

Non-Isometric Shape Matching via Functional Maps on Landmark-Adapted Bases

M. Panine¹, formerly² , M. Kirgo^{2,3}  and M. Ovsjanikov² 

¹Università della Svizzera italiana, Switzerland

²LIX, École Polytechnique, IP Paris, France

³EDF R&D, France




Figure 1: Illustration of our method on a texture transfer problem, between two surfaces with significantly different mesh structure. The source model and its texture were produced by [McG17] (LPS Head) and the target model was extracted from the Faust dataset [BRLB14]. The user-specified landmark placement is shown in green, whereas the *computed* landmarks are shown below in blue. The state-of-the-art functional maps-based method “FMapZO” [MRR*19] fails to preserve landmarks exactly, whereas the hyperbolic orbifolds [AL15] (“HyperOrb”) approach leads to a map with higher distortion compared to our approach. The “reference” transfer was obtained using the commercial registration tool R3DS Wrap [Wra21] and 33 user-defined landmarks.

Abstract

We propose a principled approach for non-isometric landmark-preserving non-rigid shape matching. Our method is based on the functional maps framework, but rather than promoting isometries we focus instead on near-conformal maps that preserve landmarks exactly. We achieve this, first, by introducing a novel landmark-adapted basis using an intrinsic Dirichlet-Steklov eigenproblem. Second, we establish the functional decomposition of conformal maps expressed in this basis. Finally, we formulate a conformally-invariant energy that promotes high-quality landmark-preserving maps, and show how it can be solved via a variant of the recently proposed ZoomOut method that we extend to our setting. Our method is descriptor-free, efficient and robust to significant mesh variability. We evaluate our approach on a range of benchmark datasets and demonstrate state-of-the-art performance on non-isometric benchmarks and near state-of-the-art performance on isometric ones.

Keywords: shape matching, landmark-based correspondence, functional maps

CCS Concepts

• **Computing methodologies** → *Shape analysis*;

1. Introduction

A common scenario in shape matching is that of very sparse user-provided landmark correspondences that need to be extended to a full map between the considered shapes. The landmarks in question are often of a semantic nature, and thus are very sensitive to exact placement. Consider, for instance the position of the eyes

or the nose on a human face (see Fig. 1) that are matched by an artist, e.g., in a texture transfer scenario. In such cases, it is crucial to preserve the landmark correspondences *exactly* when extending the map. Furthermore, it is desirable for the extension process to be time-efficient and applicable to general, possibly non-isometric shape pairs.

Functional map methods [OCB*17] constitute a highly effective shape matching framework, especially when coupled with powerful recent post-processing tools such as ZoomOut and its variants [MRR*19; HRWO20]. The existing methods, however, suffer from two major limitations: first, they heavily rely upon the assumption of near-isometry, and second, they typically formulate landmark correspondence via descriptor preservation objectives, combined with other regularizers in the least squares sense. Unfortunately, this implies that the final map is not guaranteed to preserve user-provided landmark correspondences.

In this paper, we propose a novel approach that maintains the efficiency and flexibility of the functional maps pipeline, while overcoming these drawbacks. We organize our proposal in three major stages. First, we introduce a novel functional basis in which to express our map. Crucially, our basis is explicitly adapted to the landmark correspondences, unlike the commonly-used general Laplace-Beltrami eigenbasis. Intuitively speaking, this allows us to enforce landmark preservation by only considering functional maps with a particular (block-diagonal) structure. The design of this landmark-adapted basis is the most technically involved part of our proposal, and relies on solving intrinsic Dirichlet-Steklov and Dirichlet Laplacian eigenproblems. Specifically, we first construct new boundaries at the landmarks, and then formulate and solve the associated boundary value problems.

Second, we remove the assumption of near-isometry by structuring shape matching as a search for bijective *near-conformal* maps, which are significantly more general than isometries. Following the functional maps pipeline, we express this as a carefully designed energy to be minimized.

Third, we propose an iterative minimization strategy for our energy by following in the footsteps of ZoomOut [MRR*19]. In particular, we demonstrate how landmark correspondences can be promoted throughout this iterative refinement. Furthermore, we exploit the landmark-awareness of our basis to provide a simple initial guess of the correspondence.

We test our approach on various benchmark datasets, both isometric and non-isometric. We compare our results to both state-of-the-art functional maps approaches, as well as recent methods that exactly preserve landmark correspondences. We report state-of-the-art accuracy on non-isometric datasets and near state of the art on isometric ones. Meanwhile, the computation time of our approach is significantly lower than that of the competing landmark-preserving methods.

Contributions. To summarize:

1. We introduce a novel landmark-dependent functional basis by solving an intrinsic Dirichlet-Steklov eigenproblem.
2. We formulate a functional maps-based approach to near-conformal shape matching that preserves given landmarks exactly without restrictions on the topology of the shapes.
3. We propose an efficient way to both compute the basis and to solve the shape matching problem and report state-of-the-art results on difficult non-isometric benchmarks.

2. Related Work

Non-rigid shape matching is a well-established research area with a rich history of solutions. Below we review the works that are most closely related to ours, focusing on functional maps and landmark-preserving methods, and refer the interested readers to recent surveys [vKZHC11; Sah20] for a more in-depth discussion.

Functional Maps Our approach fits within the functional maps framework that was originally introduced in [OBS*12] and extended in many follow-up works, including including [KBB*13; ADK16; RCB*17; EB17; BDK17; NO17; MRR*19; RMOW20] to name a few. An early overview of many functional maps-based techniques is given in [OCB*17]. The key idea exploited in all of these techniques is to represent correspondences as linear transformations across functional spaces, which can be compactly encoded as small-sized matrices given a choice of basis. This leads to simple optimization problems that can accommodate a range of geometric objectives such as isometry [OBS*12], accurate descriptor preservation [NO17], bijectivity [ERGB16], orientation preservation [RPWO18] or even partiality [RCB*17] among others. Typically, such objectives are formulated as soft penalties on the functional map and are optimized for in the least squares sense.

Landmarks in functional maps Landmark constraints are commonly used in functional maps-based approaches, especially in an attempt to resolve symmetry ambiguity, present, e.g., when mapping between human shapes. Starting from the segment correspondences advocated in the original approach [OBS*12], and exploited in follow-up works, e.g., [KO19], several techniques also used *pointwise landmarks*, that were either user-specified [NO17], automatically computed [MMRC18], or even extended to *curve constraints* [GBKS18]. All of these techniques, however, formulate landmark correspondences via functional descriptor preservation, e.g., based on the heat kernel [SOG09; OMMG10] or wave kernel maps [ASC11], which are enforced during optimization only in a least squares sense, alongside other descriptors and regularizers. Therefore, there is no guarantee that the final recovered point-to-point map will satisfy these user-constraints. In contrast, our approach is geared towards preserving the landmark correspondences *exactly*, while computing a smooth overall map.

Landmark-based matching Landmark-preserving shape correspondence has also been studied in other matching frameworks. Early methods relied on extrinsic shape alignment, under given constraints, e.g., using thin plate splines [Boo89; CR00] or by extending non-rigid ICP, as done in [SP04] among others. Such approaches, however, rely strongly on the shape embedding and often require a significant number of landmarks to work well in practice.

Another successful class of approaches have aimed to compute correspondences by embedding shapes to a common parametrization domain. This includes powerful approaches based on mapping surfaces to the planar domain, [APL14; WZ14], Euclidean orbifolds [AL15] general flat cone manifolds [APL15] or, more recently, the hyperbolic plane [AL16], which can accommodate an arbitrary number of landmarks.

Finally, recent techniques have also allowed landmark-preserving shape correspondence by cross-parametrizing the surfaces directly. This includes exploiting direct and inverse *averages*

on surfaces [PBDS13] or finding maps that minimize various notions of distortion, e.g., harmonicity and reversibility (using, first a surrogate high-dimensional embedding) [ESB19] or a related symmetric Dirichlet energy [SCBK20] (via direct optimization on the surface). These recent techniques can lead to accurate results, but are often computationally expensive, and typically place restrictions on the topology of the shape pair, such as having the same genus. In contrast, our method does not suffer from this limitation, as topological stability is one of the features of functional map methods, which is also inherited by our technique.

Basis selection for functional maps Finally, we remark that our construction of landmark-adapted functional bases also fits within the functional map framework, aimed at developing flexible and effective basis functions. The original article and most follow-up works [OCB*17] have advocated using the eigenfunctions of the Laplace-Beltrami operator (LBO), which are optimal for representing smooth functions with bounded variation [ABK15]. However, the Laplace-Beltrami basis has global support and may not be fully adapted to non-isometric shape changes.

The compressed manifold modes [NVT*14; OLCO13; KGB16] have been introduced to offset the global nature of the LBO by promoting sparsity and locality in the basis construction. In a related effort, Choukroun *et al.* [CSBK18] have proposed to modify the LBO through a potential function, thus defining a Hamiltonian operator, whose eigenfunctions have better localization properties. In [MRCB18], a similar approach was introduced to obtain a basis that is also orthogonal to a given set of functions. The ‘‘Coordinate Manifold Harmonics’’ used in [MMM*20], complement the LBO eigenfunctions with the coordinates of the 3D embedding, allowing to capture both extrinsic and intrinsic information. Finally, a rich family of diffusion and harmonic bases have been proposed in [Pat18], by exploiting the properties of the heat kernel.

While these basis constructions offer more flexibility and have been shown to improve the functional map pipeline in certain cases, e.g., [NVT*14; MMM*20], they nevertheless are typically still geared towards approximate isometries, and only enable approximate constraint satisfaction. In contrast our basis is geared towards landmark-preserving maps during functional map optimization, as well as during refinement.

Dirichlet-Steklov basis Finally, we note that Steklov eigenproblems have been considered within geometry processing [WBPS18] as tool for *extrinsic* shape analysis. This is achieved by considering the (two-dimensional) surface as the Steklov boundary of its (three-dimensional) interior. In contrast, we consider a fully *intrinsic* problem by using (one-dimensional) boundaries of small disks centered around the landmarks as the boundary of the remainder of the surface.

3. Method Overview

In this section, we provide a high-level overview of our approach. Our method takes as input a pair of shapes \mathcal{M}, \mathcal{N} represented as triangle meshes along with two sets of k landmark vertices $\{\gamma_i^{\mathcal{M}}\}_{i=1}^k \subset \mathcal{M}, \{\gamma_i^{\mathcal{N}}\}_{i=1}^k \subset \mathcal{N}$. We then aim to compute a high-quality vertex-to-vertex correspondence $\phi: \mathcal{N} \rightarrow \mathcal{M}$ that preserves the given landmarks exactly. I.e., $\phi(\gamma_i^{\mathcal{N}}) = \gamma_i^{\mathcal{M}}$ for all i .

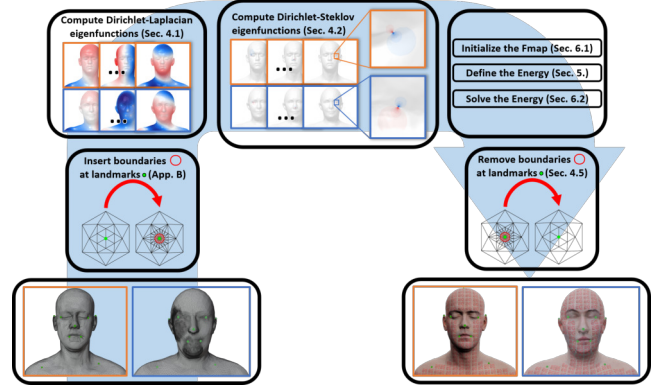


Figure 2: Schematic of the main steps involved in our method to map a source shape (orange) to a target shape (blue) as described in Sec. 3.

Our overall strategy, illustrated in Fig. 2, consists of the following major steps:

1. Remove small disks from the mesh surface, centered at each landmark point $\gamma_i, i = 1 \dots k$. This creates k circular boundary components $\{\Gamma_i\}_{i=1}^k$, which are fully contained in the one-ring neighborhood of each landmark.
2. Compute the set of the first N_{LB} Laplace-Beltrami eigenfunctions $\{\psi_j\}_{j=1}^{N_{LB}}$ with *Dirichlet boundary conditions* at the boundaries of landmark circles. I.e., $\psi_j|_{\Gamma_i} = 0$ for all i, j .
3. Add to this basis another k sets of N_{DS} basis functions each $\{u_j^{(i)}\}_{j=1}^{N_{DS}}, i = 1 \dots k$, one for each landmark circle, consisting of eigenfunctions of the intrinsic Dirichlet-Steklov eigenproblem. Each of these basis sets is well-suited to describing smooth functions in the vicinity of its corresponding landmark circle. Intuitively these functions complement the Laplace-Beltrami eigenbasis, are harmonic on the interior of the shapes, and are zero at all but one disk boundary: $u_j^{(i)}|_{\Gamma_l} = 0$ for $l \neq i$ and all j .
4. Compute an optimal functional map by minimizing an energy that promotes near-conformal maps, via an iterative refinement strategy. We split the functional map into $k + 1$ parts, and separately align the Laplacian eigenbasis and each set of Dirichlet-Steklov ones.
5. Convert the computed functional map to a vertex-to-vertex map between the shapes with the disks cut out.
6. Reinsert the landmark correspondences to obtain a landmark-preserving vertex-to-vertex map between the original meshes.

Our general strategy follows the standard functional map pipeline, especially in its recent variants based on iterative refinement [MRR*19; PRM*21; XLZ21], with several crucial changes.

First, our main motivation for introducing disks to represent landmarks in Step 1 is to associate to each landmark a well-defined functional space. In this, we are inspired by techniques that represent landmarks or seed points on a surface via associated harmonic functions [ZRKS05; Pat18] on a mesh. Unlike such harmonic functions, however, our construction is fully justified in the smooth setting. This is because it is impossible to impose boundary conditions on isolated points on a smooth surface. Furthermore, as we




Figure 3: Several Dirichlet Laplacian eigenfunctions of the annulus in 2D with external radius 1 and internal radius $1/2$. Notice that the eigenfunctions concentrate away from the boundary.

show below, the Dirichlet-Steklov eigenfunctions that we compute in Step 3, are *orthogonal to the Dirichlet Laplacian eigenbasis* and jointly form a *complete* basis for the underlying functional space.

Secondly, instead of computing a single functional map across Laplace-Beltrami eigenfunctions, we estimate a *block-diagonal* functional map that aligns each of the $k+1$ components of the functional space separately. This both improves efficiency and promotes desirable structural map properties. Indeed, we prove that this splitting must hold for conformal maps in the smooth setting, and we observe that it promotes preservation of landmark neighborhoods across a wide range of shape deformations in practice.

Finally, rather than focusing on near-isometries, we build a functional map energy that aims at computing near-conformal maps, and fully avoids the use of descriptor functions. Furthermore, we propose an efficient initialization and an iterative strategy for optimizing this energy, while promoting desirable map properties. This ensures high-quality correspondences even in challenging cases, in which existing functional maps-based methods tend to fail.

In the following sections, we discuss each step of this pipeline in detail. Throughout our discussion related to the basis construction and the structure of pointwise and functional maps, we focus mainly on derivations in the smooth setting, to highlight the theoretically justified nature of our approach.

4. Functional Basis

Central to our proposal is a careful choice of functional basis for a convenient functional space over the considered shapes. Our basis is built as the union of the solutions to the Dirichlet Laplacian and Dirichlet-Steklov eigenproblems, which we describe in Secs. 4.1 and 4.2, respectively. In Sec. 4.3 we define the functional space that we use in the rest of the proposal. The constructions described in these sections pertain to manifolds with boundaries and are not yet specialized to our shape matching method, which can be used both for shapes with and without boundaries. The specialization to our case is carried out in Sec. 4.4. There, we describe how to create a landmark adapted functional basis by, roughly speaking, treating the landmarks as boundaries. All the constructions discussed in this section are carried out in the smooth setting. Their discretization is discussed in App. B.

