

Hybrid Functional Maps for Crease-Aware Non-Isometric Shape Matching

Lennart Bastian^{1*} Yizheng Xie^{1*} Nassir Navab¹ Zorah Lähler^{2,3,4}

¹ Technical University of Munich, Germany ² University of Siegen, Germany
³ University of Bonn, Germany ⁴ Lamarr Institute, Germany

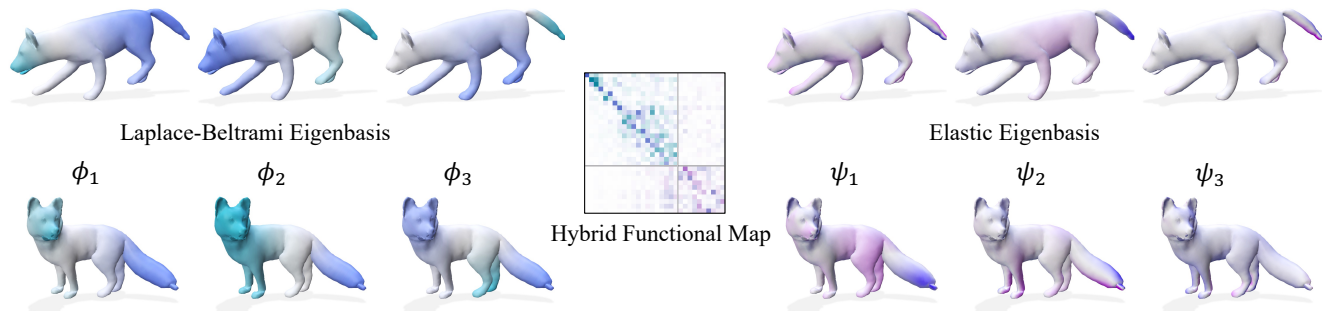


Figure 1. We propose a novel approach of hybridizing the eigenbases originating from different operators for mapping between function spaces in deformable shape correspondence. While the Laplace-Beltrami operator (LBO) eigenbasis is robust to coarse isometric deformations, it fails to encapsulate extrinsic characteristics between shapes. In contrast, elastic basis functions [21] align with high curvature details but lack the robustness for coarse isometric matching. The proposed hybrid basis can be used as a drop-in replacement for the LBO basis functions in modern functional map pipelines, improving performance in near-isometric, non-isometric, and topologically noisy settings.

Abstract

Non-isometric shape correspondence remains a fundamental challenge in computer vision. Traditional methods using Laplace-Beltrami operator (LBO) eigenmodes face limitations in characterizing high-frequency extrinsic shape changes like bending and creases. We propose a novel approach of combining the non-orthogonal extrinsic basis of eigenfunctions of the elastic thin-shell hessian with the intrinsic ones of the LBO, creating a hybrid spectral space in which we construct functional maps. To this end, we present a theoretical framework to effectively integrate non-orthogonal basis functions into descriptor- and learning-based functional map methods. Our approach can be incorporated easily into existing functional map pipelines across varying applications and can handle complex deformations beyond isometries. We show extensive evaluations across various supervised and unsupervised settings and demonstrate significant improvements. Notably, our approach achieves up to 15% better mean geodesic error for non-isometric correspondence settings and up to 45% improvement in scenarios with topological noise. Code is available at: <https://hybridfmaps.github.io/>

*Equal Contribution

✉ : {lennart.bastian,yizheng.xie}@tum.de

1. Introduction

Establishing dense correspondences between 3D shapes is a cornerstone for numerous computer vision and graphics tasks such as object recognition, character animation, and texture transfer. The complexity of this task varies significantly depending on the nature of the transformation a shape undergoes. Many classic correspondence methods leverage that rigid transformations can be represented in six degrees of freedom in \mathbb{R}^3 and preserve the Euclidean distance between pairs of points. Iterative closest point (ICP) [4], and its variations [25, 40], which alternate between transformation and correspondence estimation, are very popular due to their simplicity. In this setting, *local* extrinsic surface properties in the embedding space stay invariant under rigid transformations such that they can be used as features during optimization, for example, the change of normals. For the wider class of isometric deformations (w.r.t. the geodesic distance), the relative embedding of the shape can change significantly, and Euclidean distances between points may not be preserved. In this class, only *intrinsic* properties – those that do not depend on a specific embedding of the surface – stay invariant, and the correspondence problem becomes much harder due to the quadratic size of the solution space. For example, solving a quadratic assignment problem preserving geodesic distances [24] or heat kernel

[52] is *intrinsic* by nature, but it is also an NP-hard problem.

In this context, spectral shape analysis, a generalization of Fourier analysis to Riemannian manifolds, has emerged as a powerful tool for non-rigid correspondence by leveraging intrinsic shape structure. One popular method that takes advantage of this tool is functional maps, introduced by Ovsjanikov et al. [37], which synchronizes the eigenfunctions of the Laplace-Beltrami operator (LBO) through a low-dimensional linear change of basis. Numerous adaptations have led to advances in shape correspondence in recent years, for example, in both the learned supervised [15, 28] and unsupervised settings [11, 12, 26, 45, 51]; and while other basis choices have been proposed [21, 36, 39], almost all of these methods use the eigenfunctions of the LBO to span the to-be-mapped function spaces. One reason is that the LBO has been extensively studied, and the behavior of its eigenfunctions is well understood. For instance, the LBO’s eigenfunctions have a relatively consistent ordering and general invariance under isometric deformations. These understandings have been leveraged for efficient regularization [44] and coarse-to-fine optimization [18, 35]. Other basis sets have been studied and shown to be effective in specific cases [13, 39], but none are so generally applicable and flexible as the LBO eigenfunctions.

A known weakness of the LBO basis, which at the same time comes from its biggest strength, is the reduction to low-frequency information. This leads to efficient optimization and robustness to noise but also inaccuracy in small details. To counter this challenge, Hartwig et al. [21] proposed to utilize a basis derived from the spectral decomposition of an elastic thin-shell energy for functional mapping. These bases are particularly suitable for aligning extrinsic features of non-isometric deformations, for example, bending and creases [21]. However, due to the non-orthogonality of these basis functions, careful mathematical treatment is required to construct the appropriate optimization problem. Furthermore, the elastic basis functions do not exhibit the isometric invariance and robustness of the LBO basis functions, limiting their applicability (see Sec. 5.3).

