Filippo Maggioli, Or Litany, Maks Ovsjanikov, Leonidas Guibas, and Emanuele Rodolà. Learning spectral unions of partial deformable 3d shapes. In *Computer Graphics Forum*, pages 407–417. Wiley Online Library, 2022. [14](#)
- [82] Patrick Mullen, Pooran Memari, Fernando de Goes, and Mathieu Desbrun. Hot: Hodge-optimized triangulations. In *ACM Siggraph 2011 papers*, pages 1–12. 2011. [14](#)
- [83] Valentin Zenon Nigolian, Marcel Campen, and David Bommes. Expansion cones: A progressive volumetric mapping framework. *ACM Transactions on Graphics (TOG)*, 42(4):1–19, 2023. [2](#), [6](#), [7](#), [17](#)
- [84] Valentin Zénon Nigolian, Marcel Campen, and David Bommes. A progressive embedding approach to bijective tetrahedral maps driven by cluster mesh topology. *ACM Transactions on Graphics (TOG)*, 43(6):1–14, 2024. [2](#)
- [85] Dorian Nogneng and Maks Ovsjanikov. Informative descriptor preservation via commutativity for shape matching.

- In *Computer Graphics Forum*, pages 259–267. Wiley Online Library, 2017. [2](#), [3](#), [5](#), [8](#), [19](#)
- [86] Dorian Nogneng, Simone Melzi, Emanuele Rodola, Umberto Castellani, Michael Bronstein, and Maks Ovsjanikov. Improved functional mappings via product preservation. *Computer Graphics Forum*, 37(2):179–190, 2018. [2](#), [3](#)
- [87] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (ToG)*, 31(4):1–11, 2012. [1](#), [2](#), [3](#), [14](#), [15](#)
- [88] Maks Ovsjanikov, Etienne Corman, Michael Bronstein, Emanuele Rodola, Mirela Ben-Chen, Leonidas Guibas, Frederic Chazal, and Alex Bronstein. Computing and processing correspondences with functional maps. In *SIGGRAPH ASIA 2016 Courses*, pages 1–60. 2016. [2](#), [3](#)
- [89] G. Pai, J. Ren, S. Melzi, P. Wonka, and M. Ovsjanikov. Fast sinkhorn filters: Using matrix scaling for non-rigid shape correspondence with functional maps. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 384–393, Los Alamitos, CA, USA, 2021. IEEE Computer Society. [2](#)
- [90] Mikhail Panine, Maxime Kirgo, and Maks Ovsjanikov. Non-isometric shape matching via functional maps on landmark-adapted bases. In *Computer graphics forum*, pages 394–417. Wiley Online Library, 2022. [15](#)
- [91] Nico Pietroni, Marcel Campen, Alla Sheffer, Gianmarco Cherchi, David Bommes, Xifeng Gao, Riccardo Scateni, Franck Ledoux, Jean Remacle, and Marco Livesu. Hex-mesh generation and processing: a survey. *ACM transactions on graphics*, 42(2):1–44, 2022. [1](#), [2](#), [6](#)
- [92] Tharinda Piyadasa, Joan Glaunès, Amelia Gully, Michael Proctor, Kirrie Ballard, Tuende Szalay, Naeim Sanaei, Sheryl Foster, David Waddington, and Craig Jin. Constrained lddmm for dynamic vocal tract morphing: Integrating volumetric and real-time mri. In *Interspeech 2025*, 2025. [1](#)
- [93] Cédric Portaneri, Mael Rouxel-Labbé, Michael Hemmer, David Cohen-Steiner, and Pierre Alliez. Alpha wrapping with an offset. *ACM Transactions on Graphics (TOG)*, 41(4):1–22, 2022. [18](#)
- [94] Emil Praun and Hugues Hoppe. Spherical parametrization and remeshing. *ACM transactions on graphics (TOG)*, 22(3):340–349, 2003. [2](#)
- [95] Emil Praun, Wim Sweldens, and Peter Schröder. Consistent mesh parameterizations. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 179–184, 2001. [2](#)