4.1. Dirichlet Laplacian Eigenproblem

Let \mathcal{M} be a smooth, connected, oriented compact Riemannian manifold with metric g and a boundary $\partial\mathcal{M}$. The Dirichlet Laplacian




Figure 4: Some Dirichlet-Steklov eigenfunctions of the same annulus from Fig. 3. The Steklov boundary condition is imposed in turn on the external (top row) and internal (bottom row) boundaries. Notice that the eigenfunctions concentrate on the Steklov boundary.

ean eigenproblem is:

$$\begin{aligned} \Delta\psi_i &= \lambda_i\psi_i \\ \psi_i|_{\partial\mathcal{M}} &= 0, \end{aligned} \quad (1)$$

where Δ denotes the non-negative Laplace-Beltrami operator. Despite the vanishing boundary condition, it can be shown (see [Cha84]) that the eigenfunctions $\{\psi_i\}_{i=1}^{\infty}$ can be chosen to form an orthonormal basis for $L_2(\mathcal{M})$ (square integrable functions of \mathcal{M}). Moreover, the eigenvalues and eigenfunctions can be ordered such that $0 < \lambda_1 \leq \lambda_2 \leq \dots \rightarrow \infty$.

In Fig. 3, we illustrate the first few Dirichlet Laplacian eigenfunctions for an annulus in the plane with external radius 1 and internal radius $1/2$. We will return to this example of the annulus in our discussion to compare the properties of different bases that we consider. We very briefly discuss the discretization of the Dirichlet Laplacian problem on triangle meshes in App. B.

4.2. Dirichlet-Steklov Eigenproblem

Let \mathcal{M} be a smooth, connected, oriented compact Riemannian manifold with metric g and a Lipschitz continuous boundary $\partial\mathcal{M}$. Suppose that, up to sets of measure 0, $\partial\mathcal{M}$ consists of two disjoint nonempty open sets, denoted \mathcal{D} and \mathcal{S} . The (mixed) Dirichlet-Steklov eigenproblem is posed as follows:

$$\begin{aligned} \Delta u_i &= 0 \\ u_i|_{\mathcal{D}} &= 0 \\ \partial_n u_i|_{\mathcal{S}} &= \sigma_i u_i, \end{aligned} \quad (2)$$

where ∂_n denotes the interior normal derivative. The second and third line of the above are respectively known as the Dirichlet and Steklov boundary conditions, explaining the name Dirichlet-Steklov.

As hinted at in Sec. 3 and explained in detail in Sec. 4.4, in our approach, \mathcal{S} will be the boundary corresponding to a given landmark, while \mathcal{D} will be the union of all other landmark boundaries.

The general theory of the Dirichlet-Steklov and many other similar problems can be found in [Neč12]. For a gentle introduction to Steklov eigenproblems, see [Lab17] (in French).

The eigenvalues $\{\sigma_i\}_{i=1}^\infty$ can be ordered such that $0 < \sigma_1 \leq \sigma_2 \leq \dots \rightarrow \infty$. Unlike the eigenfunctions of the Laplacian eigenproblem, the Dirichlet-Steklov eigenfunctions do not form an orthonormal basis for $L_2(\mathcal{M})$. Instead, the restriction of $\{u_i\}_{i=1}^\infty$ to the boundary \mathcal{S} form such a basis for $L_2(\mathcal{S})$ (see [Neč12]).

We emphasize that, in contrast to previous uses of the Steklov eigenproblem in [WBPS18], we consider a purely *intrinsic* problem on the shape surface. I.e., as described in detail in Sec. 4.4 our boundaries are one-dimensional, being the boundaries of disks centered at the landmarks.

As it is written above, the Dirichlet-Steklov problem may seem a bit mysterious. However, it becomes much more approachable when written in weak form:

$$\int_{\mathcal{M}} \nabla f \cdot \nabla u_i \, d\mathcal{M} = \sigma_i \int_{\mathcal{S}} f u_i \, d(\partial\mathcal{M}). \quad (3)$$

In this form, it can be compared to the standard Laplacian eigenproblem: $\int_{\mathcal{M}} \nabla f \cdot \nabla \psi_i \, d\mathcal{M} = \lambda_i \int_{\mathcal{M}} f \psi_i \, d\mathcal{M}$. Intuitively, and as we demonstrate in practice, the Dirichlet-Steklov eigenfunctions “focus” on the boundary \mathcal{S} and provide detailed information in the vicinity of this boundary. As discussed below, in our method we establish one set of Dirichlet-Steklov eigenfunctions for each landmark, and align those functional spaces across the pair of shapes.

A derivation of the weak form of the Dirichlet-Steklov problem is provided in App. A. The discretization of the resulting problem on triangle meshes is discussed in App. B.

We illustrate some Dirichlet-Steklov eigenfunctions for the annulus in Fig. 4. Notice that the eigenfunctions concentrate on the boundary on which the Steklov boundary condition is imposed.

4.3. Functional Space $W(\mathcal{M})$

In this section, we specify the functional space used for the remainder of our proposal. Recall that our goal is to obtain a variant of the functional maps method suitable for non-isometric shape matching. We propose to search for near-conformal maps.

We thus need to translate the search for near-conformality to the functional maps setting. We do so by building upon the foundations laid in the context of conformal shape differences [ROA*13; CSB*17]. Given a pair of surfaces \mathcal{M} and \mathcal{N} , in [ROA*13], it is observed that to study the deviation from conformality of a map $\varphi : \mathcal{N} \rightarrow \mathcal{M}$, it is useful to consider its pullback $F_{\mathcal{M}\mathcal{N}}$ as a map between spaces of functions equipped with the Dirichlet form:

$$\langle f, u \rangle_{W(\mathcal{M})} = \int_{\mathcal{M}} \nabla f \cdot \nabla u \, d\mathcal{M}. \quad (4)$$

The Dirichlet form becomes an inner product on the space of smooth functions modulo constant functions. A Hilbert space is then obtained by taking the completion of the space in the induced topology. We denote the space thus obtained by $W(\mathcal{M})$. We remark that this space is different from the standard L_2 space of square integrable functions, due to the additional *smoothness* conditions. Below we describe both the properties and the utility of this space in the context of our landmark-based shape matching approach.

4.4. Landmark-Adapted Basis for $W(\mathcal{M})$

As highlighted above, a key aspect of our approach is the construction of a novel functional basis that is *adapted to the landmarks*.

Our main idea is to treat the landmarks as boundaries at which the functional bases satisfy certain boundary conditions. For this, we first slightly modify the shapes under study. Indeed, while advocated in several prior works [ZRKS05; Pat18] in geometry processing, it is not strictly speaking possible to impose boundary conditions at isolated points in the smooth setting.

Let \mathcal{M} and \mathcal{N} be compact, connected, oriented Riemannian surfaces. For simplicity, we also temporarily assume them to be without boundary. This last assumption is removed later. Let $\{\gamma_i^{\mathcal{M}}\}_{i=1}^k \subset \mathcal{M}$ and $\{\gamma_i^{\mathcal{N}}\}_{i=1}^k \subset \mathcal{N}$ be k landmarks in one-to-one correspondence. That is, we will be looking for maps $\varphi : \mathcal{N} \rightarrow \mathcal{M}$ such that $\varphi(\gamma_i^{\mathcal{N}}) = \gamma_i^{\mathcal{M}}$ for all $i = 1 \dots k$. Such φ are said to be landmark preserving. The functional map representation of φ , that is its pullback on functions, will be denoted $F_{\mathcal{M}\mathcal{N}}$, as before.

Our first step is to convert the landmarks into proper boundaries. We do so by removing small disks centered at the landmarks and treat the boundaries of these disks as boundaries of the shapes. We make sure that none of the disks intersect. Thus, we end up with a new shape that has k boundary components, one for each landmark. We denote the boundary corresponding to the landmark $\gamma_i^{\mathcal{M}}$ by $\Gamma_i^{\mathcal{M}}$. By abuse of notation, we denote the shapes thus modified by \mathcal{M} and \mathcal{N} , same as their original versions. On triangle meshes, we create boundaries that are fully contained in a one-ring neighborhood of each landmark. This operation is described in detail in App. C.

We now use the newly created boundaries to split $W(\mathcal{M})$ into convenient subspaces. These subspaces will be composed of functions satisfying carefully chosen eigenvalue problems and boundary conditions.

We begin by considering the span of Laplace-Beltrami eigenfunctions satisfying Dirichlet boundary conditions on the $\{\Gamma_i\}_i$:

$$\begin{aligned} \Delta \psi_i &= \lambda_i \psi_i, \\ \psi_i|_{\Gamma_j} &= 0, \quad \forall i, j. \end{aligned} \quad (5)$$

Recall that the eigenfunctions $\{\psi_i\}_{i=1}^\infty$ form an orthonormal basis for $L_2(\mathcal{M})$. They remain mutually orthogonal in $W(\mathcal{M})$, but interestingly fail to form a full basis for that space. This counter-intuitive behavior is due to the change of topology from $L_2(\mathcal{M})$ to $W(\mathcal{M})$ and the infinite dimensionality of the functional spaces under consideration.

Let the $W(\mathcal{M})$ closure of the subspace spanned by the $\{\psi_i\}_{i=1}^\infty$ be denoted by $\mathcal{G}(\mathcal{M})$.

Naturally, our next step is to find functions that span the remainder of $W(\mathcal{M})$. This is where the Dirichlet-Steklov eigenfunctions of Sec. 4.2 come in. We pose k Dirichlet-Steklov problems, with

the j^{th} problem being:

$$\begin{aligned} \Delta u_i^{(j)} &= 0 \\ u_i^{(j)}|_{\Gamma_q} &= 0, & q \neq j \\ \partial_n u_i^{(j)}|_{\Gamma_j} &= \sigma_i^{(j)} u_i^{(j)}. \end{aligned} \quad (6)$$

This results in k Dirichlet-Steklov eigenbases and spectra denoted $\{u_i^{(j)}\}_{i=1}^{\infty}$ and $\{\sigma_i^{(j)}\}_{i=1}^{\infty}$, respectively. Recall that the $\{u_i^{(j)}\}_{i=1}^{\infty}$ form an orthonormal basis for $L_2(\Gamma_j)$. These functions remain mutually orthogonal in $W(\mathcal{M})$. This follows directly from the weak form of the Dirichlet-Steklov problem (Eq. (3)). We denote the $W(\mathcal{M})$ closed span of the $\{u_i^{(j)}\}_{i=1}^{\infty}$ by $\mathcal{H}_j(\mathcal{M})$.

Our key result is that, once put together, the Dirichlet Laplacian eigenfunctions and the k sets of Dirichlet-Steklov eigenfunctions span all of $W(\mathcal{M})$.

Lemma 1. *The function space $W(\mathcal{M})$ admits the following decomposition:*

$$W(\mathcal{M}) = \mathcal{G}(\mathcal{M}) \oplus \overline{\left(\bigoplus_{j=1}^k \mathcal{H}_j(\mathcal{M}) \right)}, \quad (7)$$

where \oplus denotes direct sums and $\overline{\oplus}$ denotes orthogonal direct sums, and the overline denotes the $W(\mathcal{M})$ closure of the spanned functional space.

Proof. See App. D. \square

Intuitively, the above lemma says that $W(\mathcal{M})$ can be split into a non-harmonic part and k harmonic landmark-associated subspaces, with each landmark getting its own subspace of harmonic functions that are nonvanishing on it. In practice, we always $W(\mathcal{M})$ -normalize all of the considered eigenfunctions by dividing each function by its W norm. In all of the following, we use $W(\mathcal{M})$ -normalized bases.

The resultant basis is thus normalized. However, it is not quite $W(\mathcal{M})$ -orthogonal, as suggested by the notation used Lemma 1. Specifically, the problem lies in the mutual non-orthogonality of the subspaces $\mathcal{H}_j(\mathcal{M})$. This is discussed in App. D.

In principle, the energy that we are to minimize (see Sec. 5 below) can be expressed in any basis, even if it is not orthogonal. For our purposes, however, the non-orthogonality of our basis poses a few challenges, which will be detailed later. Fortunately, in practice, *our basis can be accurately approximated as orthonormal*. A typical matrix of $W(\mathcal{M})$ inner products is shown in Fig. 5 (see App. H for an extended evaluation of this approximation). In Fig. 6 we evaluate the orthonormality in the case of the 2D annulus and observe that it becomes more and more valid as the radius of the inner disk becomes smaller. We will call attention to this approximation when we use it in the implementation of our proposal.

Before proceeding further, we note briefly that on shapes with pre-existing boundaries we impose Neumann boundary conditions (vanishing normal derivatives). The above discussion remains unchanged. Note that imposing Neumann boundary conditions requires no special effort in the discrete setting.




Figure 5: W -inner products for the first 20 Dirichlet-Steklov eigenfunctions corresponding to six landmark circles on a sphere mesh (left). The first three landmarks are on the top left and the remaining three are on the bottom right. Notice that the different \mathcal{H}_i subspaces are almost orthogonal.




Figure 6: W -inner products for the first 30 Dirichlet-Steklov eigenfunctions corresponding to the two boundaries of the annulus. The external radius of the annulus is 1, while different values of the internal radius r_i are considered. \mathcal{H}_1 and \mathcal{H}_2 correspond to the internal and external boundaries of the annulus, respectively. In this case, approximation of orthogonality of \mathcal{H}_1 and \mathcal{H}_2 fails for a large r_i , but becomes more and more valid as r_i decreases.

We now illustrate our functional basis using the landmark circles and Neumann boundary conditions on both the inner and outer boundary of the annulus in Fig. 7. Notice that as their eigenvalue increases, the Dirichlet-Steklov eigenfunctions rapidly concentrate on the landmark circles. In fact, the eigenfunctions of the closely related Steklov eigenproblem (i.e. without the Dirichlet boundary) are known to decay exponentially with distance from the Steklov boundary, the rate of decay being proportional to the corresponding eigenvalue [PST19]. In contrast, the Dirichlet-Laplacian eigenfunctions remain evenly spread in the bulk of the shape. Thus, high eigenvalue Dirichlet-Steklov eigenfunctions are uninformative regarding the bulk of the manifold. Meanwhile, the high eigenvalue Dirichlet Laplacian eigenfunctions remain informative in the bulk even at high eigenvalues.

So far, we assumed that the considered shapes were connected. Our discussion remains unchanged on general shapes, as long as each connected component has at least two landmarks on it, as this is necessary to impose both boundary conditions of the Dirichlet-Steklov eigenproblem. If this is not satisfied for some connected component, at least some of the considered eigenproblems will have eigenfunctions that are piecewise constant per component and correspond to eigenvalue 0. These should not be included in a basis




Figure 7: Dirichlet-Steklov (top row) and Dirichlet Laplacian (bottom row) eigenfunctions on an annulus with three landmark circles and Neumann conditions on inner and outer boundaries. The Dirichlet-Steklov eigenfunctions correspond to the landmark on the right. Notice that as the eigenvalues increase, the Dirichlet-Steklov eigenfunctions quickly concentrate around the corresponding landmark, unlike the Dirichlet Laplacian eigenfunctions that remain distributed in the bulk of the annulus.

for $W(\mathcal{M})$, as they have vanishing W -norm. We avoid this issue by rejecting eigenfunctions with eigenvalues below a certain small threshold. Note that for components with *one* landmark we only impose the Steklov condition on the corresponding circle, omitting the second line of Eq. (2).

4.5. Structure of the Functional Map

Recall that our ultimate goal is to compute a near-conformal diffeomorphism $\varphi : \mathcal{N} \rightarrow \mathcal{M}$ that preserves the landmarks. Recall also that we propose to use functional map methods to find it. In this section we translate the structural properties of φ into properties of its pullback $F_{\mathcal{M}\mathcal{N}}$ which helps us to restrict the space of admissible functional maps, which is crucial for our approach.

We begin on a technical note. Since we have replaced landmark points with landmark circles, the notion of landmark preservation has to be slightly adjusted. We no longer can claim something as simple as $\varphi(\gamma_i^{\mathcal{N}}) = \gamma_i^{\mathcal{M}}$ for all i , as the landmark points are no longer part of the considered shapes. Instead we impose that φ restricts to a diffeomorphism on corresponding landmark circles. That is, $\varphi : \mathcal{N} \rightarrow \mathcal{M}$ is a diffeomorphism and for each i , $\varphi|_{\Gamma_i^{\mathcal{N}}} : \Gamma_i^{\mathcal{N}} \rightarrow \Gamma_i^{\mathcal{M}}$ is also a diffeomorphism.