To address the shortcomings of the bases on their own, we propose to estimate functional maps in a *hybrid* basis representation. We achieve this by constructing a joint vector space between the LBO basis functions and those of the thin shell hessian energy [21, 54]. We demonstrate that combining intrinsic and extrinsic features in this manner provides several advantages for both near-isometric and non-isometric shape-matching problems, promoting robust functional maps that can represent fine creases in the shapes as well as large topological changes. Due to the principled nature of our approach, the combined basis representation can be used in place of pure LBO basis functions in many functional map-based methods. We demonstrate this on several of the most strongly performing axiomatic and learning-based pipelines,

leading to considerable performance improvements on various challenging shape-matching datasets.

Contributions. Our contributions are as follows:

- We introduce a theoretically grounded framework to estimate functional maps between non-orthogonal basis sets using descriptor-based linear systems, a foundational element of nearly all functional map-based learning methods.
- We propose a hybrid framework for estimating functional maps that leverage the strengths of basis functions originating from different operators. We employ this framework to construct functional maps robust to coarse deformations and topological variations while capturing fine extrinsic details on the shape surface.
- We conduct an extensive experimental validation establishing a strong case for the proposed hybrid mapping framework in various challenging problem settings, achieving notable improvements upon state-of-the-art methods for deformable correspondence estimation.

2. Related Work

Shape understanding has been studied extensively; a comprehensive background is beyond the scope of this work. We refer the reader to one of several recent surveys [14, 46]. This section provides an overview of the works most closely related to ours.

Intrinsic-Extrinsic Methods. Both intrinsic and extrinsic approaches have advantages and disadvantages, and an optimal method probably uses both. Several works combining the functional maps framework with extrinsic features exist, for example, with SHOT descriptors [47], including surface orientation information [16, 41], anisotropic information [2], or spatial smoothness of the point map [51]. SmoothShells [18] uses extrinsic information as a deformation field, aligning the surfaces in a coarse-to-fine approach guided by the frequency information of the LBO eigenfunctions. These approaches still use the purely intrinsic LBO eigenfunctions to define the functional maps basis, adding extrinsic information through regularization or additional steps.

Functional Maps. The functional map framework proposed in [37] uses the eigenfunctions of the LBO to pose the correspondence problem as a low-dimensional linear system by rephrasing it as a correspondence of basis functions instead of vertices. The frequency-ordering of the LBO eigenfunctions, as well as their invariance to isometries, allow them to span a comparable but expressive space of smooth functions, which can be efficiently matched by using point descriptors, for example HKS [50], WKS [6] or SHOT [47].

Follow-up work has been proposed to improve the correspondence quality [35, 38], extend it to more general settings [23, 44], and learn to generate optimal descriptors [20, 28, 48]. These methods are particularly powerful as they exploit the structure of the geometric manifolds through the functional correspondence of eigenfunctions on

the shapes but still incorporate a learned descriptor to more accurately represent nuances in the shape surface topology. Unsupervised learning-based approaches have been proven highly effective in recent years [5, 10–12, 26, 45], even surpassing the performance of supervised methods. Such approaches have not only succeeded on a wide range of computer vision benchmarks but have recently proven effective in the medical domain [7, 8, 11, 31].

Basis Functions. Many improvements have been proposed for the functional map framework, but most methods still use the Laplace-Beltrami eigenfunctions as the underlying basis. Despite this, other basis types have been proposed for shape analysis, for example, the L1-regularized spectral basis [36], the landmark-adapted basis [39], a basis derived from Gaussians [13], or localized manifold harmonics [33]. The latter proposed to “mix” a localized basis with the normal LBO eigenfunctions. DUO-FMNet [17] proposes calculating an additional functional map for the complex-valued connection Laplacian basis. However, the basis functions in both cases are orthogonal and purely intrinsic. Another approach is to learn the optimal basis set for functional maps [22, 32, 49], but these tend to not generalize to new applications and, thus, cannot be used out of the box. Various other extrinsic bases have been proposed [1, 29, 53], but none of these have been demonstrated suitable for functional correspondence.

Recently, Hartwig et al. introduced an elastic basis based on the eigendecomposition of the Hessian of the thin-shell deformation energy for functional maps [21]. While it preserves some desirable properties of the LBO (like frequency information) and is better suited for detail alignment, our results indicate it does not perform well in learned functional map-based pipelines (c.f. Sec. 5.3). In this work, we analyze the reasons for this and propose a novel way to preserve the advantages of the elastic basis while joining it with the performance of LBO-based approaches.

3. Background: Functional Maps

Functional maps [37] offer a compelling framework for shape matching by abstracting point-to-point correspondences $S_1 \rightarrow S_2$ to a functional representation between function spaces on manifolds $\mathcal{F}(S_1) \rightarrow \mathcal{F}(S_2)$. This paradigm simplifies the map optimization problem to a linear and compact (low-rank) form, enabling additional regularization.

Until now, the Laplace-Beltrami eigenfunctions have been used almost exclusively as the basis to span the to-be-matched function spaces due to their desirable properties, for example, orthogonality, isometry invariance, and allowing a significant dimensionality reduction. In Sec. 3.1, we will study the more general setting of computing functional maps for non-orthogonal basis sets, an extension of the non-orthogonal ZoomOut [35] that has been proposed in [21]. But first, we introduce the default functional map framework. **Spectral Decomposition.** A positive semidefinite (p.s.d.)

linear operator \mathcal{T} (in most cases the LBO, Δ) is computed on the mesh representation of each shape, followed by solving the generalized eigenvalue problem:

$$\mathcal{T}\phi_i = \lambda_i M\phi_i. \quad (1)$$

Ordered by eigenvalues, the first k eigenfunctions Φ_k can be used as a truncated basis for each shape. As both \mathcal{T} and the mass matrix of lumped area elements for each shape M are p.s.d., the eigenfunctions are orthogonal w.r.t the norm induced on the vector space by M : $\Phi_k^T M \Phi_k = I$.