- [96] Adil Rasheed, Omer San, and Trond Kvamsdal. Digital twin: Values, challenges and enablers from a modeling perspective. *IEEE access*, 8:21980–22012, 2020. [1](#)
- [97] Dan Raviv, Michael M Bronstein, Alexander M Bronstein, and Ron Kimmel. Volumetric heat kernel signatures. In *Proceedings of the ACM workshop on 3D object retrieval*, pages 39–44, 2010. [5](#)
- [98] Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics (ToG)*, 37(6):1–16, 2018. [2](#), [17](#)
- [99] Jing Ren, Simone Melzi, Maks Ovsjanikov, Peter Wonka, et al. Maptree: recovering multiple solutions in the space of maps. *ACM Transactions on Graphics (TOG)*, 39(6):264–1, 2020. [2](#)
- [100] Jing Ren, Simone Melzi, Peter Wonka, and Maks Ovsjanikov. Discrete optimization for shape matching. In *Computer Graphics Forum*, pages 81–96. Wiley Online Library, 2021. [2](#)
- [101] Peter A Rinck. *Magnetic resonance in medicine: a critical introduction*. BoD–Books on Demand, 2019. [7](#)
- [102] Emanuele Rodola, Michael Moeller, and Daniel Cremers. Regularized pointwise map recovery from functional correspondence. In *Computer Graphics Forum*, pages 700–711. Wiley Online Library, 2017. [2](#), [6](#)
- [103] Yusuf Sahillioğlu. Recent advances in shape correspondence. *The Visual Computer*, 36(8):1705–1721, 2020. [2](#), [3](#)
- [104] Patrick Schmidt, Janis Born, Marcel Campen, and Leif Kobbelt. Distortion-minimizing injective maps between surfaces. *ACM Transactions on Graphics (TOG)*, 38(6):1–15, 2019. [2](#)
- [105] Patrick Schmidt, Marcel Campen, Janis Born, and Leif Kobbelt. Inter-surface maps via constant-curvature metrics. *ACM Transactions on Graphics (TOG)*, 39(4):119–1, 2020.
- [106] Patrick Schmidt, Dörte Pieper, and Leif Kobbelt. Surface maps via adaptive triangulations. In *Computer Graphics Forum*, pages 103–117. Wiley Online Library, 2023.
- [107] John Schreiner, Arul Asirvatham, Emil Praun, and Hugues Hoppe. Inter-surface mapping. In *ACM SIGGRAPH 2004 Papers*, pages 870–877. 2004. [2](#)
- [108] Nicholas Sharp, Souhaib Attaiki, Keenan Crane, and Maks Ovsjanikov. Diffusionnet: Discretization agnostic learning on surfaces, 2022. [2](#)
- [109] Hanxiao Shen, Zhongshi Jiang, Denis Zorin, and Daniele Panozzo. Progressive embedding. *ACM Transactions on Graphics*, 38(4), 2019. [2](#)
- [110] Rui Shi, Wei Zeng, Zhengyu Su, Jian Jiang, Hanna Damasio, Zhonglin Lu, Yalin Wang, Shing-Tung Yau, and Xianfeng Gu. Hyperbolic harmonic mapping for surface registration. *IEEE transactions on pattern analysis and machine intelligence*, 39(5):965–980, 2016. [2](#)
- [111] Mikkel B Stegmann and David Delgado Gomez. A brief introduction to statistical shape analysis. *Informatics and mathematical modelling, Technical University of Denmark, DTU*, 15(11), 2002. [7](#)
- [112] Kaitlin M Stouffer, Menno P Witter, Daniel J Tward, and Michael I Miller. Projective diffeomorphic mapping of molecular digital pathology with tissue mri. *Communications engineering*, 1(1):44, 2022. [1](#)
- [113] Jian-Ping Su, Xiao-Ming Fu, and Ligang Liu. Practical foldover-free volumetric mapping construction. In *Computer Graphics Forum*, pages 287–297. Wiley Online Library, 2019. [5](#), [6](#), [7](#), [14](#), [15](#), [16](#), [17](#)
- [114] Robert W. Sumner and Jovan Popović. Deformation transfer for triangle meshes. *ACM Trans. Graph.*, 23(3):399–405, 2004. [2](#)