Now, suppose that φ is indeed a conformal map. Then, $F_{\mathcal{M}\mathcal{N}}$ satisfies the following lemma.

Lemma 2 (Structure of $F_{\mathcal{M}\mathcal{N}}$). *Let $F_{\mathcal{M}\mathcal{N}} : W(\mathcal{M}) \rightarrow W(\mathcal{N})$ be the pullback of a conformal diffeomorphism that preserves the landmark circles in the sense described above. Then, $F_{\mathcal{M}\mathcal{N}}$ maps*

1. $\mathcal{G}(\mathcal{M})$ to $\mathcal{G}(\mathcal{N})$,
2. $\mathcal{H}_j(\mathcal{M})$ to $\mathcal{H}_j(\mathcal{N})$ for all j .

Proof. See App. E. □

The above lemma provides necessary, but not sufficient conditions for $F_{\mathcal{M}\mathcal{N}}$ to be the pullback of a diffeomorphism preserving the landmark circles. Nonetheless, we will use properties (1) and (2) of Lemma 2 to structure our search for $F_{\mathcal{M}\mathcal{N}}$.

From now on, we only consider functional maps that satisfy statements (1) and (2) of Lemma 2. This can be seen as $k+1$ separate maps, one for each landmark subspace \mathcal{H}_i and one for the orthogonal complement \mathcal{G} , assembled into one block-diagonal functional map. Intuitively, this keeps the overall map tethered to the landmarks.

Landmark preservation At this point, it is worth explaining what we mean when we say that our method preserves the landmark correspondences *in the discrete setting*. Indeed, the challenge of landmark preservation is to not merely enforce the condition $\varphi(\gamma_i^{\mathcal{N}}) = \gamma_i^{\mathcal{M}}$, but to also obtain a smooth (or at least continuous) map in the neighborhood of the landmarks (notice that we required φ to be a diffeomorphism when discussing the smooth setting). Our method achieves this by using a functional basis whose elements are well suited to describe smooth functions near the landmarks (recall the decay of the Dirichlet-Steklov eigenfunctions away from the Steklov boundary depicted in Fig. 7). By enforcing the functional map structure of Lemma 2 during the entire solution process, we promote vertex-to-vertex maps that smoothly map the neighborhoods of the landmarks of \mathcal{N} to the corresponding neighborhoods on \mathcal{M} , the smoothness of the map reflecting the smoothness of the functional basis. Furthermore, we reinsert the original pointwise landmarks at the end of the solution process to preserve the initial landmarks exactly. Recall that the landmark vertices are excluded from the meshes the moment the landmark circles are introduced.

5. Functional Map Energy

The previous section describes our landmark adapted basis construction, and the block-diagonal structure of landmark-preserving conformal maps when expressed in this basis. In this section we specify the optimization problem that we will solve in order to obtain landmark-preserving maps between triangle meshes.

Recall that we propose to look for conformal maps, which can be characterized in terms of the Dirichlet form ($W(\mathcal{M})$ inner product).

Theorem 3. *Let $\varphi : \mathcal{N} \rightarrow \mathcal{M}$ be a diffeomorphism between oriented Riemannian surfaces with pullback $F_{\mathcal{M}\mathcal{N}} : W(\mathcal{M}) \rightarrow W(\mathcal{N})$. Then, φ is conformal if and only if*

$$\langle u, v \rangle_{W(\mathcal{M})} = \langle F_{\mathcal{M}\mathcal{N}} u, F_{\mathcal{M}\mathcal{N}} v \rangle_{W(\mathcal{N})}, \quad \forall u, v \in W(\mathcal{M}). \quad (8)$$

Proof. See [ROA*13]. □

In practice we do not expect to obtain an exact equality of the inner products as described in the previous theorem. Instead, we will search for φ and $F_{\mathcal{M}\mathcal{N}}$ by relaxing the above equality to a minimization problem. Let $\Phi^{\mathcal{M}}$ and $\Phi^{\mathcal{N}}$ denote reduced (finite dimensional) functional bases for $W(\mathcal{M})$ and $W(\mathcal{N})$, respectively. These bases consist of the eigenfunctions of the Dirichlet Laplacian and Dirichlet-Steklov eigenproblems corresponding to small eigenvalues. The precise size of the bases is discussed in App. H.

From now on, we concentrate our attention on the discrete case.

Namely, \mathcal{M} and \mathcal{N} will now denote oriented manifold triangle meshes. Letting $\langle \Phi^{\mathcal{M}}, \Phi^{\mathcal{M}} \rangle_{W(\mathcal{M})}$ be the matrix of all inner products of the normalized basis vectors of $\Phi^{\mathcal{M}}$, we relax the equality of Theorem 3 to the minimization of the following energy term:

$$E_c(F_{\mathcal{M}\mathcal{N}}) = \left\| \langle \Phi^{\mathcal{M}}, \Phi^{\mathcal{M}} \rangle_{W(\mathcal{M})} - \langle F_{\mathcal{M}\mathcal{N}} \Phi^{\mathcal{M}}, F_{\mathcal{M}\mathcal{N}} \Phi^{\mathcal{M}} \rangle_{W(\mathcal{N})} \right\|_F^2. \quad (9)$$

We call this the *conformal term* of the energy. Here, as well as everywhere else in this text, $\|\cdot\|_F$ denotes the Frobenius norm.

Having covered the conformality of the map, it remains to rephrase the restriction of $F_{\mathcal{M}\mathcal{N}}$ to pullbacks of landmark-preserving diffeomorphisms. This assumption cannot be exactly imposed in the discrete case. Still, we would like $F_{\mathcal{M}\mathcal{N}}$ to exhibit the properties of such a map. In order to do so, we complete our energy by specifying two *structural terms*. Specifically, the first term promotes $F_{\mathcal{M}\mathcal{N}}$ being a *proper* functional map (i.e., the pullback of a vertex-to-vertex map), as recently defined in [RMWO21], and the second promotes the invertibility of $F_{\mathcal{M}\mathcal{N}}$ [ERGB16].

Let $\Pi_{\mathcal{N}\mathcal{M}}$ denote the vertex-to-vertex map from \mathcal{N} to \mathcal{M} expressed as a matrix (i.e. a binary matrix that contains exactly one 1 per row). Then, $F_{\mathcal{M}\mathcal{N}}$ should satisfy:

$$F_{\mathcal{M}\mathcal{N}} = (\Phi^{\mathcal{N}})^+ \Pi_{\mathcal{N}\mathcal{M}} \Phi^{\mathcal{M}}, \quad (10)$$

where $(\Phi^{\mathcal{N}})^+$ denotes the pseudoinverse of $\Phi^{\mathcal{N}}$, or in other words, the $W(\mathcal{N})$ projection onto the reduced basis $\Phi^{\mathcal{N}}$. As before, we relax the equality into an energy to be optimized:

$$E_p(F_{\mathcal{M}\mathcal{N}}, \Pi_{\mathcal{N}\mathcal{M}}) = \left\| (\Phi^{\mathcal{N}})^+ \Pi_{\mathcal{N}\mathcal{M}} \Phi^{\mathcal{M}} - F_{\mathcal{M}\mathcal{N}} \right\|_F^2. \quad (11)$$

We call this the *properness term* of the energy. Notice that we have expressed the energy as a function of both F and Π . We do so as we will have to consider these two objects as independent variables when minimizing the energy. The exact meaning of this is discussed in Sec. 6.

In addition to $F_{\mathcal{M}\mathcal{N}}$ arising from a point-to-point map, we would also like for it to be invertible. For this, we consider two maps $F_{\mathcal{M}\mathcal{N}} : W(\mathcal{M}) \rightarrow W(\mathcal{N})$ and $F_{\mathcal{N}\mathcal{M}} : W(\mathcal{N}) \rightarrow W(\mathcal{M})$, the latter arising from a vertex-to-vertex map $\Pi_{\mathcal{M}\mathcal{N}} : \mathcal{M} \rightarrow \mathcal{N}$. Thus, in what follows, we will be simultaneously optimizing for maps going in both directions between the shapes. With I being the identity matrix, the invertibility condition is, of course:

$$\begin{aligned} F_{\mathcal{N}\mathcal{M}} F_{\mathcal{M}\mathcal{N}} &= I, \\ F_{\mathcal{M}\mathcal{N}} F_{\mathcal{N}\mathcal{M}} &= I. \end{aligned} \quad (12)$$

Once again, we convert the above into minimization form. The *invertibility term* corresponding to the first line above is

$$E_{I,\mathcal{M}\mathcal{N}}(F_{\mathcal{M}\mathcal{N}}, F_{\mathcal{N}\mathcal{M}}) = \|F_{\mathcal{M}\mathcal{N}} F_{\mathcal{N}\mathcal{M}} - I\|_F^2. \quad (13)$$

The invertibility term $E_{I,\mathcal{N}\mathcal{M}}$ is defined analogously.

In sum, our search for the correspondence between \mathcal{M} and \mathcal{N} will involve the joint minimization of the energy

$$\begin{aligned} E_{\mathcal{M}\mathcal{N}} &= a_C E_c(F_{\mathcal{M}\mathcal{N}}) + \\ & a_P E_p(F_{\mathcal{M}\mathcal{N}}, \Pi_{\mathcal{N}\mathcal{M}}) + \\ & a_I E_{I,\mathcal{M}\mathcal{N}}(F_{\mathcal{M}\mathcal{N}}, F_{\mathcal{N}\mathcal{M}}) \end{aligned} \quad (14)$$

and an analogously defined energy $E_{\mathcal{N}\mathcal{M}}$. Here, a_C, a_P and a_I are nonnegative tunable weights controlling the relative strength of the conformality, properness and invertibility terms, respectively. Different values of these parameters are explored in App. H.

The above energy is conformally invariant in the following sense.

Lemma 4 (Energy Invariance). *The conformality, properness and invertibility terms of the energy (Eqs. (9), (11) and (13)), as well as the energy (their weighted sum, Eq. (14)) are invariant under (combinations of) the following transformations:*

1. *Conformal transformations of the meshes keeping the reduced bases fixed.*
2. *Orthogonal transformations of the reduced bases.*

Proof. Let the functional and vertex-to-vertex maps be fixed. Since conformal transformations leave the W inner product invariant, the energy terms are conformally invariant for a fixed choice of functional basis. Statement 1. is now proven. Statement 2. follows from the fact the Frobenius norm is invariant under orthogonal transformations. \square

Note that, in the lemma above, conformal transformations and changes of basis are treated as independent. In practice, they are not, as reduced bases are usually mesh-dependent. Thus, the change of basis induced by a conformal transformation may fail to be orthogonal and then Lemma 4 will not apply. As long as one works with reduced rather than full bases (i.e. spanning all functions of the mesh), the invariance of the energy under conformal transformations is therefore only approximately guaranteed.

The conformal invariance of the energy can be violated in another way. Suppose that the recipe for constructing the reduced bases produces non-orthogonal bases. Then, the change of basis induced by a conformal transformation may fail to be orthogonal even when full bases are used.

Of course, we raise the previous two issues precisely because our method uses non-orthogonal reduced bases. Thus, Lemma 4 does not offer a full guarantee of conformal invariance for our energy. Still, the bases that we use turn out to be very nearly orthogonal and thus the energy that we employ remains approximately conformally invariant. Obtaining a truly conformally invariant energy (at least up to basis truncation) is a subject for future work.

6. Solving the Problem

In this section, we propose an efficient approach for the optimization problem posed in Eq. (14). Our approach is inspired by a discrete optimization strategy, first suggested in [MRR*19] and recently extended to other general energies [RMWO21]. The general idea is to recast the problem in a way that makes every iteration of the optimization into a nearest-neighbor search. The overall process

then consists of two qualitatively different parts. First, an initial guess of the correspondence is obtained. Then, the correspondence is refined via the iterative process mentioned above. These steps are explained in Secs. 6.1 and 6.2, respectively.

6.1. Initial Correspondence

In this section we explain how we obtain an initial guess of the functional maps $F_{\mathcal{M}\mathcal{N}}$ and $F_{\mathcal{N}\mathcal{M}}$. A common way to initialize functional maps with landmarks is via descriptor preservation [OBS*12; RPWO18]. However, common descriptors such as HKS or WKS [SOG09; ASC11] strongly rely on the isometry assumption and moreover the initial functional map is not guaranteed to respect landmark correspondences exactly. To overcome this, we propose a simple and lightweight initialization scheme.

Recall that our approach upgrades landmark correspondence to landmark circle correspondence. Moreover, in the smooth setting, we require this correspondence to be a diffeomorphism. We now make the assumption that the correspondence between landmark circles can be seen as a rotation of one circle to match the other.

Specifically, we label the vertices of each landmark circle in counter-clockwise order using the outward-facing normal orientation. We can then assign each vertex in a landmark circle coordinates in $[0, 1)$. All that remains is to ensure that the origin of this coordinate system is placed consistently on both shapes. In other words, the matching of two corresponding landmark circles reduces to finding an appropriate shift of one of the parametrizations.

We propose to align the parametrizations of the boundary circles such that the landmark circles are consistently oriented relative to the other landmarks. In order to do so, we construct functions on the landmark circles that have maxima in directions roughly pointing towards other landmarks. Consider the following problem:

$$\begin{aligned} \Delta h_i &= 0, \quad i = 1 \dots k, \\ h_i|_{\Gamma_j} &= \delta_{ij}. \end{aligned} \quad (15)$$

This results in k harmonic functions, one for each landmark, where each function h_i equals 1 on the boundary of landmark circle i , and zeros on the boundaries of other circles. As they stand, these functions are constant on each landmark circle. Their normal derivatives, however, are not. In essence, we use $\partial_n|_{\Gamma_i} h_j$ as an indication of the direction one should take from Γ_i to reach Γ_j . See Fig. 8 for an illustration. This is similar in spirit to the Geodesics in Heat construction [CWW13], where gradients of solutions to the heat transfer problem are used to construct approximate geodesics.

Denoting the landmark circle coordinates on $\Gamma_i^{\mathcal{M}}$ by θ_i , we select the optimal shift α_i by solving:

$$\alpha_i = \arg \min_{\alpha} \left(\sum_{j \neq i} \left\| \partial_n|_{\Gamma_i^{\mathcal{M}}} h_j^{\mathcal{M}}(\theta_i) - \partial_n|_{\Gamma_i^{\mathcal{N}}} h_j^{\mathcal{N}}(\text{mod}(\theta_i - \alpha, 1)) \right\|_{L_2(\Gamma_i^{\mathcal{M}})}^2 \right). \quad (16)$$

In this problem, we consider each landmark i and examine directions to all other landmarks (via normal derivatives). We then align the coordinates of landmarks on \mathcal{M} and \mathcal{N} so that these directions align in the best possible way. This problem can be solved simply




Figure 8: Harmonic function satisfying Eq. (15) corresponding to the top landmark on a disk with three landmark circles. Notice that the gradient of the function roughly points towards the top landmark. Hence, its normal derivatives at the bottom landmark circles can be seen as specifying the direction towards the top landmark.

by directly examining all possible shifts and taking the optimum. In order to gain robustness to changes in triangulation, we first project the normal derivatives (as circular functions $\partial_n|_{\Gamma_i^{\mathcal{M}}} h_j^{\mathcal{M}}(\theta)$) onto the reduced basis in order to remove spurious high frequency components. Recall that this makes sense as the Dirichlet-Steklov eigenfunctions belonging to landmark Γ_i form a basis for $L_2(\Gamma_i)$.

Converting the optimal shifts α_i into vertex-to-vertex correspondences on the landmark circles is a matter of a nearest-neighbor search between the circular coordinates of the vertices of $\Gamma_i^{\mathcal{M}}$ and the α_i -shifted coordinates of the vertices of $\Gamma_i^{\mathcal{N}}$.