Functional Map Estimation. Given two point descriptors functions $D_1 \in \mathcal{F}(S_1), D_2 \in \mathcal{F}(S_2)$ which are known to be corresponding, the functional map between two basis sets can be computed via a least-squares problem. Let $D_{\Phi_i} := \Phi_i^\dagger D_i$ denote the descriptor functions projected into the LBO eigenfunctions Φ_i using the Moore-Penrose pseudo inverse Φ_i^\dagger . We can then compute an optimal functional map by solving the following optimization problem [15, 37]:

$$C^* = \arg \min_C E(C) = E_{\text{data}}(C) + \lambda E_{\text{reg}}(C) \quad (2)$$

$$E_{\text{data}}(C) = \|CD_{\Phi_1} - D_{\Phi_2}\|_F^2$$

$$E_{\text{reg}}(C) = \|C\Lambda_1 - \Lambda_2 C\|_F^2$$

where Λ_1 a diagonal matrix of the eigenvalues of \mathcal{T} [15] or the resolvent [42]. This energy can be solved in closed form row-by-row with k least squares problems when defined in the Frobenius norm [15].

Learned features have proven robust for a wide variety of surface representations. Unless mentioned otherwise, we use deep features from DiffusionNet [48] and denote these as $D_i \in \mathbb{R}^{n_i \times d}$ for shapes S_1 and S_2 .

Map Regularization. The estimated map can be interpreted as a change of basis between shapes. In case of an underdetermined linear system in Eq. (2) or noisy descriptor function, C can be further regularized with losses that promote orthogonality, bijectivity, isometry, or additional pointwise descriptor preservation [12, 15, 45]. If the regularizer is in a simple quadratic form, it can be backpropagated through and used to train the descriptor functions.

3.1. Non-Orthogonal Basis Functions

Wirth et al. [54] originally proposed an elastic thin-shell energy for spectral analysis. Hartwig et al. [21] then recently demonstrated how the spectral decomposition of this elastic deformation energy can be used for functional mapping despite being non-orthogonal [21]. The elastic energy $\mathcal{W}_S[\mathbf{f}]$ consists of a membrane contribution \mathcal{W}_{mem} , which measures the local distortion of the surface, and bending energy $\mathcal{W}_{\text{bend}}$ encapsulating curvature (c.f. appendix for a complete definition). By construction, the semi-positive definite hessian of the elastic deformation energy can be decomposed at the identity as in Eq. (1), yielding a set of eigenfunctions.




Figure 2. A Percentage-Correct-Keypoint ablation between the pure LB basis, pure elastic basis (orthogonalized), and our hybrid approach at the same spectral resolution ($k = 200$). The elastic basis attains better detail alignment than the LB basis but yields inferior overall global correspondences. The proposed hybrid approach achieves the best of both worlds. Experiments are conducted with the ULRSSM [12] framework on SMAL.

These vector-valued eigenmodes are suitable for functional mapping after projection onto the vertex-wise normals of the surface and selecting the first k non-orthogonal basis functions $\Psi = [\psi_1, \dots, \psi_k] \in \mathbf{R}^{n,k}$ [21].

Much of the simplicity of the functional maps framework can be attributed to the orthogonality of the basis functions w.r.t the mass matrices on each shape. The mass matrix accounts for the anisotropic metric on the non-Euclidean shape manifolds which must be observed for common operators such as the inner product $\langle \cdot, \cdot \rangle_M$ and norm $\| \cdot \|_M$. The reduced mass representation $M_k = \Psi^T M \Psi \in \mathbf{R}^{k \times k}$ can be used to construct a metric in the spectral space of each shape. Notably, these operations reduce to the standard inner product when represented in the orthogonal LBO basis.

However, this is not the case for non-orthogonal basis functions, and careful treatment must be taken to avoid neglecting the anisotropic metric. Hartwig et al. [21] derive the necessary operations, such as the orthogonal projector, and reformulate optimization problems to use the elastic basis in the ZoomOut [35] framework for functional map refinement. For a thorough treatment of these fundamental definitions, we refer the reader to the relevant literature [21, 54]. Our method requires several additional operations and losses to utilize the elastic basis in a learned setting, including the formulation in Eq. (2), which we will derive in Sec. 4.

4. Method: A Hybrid Approach

The LBO eigenbasis is the predominant choice in functional map-based [37] approaches due to their robustness and invariance to isometric deformations, but they tend to struggle with aligning high-frequency details. On the other hand, the recently proposed elastic basis functions have proven effective at representing extrinsic creases and bending [21] (see

Fig. 2). However, we observed that naively replacing the LBO basis with the elastic basis does not always improve performance, particularly in learning-based frameworks (see Tab. 2).

To overcome the deficiencies of both basis choices, we propose constructing functional maps between *hybrid spaces consisting of the LBO and elastic basis functions*. This attains the best of both worlds: a stable, isometric functional map at low frequencies and sensitivity to extrinsic creases and high-curvature details. To achieve this, we generalize the deep functional maps framework outlined in Sec. 3 to non-orthogonal basis functions in Sec. 4.1, then introduce the hybrid functional map estimation in Sec. 4.2, and discuss necessary adjustments for learning pipelines in Sec. 4.3.

4.1. Generalization to the Hilbert-Schmidt Norm

In this section, we will generalize Eq. (2) to functional maps between non-orthogonal basis sets, for example, the elastic basis [21]. For an orthogonal basis, Eq. (2) can be written with the Frobenius norm in spectral space. For non-orthogonal basis functions, this requires using an inner product induced by the mass matrices on each shape [21]; norms to measure distances in each Hilbert space or the magnitude of linear operators must be scaled similarly.

Data Term. The original formulation of Eq. (2) takes the difference of the descriptors D_1, D_2 as functions on the surface using the S_2 inner product; this reduces to the standard inner product in spectral space for the LBO eigenfunctions. For non-orthogonal basis sets, the spectral space is a Hilbert space equipped with an inner product induced by the reduced mass matrix $M_{k,2} = \Psi_2^T M_2 \Psi_2$. The data term then reads:

Lemma 4.1. *The descriptor preservation term E_{data} can be represented in the norm induced by $M_{k,2}$ as:*

$$\|CD_{\Psi_1} - D_{\Psi_2}\|_{M_{k,2}} = \|\sqrt{M_{k,2}}(CD_{\Psi_1} - D_{\Psi_2})\|_F \quad (3)$$

We include a derivation in the appendix for completeness.