- [115] Ramana Sundararaman, Nicolas Donati, Simone Melzi, Etienne Corman, and Maks Ovsjanikov. Deformation recovery: Localized learning for detail-preserving deformations. *ACM Trans. Graph.*, 43(6), 2024. [2](#)
- [116] Gabriel Taubin. A signal processing approach to fair surface design. In *Proc. CGIT*, pages 351–358, New York, NY, 1995. ACM. [2](#)
- [117] Roberto Triggiani. Linear independence of boundary traces of eigenfunctions of elliptic and stokes operators and applications. *Applicationes Mathematicae*, 35(4):481–512, 2008. [4](#), [16](#)
- [118] William Thomas Tutte. How to draw a graph. *Proceedings of the London Mathematical Society*, 3(1):743–767, 1963. [2](#)
- [119] Max Wardetzky, Saurabh Mathur, Felix Kälberer, and Eitan Grinspun. Discrete laplace operators: no free lunch. In *Symposium on Geometry processing*, pages 33–37. Aire-la-Ville, Switzerland, 2007. [14](#)
- [120] Ofir Weber and Denis Zorin. Locally injective parametrization with arbitrary fixed boundaries. *ACM Transactions on Graphics (TOG)*, 33(4):1–12, 2014. [2](#)
- [121] Hermann Weyl. Über die asymptotische verteilung der eigenwerte. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 1911:110–117, 1911. [14](#)
- [122] Tianyu Zhang, Guoxin Fang, Yuming Huang, Neelotpal Dutta, Sylvain Lefebvre, Zekai Murat Kilic, and Charlie CL Wang. S3-slicer: A general slicing framework for multi-axis 3d printing. *ACM Transactions on Graphics (TOG)*, 41(6):1–15, 2022. [6](#)
- [123] Silvia Zuffi, Angjoo Kanazawa, David W Jacobs, and Michael J Black. 3d menagerie: Modeling the 3d shape and pose of animals. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6365–6373, 2017. [18](#)

A. Choice of the Laplacian

The original functional maps paper introduced the use of the LBO eigenbasis because it satisfies the desired properties of compactness of representation and stability under near-isometric deformations [87, §5.1]. Furthermore, later works from Aflalo et al. [3, 4] have shown that, in a general setting, and especially when working with smooth signals, the LBO eigenbasis of a manifold $\mathcal{M}$ (of arbitrary dimension) is the optimal orthonormal basis for representing a finite-dimensional subspace of $L^2(\mathcal{M})$. In this setting, it is reasonable to extend the choice of the LBO to the volumetric domain, especially considering that volumetric deformations do not induce a larger geodesic distortion with respect to near-isometric surface deformations (see Sec. C for an empirical study in this regard).

On the other hand, it should be noted that discretizing the LBO operator in a way that fulfills all the algebraic properties of the same operator in the smooth setting is often difficult [119]. In literature there exist multiple alternative discretizations [8, 45, 82], each one with its own pros and cons. Among the various alternatives, we chose the cotangent formula due to its popularity, ease of implementation, and efficiency. This choice is not restrictive. Any alternative discretization can be used equivalently. Nonetheless, since most applications that we address involve analyzing information at the surface (which is the boundary of a tetrahedral mesh), we are forced to a discretization that imposes Neumann’s boundary conditions to avoid zero-valued eigenfunctions on the surface.

Volumetric Cotangent Formula We discretize the LBO of a tetrahedral mesh $\mathcal{M} = (V_{\mathcal{M}}, T_{\mathcal{M}})$ with the following per vertex relation

$$\Delta(v_i) = \sum_{j \in N(v_i)} w_{ij}(v_i - v_j), \quad (10)$$

using the n -dimensional cotangent formula [20, 57] to compute per edge weights w_{ij} . These weights consider all tetrahedra $ijkl$ incident to edge ij , according to the following formula

$$w_{ij} = \frac{1}{6} \sum_{ijkl} |v_k - v_l| \cot \theta_{kl}, \quad (11)$$

where angle θ_{kl} denotes the dihedral angle at the edge kl opposite to edge ij w.r.t. tetrahedron $ijkl$.

B. Spectral Comparison

LBO Eigenbasis. On continuous Riemannian manifolds, the LBO and its eigenbasis are invariant to isometric deformations. However, in the discrete setting, the position and connectivity of the vertices have a significant impact on the equivalence of the operator across different meshes, and




Figure 9. Eigenvalue errors distribution for the volumetric LBO (blue) and the surface LBO (orange).

changing the triangulation of a surface or its vertex density can strongly affect the spectral decomposition [53].

Moving from discrete surfaces to discrete volumes makes this problem even more evident. In a regular triangular mesh, each vertex has on average 6 neighboring vertices and 6 neighboring triangles, while in regular tetrahedral meshes, the average size of a vertex neighborhood is 12 vertices and 20 tetrahedra.

To ensure that the LBO has a consistent behavior across isometric or similar volumetric shapes, we compute the spectral decomposition of the LBO on 40 pairs of shapes with identical connectivity released by Su et al. [113] and compare both the eigenvalues and the eigenfunctions of the LBO. Our results are summarized in Fig. 9. The first two rows of the figure show that two humanoid shapes share the same LBO eigenfunctions (up to sign flips), even if they are not isometric and do not represent the same subject; the overall similarity is sufficient to generate the same low-frequency spectrum. To compare the eigenvalues, we first compute the vectors $\lambda_{\mathcal{M}}, \lambda_{\mathcal{N}}$ of the first 100 LBO eigenvalues on two tetrahedral meshes. Since the absolute difference between LBO eigenvalues is notoriously not informative, we rely on both the relative difference and the offset difference proposed by Moschella et al. [81] (see below).