It remains to convert the resulting vertex-to-vertex map into a functional map. Once again, recall that our reduced basis contains an L_2 basis for each landmark circle. Thus, by using an expression of the form of Eq. (10) we can construct functional maps between $\mathcal{H}_i(\mathcal{M})$ and $\mathcal{H}_i(\mathcal{N})$. Assembling the resulting maps into block-diagonal matrices gives our initial guesses of $F_{\mathcal{M}\mathcal{N}}$ and $F_{\mathcal{N}\mathcal{M}}$.

In App. G, we compare this approach to two alternative initialization strategies, and demonstrate its relative advantages.

6.2. Energy Minimization via Nearest Neighbor Search

The procedure described in Sec. 6.1 provides a descriptor-free initial guess for the functional map. In this section we describe a refinement method that significantly improves the map.

Recall that we are looking for a vertex-to-vertex map by minimizing an energy that depends on both the point-to-point and the associated functional map (pullback). In [MRR*19] it is observed that a particular case of such problems can be efficiently solved by considering the two maps as being independent variables. This observation was recently extended to a wide range of energies in [RMWO21]. Following this line of work we will move all of the difficult optimization on the side of the vertex-to-vertex map and use Eq. (10) to restore the relationship between the maps.

Our main tool is the following result, standard in functional maps literature [EB17; RMWO21], which allows one to reduce optimization problems of a certain form to nearest neighbor searches.

Lemma 5. *Let A be a symmetric positive-definite matrix inducing the matrix norm $\|M\|_A^2 = \text{Tr}(M^T A M)$. Let Φ be a reduced basis orthogonal with respect to A , that is $\Phi^T A \Phi = I$. Then, given n pairs of matrices X_i and Y_i , the following two expressions are equal:*

1. $\sum_i^n \left(\|\Phi^T A \Pi X_i - Y_i\|_F^2 + \|(I - \Phi \Phi^T A) \Pi X_i\|_F^2 \right)$
2. $\sum_i^n \|\Pi X_i - \Phi Y_i\|_F^2$

Moreover, if A is diagonal, minimizing the above expressions over matrices Π that reflect point-to-point maps (i.e., binary matrices that contain exactly one 1 per row) is equivalent to

$$\min_{\Pi} \sum_i^n \|\Pi X_i - \Phi Y_i\|_F^2. \quad (17)$$

This problem can be solved via nearest-neighbor search between the rows of the concatenated matrices $[X_1 \dots X_n]$ and $[\Phi Y_1 \dots \Phi Y_n]$.

Proof. See [EB17] for a proof of a special case and [RMWO21] (Lemma 4.1) for the general statement. \square

We first convert the energy of Eq. (14) into the form used in the above lemma. For brevity's sake, we will only develop the expression for $F_{\mathcal{M}\mathcal{N}}$ and $\Pi_{\mathcal{N}\mathcal{M}}$. The expression for the pair $F_{\mathcal{N}\mathcal{M}}$ and $\Pi_{\mathcal{M}\mathcal{N}}$ is analogous. As mentioned in Sec. 4.4, we approximate the functional bases $\Phi^{\mathcal{M}}$ and $\Phi^{\mathcal{N}}$ to be orthonormal with respect to the Dirichlet form. Then, the energy minimized by the desired $F_{\mathcal{M}\mathcal{N}}$ and $\Pi_{\mathcal{N}\mathcal{M}}$ becomes:

$$\begin{aligned} E_{\mathcal{M}\mathcal{N}} = & a_c \left\| I - (F_{\mathcal{M}\mathcal{N}})^T F_{\mathcal{M}\mathcal{N}} \right\|_F^2 \\ & + a_p \left\| (\Phi^{\mathcal{N}})^T W^{\mathcal{N}} \Pi_{\mathcal{N}\mathcal{M}} \Phi^{\mathcal{M}} - F_{\mathcal{M}\mathcal{N}} \right\|_F^2 \\ & + a_l \|F_{\mathcal{M}\mathcal{N}} F_{\mathcal{N}\mathcal{M}} - I\|_F^2. \end{aligned} \quad (18)$$

Here, we used the approximation of basis orthonormality in two ways. First, we used it to evaluate the inner products in the conformality term (first line of the above equation). Second, we used it to express $(\Phi^{\mathcal{N}})^+ = (\Phi^{\mathcal{N}})^T W^{\mathcal{N}}$, where $W^{\mathcal{N}}$ is the so-called cotangent Laplacian on \mathcal{N} , which also corresponds to the piecewise linear finite element discretization of the Dirichlet form. We are still a few modifications away from being able to apply Lemma 5 to this problem.

We obtain the desired form for the expression by replacing certain instances of $F_{\mathcal{M}\mathcal{N}}$ with its expression in terms of the vertex-to-vertex map $\Pi_{\mathcal{N}\mathcal{M}}$: $(\Phi^{\mathcal{N}})^T W^{\mathcal{N}} \Pi_{\mathcal{N}\mathcal{M}} \Phi^{\mathcal{M}}$. By using the fact that $\|I - F^T F\|_F^2 = \|F F^T - I\|_F^2$ and making this replacement, we obtain:

$$\begin{aligned} E_{\mathcal{M}\mathcal{N}} = & a_c \left\| (\Phi^{\mathcal{N}})^T W^{\mathcal{N}} \Pi_{\mathcal{N}\mathcal{M}} \Phi^{\mathcal{M}} (F_{\mathcal{M}\mathcal{N}})^T - I \right\|_F^2 \\ & + a_p \left\| (\Phi^{\mathcal{N}})^T W^{\mathcal{N}} \Pi_{\mathcal{N}\mathcal{M}} \Phi^{\mathcal{M}} - F_{\mathcal{M}\mathcal{N}} \right\|_F^2 \\ & + a_l \left\| (\Phi^{\mathcal{N}})^T W^{\mathcal{N}} \Pi_{\mathcal{N}\mathcal{M}} \Phi^{\mathcal{M}} F_{\mathcal{N}\mathcal{M}} - I \right\|_F^2. \end{aligned} \quad (19)$$

Now, all of the terms of the above are of the form $\|\Phi^T A \Pi X_i - Y_i\|_F^2$, with $W^{\mathcal{N}}$ playing the role of the matrix A . Our energy is thus of the form of line (1) of Lemma 5, up to three terms of the form $\|(I - \Phi \Phi^T A) \Pi X_i\|_F^2$. Notice that $(I - \Phi \Phi^T A)$ is the orthogonal projection onto the orthogonal complement of the reduced (approximately) orthonormal basis Φ . Thus, this term can be seen as a regularizer penalizing solutions lying outside of the considered reduced basis. Indeed, this is how this term was originally introduced

in [EB17]. Consequently, by implicitly introducing the appropriate regularizers we can use the first part of Lemma 5 to obtain the following expression for the energy:

$$\begin{aligned} E_{\mathcal{M}\mathcal{N}} = & a_c \left\| \Pi_{\mathcal{N}\mathcal{M}} \Phi^{\mathcal{M}} (F_{\mathcal{M}\mathcal{N}})^T - \Phi^{\mathcal{N}} \right\|_{W^{\mathcal{N}}}^2 \\ & + a_p \left\| \Pi_{\mathcal{N}\mathcal{M}} \Phi^{\mathcal{M}} - \Phi^{\mathcal{N}} F_{\mathcal{M}\mathcal{N}} \right\|_{W^{\mathcal{N}}}^2 \\ & + a_l \left\| \Pi_{\mathcal{N}\mathcal{M}} \Phi^{\mathcal{M}} F_{\mathcal{N}\mathcal{M}} - \Phi^{\mathcal{N}} \right\|_{W^{\mathcal{N}}}^2. \end{aligned} \quad (20)$$

At this point we are forced to make an approximation. Namely, we assume that the second part of the lemma applies, which would normally require $W^{\mathcal{N}}$ to be diagonal. In other words, we convert the problem into a nearest neighbor search without having the guarantee of the equivalence of solutions. Despite this approximation, we have observed that the resulting approach works remarkably well in practice.

This finally brings us to the procedure that we use to minimize the energy. As mentioned above, we will consider the functional and vertex-to-vertex maps as independent variables. Thus, given functional maps $F_{\mathcal{M}\mathcal{N}}$ and $F_{\mathcal{N}\mathcal{M}}$, the point-to-point map $\Pi_{\mathcal{N}\mathcal{M}}$ can be found by solving the nearest-neighbor search problem:

$$\Pi_{\mathcal{N}\mathcal{M}} = \text{NNS} \left(\begin{bmatrix} \Phi^{\mathcal{M}} (F_{\mathcal{M}\mathcal{N}})^T & \Phi^{\mathcal{M}} & \Phi^{\mathcal{M}} F_{\mathcal{N}\mathcal{M}} \\ \Phi^{\mathcal{N}} & \Phi^{\mathcal{N}} F_{\mathcal{M}\mathcal{N}} & \Phi^{\mathcal{N}} \end{bmatrix} \right). \quad (21)$$

Here $\text{NNS}(A, B)$ denotes a set of nearest neighbor problems: for each row of B among the rows of A . The vertex-to-vertex map $\Pi_{\mathcal{M}\mathcal{N}}$ can be obtained analogously. In sum, minimizing the energy with respect to the vertex-to-vertex maps is also a recipe for converting functional maps into vertex-to-vertex maps, while taking into account the original functional map energy.

We are now ready to formulate the optimization algorithm. Following [MRR*19], the overall procedure is based on an iterative spectral upsampling of the functional map. Specifically, we iteratively convert the functional map into a vertex-to-vertex map while increasing the size of the reduced basis. As explained earlier (see Fig. 7), the Dirichlet-Steklov functions are concentrated near the landmark circles. Thus, increasing their number does not provide much additional information about the map in the bulk of the shapes. Thus, we only increase the number of Dirichlet Laplacian eigenfunctions.

Beginning from the initial functional maps $F_{\mathcal{M}\mathcal{N}}$ and $F_{\mathcal{N}\mathcal{M}}$ obtained in Sec. 6.1, we proceed as follows.

1. Convert $F_{\mathcal{M}\mathcal{N}}$ and $F_{\mathcal{N}\mathcal{M}}$ into $\Pi_{\mathcal{N}\mathcal{M}}$ and $\Pi_{\mathcal{M}\mathcal{N}}$ via Eq. (21).
2. Increase the reduced bases $\Phi^{\mathcal{M}}$ and $\Phi^{\mathcal{N}}$ by including k_{step} additional Dirichlet Laplacian eigenfunctions.
3. Update the functional maps to the new basis size via $F_{\mathcal{M}\mathcal{N}} = (\Phi^{\mathcal{N}})^+ \Pi_{\mathcal{N}\mathcal{M}} \Phi^{\mathcal{M}}$ and $F_{\mathcal{N}\mathcal{M}} = (\Phi^{\mathcal{M}})^+ \Pi_{\mathcal{M}\mathcal{N}} \Phi^{\mathcal{N}}$.
4. Iterate steps (1) to (3) until the desired basis size is reached.
5. Repeat step (1) using only the original non-landmark vertices. This produces a vertex-to-vertex map between the original meshes, landmarks excluded.
6. Insert the landmark correspondence into the vertex-to-vertex map.

A Fast Approximation. We conclude this section by proposing an

acceleration strategy to perform the nearest-neighbor search. The method proposed here is unprincipled, but is validated by both the overall quality of our results and explicit tests found in App. G. The method proposed below is the only one used in the main text of this paper.

In the language of Lemma 5, we propose to replace the nearest neighbor search between the concatenated matrices $[X_1 \dots X_n]$ and $[\Phi Y_1 \dots \Phi Y_n]$ by a nearest neighbor search between the summed matrices $X_1 + \dots + X_n$ and $\Phi Y_1 + \dots + \Phi Y_n$. This corresponds to solving the following problem:

$$\min_{\Pi} \left\| \sum_i^n \Pi X_i - \Phi Y_i \right\|_F^2. \quad (22)$$

This reformulation helps to decrease the dimensionality of the nearest neighbor searches. Essentially, we assume that the different energy terms will not cancel each other. The payoff for this approximation is that the matrices involved in the nearest-neighbor search become n times smaller. In our case, there are $n = 3$ energy terms. The experiments in App. G show that this reduction in matrix size results in a slightly more than threefold speed-up.

7. Evaluation

We evaluate our method[†] on standard shape matching datasets, which we describe in Sec. 7.1. We first analyze the parameters involved in our computations (Sec. 7.2). Second, we conduct an in-depth evaluation to compare our method to state-of-the-art approaches on shape matching benchmarks (Sec. 7.3).

For our quantitative evaluation in Fig. 12 (right), Fig. 14, Fig. 16 and Fig. 18, we follow the commonly-used protocol, introduced in [KLF11] by plotting the percentage of correspondences below a certain geodesic distance threshold from the ground truth.

7.1. Datasets

We perform all our experiments on the following datasets.

FAUST [BRLB14]. This dataset contains models of ten different humans in ten poses each. Despite the variability in the body types of said humans, this dataset is typically considered as near-isometric. We remesh the shapes of the dataset to shapes with approximately 5K vertices and use 300 shape pairs following the procedure of the authors of [RPWO18]. Note that the shapes in question are remeshed independently and do not share the same connectivity.

TOSCA [BBK08]. This dataset consists of meshes of humans and animals. Following [RPWO18], we split this dataset into isometric and non-isometric shape pairs. We call the resulting datasets TOSCA isometric (284 shape pairs) and TOSCA non-isometric (95 shape pairs) respectively. The shapes of these datasets are remeshed independently to count around 5K vertices per shape. Once again, the remeshed shapes have distinct connectivity.

[†] Our code is available at <https://github.com/mpanine/DirichletSteklovLandmarkMatching>




Figure 9: Impact of the r_f parameter on the shape matching quality. The mean geodesic error is averaged on the 95 shape pairs of the TOSCA non-isometric dataset (remeshed to 5K vertices). Notice how stable our method remains, even for extreme values of r_f .

SHREC’19 [MMR*19b]. This challenging dataset is composed of human shapes with high variability in pose, vertex count (ranging from 5K to 200K vertices) and topology (some shapes are watertight manifold meshes whereas other have holes and other surface noise sources).

FAUST “Wild” [SACO20]. This dataset is a variant of FAUST in which challenging differences in connectivity are introduced via remeshing. We use the following types of remeshing of the dataset: a uniform isotropic remeshing (*iso*), a remeshing where randomly sampled regions are refined (*dense*), and the remeshing proposed in [GH97] (*qes*). Finally, we consider correspondences *across* the 20 template models of the dataset instead of solely considering the initial template shape as the source shape.

SHREC’20 [DLR*20]. This dataset proposes a collection of 14 animal shapes with a set of landmarks determined by experts. The animal pairs contain parts in correspondence with highly non-isometric deformations. We only consider the correspondences between full shapes for our experiments (test sets 1 to 4).

7.2. Parameter Study

We present here the main results concerning the parameters of our method. Other minor experiments on this topic are presented in App. H (influence of the weights in the energy, qualitative illustration of the impact of landmark placement, near-orthogonality assessment for our basis and study of the effect of basis size).

7.2.1. Radius r_f

The construction of the landmark boundaries Γ_i explained in App. C relies on the user-defined scalar parameter $r_f \in (0, 1)$. In Fig. 9, we study the influence of r_f on the geodesic matching error averaged on the TOSCA non-isometric dataset, with 7 landmark correspondences at their standard locations (see App. I). It demonstrates empirically that this parameter has no significant impact on the matching performance. We therefore set $r_f = 0.5$ in all our other experiments.

7.2.2. Landmark placement

In order to study the influence of landmark placement on our method, we conduct the following experiment on 10 shapes of the TOSCA Isometric dataset (cat category). We consider an increasing number of landmark correspondences, ranging from 3 to 100,




Figure 10: Error summary when increasing the number of landmarks k for different surface sampling strategies. The mean geodesic error on 10 cat shapes of the TOSCA Isometric dataset is reported. “FPS” stands for Farthest Point Sampling.