Regularizer. Next, we derive E_{reg} which ensures the functional map C commutes with the diagonal matrix of eigenvalues Λ_i of the respective linear operator [15, 37] or its resolvent [42]. A key to functional map formulation of Hartwig et al. [21] is the use of the Hilbert-Schmidt norm for linear operators between Hilbert spaces, as it considers the geometry on both the domain *and* range of the operator as opposed to the Frobenius norm. We note that the term E_{reg} measures the magnitude of the operator $(C\Lambda_1 - \Lambda_2 C) : \mathcal{F}(S_1) \rightarrow \mathcal{F}(S_2)$, and should therefore take into account the anisotropic metrics on each space.

Proposition 4.2. *The regularization term E_{reg} can be formulated in the Hilbert-Schmidt norm as:*

$$\|C\Lambda_1 - \Lambda_2 C\|_{HS} = \|\sqrt{M_{k,2}}(C\Lambda_1 - \Lambda_2 C)\sqrt{M_{k,1}^{-1}}\|_F$$

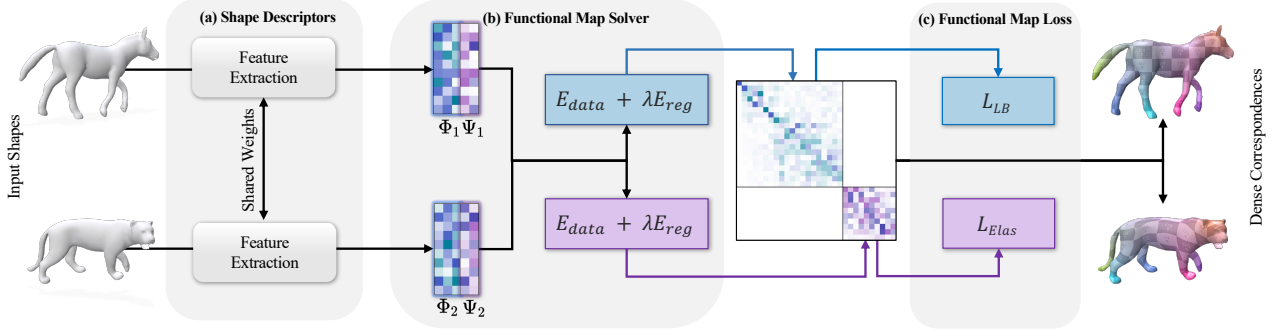


Figure 3. **Hybrid Functional Maps** in a typical pipeline. Features are first extracted from a pair of shapes with a Siamese network (a). They are then projected onto eigenbasis sets from different linear operators (b). We then solve for a block diagonal functional map spanning the constructed hybrid function space (b). Additional regularization can be used to impose structure on parts of the hybrid functional map (c).

This problem can be expanded to a $k^2 \times k^2$ linear system:

$$\|((\Lambda_1 \sqrt{M_{k,1}^{-1}}) \otimes \sqrt{M_{k,2}} - \sqrt{M_{k,1}^{-1}} \otimes (\sqrt{M_{k,2}} \Lambda_2)) \vec{C}\|_2$$

with the Kronecker product \otimes , and $\vec{C} = \text{vec}(C)$ the column stacked vectorization of C , and using Lemma 4.1. This system can be solved in closed form.

Proof. The first statement follows from the definition of the HS-norm, using the cyclicity of the trace and equivalence with the scaled Frobenius norm. A detailed discussion regarding how to reformulate this optimization problem in the expanded form can be found in the appendix. \square

It was previously shown that the formulation of $E(C)$ in the Frobenius norm admits a closed-form solution [15, 42]. This is crucial to the deep functional maps pipeline; we found that solving iteratively with a differentiable convex optimizer proves prohibitively expensive when solved to sufficient accuracy for the outer SGD iteration to converge. We therefore solve Eq. (2) in the expanded form to consider the anisotropic metrics on $\mathcal{F}(S_1)$ and $\mathcal{F}(S_2)$. However, the expanded $k^2 \times k^2$ system becomes prohibitively large at $k = 200$, the spectral resolution typically used in advanced functional maps pipelines. In the next section, we show that separating the functional map optimization in Eq. (2) into two problems under mild assumptions effectively resolves this issue and enables proper regularization for practical applications.

4.2. Hybrid Functional Map Estimation

Our experiments suggest that although the elastic basis performs sub-optimally compared to the LB basis in deep-learning settings, it achieves a higher percentage of matches at a low geodesic error threshold, suggesting superior alignment of fine details and creases (see Fig. 2). Motivated by this, we propose constructing a hybrid basis by combining basis functions from both operators. Intuitively, the low-frequency LBO eigenfunctions approximate the shape and

enable coarse alignment, while the elastic eigenfunctions conform to creases and regions of high curvature. In this hybrid function space, a functional map C is articulated as a block matrix, with each entry C^{ij} encoding the correspondence between two basis sets (see Fig. 1).

$$C = \begin{pmatrix} C^{11} & C^{12} \\ C^{21} & C^{22} \end{pmatrix} \quad (4)$$

C^{11} and C^{22} correspond to intra-basis maps and the off-diagonal blocks C^{12} and C^{21} to inter-basis maps. The resulting hybrid map can be used directly to obtain dense point-to-point correspondences via nearest neighbor search in the hybrid basis or via map refinement strategies [35].

We observe that, while mutually non-orthogonal, the LBO and elastic eigenbasis exhibit very different behaviors on the shape, and therefore assume that inter-basis maps are undesirable (enforcing $C^{12} = 0$ and $C^{21} = 0$). Eq. (4) then separates into two optimization problems (one for each basis type), resulting in a block diagonal functional map (see Fig. 3). In the appendix, we support this with mathematical intuition and empirically show that inter-basis matchings adversely affect the convergence of the map.

The combined framework enables the application of the hybrid basis to large functional map systems (e.g. $k = 200$) with proper adaptation for a non-orthogonal basis, combining the benefits of different basis types and exceeding the performance of each basis when used individually (see Fig. 2).

4.3. Learning in a Hybrid Basis

Deep functional map pipelines regularize the FM obtained from Eq. (2) with additional loss functions. These are described in detail in Sec. 5 for each specific pipeline. Similar to Sec. 4.2, we note that each loss can be separated in our block diagonal hybrid formulation (c.f. appendix for details).

The isometric invariance of the LBO loss functions provides a strong supervision signal for unsupervised FM methods. As the elastic basis functions lack these properties, training from scratch with various architectures leads to

Table 1. **Shape correspondence estimation** under various conditions, including isometric, non-isometric, and settings with topological noise. The proposed hybrid approach yields performance improvements in axiomatic, supervised, and unsupervised settings.