For each pair of tetrahedral meshes, we also extract the triangular meshes representing their boundary surface and compute the same quantities. For a more meaningful comparison, all the tetrahedral meshes are rescaled to have unitary volume, and all the surface meshes are rescaled to have unitary surface area. The error distributions are shown at the bottom of Fig. 9, where we can see that the spectral similarity between pairs of tetrahedral and surface meshes is comparable.

Offset Error. The eigenvalues of the Laplace-Beltrami operator form a non-decreasing sequence that grows with a rate $\lambda_k \in \mathcal{O}(\sqrt[k]{k^2})$, where d is the dimensionality of the manifold [121]. A direct comparison of the eigenvalues is thus misleading, as the difference will be dominated by the eigenvalues with larger index. One possibility for handling this issue is to introduce relative differences in the comparison. Alternatively, Moschella et al. [81] introduced the offset representation for comparing and handling Laplacian spectra. Given two manifolds $\mathcal{M}$ and $\mathcal{N}$, the offset




Figure 10. Distribution of the orthogonality (left) and Laplacian commutativity (right) metrics for pairs of volumetric shapes (blue) and surface meshes (orange).

difference can be defined as follows

$$\text{offset}(\lambda_{\mathcal{M}}, \lambda_{\mathcal{N}}) = |\text{off}(\lambda_{\mathcal{M}}) - \text{off}(\lambda_{\mathcal{N}})|, \quad (12)$$

$$\text{off}(\lambda_{\mathcal{M}}) = (\text{off}(\lambda_{\mathcal{M},1}), \text{off}(\lambda_{\mathcal{M},2}), \dots), \quad (13)$$

$$\text{off}(\lambda_{\mathcal{M},i}) = \lambda_{\mathcal{M},i} - \lambda_{\mathcal{M},i-1}. \quad (14)$$

This method for comparing eigenvalues is designed to compensate for the Weyl’s law on 2-dimensional manifolds: since $\lambda_k \in \mathcal{O}(k)$, comparing consecutive differences (which asymptotically behave like a constant function) is more informative than comparing absolute values. However, when dealing with 3-dimensional manifolds, the Weyl’s law changes to $\lambda_k \in \mathcal{O}(\sqrt[3]{k^2})$ and such consecutive differences are not asymptotically constant anymore. This change to the asymptotic behavior also affects the relative difference: if we consider that the k -th eigenvalue can be written as the Weyl’s asymptotic relation plus a correction term, namely $\lambda_k = c\sqrt[3]{k^2} + r$, we have that

$$\frac{c_1 k + r_1 - c_2 k - r_2}{c_1 k + r_1} \not\approx \frac{c_1 \sqrt[3]{k^2} + r_1 - c_2 \sqrt[3]{k^2} - r_2}{c_1 \sqrt[3]{k^2} + r_1}. \quad (15)$$

As a consequence neither the relative difference or the offset difference have the same scale for 2-dimensional and 3-dimensional manifolds.

However, we can use the Weyl’s law to compensate the non-linear growth of the eigenvalues. Since $\lambda_k \in \mathcal{O}(\sqrt[3]{k^2})$, before computing the offsets we can transform the eigenvalues as $\lambda'_k = \sqrt{\lambda_k^3} \in \mathcal{O}(k)$. With this transformation, the error measures has the same scale for both the surfaces and the volumes. Clearly, this holds under the assumption that all the 2-dimensional manifolds have unitary area and all the 3-dimensional manifolds have unitary volume, which is easy to enforce as a preprocessing step.

Volumetric Functional Maps. Given a definition for the discrete LBO, the definition of the functional maps framework in tetrahedral meshes follows directly from its 2-dimensional counterpart (see Sec. 3). Nonetheless, in order to ensure a seamless integration of the existing approaches in the volumetric setting, we study how the usual properties of a functional map behave on tetrahedral meshes.