Figure 11: Error summary when increasing the number of landmarks k for two surface sampling strategies. The mean geodesic error on 95 shape pairs of the TOSCA non-isometric dataset with 3 different seed initializations for each pair is displayed. “FPS” and “std dev.” respectively stand for Farthest Point Sampling and standard deviation.

placed according to four standard surface sampling strategies: (i) random, (ii) euclidean farthest point (iii) geodesic distance farthest point (iv) Poisson disk (as implemented in [Jac*18]). The outcome of these experiments is illustrated in Fig. 10. The farthest point sampling strategies result in the fastest decrease of the error, Poisson disk is slightly slower and random placement is predictably the slowest. This indicates that our method performs best when the extremities of the shapes are prioritized for landmark placement. The landmark placement used in the benchmarks of Sec. 7.3 makes use of this observation (see App. I for details).

To complement the above experiment, we show the variance of our method when initializing two sampling strategies with 3 different seeds in Fig. 11 on the full TOSCA non-isometric dataset.

7.2.3. Remeshing invariance

In order to show that our method remains applicable on shapes with different triangulations, we remesh independently the target pair of each FAUST shape pair and compute the mean geodesic error in Fig. 12 (left). We additionally experiment with the FAUST “Wild” dataset created in [SACO20] to assess invariance to the remeshing proposed by the authors. Fig. 12 (right) and Tab. 1 present the output of this experiment. We observe marginal difference when considering the various remeshing approaches tested, which highlights the insensitivity of the proposed approach to the shape connectivity. Fig. 13 illustrates qualitatively the median transfer obtained on this dataset.




Figure 12: **Left:** remeshing stability when varying the triangle reduction factor r of the target shape. The geodesic error, averaged over 300 test pairs of the FAUST data set, slightly increases when the target mesh becomes coarse (low value of r). **Right:** stability of our method when performing resmeshings on the FAUST dataset (Remeshed to 5K vertices and FAUST “Wild” (see Sec. 7.1)). The geodesic error is measured in mean geodesic distance $\times 100$ after normalizing by the geodesic diameter. The mean values, mean execution times and vertex counts for each remeshing is presented in Tab. 1.




Figure 13: Qualitative illustration of the median map quality obtained with our method on three types of remeshing in the FAUST “Wild” dataset (see Sec. 7.1). Despite the great disparity of the underlying meshes, our method provides smooth transfers.

7.3. Benchmarks

In this section, we describe the competing state-of-the-art methods that we employ (Sec. 7.3.1) and present our main results for shape matching (Sec. 7.3.2).

7.3.1. Setup

We compare our method against three competitors that leverage landmark information to compute correspondences between shapes. The detailed setup for each method, including the landmark placement is provided in App. I. The competing methods are:

	<i>vtx5k</i>	<i>iso</i>	<i>dense</i>	<i>qes</i>
Geo. Err.	13.7	14.3	14.1	14.2
n_v	5001	7117	13399	14002
Exec. t. (s)	7.3	8.35	13.75	14.1

Table 1: Stability of our method when performing resmeshings on the FAUST dataset. The geodesic error (geo. err.) is measured in mean geodesic distance $\times 100$ after normalizing by the geodesic diameter. The corresponding error curves are displayed in Fig. 12 (right). The execution time (exec. t.) is also reported, along with the mean number of vertices for each remeshing type (n_v).




Figure 14: Error summary on the FAUST (**left**) and TOSCA Isometric dataset (**right**). The geodesic error is measured in mean geodesic distance $\times 100$ after normalizing by the geodesic diameter.

Hyperbolic Orbifold Tutte Embeddings (hyperOrb) [AL16] constructs a parameterization of each surface by embedding the points to the hyperbolic plane. The surfaces are cut along the input correspondences, which are *de facto* preserved.

Weighted Averages (WA) [PBDS13] also defines a parameterization of the input surfaces that preserves landmarks exactly: each point at the surface is expressed as a weighted average of its distance to a set of landmarks.

Functional Maps With ZoomOut Refinement (FMap ZO) [MRR*19] computes correspondences between shapes by leveraging a functional basis defined on the source and target shapes. While the method does not allow to retrieve exact correspondence between user-specified landmarks, it constitutes the current state-of-the-art method for isometric shape matching.

7.3.2. Results

In this section, we present our main results on shape matching.

Isometric shape matching. The evaluation on FAUST and TOSCA Isometric are illustrated in Fig. 14, with averaged errors and runtimes displayed in Tab. 2. On the FAUST data set, our approach remains competitive with a mean geodesic error of 1.40×10^{-2} and a mean computation time of 8.83 s. On the TOSCA isometric data set, we obtain a slightly better average geodesic error score than competitors. Qualitatively, our method produces smooth texture transfers on both data sets, as highlighted in Fig. 15.

Non-isometric shape matching. We run an evaluation of our method on the TOSCA non-isometric and the SHREC’20 datasets (Fig. 16). The mean error values and timings are showed in Tab. 3. In this challenging setup, our method has the best results in terms of mean geodesic error, while being the second best in terms of computation time. Fig. 17 presents a qualitative evaluation using a texture transfer on a pair of shapes for each data set.

SHREC’19 benchmark. The quantitative evaluation is reported in Fig. 18, with the associated averaged geodesic errors on the right of the figure. Our method obtains the best mean geodesic error score for this difficult benchmark. In addition, a qualitative evaluation via texture transfer is depicted in Fig. 19. Our method’s strong

Method	Data Set	Av. Geo. Err.	Av. Time (in s.)
FMap ZO	FAUST	1.23×10^{-2}	5.93
	TOSCA Iso.	1.95×10^{-2}	6.27
HyperOrb	FAUST	2.19×10^{-2}	26.8
	TOSCA Iso.	2.10×10^{-2}	10.5
WA	FAUST	4.08×10^{-2}	59.3
	TOSCA Iso.	5.26×10^{-2}	81.0
Ours	FAUST	1.40×10^{-2}	8.83
	TOSCA Iso.	1.90×10^{-2}	11.3

Table 2: Quantitative evaluation results on the remeshed FAUST and TOSCA Isometric (TOSCA Iso.) data sets. The average geodesic error (Av. Geo. Err.) and average execution time (Av. Time) on both data sets are displayed for our method and competing approaches.

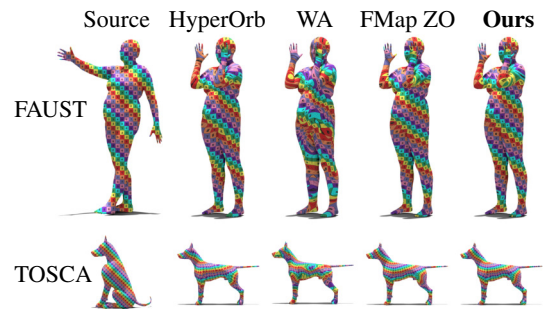


Figure 15: Qualitative evaluation of our method and competing approaches on isometric shapes. The first row corresponds to shapes from the FAUST data set. The bottom row consists of shapes from the TOSCA isometric data set. The shape pair is selected such that the geodesic error of our method is **median** over the dataset. The best and worst cases are illustrated in App. G.




Figure 16: Error summary on the TOSCA non-isometric (**left**) and on the SHREC’20 lores dataset (**right**).

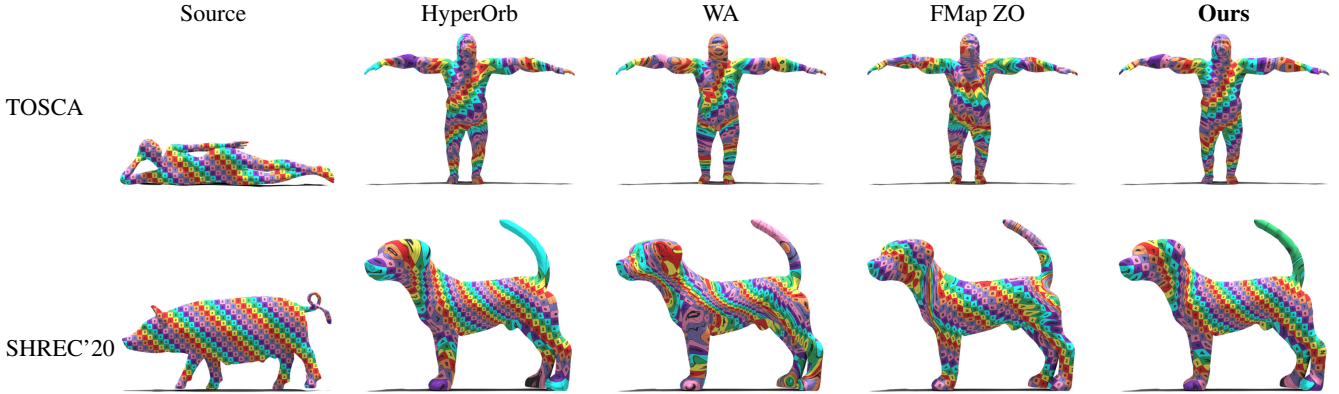


Figure 17: Qualitative evaluation of our method and competing approaches on non-isometric shapes. The first row corresponds to shapes from the TOSCA non-isometric data set. The bottom row consists of shapes from the SHREC’20 lores data set. Each shape pair is selected such that the geodesic error of our method is **median** over the dataset. The best and worst cases are illustrated in App. G.

Method	Data Set	Av. Geo. Err.	Av. Time (in s.)
FMap ZO	TOSCA n-i.	1.10×10^{-1}	7.78
	SHREC’20	7.86×10^{-2}	27.9
HyperOrb	TOSCA n-i.	4.33×10^{-2}	17.8
	SHREC’20	5.78×10^{-2}	270
WA	TOSCA n-i.	6.50×10^{-2}	79.7
	SHREC’20	7.62×10^{-2}	140
Ours	TOSCA n-i.	4.11×10^{-2}	13.5
	SHREC’20	5.09×10^{-2}	63.8

Table 3: Quantitative evaluation results on the TOSCA non-isometric (n-i.) and the SHREC’20 lores (without partial shapes) data sets. The average geodesic error (Av. Geo. Err.) and average execution time (Av. Time) on both data sets are displayed for competing approaches and our method.




Figure 18: Error summary on 165 shapes of the SHREC’19 data set. The average geodesic error (Av. Geo. Err.) is displayed for our method and competing approaches.

performance on this dataset is indicative of its stability and applicability across diverse changes in shape topology, such as the introduction of small holes. This is a general feature of the functional maps methods, which our approach inherits.




Figure 19: Qualitative evaluation of our method and competing approaches on a shape pair from the SHREC’19 data set, selected such that the geodesic error of our method is **median** over the dataset. The best and worst cases are illustrated in App. G.

8. Conclusion, Limitations and Outlook

We have proposed an efficient functional maps-based shape matching approach that promotes conformal maps and exactly preserves landmark correspondences. This was achieved via the introduction of a novel functional basis and an energy promoting bijective conformal maps. The efficiency of our solution comes from an adaptation of the ZoomOut procedure [MRR*19; RMWO21] using our energy and novel basis. The resulting method exhibits state-of-the-art performance on non-isometric benchmark datasets and near state of the art performance on isometric ones.

Recall, however, that our usage of the ZoomOut procedure was not fully principled. Indeed, we needed to make some approximations in order to use Lemma 5, which converts certain optimization problems into nearest-neighbor searches. The quality of our results indicates that our approximations were justified, suggesting that Lemma 5 could likely be rigorously extended to suit our needs. In fact, extending Lemma 5 would be of general interest to the functional maps community, as it would enable the efficient minimization of various other energies.

The construction of our landmark-adapted basis required us to upgrade the landmarks to proper boundaries. We did so by cutting

out small disks centered at the landmarks, resulting in the introduction of landmark circles. The landmark circles offer an intriguing possibility that we have not explored here. Namely, one could augment landmark correspondence to include a user-specified matching of the landmark circles. This could allow for greater semantic or artistic control of the resulting map. Our initialization procedure of Sec. 6.1 can be seen as an automated implementation of a similar idea.

Furthermore, since our present work has demonstrated the fruitfulness of landmark-adapted bases, it is natural to ask whether better performance can be achieved by improving upon basis construction. In particular, we have noted that the Dirichlet-Steklov eigenfunctions have their amplitude intensely concentrated near the landmark circle equipped with the Steklov boundary condition (see Fig. 7). It seems likely that an analogous basis with less concentrated functions could be better suited to describe the behavior of the functional map near the landmarks. Notice that this dovetails with the idea of user-specified landmark circle correspondence, as the user-provided information would have impact further away from the landmarks.

9. Acknowledgements

The authors would like to acknowledge the anonymous reviewers for their helpful feedback and suggestions. Parts of this work were supported by the ERC Starting Grants No. 758800 (EXPROTEA), the ANR AI Chair AIGRETTE, the Swiss National Science Foundation (SNSF) under project number 188577 and the Association Nationale de la Recherche et de la Technologie (ANRT) via the Convention industrielle de formation par la recherche (CIFRE) grant No. 2019/0433. Finally, the authors wish to thank Prof. Alexandre Girouard for his assistance in navigating the literature on the Steklov eigenproblem, Jing Ren and Simone Melzi for providing tools for shape analysis via functional maps and the TOSCA/FAUST/SHREC'19 datasets, Nicholas Sharp for releasing the "FAUST Wild" dataset and Patrick Schmidt for computing a baseline to his work [SCBK20].

References

- [ABK15] AFLALO, Y., BREZIS, H., and KIMMEL, R. "On the optimality of shape and data representation in the spectral domain". *SIAM Journal on Imaging Sciences* 8.2 (2015), 1141–1160 3.
- [ADK16] AFLALO, YONATHAN, DUBROVINA, ANASTASIA, and KIMMEL, RON. "Spectral Generalized Multi-dimensional Scaling". *International Journal of Computer Vision* 118.3 (July 2016), 380–392. ISSN: 1573-1405. DOI: 10.1007/s11263-016-0883-8 2.
- [AL15] AIGERMAN, NOAM and LIPMAN, YARON. "Orbifold Tutte embeddings". *ACM Transactions on Graphics* 34.6 (2015), 190–1 1, 2.
- [AL16] AIGERMAN, NOAM and LIPMAN, YARON. "Hyperbolic Orbifold Tutte Embeddings". *ACM Transactions on Graphics* 35.6 (Nov. 2016), 217:1–217:14 2, 13.
- [APL14] AIGERMAN, NOAM, PORANNE, ROI, and LIPMAN, YARON. "Lifted bijections for low distortion surface mappings". *ACM Transactions on Graphics (TOG)* 33.4 (2014), 1–12 2.
- [APL15] AIGERMAN, NOAM, PORANNE, ROI, and LIPMAN, YARON. "Seamless surface mappings". *ACM Transactions on Graphics (TOG)* 34.4 (2015), 1–13 2.
- [ASC11] AUBRY, MATHIEU, SCHLICKEWEI, ULRICH, and CREMERS, DANIEL. "The wave kernel signature: A quantum mechanical approach to shape analysis". *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on*. IEEE, 2011, 1626–1633 2, 9.
- [BBK08] BRONSTEIN, ALEXANDER M, BRONSTEIN, MICHAEL M, and KIMMEL, RON. *Numerical geometry of non-rigid shapes*. Springer Science & Business Media, 2008 11.
- [BDK17] BURGHARD, OLIVER, DIECKMANN, ALEXANDER, and KLEIN, REINHARD. "Embedding shapes with Green's functions for global shape matching". *Computers & Graphics* 68 (2017), 1–10 2.
- [BHK20] BUNGE, ASTRID, HERHOLZ, PHILIPP, KAZHDAN, MISHA, and BOTSCH, MARIO. "Polygon Laplacian Made Simple". *Computer Graphics Forum* 39.2 (2020), 303–313. DOI: <https://doi.org/10.1111/cgfm.13931> 17.
- [Boo89] BOOKSTEIN, FRED L. "Principal warps: Thin-plate splines and the decomposition of deformations". *IEEE Transactions on pattern analysis and machine intelligence* 11.6 (1989), 567–585 2.
- [BRL14] BOGO, FEDERICA, ROMERO, JAVIER, LOPER, MATTHEW, and BLACK, MICHAEL J. "FAUST: Dataset and Evaluation for 3D Mesh Registration". *2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, June 2014, 3794–3801. DOI: 10.1109/CVPR.2014.491 1, 11.
- [Cha84] CHAVEL, ISAAC. *Eigenvalues in Riemannian geometry*. Orlando, FL: Academic Press, 1984 4.
- [CR00] CHUI, HAILI and RANGARAJAN, ANAND. "A new algorithm for non-rigid point matching". *Proceedings IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2000 (Cat. No. PR00662)*. Vol. 2. IEEE, 2000, 44–51 2.
- [CSB*17] CORMAN, ETIENNE, SOLOMON, JUSTIN, BEN-CHEN, MIRELA, et al. "Functional characterization of intrinsic and extrinsic geometry". *ACM Transactions on Graphics (TOG)* 36.2 (2017), 1–17 5.
- [CSBK18] CHOUKROUN, YONI, SHTERN, ALON, BRONSTEIN, ALEX, and KIMMEL, RON. "Hamiltonian operator for spectral shape analysis". *IEEE transactions on visualization and computer graphics* 26.2 (2018), 1320–1331 3.
- [CWW13] CRANE, KEENAN, WEISCHEDER, CLARISSE, and WARDETZKY, MAX. "Geodesics in heat: A new approach to computing distance based on heat flow". *ACM Transactions on Graphics (TOG)* 32.5 (2013), 1–11 9.
- [DLR*20] DYKE, ROBERTO M, LAI, YU-KUN, ROSIN, PAUL L, et al. "SHREC'20: Shape correspondence with non-isometric deformations". *Computers & Graphics* 92 (2020), 28–43 11.
- [EB17] EZUZ, DANIELLE and BEN-CHEN, MIRELA. "Deblurring and Denoising of Maps between Shapes". *Computer Graphics Forum* 36.5 (2017), 165–174 2, 9, 10.
- [ERGB16] EYNARD, DAVIDE, RODOLA, EMANUELE, GLASHOFF, KLAUS, and BRONSTEIN, MICHAEL M. "Coupled functional maps". *2016 Fourth International Conference on 3D Vision (3DV)*. IEEE, 2016, 399–407 2, 8.
- [ESB19] EZUZ, DANIELLE, SOLOMON, JUSTIN, and BEN-CHEN, MIRELA. "Reversible Harmonic Maps between Discrete Surfaces". *ACM Trans. Graph.* 38.2 (2019). ISSN: 0730-0301. DOI: 10.1145/3202660 3, 20.
- [GBKS18] GEHRE, ANNE, BRONSTEIN, MICHAEL, KOBELT, LEIF, and SOLOMON, JUSTIN. "Interactive curve constrained functional maps". *Computer Graphics Forum* 37.5 (2018), 1–12 2.
- [GH97] GARLAND, MICHAEL and HECKBERT, PAUL S. "Surface simplification using quadric error metrics". *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co. New York, NY, USA, 1997, 209–216 11.
- [HRWO20] HUANG, RUQI, REN, JING, WONKA, PETER, and OVSIANIKOV, MAK. "Consistent ZoomOut: Efficient Spectral Map Synchronization". *Computer Graphics Forum*. Vol. 39. 5. Wiley Online Library, 2020, 265–278 2.