† SHREC’19 methods are trained on FAUST and SCAPE as in recent methods [12, 26].

	Geodesic Error ($\times 100$)	FAUST	SCAPE	SHREC’19†	SMAL	DT4D-H		TOPKIDS
						intra-class	inter-class	
Axiomatic	ZoomOut [35]	6.1	7.5	-	38.4	4.0	29.0	33.7
	DiscreteOp [43]	5.6	13.1	-	38.1	3.6	27.6	35.5
	Smooth Shells [18]	2.5	4.2	-	30.0	1.2	6.4	10.8
	Hybrid Smooth Shells (ours)	2.6	4.2	-	28.4	1.3	5.7	7.5
Sup.	FMNet [28]	11.0	33.0	-	42.0	9.6	38.0	-
	GeomFMaps [15]	2.6	3.0	7.9	8.4	1.9	4.2	-
	Hybrid GeomFMaps (ours)	2.4	2.8	5.6	7.6	2.2	4.1	-
Unsupervised	Deep Shells [19]	1.7	2.5	21.1	29.3	3.4	31.1	13.7
	DUO-FMNet [17]	2.5	4.2	6.4	6.7	2.6	15.8	-
	AttentiveFMaps-Fast [26]	1.9	2.1	6.3	5.8	1.2	14.6	28.5
	AttentiveFMaps [26]	1.9	2.2	5.8	5.4	1.7	11.6	23.4
	SSCDFM [51]	1.7	2.6	3.8	5.4	1.2	6.1	-
	ULRSSM [12]	1.6	1.9	4.6	3.9	0.9	4.1	9.2
	Hybrid ULRSSM (ours)	1.5	1.8	3.6	3.3	1.0	3.5	5.0

suboptimal convergence. We, therefore, parameterize the optimization in Eq. (5) through linear annealing during training:

$$\mathcal{L}(C) = \mathcal{L}_{\text{LB}}(C) + \mu\mathcal{L}_{\text{Elas}}(C) \quad (5)$$

Intuitively, this favors coarse isometric matching early on during training with the LBO eigenfunctions and leverages the tendency of the elastic basis to align creases and details later for optimal convergence. Empirically, we find this achieves superior performance compared to fine-tuning from LBO pre-trained descriptors, which likely converges to local minima near the LBO optimum.

5. Experimental Results

This section provides a summary of the datasets used and our experimental setup. We refer to the appendix for a complete description of the datasets, splits, hyperparameters, and reformulation of method-specific losses in the HS-norm. We use $k = k_{\text{LB}} + k_{\text{Elas}}$ to signify the total spectral resolution, the number of LBO and elastic basis functions, respectively.

5.1. Datasets

We evaluate our method on several challenging benchmarks encompassing *near-isometric* (FAUST [9], SCAPE [3], SHREC [34], DeformingThings4D intra- [27]), *non-isometric* (SMAL [55], DeformingThings4D inter- [27]), and *topologically noisy* (TOPKIDS [30]) settings. We use the more challenging re-meshed versions as established [12, 17].

5.2. Hybrid Basis in Different Frameworks

To understand the efficacy of the proposed hybrid basis in various methodological settings, we use it instead of the

LBO basis in three different methods spanning supervised (GeomFMaps [15]), unsupervised (ULRSSM [12]), and axiomatic settings (Smooth Shells [18]). Due to inherent variability, we reproduce each experiment 5 times in both the baseline (LBO) and hybrid configuration, reporting the best results consistent with standard practices. The total number of basis elements k is kept fixed per method for all experiments; we replace only the highest-frequency LBO eigenfunctions with the elastic basis functions corresponding to the smallest eigenvalues. Quantitative experimental results (c.f. Tab. 1) are organized into sections (supervised, unsupervised, axiomatic), where we compare to competitive methods in the same category. Qualitative results are shown in Fig. 4 and in the supplementary.

GeomFMaps [15] originally proposed the addition of a Laplacian regularization term to the FMNet framework, which has proven effective at enforcing isometric characteristics of the map calculated from Eq. (2). We replace the LBO basis functions with the hybrid formulation, solving them separately as proposed in Sec. 4. For the elastic part of the functional map, we replace both the E_{data} and E_{reg} terms in the map optimization problem with our weighted variations. We also regularize the ground truth supervision loss $\mathcal{L}_{\text{gt}} = (C - C_{\text{gt}})$ with the weighted HS-norm. The hybrid functional map is refined during inference to obtain dense point-to-point correspondences by performing a nearest-neighbor search in the hybrid vector space. Following the recommendations of the original authors [15], all final results in Tab. 1 are run at a spectral resolution of $k_{\text{LB}} = 20$, $k_{\text{Elas}} = 10$.

Results. We compare our results with those of GeomFMaps under the LBO basis and the supervised method FMNet [28]. Notably, the proposed hybrid basis outperforms LBO Ge-




Figure 4. **Qualitative Results on SMAL, DT4D-H, and TOPKIDS.** Comparison of ULRSSM in the LBO basis and in the proposed hybrid basis. Hybrid functional maps yield higher-quality correspondences, particularly under topological noise. ULRSSM in the LBO basis frequently creates coarse mismatches such as incorrectly assigning appendages, whereas the elastic basis better represents these details. The first six columns show texture transfer. The last columns transfer normals making the less accurate alignment of creases in ULRSSM visible.

omFMaps in most settings, spanning near-isometric and non-isometric shape matching, where a particular benefit can be seen for SHREC’19 and SMAL, with a 2.3 and 0.8 improvement in mean geodesic error, respectively.

ULRSSM [12] has recently achieved SoTA performance in various challenging shape-matching settings. We evaluate our proposed hybrid basis when used in ULRSSM instead of the pure LBO functional basis. ULRSSM uses the functional map computation term described in Eq. (2). Hence, we proceed to split the optimization problem as described in Sec. 4.2 and adapt the elastic part with the proposed weighted formulation. The authors of ULRSSM additionally regularize the functional map C to preserve bijectivity \mathcal{L}_{bij} , orthogonality $\mathcal{L}_{\text{orth}}$, and a loss coupling functional and point-to-point maps $\mathcal{L}_{\text{couple}}$ in a differentiable manner. For the elastic optimization, these are all reformulated in the HS-norm. We use the same overall spectral resolution $k = 200$ as the original implementation[12], with $k_{\text{LB}} = 140$, $k_{\text{Elas}} = 60$. This choice of basis ratio is discussed in the appendix.