Given two meshes $\mathcal{M}$ and $\mathcal{N}$ and a functional map $\mathbf{C} : \mathcal{F}(\mathcal{M}) \rightarrow \mathcal{F}(\mathcal{N})$, Lescoat et al. [53] propose two metrics

for evaluating the ideal qualities of $\mathbf{C}$: the orthogonality $\|\mathbf{C}\|_O$ and the Laplacian commutativity $\|\mathbf{C}\|_L$

$$\|\mathbf{C}\|_O = \frac{\|\mathbf{C}^\top \mathbf{C} - \mathbf{I}\|}{\|\mathbf{I}\|}, \quad \|\mathbf{C}\|_L = \frac{\|\mathbf{C} \Lambda_{\mathcal{M}} - \Lambda_{\mathcal{N}} \mathbf{C}\|}{\|\mathbf{C} \Lambda_{\mathcal{M}}\|}, \quad (16)$$

where $\mathbf{I}$ is the identity matrix, $\|\cdot\|$ is the Frobenius norm, and $\Lambda_{\mathcal{M}}, \Lambda_{\mathcal{N}}$ are the diagonal matrices containing the LBO eigenvalues of, respectively, $\mathcal{M}$ and $\mathcal{N}$.

Again, we compute the same quantities for tetrahedral and triangular meshes across the same 40 pairs of shapes, summarizing the results in Fig. 10. As shown in the figure, the quality of volumetric functional maps is comparable to the surface cases.

C. Volumetric Isometries

The functional maps framework relies on the assumption that the two shapes are isometric or quasi-isometric. In practice, the framework proposed by Ovsjanikov et al. [87] works well even under mild non-isometric deformations and meshes at a different scale. However, when dealing with strong non-isometric distortion, more sophisticated approaches are needed, which cannot be easily lifted to manifolds of higher dimensions [90].

In order to verify that the functional maps framework can be meaningfully transported to a volumetric setting, we need to verify that the level of isometric distortion does not change when a 3D shape is represented as a 2-dimensional surface or a 3-dimensional volume. The dataset introduced by Su et al. [113] contains 40 pairs of tetrahedral meshes in bijective correspondence, and thus it is easy to compute the distortion of geodesic distances induced by the correspondence. For each pair of shapes $\mathcal{M}$ and $\mathcal{N}$, we consider 100k random pairs of points $x, y \in \partial\mathcal{M}$. We then compute the geodesic distortion induced by the ground truth correspondence $\pi : \mathcal{M} \rightarrow \mathcal{N}$ for both the surface and the volumetric setting, respectively as

$$\text{SurfErr}(x, y) = \frac{|d_{\partial\mathcal{M}}(x, y) - d_{\partial\mathcal{N}}(\pi(x), \pi(y))|}{d_{\partial\mathcal{M}}(x, y)}, \quad (17)$$

$$\text{VolErr}(x, y) = \frac{|d_{\mathcal{M}}(x, y) - d_{\mathcal{N}}(\pi(x), \pi(y))|}{d_{\mathcal{M}}(x, y)}, \quad (18)$$

where d_S denotes the approximation of the geodesic distance in the volume S provided by the Dijkstra distance, and $d_{\partial S}$ denotes the analogous approximation over the boundary surface ∂S .

The results are summarized in Fig. 11, where we report for each pair of shapes the average distortion over the surface (in blue) and the volume (in orange), together with the standard deviation. We can see that, in all cases, the distortions are comparable, with the tetrahedral meshes generally achieving a slightly better preservation of the geodesic distances. As a consequence, we guess that, for most practical




Figure 11. For each subject, the relative geodesic distortion between pairs of volumes (orange) and pairs of surfaces (blue). The plot provides both the average distortion and the standard deviation.




Figure 12. For a varying basis size, we try to reconstruct the LBO eigenfunctions of the surface using the restriction of the volumetric LBO eigenfunctions at the boundary. The reconstruction error lies in the range $[0, 1]$.

applications, the property of a pair of shapes of being isometric is preserved between a surface and volumetric representation, thus allowing us to safely extend the functional maps framework from surfaces to volumes.

D. Boundary Trace of Eigenfunctions

As mentioned in Sec. 4.1, given the volumetric LBO eigenfunctions $\Phi_{\mathcal{M}}$ of a 3-dimensional manifold $\mathcal{M}$, the restriction $\Phi_{\partial\mathcal{M}}$ of such basis at the boundary surface $\partial\mathcal{M}$ (often also called *boundary trace*) forms a set of linearly independent functions over $\partial\mathcal{M}$ [51, 117].

However, $\Phi_{\partial\mathcal{M}}$ is not necessarily equivalent to the LBO eigenbasis $\Psi_{\partial\mathcal{M}}$ of the surface, meaning that it could be a sub-optimal basis [3, 4]. In order to evaluate the expressive power of the boundary trace of the volumetric eigenfunctions, we try to reconstruct the basis $\Psi_{\partial\mathcal{M}}$ using the basis $\Phi_{\partial\mathcal{M}}$. For our experiment, we use the 80 shapes from the dataset presented by Su et al. [113] and we try to reconstruct the first k eigenfunctions of the surface with the first k boundary traces. In particular, we first compute the reconstruction $\Psi' = \Phi_{\partial\mathcal{M}}\Phi_{\partial\mathcal{M}}^\dagger\Psi_{\partial\mathcal{M}}$. Then, for each surface eigenfunction ψ_i and its reconstruction ψ'_i , we compute the reconstruction error as

$$\text{err}(\psi_i) = \int_{\partial\mathcal{M}} (\psi_i(x) - \psi'_i(x))^2 dx. \quad (19)$$

This error measure lies in the range $[0, 1]$.