- [Jac*18] JACOBSON, ALEC et al. *gptoolbox: Geometry Processing Toolbox*. <http://github.com/alecjacobson/gptoolbox>. 2018 12.
- [Jos08] JOST, JÜRGEN. *Riemannian Geometry and Geometric Analysis*. fifth. Springer, 2008 20.
- [KBB*13] KOVNATSKY, A., BRONSTEIN, M. M., BRONSTEIN, A. M., et al. “Coupled quasi-harmonic bases”. *Computer Graphics Forum* 32.2pt4 (2013), 439–448. DOI: <https://doi.org/10.1111/cgf.12064> 2.
- [KGB16] KOVNATSKY, ARTIOM, GLASHOFF, KLAUS, and BRONSTEIN, MICHAEL M. “MADMM: A Generic Algorithm for Non-smooth Optimization on Manifolds”. *Computer Vision – ECCV 2016*. Ed. by LEIBE, BASTIAN, MATAS, JIRI, SEBE, NICU, and WELLING, MAX. Cham: Springer International Publishing, 2016, 680–696. ISBN: 978-3-319-46454-1 3.
- [KLF11] KIM, VLADIMIR G, LIPMAN, YARON, and FUNKHOUSER, THOMAS. “Blended intrinsic maps”. *ACM Transactions on Graphics (TOG)* 30.4 (2011), 79 11.
- [KO19] KLEIMAN, YANIR and OVSJANIKOV, MAKS. “Robust Structure-Based Shape Correspondence”. *Computer Graphics Forum* 38.1 (2019), 7–20. DOI: <https://doi.org/10.1111/cgf.13389> 2.
- [Kob00] KOBBELT, LEIF. “ $\sqrt{3}$ -subdivision”. *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley, 2000, 103–112 18.
- [Lab17] LABRIE, MARC-ANTOINE. “Le théorème spectral pour le problème de Steklov sur un domaine euclidien”. MA thesis. Québec: Université Laval, 2017 4.
- [McG17] MCGUIRE, MORGAN. *Computer Graphics Archive*. July 2017. URL: <https://casual-effects.com/data> 1.
- [MMM*20] MELZI, SIMONE, MARIN, RICCARDO, MUSONI, PIETRO, et al. “Intrinsic/extrinsic embedding for functional remeshing of 3D shapes”. *Computers & Graphics* 88 (2020), 1–12 3.
- [MMR*19a] MELZI, SIMONE, MARIN, RICCARDO, RODOLÀ, EMANUELE, et al. “SHREC 2019: Matching Humans with Different Connectivity”. *Eurographics Workshop on 3D Object Retrieval*. The Eurographics Association, 2019 22.
- [MMR*19b] MELZI, SIMONE, MARIN, RICCARDO, RODOLÀ, EMANUELE, et al. “SHREC 2019: Matching Humans with Different Connectivity”. *Eurographics Workshop on 3D Object Retrieval*. Vol. 7. The Eurographics Association, 2019 11.
- [MMRC18] MARIN, RICCARDO, MELZI, SIMONE, RODOLÀ, EMANUELE, and CASTELLANI, UMBERTO. *FARM: Functional Automatic Registration Method for 3D Human Bodies*. 2018 2.
- [MRCB18] MELZI, S., RODOLÀ, E., CASTELLANI, U., and BRONSTEIN, M. M. “Localized Manifold Harmonics for Spectral Shape Analysis”. *Computer Graphics Forum* 37.6 (2018), 20–34. DOI: <https://doi.org/10.1111/cgf.13309> 3.
- [MRR*19] MELZI, SIMONE, REN, JING, RODOLÀ, EMANUELE, et al. “Zoomout: Spectral upsampling for efficient shape correspondence”. *ACM Transactions on Graphics (TOG)* 38.6 (2019), 155 1–3, 8–10, 13, 14, 23.
- [Neč12] NEČAS, JINDŘICH. *Direct Methods in the Theory of Elliptic Equations*. Berlin: Springer-Verlag, 2012 4, 5.
- [NO17] NOGNENG, DORIAN and OVSJANIKOV, MAKS. “Informative Descriptor Preservation via Commutativity for Shape Matching”. *Computer Graphics Forum* 36.2 (2017), 259–267. DOI: <https://doi.org/10.1111/cgf.13124> 2.
- [NVT*14] NEUMANN, T., VARANASI, K., THEOBALT, C., et al. “Compressed Manifold Modes for Mesh Processing”. *Computer Graphics Forum* 33.5 (2014), 35–44. DOI: <https://doi.org/10.1111/cgf.12429> 3.
- [OBS*12] OVSJANIKOV, MAKS, BEN-CHEN, MIRELA, SOLOMON, JUSTIN, et al. “Functional maps: a flexible representation of maps between shapes”. *ACM Transactions on Graphics (TOG)* 31.4 (2012), 30:1–30:11 2, 9.
- [OCB*17] OVSJANIKOV, MAKS, CORMAN, ETIENNE, BRONSTEIN, MICHAEL, et al. “Computing and Processing Correspondences with Functional Maps”. *ACM SIGGRAPH 2017 Courses*. 2017, 5:1–5:62 2, 3.
- [OLCO13] OZOLIŅŠ, VIDVUDS, LAI, RONGJIE, CAFLISCH, RUSSEL, and OSHER, STANLEY. “Compressed modes for variational problems in mathematics and physics”. *Proc. of the National Academy of Sciences* 110.46 (2013), 18368–18373 3.
- [OMMG10] OVSJANIKOV, MAKS, MÉRIGOT, QUENTIN, MÉMOLI, FACUNDO, and GUIBAS, LEONIDAS. “One Point Isometric Matching with the Heat Kernel”. *Computer Graphics Forum* 29.5 (2010), 1555–1564. DOI: <https://doi.org/10.1111/j.1467-8659.2010.01764.x> 2.
- [Pat18] PATANÈ, G. “Laplacian spectral basis functions”. *Computer-Aided Geometric Design* 65 (2018), 31–47. ISSN: 0167-8396 3, 5.
- [PBDS13] PANOZZO, DANIELE, BARAN, ILYA, DIAMANTI, OLGA, and SORKINE-HORNUNG, OLGA. “Weighted averages on surfaces”. *ACM Transactions on Graphics (TOG)* 32.4 (2013), 1–12 3, 13.
- [PRM*21] PAI, GAUTAM, REN, JING, MELZI, SIMONE, et al. “Fast Sinkhorn Filters: Using matrix scaling for non-rigid shape correspondence with functional maps”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 2021, 384–393 3.
- [PST19] POLTEROVICH, IOSIF, SHER, DAVID A, and TOTH, JOHN A. “Nodal length of Steklov eigenfunctions on real-analytic Riemannian surfaces”. *Journal für die reine und angewandte Mathematik (Crelles Journal)* 2019.754 (2019), 17–47 6.
- [RCB*17] RODOLÀ, EMANUELE, COSMO, LUCA, BRONSTEIN, MICHAEL, et al. “Partial functional correspondence”. *Computer Graphics Forum* 36.1 (2017), 222–236 2.
- [RMOW20] REN, JING, MELZI, SIMONE, OVSJANIKOV, MAKS, and WONKA, PETER. “MapTree: Recovering Multiple Solutions in the Space of Maps”. *ACM Trans. Graph.* 39.6 (Nov. 2020) 2.
- [RMWO21] REN, JING, MELZI, SIMONE, WONKA, PETER, and OVSJANIKOV, MAKS. “Discrete Optimization for Shape Matching”. *Computer Graphics Forum*. Vol. 40. 5. Wiley Online Library, 2021, 81–96 8–10, 14.
- [ROA*13] RUSTAMOV, RAIF M, OVSJANIKOV, MAKS, AZENCOT, OMRI, et al. “Map-based exploration of intrinsic shape differences and variability”. *ACM Transactions on Graphics (TOG)* 32.4 (2013), 1–12 5, 7.
- [RPWO18] REN, JING, POULENARD, ADRIEN, WONKA, PETER, and OVSJANIKOV, MAKS. “Continuous and orientation-preserving correspondences via functional maps”. *ACM Transactions on Graphics (TOG)* 37.6 (2018), 1–16 2, 9, 11.
- [SACO20] SHARP, NICHOLAS, ATTAIKI, SOUHAIB, CRANE, KEENAN, and OVSJANIKOV, MAKS. “Diffusion is all you need for learning on surfaces”. *arXiv preprint arXiv:2012.00888* (2020) 11, 12.
- [Sah20] SAHILIOĞLU, YUSUF. “Recent advances in shape correspondence”. *The Visual Computer* 36.8 (2020), 1705–1721 2.
- [SCBK20] SCHMIDT, PATRICK, CAMPEN, MARCEL, BORN, JANIS, and KOBBELT, LEIF. “Inter-surface maps via constant-curvature metrics”. *ACM Transactions on Graphics (TOG)* 39.4 (2020), 119–1 3, 15, 21.
- [SOG09] SUN, JIAN, OVSJANIKOV, MAKS, and GUIBAS, LEONIDAS. “A concise and provably informative multi-scale signature based on heat diffusion”. *Computer graphics forum* 28.5 (2009), 1383–1392 2, 9.
- [SP04] SUMNER, ROBERT W and POPOVIĆ, JOVAN. “Deformation transfer for triangle meshes”. *ACM Transactions on Graphics (TOG)* 23.3 (2004), 399–405 2.
- [vKZHC11] Van KAICK, OLIVER, ZHANG, HAO, HAMARNEH, GHASAN, and COHEN-OR, DANIEL. “A Survey on Shape Correspondence”. *Computer Graphics Forum* 30.6 (2011), 1681–1707. DOI: <https://doi.org/10.1111/j.1467-8659.2011.01884.x> 2.

- [WBPS18] WANG, YU, BEN-CHEN, MIRELA, POLTEROVICH, IOSIF, and SOLOMON, JUSTIN. “Steklov Spectral Geometry for Extrinsic Shape Analysis”. *ACM Transactions on Graphics (TOG)* 38.1 (2018), 1–21 [3](#), [5](#).
- [Wra21] WRAP. *RUSSIAN3DSCANNER: Wrap 2020.12.2*. Oct. 2021. URL: <https://www.russian3dscanner.com/> 1.
- [WZ14] WEBER, OFIR and ZORIN, DENIS. “Locally injective parametrization with arbitrary fixed boundaries”. *ACM Transactions on Graphics (TOG)* 33.4 (2014), 1–12 [2](#).
- [XLZ21] XIANG, RUI, LAI, RONGJIE, and ZHAO, HONGKAI. “A Dual Iterative Refinement Method for Non-Rigid Shape Matching”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2021, 15930–15939 [3](#).
- [ZRKS05] ZAYER, RHALEB, RÖSSL, CHRISTIAN, KARNI, ZACHI, and SEIDEL, HANS-PETER. “Harmonic Guidance for Surface Deformation”. *Computer Graphics Forum* 24.3 (2005), 601–609. DOI: <https://doi.org/10.1111/j.1467-8659.2005.00885.x> [3](#), [5](#).

Appendix A: Weak Form of the Dirichlet-Steklov Eigenproblem

In this appendix, we derive the weak form of the Dirichlet-Steklov eigenproblem (Eq. (3)), in which it becomes very similar to the weak form of the more familiar Laplacian eigenproblem. For sufficiently smooth functions f and u , Stokes’ theorem implies that

$$\int_{\mathcal{M}} f(\Delta u) d\mathcal{M} = \int_{\mathcal{M}} \nabla f \cdot \nabla u d\mathcal{M} - \int_{\partial\mathcal{M}} f(\partial_n u) d(\partial\mathcal{M}). \quad (23)$$

Applying this to a u_i satisfying the Dirichlet-Steklov eigenproblem (Eq. (2)) and a smooth test function f vanishing on \mathcal{D} yields

$$\begin{aligned} \int_{\mathcal{M}} \nabla f \cdot \nabla u_i d\mathcal{M} &= \int_{\mathcal{M}} f(\Delta u_i) d\mathcal{M} \\ &+ \int_{\partial\mathcal{M}} f(\partial_n u_i) d(\partial\mathcal{M}) \\ &= \int_{\mathcal{D}} f(\partial_n u_i) d(\partial\mathcal{M}) \\ &+ \int_S f(\partial_n u_i) d(\partial\mathcal{M}), \end{aligned} \quad (24)$$

where the first cancellation arises from the harmonicity of u_i and the second one from f vanishing on \mathcal{D} . Finally, using the third line of Eq. (2) results in the weak form of the Dirichlet-Steklov problem:

$$\int_{\mathcal{M}} \nabla f \cdot \nabla u_i d\mathcal{M} = \sigma_i \int_S f u_i d(\partial\mathcal{M}). \quad (25)$$

This can be readily discretized on triangle meshes, as discussed in App. B.

Appendix B: Discretization of the Eigenproblems

In this appendix, we briefly discuss the discretization on triangle meshes of the eigenproblems used in our approach.

Discretization of the Dirichlet Laplacian eigenproblem We begin with the familiar Dirichlet Laplacian eigenproblem (Eq. (5)). We discretize this problem using the well-known cotangent scheme (piecewise-linear finite elements). The problem then becomes

$$\begin{aligned} W^{\mathcal{M}} \psi_i &= \lambda_i A^{\mathcal{M}} \psi_i, \\ \psi_i|_{\partial\mathcal{M}} &= 0, \end{aligned} \quad (26)$$

where $W^{\mathcal{M}}$ denotes the so-called cotangent Laplacian and $A^{\mathcal{M}}$ denotes the lumped mass matrix. See [BHKB20], among many others, for a definition of these objects.