Results. Using the hybrid basis instead of the LBO basis in ULRSSM results in notable performance improvements, even in near-isometric matching settings such as FAUST and SCAPE Tab. 1. Improvements are most significant in the non-isometric settings, including SMAL and inter-class DT4D-H, where the hybrid basis outperforms LBO with a geodesic error of 0.6. The most notable performance increase can be observed for TOPKIDS, where the hybrid basis yields a 45% improvement in geodesic error. Percentage-correct-keypoints (PCK) plots underscore these results (see Fig. 5).

Smooth Shells [18] remains one of the most strongly performing axiomatic methods for spectral shape matching. The method generates initial hypotheses for aligning a shape pair through a Markov-Chain Monte-Carlo (MCMC) step in a low-dimensional spectral basis ($k = 20$). The algorithm then proceeds with an alternating optimization using both extrinsic and intrinsic information. Following the principle that Laplacian eigenfunctions capture coarse shape features well, we perform the MCMC initialization in the LBO basis. During the hierarchical matching step, we extend the product manifold with an additional dimension consisting of the elastic basis. We use a spectral resolution of $k_{\text{LB}} = 300$, $k_{\text{Elas}} = 200$, while the original implementation uses $k = 500$.

Results. We observe that the performance with the proposed hybrid basis also leads to improved performance of Smooth Shells’ over pure LBO, particularly for non-isometric and topologically noisy settings. Notable improvements can be seen for the TOPKIDS and SMAL datasets, with a 3.3 and 1.6 improvement in mean geodesic error, respectively.

5.3. Ablations and Implementation

We conduct two ablations to support the design choices regarding the generalization to the Hilbert-Schmidt (HS) norm in Sec. 4, and the hybrid formulation. All experiments are conducted on ULRSSM with $k = 200$ basis functions.

To motivate the generalization to the HS norm, we consider two alternatives. The first involves using the standard functional map solver without making any adjustments to




Figure 5. **Percentage-Correct-Keypoint Plots** depicting the geodesic error for state-of-the-art unsupervised methods on the datasets SMAL, DT4D-H inter, and TOPKIDS. We compare AttentiveFMaps, ULRSSM, and Hybrid ULRSSM(Ours).

the non-orthogonal elastic basis. The second alternative, we orthogonalize the basis using Gram-Schmidt under the inner product induced by M , making it directly usable in a standard FMap framework. Next, we compare the proposed hybrid formulation against using either pure LB or pure elastic basis functions. Due to the complexity of the $k^2 \times k^2$ expansion under the HS-norm in Sec. 4, ULRSSM with 200 elastic basis becomes computationally intractable. We therefore orthogonalize the elastic basis functions as an approximation to the proper adaptation of a fully elastic functional map.

Results. The results can be seen in Tab. 2 and Fig. 2. We observe that using the standard Frobenius norm or orthogonalizing the elastic basis yields inferior results compared to the HS-norm adaptation. Furthermore, both pure LB and pure elastic basis adaptations perform worse than the proposed hybrid framework (c.f. Fig. 2). Interestingly, we observe that while the elastic basis functions achieve a superior detail alignment, and the LBO a better coarse alignment, the hybrid basis surpasses the performance in both regimes. We conclude that both basis hybridization and the HS norm adaptation contribute to notable performance improvements.

Implementation Details. Experiments are carried out in Pytorch 2.1.0 with CUDA version 12.1, except for Smooth Shells, which is run in Matlab based on the implementa-

tion provided by the authors. Supervised and unsupervised methods are trained and evaluated on an NVIDIA A40. A complete list of hyperparameters for each of the methods used is provided in the appendix.

6. Limitations and Conclusion

This work explores the efficacy of combining basis functions originating from different operators for deformable shape correspondence. Our findings highlight the importance of accurately treating non-orthogonal basis functions to reflect the anisotropic metric on each shape. Imposing orthogonality on the basis functions shows improvement over naive adaptation but does not supplant proper mathematical adaptation of the optimization objectives. Additionally, the elastic basis functions underperform when used independently in a learned context; integrating it with low-frequency LBO basis functions significantly enhances spectral matching accuracy.

Solving the expanded $k^2 \times k^2$ system from Sec. 4 leads to computational overhead; however, this is tractable for the elastic basis size of 60. Performance gains in the expanded form justify this trade-off. Future research could potentially address partial shapes or noisy point clouds with non-orthogonal basis functions as these are active areas of interest [5, 11, 22].

Overall, the proposed hybrid functional mapping approach, leveraging both elastic and LBO eigenfunctions, exhibits notable performance in diverse settings, including isometric and non-isometric deformations and under topological noise. Our findings open new avenues for integrating various non-orthogonal basis functions into deep functional mapping frameworks, paving the way for further advances in spectral shape matching for challenging settings.

Acknowledgements. We thank the reviewers, Lei Li, and Dongliang Cao for their support and valuable feedback. Zorah Löhner is funded by a KI-Starter grant from the Ministry of Culture and Science of the State of North Rhine-Westphalia.

LB	Elastic	Adaptation	Geo. error ($\times 100$)
✗	✓	◆	intractable
✗	✓	-	40.2 ± 0.80
✗	✓	+	5.75 ± 1.20
✓	✗	-	5.15 ± 0.99
✓	✓	-	4.37 ± 1.57
✓	✓	+	4.33 ± 0.56
✓	✓	◆	3.83 ± 0.74

Table 2. Ablation study of the proposed hybrid basis and the effect of two adaptations to the non-orthogonal elastic basis: generalization to the HS norm (◆) as proposed in Sec. 4 and orthogonalization (+). Experiments are conducted with ULRSSM [12] on the SMAL dataset at spectral resolution $k = 200$. The green row represents our approach, and the orange row the original ULRSSM [12]. Experiments are conducted 5 times; mean \pm stdev. is reported.