In Fig. 12 we show the distributions of the reconstruction errors across the entire dataset for different values of k . As shown in the figure, very few of the surface eigenfunctions are close to be orthogonal to the boundary traces. However, the boundary restriction of the volumetric eigenfunctions is able to faithfully represent a significant portion of the surface eigenbasis, and provides a reasonable approximation for most of them. Therefore, while not being the optimal basis, the boundary trace of the volumetric LBO eigenfunctions still provides a good candidate for representing a k -dimensional function space over the surface.




Figure 13. Number of flipped tetrahedra obtained by transferring the mesh connectivity considering a growing number of basis functions. From left to right: coordinate transfer using the LBO basis; coordinate transfer using the CMH basis; coordinate extrapolation using the LBO basis; coordinate extrapolation using the CMH basis. The smaller colored lines show the trend for each pair of shapes in our reference dataset [113]. The bold black line is the average trend across the entire dataset. The two pairs that take longer to converge with the LBO basis functions (yellow and orange lines) are the ones where the source and target differ the most and are strongly non-isometric. This explains why using only the intrinsic information from the LBO eigenbasis makes the method converge slowly and adding extrinsic information from CMH aligns the pairs with the other results.

Method	Continuity	Coverage	Dirichlet	Flips
FMaps	1.32	91.48%	3.82%	10.03%
ZoomOut	1.26	92.63%	3.80%	9.72%
ZoomOut Fast	1.32	92.88%	3.79%	10.80%
Orthoprods	1.29	87.78%	3.81%	5.80%

Table 3. Additional metrics for the volumetric shape matching experiments from Sec. 5.1 on the dataset from Su et al. [113].

E. Volumetric Functional Maps Metrics

Finally, we consider additional metrics to better evaluate the quality of volumetric shape matching discussed in Sec. 5.1. In particular, we report in Tab. 3 the continuity and coverage of the correspondence, as well as its Dirichlet energy, which provide a further characterization of maps in general [98, §5]. We also report the average percentage of flipped tetrahedra, which is an intrinsically volumetric measure. All the values are averaged across the entire dataset from Su et al. [113]. We stress that all reported values (including those in Sec. 5.1) refer to volumetric estimates of the metrics.

F. Connectivity Transfer and Basis Size

While the average behavior of our connectivity transfer approaches is clear from the results shown in Tab. 1, there are additional information that can be better appreciated by examining the behavior of the methods on the single pairs. In particular, we can see from Fig. 13 that some pairs are particularly challenging and results in a number of flips that is

significantly higher than the average, when using the LBO eigenbasis. However, CMH seems to be a strong regularizer: the basis extended with the vertex coordinates does not seem to have a significant impact on the other curves, but it is definitely able to align the behavior of these outliers to the average.

On the other hand, it seems that the introduction of the coordinate harmonics negatively affects smaller bases. Indeed, on the left end of the curves, the use of CMH as a basis seems to slow down the decrease rate of the curve.

In summary, by inspecting these results we educatedly guess that a good rule of thumb could be to use CMH when we have access to a larger basis, and stick to the LBO eigenbasis if we can only afford to compute a few basis functions.

We stress that using 20% of the entire basis, in general, is computationally expensive, and we report the values only for analysis purposes. Indeed, for our experiments on the datasets from Nigolian et al. [83] (see Sec. 5.2), we consider a significantly smaller basis with 500 eigenfunctions, covering from 0.5% to 5% of the total basis, depending on the resolution of the tetrahedral mesh.

G. Volumetric Shape Matching Dataset

As mentioned in Sec. 5.4, well-established datasets for shape matching only provide surfaces. On top of that, most of the available recent datasets try to highlight some challenges in the problem of finding a correspondence; this includes partialities, holes, heavy topological noise, self-




Figure 14. The two partial surfaces from the SHREC’20 dataset that we discarded from our experiments because once turned into tetrahedral meshes they do not correctly convey the intended shape, introducing an unwanted bias in the comparative analysis with surface-based methods.

intersections, clutter, and so on.