Discretization of the Dirichlet-Steklov eigenproblem We use piecewise linear finite elements to discretize the weak form of the Dirichlet-Steklov eigenproblem (Eq. (25)). The left-hand side of the expression becomes the familiar cotangent Laplacian, denoted by $W^{\mathcal{M}}$. The discretization of the integral on the right-hand side requires a mass matrix defined *strictly on the boundary*. Similarly to the mass matrix used in the Laplacian eigenproblem, it can be discretized either according to a piecewise-linear finite element scheme, or as a lumped mass matrix. Regardless of the chosen discretization, we call this mass matrix $S^{\mathcal{M}}$. Note that $S^{\mathcal{M}}$ is of the same size as $W^{\mathcal{M}}$.

We begin by the lumped discretization. The boundary is one-dimensional. Thus, a vertex $p \in \partial\mathcal{M}$, has (at most) two neighbors that are also in $\partial\mathcal{M}$, which we denote $p-1$ and $p+1$. The length of the edges $(p-1, p)$ and $(p, p+1)$ are denoted r_{p-1} and r_{p+1} , respectively. The lumped Steklov mass matrix is given by

$$S_{pq}^{\mathcal{M}} = \begin{cases} \frac{1}{2}(r_{p-1} + r_{p+1}) & , p = q \text{ and } p, q \in \partial\mathcal{M} \\ 0 & , \text{elsewhere.} \end{cases} \quad (27)$$

The non-lumped mass matrix is computed from a piecewise linear finite element discretization on the boundary. This discretization corresponds to the restriction of the piecewise linear finite elements of the mesh to the boundary edges. Whenever vertices p and q are distinct endpoints of the same edge, we write $p \sim q$. The length of the edge connecting p and q is denoted r_{pq} . After a straightforward computation which we omit, the non-lumped Steklov mass matrix is given by

$$S_{pq}^{\mathcal{M}} = \begin{cases} \frac{1}{3}(r_{p-1} + r_{p+1}) & , p = q \text{ and } p, q \in \partial\mathcal{M} \\ \frac{1}{6}r_{pq} & , p \sim q \text{ and } p, q \in \partial\mathcal{M} \\ 0 & , \text{elsewhere.} \end{cases} \quad (28)$$

In sum, no matter the version of $S^{\mathcal{M}}$ chosen, the discretization of the Dirichlet-Steklov problem becomes

$$\begin{aligned} W^{\mathcal{M}} u_i &= \sigma_i S^{\mathcal{M}} u_i, \\ u_i|_{\mathcal{D}} &= 0, \end{aligned} \quad (29)$$

which is quite similar to the more familiar Laplacian eigenvalue problem with Dirichlet boundary conditions (Eq. (26)).

A Word of Warning As a final note on the discretization of the considered eigenproblems, we would like to warn the reader of a small issue one may encounter when numerically solving them. Recall that we want the Dirichlet–Steklov eigenfunctions to be normalized with respect to the boundary mass matrix S^M . Solvers for generalized eigenvalue problems, such as Matlab’s `eigs` routine, which we use in our implementation, will typically do so automatically. However, according to our observations, sometimes this automated process will not happen. This seems to be related to the fact that S^M is a positive semi-definite matrix rather than a positive definite one. Thus, one needs to explicitly normalize the solutions with respect to S^M . In fact, we suggest explicitly normalizing even the Laplacian eigenfunctions, despite the fact that there the mass matrix A^M is positive definite on (good quality) triangle meshes. Indeed, A^M can fail to be positive-definite on pathological inputs. Consider for instance an otherwise good mesh with an isolated vertex belonging to no triangle. Functions vanishing everywhere except on said vertex have norm 0 with respect to A^M , despite being nonzero.

Appendix C: Boundary Circles on Triangle Meshes

In Sec. 4.4, small disks centered at the landmarks are removed in order to create new boundaries for the shapes under study. Here, we describe in detail how this is achieved on triangle meshes. Crucially, we do not want to unduly disturb the geometry of the shapes. In order to achieve this we construct the new boundaries entirely within the triangles adjacent to the landmarks.

Let’s say that we are constructing the boundary circle for the landmark γ_i . We begin by selecting the radius r_i of the disk to be removed. This is done by finding the length s_i of the shortest edge connected to γ_i . The minimum is taken over both shapes, which are scaled to be of identical surface area and thus of comparable size. Then, we set $r_i = r_f \cdot s_i$, where $r_f \in (0, 1)$ is a user-set parameter. The (surprisingly low) impact of this parameter is studied in Sec. 7.2.1.

We are now ready to construct the boundary Γ_i . This process is best understood by looking at its illustration in Fig. 20. First, we split each triangle adjacent to the landmark into n_s wedges of equal angle, which introduces $n_s - 1$ new vertices at the opposite edge of the original triangle, as well as edges connecting them to the landmark. Then, we introduce $n_s + 1$ new vertices situated on the new edges at a distance r_i away from the landmark γ_i . We then connect these vertices in a way that creates an approximation of a sector of a disk of radius r_i . Doing so produces n_s quadrilaterals in the part of the original triangle far from the landmark. We split those quadrilaterals into triangles along their diagonals. This concludes the refinement of the triangles adjacent to the landmark. It remains to refine the triangles adjacent to them across the edges opposite to the landmark. There, the common edges between the triangles contains $n_s - 1$ new vertices. On each triangle, we connect these new vertices to the original vertex not on the common edge. This concludes the refinement process. Note that all of the new triangles are contained within the original ones. An example of a mesh




Figure 20: Illustration of the creation of a landmark boundary. The landmark position is indicated by a green dot. The triangles composing the landmark disk are shown in light red. The boundary circle is highlighted as a red line. Note that a gap of connectivity appears when creating the boundary around the landmark. This gap is closed when the process finishes producing the boundary.

with landmark circles constructed in this manner is shown in Fig. 5.

The construction of the boundaries associated to different landmarks is done sequentially over the landmarks. This requires some additional care if the landmarks are placed too close to each other. Indeed, during the construction of Γ_i , new faces are created in what was originally the 2–ring neighborhood of the landmark γ_i . Thus, if a different landmark γ_j is closer than 4 rings away from γ_i , there will be overlap between the newly created mesh faces. The resulting mesh will then be dependent upon the order in which the boundary circles Γ_i and Γ_j are created. In the present paper, we avoid this issue by disallowing such landmark placement. If such landmark placement becomes necessary in a given application, we suggest locally refining the mesh via, say, $\sqrt{3}$ –subdivision [Kob00] such that the landmarks are no longer closer than 4 triangle rings from one another. We do not pursue this here.

Appendix D: Proof of Lemma 1 and Discussion on its Meaning

Lemma 1. *The function space $W(\mathcal{M})$ admits the following decomposition:*

$$W(\mathcal{M}) = \mathcal{G}(\mathcal{M}) \oplus \left(\bigoplus_{j=1}^k \mathcal{H}_j(\mathcal{M}) \right), \quad (30)$$

where \oplus denotes direct sums and \bigoplus denotes orthogonal direct sums.

Proof. Recall that, by construction, $W(\mathcal{M})$ is the completion of smooth functions modulo constants with respect to the Dirichlet form. Thus, we begin our analysis on smooth functions.

Let u be smooth and $W(\mathcal{M})$ -orthogonal to all of the Dirichlet–Laplacian eigenfunctions $\{\psi_i\}_{i=1}^{\infty}$. Then, by Stokes’ theorem,

$$\begin{aligned} 0 &= \int_{\mathcal{M}} \nabla \psi_i \cdot \nabla u \, d\mathcal{M} \\ &= \int_{\mathcal{M}} \psi_i (\Delta u) \, d\mathcal{M} + \int_{\partial \mathcal{M}} \psi_i (\partial_{\pi} u) \, d(\partial \mathcal{M}), \end{aligned} \quad (31)$$

where the cancellation results from ψ_i vanishing at the boundary. Since the $\{\psi_i\}_{i=1}^\infty$ form an orthogonal basis for $L_2(\mathcal{M})$, this implies that $\Delta u = 0$. Thus, smooth functions can be $W(\mathcal{M})$ -orthogonally decomposed into a part that lies in $\mathcal{G}(\mathcal{M})$ (the closed span of $\{\psi_i\}_{i=1}^\infty$) and a harmonic part.

Recall that each $\mathcal{H}_j(\mathcal{M})$ spans harmonic functions that vanish on all landmark boundaries, but the j^{th} one. Since harmonic functions are uniquely determined by their values at the boundaries, the harmonic part of u can be naturally expressed as an element in $\oplus_{j=1}^k \mathcal{H}_j(\mathcal{M})$.

Since $W(\mathcal{M})$ is complete by construction and smooth functions are dense in $W(\mathcal{M})$, the desired result is achieved by taking the closure of the subspaces. \square

Notice that in the above lemma, the subspaces $\mathcal{H}_j(\mathcal{M})$ are not marked as $W(\mathcal{M})$ -orthogonal. Indeed, by Stokes' theorem,

$$\begin{aligned} \langle u_i^{(p)}, u_l^{(q)} \rangle_{W(\mathcal{M})} &= \int_{\mathcal{M}} \nabla u_i^{(p)} \cdot \nabla u_l^{(q)} d\mathcal{M} \\ &= \int_{\mathcal{M}} u_i^{(p)} (\Delta u_l^{(q)}) d\mathcal{M} \\ &\quad + \int_{\partial\mathcal{M}} u_i^{(p)} (\partial_n u_l^{(q)}) d(\partial\mathcal{M}) \quad (32) \\ &= \sum_{\mu=1}^k \int_{\Gamma_\mu} u_i^{(p)} (\partial_n u_l^{(q)}) d(\partial\mathcal{M}) \\ &= \int_{\Gamma_p} u_i^{(p)} (\partial_n u_l^{(q)}) d(\partial\mathcal{M}) \end{aligned}$$

The above expression yields different results depending on whether p and q coincide or not. We begin by considering $p = q$.

$$\begin{aligned} \langle u_i^{(p)}, u_l^{(p)} \rangle_{W(\mathcal{M})} &= \int_{\Gamma_p} u_i^{(p)} (\partial_n u_l^{(p)}) d(\partial\mathcal{M}) \\ &= \sigma_l^{(p)} \int_{\Gamma_p} u_i^{(p)} u_l^{(p)} d(\partial\mathcal{M}) \quad (33) \\ &= \sigma_l^{(p)} \delta_{il} . \end{aligned}$$

Here δ_{il} denotes the Kronecker delta. Thus, for every p , the Dirichlet-Steklov basis $\{u_i^{(p)}\}_{i=1}^\infty$ is composed of $W(\mathcal{M})$ -orthogonal functions. Notice that said eigenfunctions can be $W(\mathcal{M})$ -normalized by dividing them by the square root of the corresponding eigenvalue.

Now, consider $p \neq q$. In that case, Eq. (32) can no longer be evaluated by substituting the Dirichlet-Steklov eigenvalue for the normal derivative, as it is evaluated on the wrong boundary component. Moreover, the normal derivative $\partial_n u_l^{(q)}$ has no reason to vanish on Γ_p , which implies that the subspaces spanned by $\{u_i^{(p)}\}_{i=1}^\infty$ and $\{u_i^{(q)}\}_{i=1}^\infty$ are not $W(\mathcal{M})$ -orthogonal.

Appendix E: Proof of Lemma 2

Lemma 2 (Structure of $F_{\mathcal{M}\mathcal{N}}$). *Let $F_{\mathcal{M}\mathcal{N}} : W(\mathcal{M}) \rightarrow W(\mathcal{N})$ be the pullback of a conformal diffeomorphism that preserves the landmark circles. Then, $F_{\mathcal{M}\mathcal{N}}$ maps*

1. $\mathcal{G}(\mathcal{M})$ to $\mathcal{G}(\mathcal{N})$,
2. $\mathcal{H}_j(\mathcal{M})$ to $\mathcal{H}_j(\mathcal{N})$ for all j .

Proof. Since $\phi : \mathcal{N} \rightarrow \mathcal{M}$ is a diffeomorphism, we can express everything on the surface \mathcal{N} . Thus, instead of thinking of \mathcal{M} as separate manifold, we treat \mathcal{N} as being equipped with two Riemannian metrics: its original metric $g^\mathcal{N}$ and the pullback metric $g^\mathcal{M}$. In this representation, the pullback acts as the identity. In particular, this means that $F_{\mathcal{M}\mathcal{N}} : W(\mathcal{M}) \rightarrow W(\mathcal{N})$ is a bounded operator.

Since ϕ is conformal, there exists a positive function ω such that $g^\mathcal{M} = \omega g^\mathcal{N}$ and $\Delta^\mathcal{M} = (1/\omega)\Delta^\mathcal{N}$. Let u be a harmonic function on \mathcal{M} . Then, $\Delta^\mathcal{N} F_{\mathcal{M}\mathcal{N}} u = \omega \Delta^\mathcal{M} u = 0$. Thus, $F_{\mathcal{M}\mathcal{N}}$ maps harmonic functions to harmonic functions. Furthermore, since $F_{\mathcal{M}\mathcal{N}}$ is the pullback of a map that preserves the landmark circles, it maps smooth functions that vanish on all landmark circles of \mathcal{M} but $\Gamma_j^\mathcal{M}$ to smooth functions that vanish on all landmark circles of \mathcal{N} but $\Gamma_j^\mathcal{N}$ and so for any fixed j . Statement 2. then follows from the completeness of $W(\mathcal{M})$ and $W(\mathcal{N})$ and the boundedness of $F_{\mathcal{M}\mathcal{N}}$ by taking the closure of the relevant subspaces.

Now consider $f \in \mathcal{G}(\mathcal{M})$. By Lemma 1, for all harmonic u ,

$$\langle u, f \rangle_{W(\mathcal{M})} = 0 . \quad (34)$$

By Theorem 3, the conformality of ϕ allows us to replace the inner product on $W(\mathcal{M})$ with that on $W(\mathcal{N})$ up to the introduction of two functional maps:

$$\langle F_{\mathcal{M}\mathcal{N}} u, F_{\mathcal{M}\mathcal{N}} f \rangle_{W(\mathcal{N})} = 0 . \quad (35)$$

Since $F_{\mathcal{M}\mathcal{N}}$ maps harmonic functions to harmonic functions and is invertible, $F_{\mathcal{M}\mathcal{N}} u$ can be any desired harmonic function of \mathcal{N} . Thus, $F_{\mathcal{M}\mathcal{N}} f$ is $W(\mathcal{N})$ -orthogonal to harmonic functions of \mathcal{N} , that is $F_{\mathcal{M}\mathcal{N}} f \in \mathcal{G}(\mathcal{N})$. This concludes the proof of statement 1. \square

Appendix F: Definition of the Dirichlet Energy

Consider a smooth map $\phi : \mathcal{M} \rightarrow \mathcal{N}$ between two smooth Riemannian manifolds. The Dirichlet energy of the map is

$$D(\phi) = \frac{1}{2} \int_{\mathcal{M}} \|d\phi\|^2 d\mathcal{M}, \quad (36)$$

where $d\phi$ is the differential of ϕ . Informally speaking, the Dirichlet energy measures the oscillation of the map ϕ . The larger the energy, the more oscillatory the map. Maps minimizing the Dirichlet energy are known as harmonic maps. Such maps are a simultaneous generalization of geodesics and harmonic functions.

See [Jos08] for the relevant theory.

In the discrete setting, we use the same method as in [ESB19] to compute the Dirichlet energy. Namely, the expression becomes

$$D(\varphi) = \frac{1}{4} \sum_{(u,v) \in \mathcal{E}(\mathcal{M})} w_{uv}^{\mathcal{M}} D_{\mathcal{N}}^2(\varphi(u), \varphi(v)), \quad (37)$$

where $\mathcal{E}(\mathcal{M})$ denotes the edges of the mesh \mathcal{M} , $w_{uv}^{\mathcal{M}}$ denotes the cotangent weight of the edge (u,v) and $D_{\mathcal{N}}^2(\cdot, \cdot)$ is the matrix of square geodesic distances on \mathcal{N} .