References

- [1] Yonathan Aflalo, Ron Kimmel, and Dan Raviv. Scale invariant geometry for nonrigid shapes. *SIAM Journal on Imaging Sciences*, 6(3):1579–1597, 2013. [3](#)
- [2] M. Andreux, E. Rodolà, M. Aubry, and D. Cremers. Anisotropic laplace-beltrami operators for shape analysis. In *Proc. European Conference on Computer Vision Workshops (ECCV - NORDIA)*, Zurich, Switzerland, 2014. [2](#)
- [3] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. SCAPE: shape completion and animation of people. In *ACM SIGGRAPH 2005 Papers*, pages 408–416, New York, NY, USA, 2005. Association for Computing Machinery. [6](#)
- [4] Somani Arun, Thomas S. Huang, and Steven D. Blostein. Least-square fitting of two 3-d point sets. *IEEE Pattern Analysis and Machine Intelligence*, 9(5), 1987. [1](#)
- [5] Souhaib Attaiki, Gautam Pai, and Maks Ovsjanikov. DPFM: Deep Partial Functional Maps. In *2021 International Conference on 3D Vision (3DV)*, pages 175–185, 2021. ISSN: 2475-7888. [3](#), [8](#)
- [6] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 1626–1633, 2011. [2](#)
- [7] Lennart Bastian, Alexander Baumann, Emily Hoppe, Vincent Bürgin, Ha Young Kim, Mahdi Saleh, Benjamin Busam, and Nassir Navab. S3M: Scalable Statistical Shape Modeling Through Unsupervised Correspondences. In *MICCAI 2023*, pages 459–469, Cham, 2023. Springer Nature Switzerland. [3](#)
- [8] Lennart Bastian, Vincent Bürgin, Ha Young Kim, Alexander Baumann, Benjamin Busam, Mahdi Saleh, and Nassir Navab. On the localization of ultrasound image slices within point distribution models. In *International Workshop on Shape in Medical Imaging*, pages 133–144. Springer, 2023. [3](#)
- [9] Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. FAUST: Dataset and Evaluation for 3D Mesh Registration. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3794–3801, Columbus, OH, USA, 2014. IEEE. [6](#)
- [10] Dongliang Cao and Florian Bernard. Unsupervised Deep Multi-shape Matching. In *Computer Vision – ECCV 2022*, pages 55–71, Cham, 2022. Springer Nature Switzerland. [3](#)
- [11] Dongliang Cao and Florian Bernard. Self-supervised learning for multimodal non-rigid 3d shape matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17735–17744, 2023. [2](#), [3](#), [8](#)
- [12] Dongliang Cao, Paul Roetzer, and Florian Bernard. Unsupervised learning of robust spectral shape matching. *ACM Transactions on Graphics (TOG)*, 42(4):1–15, 2023. [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [13] Michele Colombo, Giacomo Boracchi, and Simone Melzi. Pcgau: Pca basis of scattered gaussians for shape matching via functional maps. In *Smart Tools and Applications in Graphics (STAG)*, 2022. [2](#), [3](#)
- [14] Bailin Deng, Yuxin Yao, Roberto M. Dyke, and Juyong Zhang. A Survey of Non-Rigid 3D Registration. *Computer Graphics Forum*, 41(2):559–589, 2022. [eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14502](#). [2](#)
- [15] Nicolas Donati, Abhishek Sharma, and Maks Ovsjanikov. Deep Geometric Functional Maps: Robust Feature Learning for Shape Correspondence. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8589–8598, Seattle, WA, USA, 2020. IEEE. [2](#), [3](#), [4](#), [5](#), [6](#)
- [16] Nicolas Donati, Etienne Corman, Simone Melzi, and Maks Ovsjanikov. Complex functional maps: A conformal link between tangent bundles. In *Computer Graphics Forum*, pages 317–334. Wiley Online Library, 2022. [2](#)
- [17] Nicolas Donati, Etienne Corman, and Maks Ovsjanikov. Deep orientation-aware functional maps: Tackling symmetry issues in shape matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 742–751, 2022. [3](#), [6](#)
- [18] Marvin Eisenberger, Zorah Lahner, and Daniel Cremers. Smooth Shells: Multi-Scale Shape Registration With Functional Maps. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12262–12271, Seattle, WA, USA, 2020. IEEE. [2](#), [6](#), [7](#)
- [19] Marvin Eisenberger, Aysim Toker, Laura Leal-Taixé, and Daniel Cremers. Deep shells: Unsupervised shape correspondence with optimal transport. *Advances in Neural information processing systems*, 33:10491–10502, 2020. [6](#)
- [20] Oshri Halimi, Or Litany, Emanuele Rodola Rodola, Alex M. Bronstein, and Ron Kimmel. Unsupervised Learning of Dense Shape Correspondence. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4365–4374, Long Beach, CA, USA, 2019. IEEE. [2](#)
- [21] Florine Hartwig, Josua Sassen, Omri Azencot, Martin Rumpf, and Mirela Ben-Chen. An Elastic Basis for Spectral Shape Correspondence. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Proceedings*, pages 1–11, Los Angeles CA USA, 2023. ACM. [1](#), [2](#), [3](#), [4](#)
- [22] Jiahui Huang, Tolga Birdal, Zan Gojcic, Leonidas J. Guibas, and Shi-Min Hu. Multiway Non-rigid Point Cloud Registration via Learned Functional Map Synchronization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2022. [3](#), [8](#)
- [23] Ruqi Huang, Panos Achlioptas, Leonidas Guibas, and Maks Ovsjanikov. Limit Shapes – A Tool for Understanding Shape Differences and Variability in 3D Model Collections. *Computer Graphics Forum*, 38(5):187–202, 2019. [2](#)
- [24] S. Kezurer, Z. Kovalsky, R. Basri, and Y. Lipman. Tight relaxation of quadratic matching. *Computer Graphics Forum (CGF)*, 34, 2015. [1](#)
- [25] Hao Li, Robert W. Sumner, and Mark Pauly. Global correspondence optimization for non-rigid registration of depth scans. *Computer Graphics Forum, Symposium on Geometry Processing*, 27(5), 2008. [1](#)
- [26] Lei Li, Nicolas Donati, and Maks Ovsjanikov. Learning multi-resolution functional maps with spectral attention for robust shape matching. *Advances in Neural Information Processing Systems*, 35:29336–29349, 2022. [2](#), [3](#), [6](#)