Dropping the assumptions of manifoldness, watertightness, and absence of self-intersections for the surfaces, the most robust tool for producing tetrahedral meshes results to be fTetWild [43]. Nonetheless, for some of the surfaces, the problem of finding a tetrahedralization is so ill-posed that any resulting volume would not provide meaningful information about the 3D shape.

Indeed, for SHREC’20, we discard two partial meshes (see Fig. 14) because the inside/outside system of fTetWild was not able to correctly reproduce a volumetric mesh that allows a meaningful comparison with surface methods. For the same reason, we are not able to consider alternative datasets such as SMAL [123], which contains numerous self-intersections that yield topological changes during tetrahedralization. Attempts to sanitize these models with various combinations of geometry processing tools [23, 24, 39, 60, 93] ended with no luck. Indeed, our difficulties in creating a dataset that allows to meaningfully compare surface-based and volume-based shape matching methods raises an open research challenge that will hopefully be addressed by the community in the near future (Sec. 6).




Figure 15. Cumulative plot of the geodesic error obtained with the standard functional maps framework (blue) on surfaces and our volumetric functional maps approach (orange) on four different datasets (columns) and using three different bases (rows). The results are averaged across each dataset.



Figure 16. The tool we use to tetrahedralize surface matching datasets [43] is robust, but occasionally introduces topological noise, closing unwanted handles that alter the mapping task (closeup). The LBO eigenbasis is not robust against such defects. Conversely, the CMH basis is robust and greatly improves the mapping quality. Both the camel and the cow are from the SHREC’20 dataset.

H. Choice of the Basis for Shape Matching

More detailed results about the comparison can be found in Fig. 15, where we show the cumulative geodesic error curves averaged across each dataset and their corresponding area under the curve (AUC). By comparing these curves against the results shown in Tab. 2, we can appreciate how the results summarized by the average geodesic error reflect on the cumulative curves, highlighting the consistently better behavior of our approach, as well as the importance in the choice of the basis.

In this regard, the use of the CMH basis seems to be particularly effective on SHREC’20, where the surface-based approach achieves better accuracy otherwise. Based

on our understanding, this depends on the basis alignment algorithm we use [85], which is known to not be robust under strong non-isometric deformations and topological changes, which occasionally occurred when tetrahedralizing our surface-based reference datasets. This behavior can be partially compensated by using more informative bases such as CMH, which provide a better tool for representing extrinsic information. This can be better appreciated in the example from Fig. 16, where we show a surface coordinate transfer in a pair from the SHREC’20 dataset with our volumetric approach using the standard LBO eigenbasis and the CMH basis. However, integrating more robust basis alignment methods is certainly worth of future investigation.

As a final note, for all our experiments we use the values suggested in the original papers when selecting the size of the bases. In particular, for the LBO eigenbasis we compute a 20×20 functional map using descriptors derived from 200 eigenfunctions for each shape. The map is extended up to the size 120×120 using ZoomOut, with a step size of 5 basis functions per iteration. In the case of CMH, the last 3 functions are replaced with the orthogonalized coordinates. For the Orthoprops basis, we use the suggested values of 40 eigenfunctions and second order polynomials. However, instead of the numerically unstable derivation of the map proposed in the original paper [68], we extend the functional map with ZoomOut.

I. Algorithmic Details

A sample MATLAB implementation of our volumetric functional maps framework is available at <https://github.com/filthynobleman/vol-fmaps>. Here we provide the pseudocode for both our connectivity transfer approaches (Sec. 4.1) and our volume-aware surface correspondence method (Sec. 4.2).

Algorithm 1 summarizes our solution for transferring connectivity through functional mapping. Firstly, we compute a tetrahedralization $\mathcal{N}$ for the interior of the target surface $\partial\mathcal{N}$ (Line 2). The volumetric meshes $\mathcal{M}$ and $\mathcal{N}$ are then put in functional correspondence via the map $\mathbf{C}$, which is approximated from the known alignment of the eigenfunctions at the surface (Lines 3-5). The source volume $\mathcal{M}$ is used to initialize the final mesh $\mathcal{R}$ (Line 6), whose coordinates are replaced from the functional transfer of the coordinates of $\mathcal{N}$.