Appendix G: Additional Experiments

Analysis of Alternative Initialization Methods

The iterative optimization procedure detailed in Sec. 6.2 requires as an input an initial guess of the functional map. In Sec. 6.1, we thus introduce an initialization procedure for this initial guess based on the landmark correspondence and the normal derivatives of certain landmark-dependent harmonic functions. In this section we compare this approach to two alternatives.

For the purposes of this discussion, the approach of Sec. 6.1 shall be referred to as the “normal derivatives” method. The two alternatives described below will be termed “trivial” and the “conformal energy”, for reasons that should soon become apparent.

The landmark circles can be seen as lists of vertices ordered counter-clockwise as seen from outside the shape. The choice of the first element of this list carries no particular meaning and is left to the whims of the indexing of the faces of the mesh. Thus, the first elements of two corresponding boundary circles need not match. The “trivial” approach consists in assuming that the first elements of the boundary circles do indeed match. This correspondence is then proportionally extended to the rest of the landmark circle.

The “conformal energy” approach stems from the observation that mapping the landmark circles $\Gamma_i^{\mathcal{N}} \rightarrow \Gamma_i^{\mathcal{M}}$ induces a restricted functional map $\mathcal{H}_i(\mathcal{M}) \rightarrow \mathcal{H}_i(\mathcal{N})$. The conformal term of the energy (Eq. (9)) can be easily evaluated on these subspaces. The “conformal energy” approach consists in choosing the shifts $\{\alpha_i\}_{i=1}^k$ (see Sec. 6.1) such that they minimize the conformal energy of the resulting $\mathcal{H}_i(\mathcal{M}) \rightarrow \mathcal{H}_i(\mathcal{N})$ map.

Fig. 21 (left) depicts the performance of the three initializations in terms of geodesic error on the SHREC’20 dataset (lores), using 7 landmarks. Tab. 4 provides quantitative evaluations for the same experiment in terms of averaged geodesic error and Dirichlet energy. The “normal derivatives” approach slightly outperforms the other two on all metrics, which is why it is the one used in the main text.




Figure 21: **Left:** comparison of initializations for our method, where “Norm. De.” and “Conf. En.” respectively stand for “Normal Derivative” and “Conformal Energy”. **Right:** comparison of the “fast” and “principled” energy formulations of our method. Both experiments are performed on the SHREC’20 lores dataset (partial shapes excluded).

Method	Av. Geo. Err.	Dir. E.	Av. Time (in s.)
Trivial	6.36×10^{-2}	16.8	41.4
Conf. En.	6.36×10^{-2}	16.7	53.2
Norm. De.	6.26×10^{-2}	16.2	40.4

Table 4: Quantitative evaluation results on the SHREC’20 lores (without partial shapes) data sets. The average geodesic error (Av. Geo. Err.), the Dirichlet energy (Dir. E.) and average execution time (Av. Time) on both data sets are displayed for the three initialization methods that we tried: Trivial, Conformal Energy (“Conf. En.”) and Normal Derivatives (“Norm. De.”). Normal Derivatives is the method used in the rest of the paper.

Comparison of the “Principled” and “Fast” Energy Optimization

At the end of Sec. 6.2, we introduced an unprincipled way to accelerate the nearest neighbor search used in the solution of our problem. In this section, we quantitatively compare this “fast” method to the “principled” one on the SHREC’20 data set (partial shapes excluded). The output of this evaluation is displayed in Fig. 21 (right) and Tab. 5. While very similar in terms of matching performance, the “fast” method is more than three times faster to compute. We therefore employ it instead of the “principled” approach. Note that the more than threefold speedup is consistent with the fact that the matrices used in the “fast” method are three times smaller.

Method	Av. Geo. Err.	Av. Time (in s.)
Principled	4.96×10^{-2}	184
Fast	5.13×10^{-2}	48.7

Table 5: Average geodesic error (Av. Geo. Err.) and average execution time (Av. Time) associated to the comparison of the “principled” and “fast” computation methods.

Complementary benchmark on SHREC'20 lores

As a complement to our main evaluation on SHREC'20 lores, we conducted an evaluation using only 8 pairs from the initial benchmark to compare against the method proposed in [SCBK20] (InterSurf). InterSurf, WA, HyperOrb FMap ZO and our approach obtain a geodesic error (scaled by a factor $\times 100$) of respectively 11.9, 5.41, 5.99, 8.69 and **5.2**. The restricted number of shapes on which we evaluate is due to the fact that InterSurf does not handle shapes with complex topologies well. In particular, the method assumes that the meshes are watertight and share the same genus, in strong contrast to our approach that does not make such assumptions. However, we note that this method was not primarily designed for shape matching.

Additional Qualitative Evaluations

We provide additional qualitative evaluations on isometric and non-isometric shape pairs in order to show best- and worst-case shape matching scenarios for our method.

For isometric shapes, the best pairs are depicted in Fig. 22 and the worst pairs in Fig 23.

For non-isometric shapes, the best pairs are illustrated in Fig. 24 and the worst pairs in Fig. 25.

Finally, in Fig. 26, we show the best and worst pairs for the SHREC'19 benchmark.

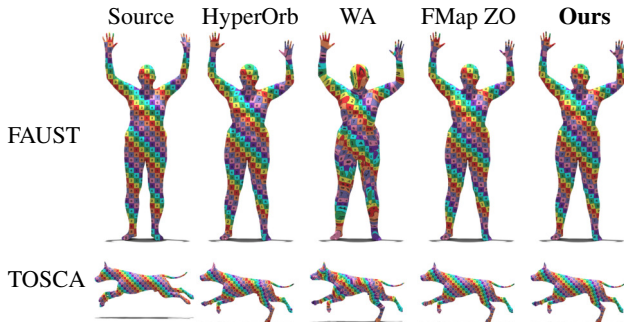


Figure 22: Qualitative evaluation of our method and competitors on **isometric shapes** from the FAUST dataset (top row) and the TOSCA isometric dataset (bottom row). The shape pair is selected such that the geodesic error of our method is **the best** over the dataset.

Appendix H: Additional Parameter Study

Study of the Weights in the Energy

We define three weights to compute a point-to-point map between two shapes based on the energy (Eq. (14)): the conformal, the properness and the invertibility weights, denoted respectively a_C , a_P and a_I . Since we normalize the weights, their absolute value is unimportant.

To study how their relative value influences the quality of the output map we conduct a dedicated experiment on the SHREC'20




Figure 23: Qualitative evaluation of our method and competitors on **isometric shapes** from the FAUST dataset (top row) and the TOSCA isometric dataset (bottom row). The shape pair is selected such that the geodesic error of our method is **the worst** over the dataset.

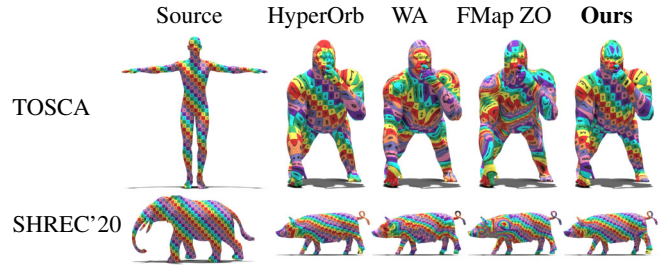


Figure 24: Qualitative evaluation of our method and competitors on **non-isometric shapes**. The first row corresponds to shapes from the TOSCA non-isometric data set. The bottom row consists of shapes from the SHREC'20 lores data set. The shape pair is selected such that the geodesic error of our method is **the best** over the dataset.

dataset, with shapes remeshed to count 1K vertices and excluding partial shapes. 8 landmarks in ground-truth correspondence are placed on each shape, in the locations described in App. I. For each set of weight values, the geodesic error and the Dirichlet energy (see App. F), averaged over all shape pairs (in both directions) in the dataset, are computed.

We first fix the conformality weight to 1 and vary the two remaining weights within a range of energy values in Fig. 27 left (geodesic error) and Fig. 28 left (Dirichlet energy). Second, we let one weight vary and fix the two remaining values either to 0 or to 1, as illustrated in Fig. 27 right (geodesic error) and Fig. 28 right (Dirichlet energy). Finally, we report in Tab. 6 the average geodesic error and Dirichlet energy on the data set, obtained when fixing one weight to 1 and setting the two others to 0. This experiment allows to measure which term carries the greatest influence on the final map quality.

These quantitative evaluations highlight the existence of a trade-off between the accuracy of the map (minimization of the geodesic error) and the smoothness of the map (minimizing the Dirichlet en-

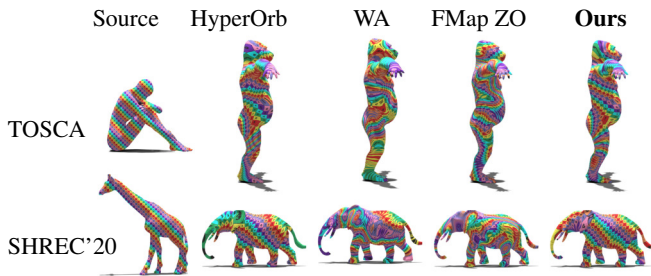


Figure 25: Qualitative evaluation of our method and competitors on **non-isometric shapes**. The first row corresponds to shapes from the TOSCA non-isometric data set. The bottom row consists of shapes from the SHREC'20 lores data set. The shape pair is selected such that the geodesic error of our method is **the worst** over the dataset.




Figure 26: Qualitative evaluation of our method and competitors on the **SHREC'19** data set. The first row corresponds to the **best** shape pair, while the bottom row corresponds to the **worst** shape pair on this data set.

ergy) when choosing the weight configuration. Roughly speaking, the invertibility and properness terms promote accuracy, while the conformality term promotes smoothness.

Since this trade-off is application-dependent, we leave the fine-tuning of the energy weights to the end-user and set all weights to 1 in the remaining of our experiments as it provides a satisfactory balance in practice.

Non-Zero Weight	Av. Geo. Err.	Dir. E.
Conformality (a_C)	5.91×10^{-2}	6.82
Properness (a_P)	7.06×10^{-2}	7.82
Invertibility (a_I)	5.42×10^{-2}	11.4

Table 6: Quantitative evaluation results on the SHREC'20 data set (full shapes remeshed to 1K vertices) when fixing one weight to 1 (Non-Zero Weight) and setting the remaining weights to 0. The average geodesic error (Av. Geo. Err.) and Dirichlet Energy (Dir. E.) is given for each.




Figure 27: Weight study on the SHREC'20 data set (full shapes remeshed to 1K vertices). The error measure is the **mean geodesic error**, averaged on the data set. a_C , a_P and a_I are the Conformality, Properness and Invertibility weights. On the **left**, we fix the conformality weight a_C and vary the properness and invertibility weights a_P and a_I . On the **right**, we vary one weight $a_C/P/I$ and fix the remaining weights either to 0 or to 1.




Figure 28: Weight study on the SHREC'20 data set (full shapes remeshed to 1K vertices). The error measure is the **Dirichlet energy**, averaged on the data set. a_C , a_P and a_I are the Conformality, Properness and Invertibility weights. On the **left**, we fix the conformality weight a_C and vary the properness and invertibility weights a_P and a_I . On the **right**, we vary one weight $a_C/P/I$ and fix the remaining weights either to 0 or to 1.

Landmark Sampling Qualitative Illustration

We visualize qualitatively the interest of introducing more landmark correspondences in Fig. 29. In this visualisation, since ‘‘HyperOrb’’ does not support less than 5 landmark correspondences, no map for 3 and 4 landmark correspondences can be computed for this method.

Note how the regions around the mouth and the eyes are accurately mapped with our approach compared to the two other approaches.

Basis Near-Orthogonality

For each shape \mathcal{M} of the SHREC'19 data set [MMR*19a], we compute the matrix with entries $m_{i,j} = \left| \langle \Phi_i^{\mathcal{M}}, \Phi_j^{\mathcal{M}} \rangle_{W(\mathcal{M})} \right|$, where $\Phi_i^{\mathcal{M}}$ designates the i -th basis vector. We use 7 landmarks, 10 Dirichlet-Steklov eigenfunctions, leading to a Dirichlet-Steklov block of size 70×70 , and 120 Dirichlet Laplacian eigenfunctions. Since we are only interested in the computation of the basis itself in this setup, the landmarks were placed at random locations to maximize the diversity of situations encountered. The average of all matrices is displayed in Fig. 30. Note the clear diagonal behavior, that is in agreement with our observations on a simple sphere shape (Fig. 5).




Figure 29: Qualitative comparison of our method to competitors when increasing the number of landmarks on the same shape pair as for our teaser (Fig. 1). The ground truth landmark locations are denoted by green dots. In the case of FMapZO (no exact landmark preservation), the blue dots indicate the location of the mapped landmarks.




Figure 30: Average of the absolute values of the inner product matrix of each shape in the SHREC'19 data set. Except for the first few Dirichlet-Steklov eigenfunctions, the off-diagonal inner products are negligible. This validates the approximation of orthogonality. We highlight that this computation also sheds light on the robustness of our basis computation to complex triangulation and partiality setups.

Number of Basis Functions

To select the number of basis functions for $\mathcal{G}(\mathcal{M})$ and each $\mathcal{H}_j(\mathcal{M})$ (see Sec. 4.4), we study their respective size N_{LB} and N_{DS} separately, as illustrated in Fig. 31.

Increasing the size of $\mathcal{G}(\mathcal{M})$ slightly increases the matching performance up to $N_{LB} = 120$. In contrast, varying N_{DS} above 10 decreases the quality of the maps. Hence, we fix the following basis sizes throughout the rest of the article: $N_{LB} = 120$ and $N_{DS} = 10$.




Figure 31: Effect of varying the size of our basis on the $\mathcal{G}(\mathcal{M})$ space (left) and on the $\mathcal{H}_j(\mathcal{M})$ space (right). Both figures are an average over all pairs of the TOSCA non-isometric dataset.

Appendix I: Evaluation Setup Details

Landmark Position

The benchmark datasets that we use contain either humanoid shapes (humans and gorillas) or four-legged animals. Depending on the type of creature, we place our landmarks at either 7 or 8 semantically compatible locations:

1. Top of the head
2. Bottom of the right (hind) leg
3. Bottom of the left (hind) leg
4. Bottom of the right front leg / extremity of the third finger on the right hand
5. Bottom of the left front leg / extremity of the third finger on the left hand
6. Middle of the belly/umbilicus
7. Middle of the back
8. Tip of the tail (Four-legged animals only)

The last landmark is only used on the TOSCA and SHREC'20 data sets. Notice that our landmark placement is reminiscent of farthest point sampling. The landmark placement is common to all considered methods. The other parameters depend on the method used.

Method Configuration

Hyperbolic Orbifold Tutte Embeddings (hyperOrb) and Weighted Averages (WA). These methods do not require any additional parameters.

Functional Maps With ZoomOut Refinement (FMap ZO). A 20×20 functional map is computed for each source-target pair in setup 1 and 2, following the setup of [MRR*19]. In particular, we use wave kernel signature and wave kernel map functions as descriptors. The descriptor functions are computed at the same ground truth landmark positions used for the other methods. At each landmark location, 12 wave kernel map functions are computed using a basis of 120 LB-eigenfunctions.

The energy employed to compute the functional map leverages the descriptor preservation, descriptor commutativity and LB-commutativity terms. Contrary to [MRR*19], we did not employ the orientation term in the energy. Indeed, with a high number of landmarks as in our setup, the symmetry ambiguities are easily solved by the functional maps pipeline.

Ours. We use the provided landmark locations together with the settings specified previously. We summarize them here for convenience.

- Energy weights: $a_C = a_P = a_I = 1$.
- Number of Dirichlet-Steklov eigenfunctions per landmark: $N_{DS} = 10$.
- Number of Dirichlet Laplacian eigenfunctions: $N_{LB} = 120$.
- Landmark circle size factor: $r_f = 0.5$.

Moreover, recall that we use the acceleration strategy described at the end of Sec. 6.2.