The approach for extrapolating the interior volume from the surface is summarized in Algorithm 1. In this case, we only need the LBO eigenfunctions of $\mathcal{M}$ (Line 2). Again, we initialize a final mesh $\mathcal{R}$ from the source volume $\mathcal{M}$ (Line 3). However, with this approach the interior volume is obtained by projecting the coordinates at the surface in the functional space and reconstructing them in the whole volumetric mesh. The projection is performed using the boundary traces of the eigenfunctions, and the reconstruction is

Algorithm 1 Functional connectivity transfer algorithm.

```

1: procedure TRANSFER( $\mathcal{M}, \partial\mathcal{N}, \pi : \partial\mathcal{M} \rightarrow \partial\mathcal{M}$ )
2:    $\mathcal{N} \leftarrow$  computed interior of  $\partial\mathcal{N}$ 
3:    $\Phi_{\mathcal{M}} \leftarrow$  LBODECOMPOSITION( $\mathcal{M}$ )
4:    $\Phi_{\mathcal{N}} \leftarrow$  LBODECOMPOSITION( $\mathcal{N}$ )
5:    $\mathbf{C} \leftarrow \Phi_{\partial\mathcal{M}}^\dagger T_\pi(\Phi_{\partial\mathcal{N}})$ 
6:    $\mathcal{R} \leftarrow \mathcal{M}$ 
7:    $\mathbf{x}_{\mathcal{R}} \leftarrow \Phi_{\mathcal{M}} \mathbf{C} \Phi_{\mathcal{N}}^\dagger \mathbf{x}_{\mathcal{N}}$ 
8:    $\mathbf{y}_{\mathcal{R}} \leftarrow \Phi_{\mathcal{M}} \mathbf{C} \Phi_{\mathcal{N}}^\dagger \mathbf{y}_{\mathcal{N}}$ 
9:    $\mathbf{z}_{\mathcal{R}} \leftarrow \Phi_{\mathcal{M}} \mathbf{C} \Phi_{\mathcal{N}}^\dagger \mathbf{z}_{\mathcal{N}}$ 
10:  return  $\mathcal{R}$ 
11: end procedure

```

Algorithm 2 Spectral coordinate extrapolation algorithm.

```

1: procedure EXTRAPOLATE( $\mathcal{M}, \partial\mathcal{N}, \pi : \partial\mathcal{M} \rightarrow \partial\mathcal{M}$ )
2:    $\Phi_{\mathcal{M}} \leftarrow$  LBODECOMPOSITION( $\mathcal{M}$ )
3:    $\mathcal{R} \leftarrow \mathcal{M}$ 
4:    $\mathbf{x}_{\mathcal{R}} \leftarrow \Phi_{\mathcal{M}} \Phi_{\partial\mathcal{M}}^\dagger T_\pi(\mathbf{x}_{\partial\mathcal{N}})$ 
5:    $\mathbf{y}_{\mathcal{R}} \leftarrow \Phi_{\mathcal{M}} \Phi_{\partial\mathcal{M}}^\dagger T_\pi(\mathbf{y}_{\partial\mathcal{N}})$ 
6:    $\mathbf{z}_{\mathcal{R}} \leftarrow \Phi_{\mathcal{M}} \Phi_{\partial\mathcal{M}}^\dagger T_\pi(\mathbf{z}_{\partial\mathcal{N}})$ 
7:   return  $\mathcal{R}$ 
8: end procedure

```

Algorithm 3 Surface restricted volume matching algorithm.

```

1: procedure VOL2SURF( $\partial\mathcal{M}, \partial\mathcal{N}$ )
2:    $\mathcal{M} \leftarrow$  computed interior of  $\partial\mathcal{M}$ 
3:    $\mathcal{N} \leftarrow$  computed interior of  $\partial\mathcal{N}$ 
4:    $\Phi_{\mathcal{M}} \leftarrow$  LBODECOMPOSITION( $\mathcal{M}$ )
5:    $\Phi_{\mathcal{N}} \leftarrow$  LBODECOMPOSITION( $\mathcal{N}$ )
6:    $\mathbf{C} \leftarrow$  OPTIMIZEVOLFMAP( $\Phi_{\mathcal{M}}, \Phi_{\mathcal{N}}$ )
7:    $\pi \leftarrow$  NNSEARCH( $\Phi_{\partial\mathcal{M}} \mathbf{C}, \Phi_{\partial\mathcal{N}}$ )
8:   return  $\pi$ 
9: end procedure

```

done with the full basis.

Finally, we summarize in Algorithm 3 our procedure to compute the surface correspondence by exploiting the volumetric information. We first compute the interior volumes $\mathcal{M}$ and $\mathcal{N}$ for the surfaces $\partial\mathcal{M}$ and $\partial\mathcal{N}$, respectively (Line 2-3). The LBO eigenbases are then computed on the entire volume (Lines 4-5) and put in functional correspondence through the optimization of the map $\mathbf{C}$ (Line 6). Finally, we compute the correspondence π via nearest neighbor, limiting the search to the boundary traces of the eigenfunctions.