- [27] Yang Li, Hikari Takehara, Takafumi Taketomi, Bo Zheng, and Matthias Nießner. 4DComplete: Non-Rigid Motion Estimation Beyond the Observable Surface. pages 12706–12716, 2021. [6](#)
- [28] Or Litany, Tal Remez, Emanuele Rodola, Alex Bronstein, and Michael Bronstein. Deep functional maps: Structured prediction for dense shape correspondence. In *Proceedings of the IEEE international conference on computer vision*, pages 5659–5667, 2017. [2](#), [6](#)
- [29] Hsueh-Ti Derek Liu, Alec Jacobson, and Keenan Crane. A dirac operator for extrinsic shape analysis. In *Computer Graphics Forum*, pages 139–149. Wiley Online Library, 2017. [3](#)
- [30] Z Löhner, E Rodolà, M M Bronstein, D Cremers, O Burghard, L Cosmo, A Dieckmann, R Klein, and Y Sahilliog. SHREC’16: Matching of Deformable Shapes with Topological Noise. *Eurographics Workshop on 3D Object Retrieval, EG 3DOR*, pages 55–60, 2016. [6](#)
- [31] Robin Magnet, Kevin Bloch, Maxime Taverne, Simone Melzi, Maya Geoffroy, Roman H. Khonsari, and Maks Ovsjanikov. Assessing craniofacial growth and form without landmarks: A new automatic approach based on spectral methods. *Journal of Morphology*, 284(8):e21609, 2023. [3](#)
- [32] Riccardo Marin, Marie-Julie Rakotosaona, Simone Melzi, and Maks Ovsjanikov. Correspondence learning via linearly-invariant embedding. In *Advances in Neural Information Processing Systems*, pages 1608–1620. Curran Associates, Inc., 2020. [3](#)
- [33] Simone Melzi, Emanuele Rodolà, Umberto Castellani, and Michael M. Bronstein. Localized manifold harmonics for spectral shape analysis. *Computer Graphics Forum (CGF)*, 2017. [3](#)
- [34] S. Melzi, R. Marín, E. Rodolà, U. Castellani, J. Ren, A. Poulenard, Peter Wonka, and M. Ovsjanikov. SHREC 2019: Matching Humans with Different Connectivity. 2019. [6](#)
- [35] Simone Melzi, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, Maks Ovsjanikov, et al. Zoomout: spectral upsampling for efficient shape correspondence. *ACM TRANSACTIONS ON GRAPHICS*, 38(6):1–14, 2019. [2](#), [3](#), [4](#), [5](#), [6](#)
- [36] Thomas Neumann, Kiran Varanasi, Christian Theobalt, Marcus Magnor, and Markus Wacker. Compressed manifold modes for mesh processing. *Computer Graphics Forum (Proc. of Symposium on Geometry Processing SGP)*, 33(5):35–44, 2014. [2](#), [3](#)
- [37] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional Maps: A Flexible Representation of Maps Between Shapes. *ACM Transactions on Graphics (ToG)*, pages 1–11, 2012. [2](#), [3](#), [4](#)
- [38] Gautam Pai, Jing Ren, Simone Melzi, Peter Wonka, and Maks Ovsjanikov. Fast Sinkhorn Filters: Using Matrix Scaling for Non- Rigid Shape Correspondence with Functional Maps. *Conference on Computer Vision and Pattern Recognition (CVPR)*, page 11, 2021. [2](#)
- [39] Mikhail Panine, Maxime Kirgo, and Maks Ovsjanikov. Non-isometric shape matching via functional maps on landmark-adapted bases. *Computer Graphics Forum (CGF)*, 41(6): 394–417, 2022. [2](#), [3](#)
- [40] Francois Pomerleau, Francis Colas, Roland Siegwart, and Stéphane Magnenat. Comparing icp variants on real-world data sets. *Autonomous robots*, 34, 2013. [1](#)
- [41] Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics*, 37(6):1–16, 2018. [2](#)
- [42] Jing Ren, Mikhail Panine, Peter Wonka, and Maks Ovsjanikov. Structured Regularization of Functional Map Computations. *Computer Graphics Forum*, 38(5):39–53, 2019. [3](#), [4](#), [5](#)
- [43] Jing Ren, Simone Melzi, Peter Wonka, and Maks Ovsjanikov. Discrete Optimization for Shape Matching. *Computer Graphics Forum*, 40(5):81–96, 2021. [6](#)
- [44] Emanuele Rodolà, Luca Cosmo, Michael M Bronstein, Andrea Torsello, and Daniel Cremers. Partial functional correspondence. In *Computer graphics forum*, pages 222–236. Wiley Online Library, 2017. [2](#)
- [45] Jean-Michel Roufousse, Abhishek Sharma, and Maks Ovsjanikov. Unsupervised Deep Learning for Structured Shape Matching. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1617–1627, Seoul, Korea (South), 2019. IEEE. [2](#), [3](#)
- [46] Yusuf Sahillioglu. Recent advances in shape correspondence. *The Visual Computer*, 36(8):1705–1721, 2020. [2](#)
- [47] Samuele Salti, Federico Tombari, and Luigi Di Stefano. SHOT: Unique signatures of histograms for surface and texture description. *Computer Vision and Image Understanding*, 125:251–264, 2014. [2](#)
- [48] Nicholas Sharp, Souhaib Attaiki, Keenan Crane, and Maks Ovsjanikov. Diffusionnet: Discretization agnostic learning on surfaces. *ACM Transactions on Graphics (TOG)*, 41(3):1–16, 2022. [2](#), [3](#)
- [49] Sharik Siddiqi and Zorah Löhner. A network analysis for correspondence learning via linearly-embedded functions. In *DAGM German Conference on Pattern Recognition*, pages 100–114. Springer, 2023. [3](#)
- [50] Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. *Computer Graphics Forum*, 28(5):1383–1392, 2009. [2](#)
- [51] Mingze Sun, Shiwei Mao, Puhua Jiang, Maks Ovsjanikov, and Ruqi Huang. Spatially and spectrally consistent deep functional maps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14497–14507, 2023. [2](#), [6](#)
- [52] Matthias Vestner, Zorah Löhner, Amit Boyarski, Or Litany, Ron Slossberg, Tal Remez, Emanuele Rodolà, Alex M. Bronstein, Michael M. Bronstein, Ron Kimmel, and Daniel Cremers. Efficient deformable shape correspondence via kernel matching. In *International Conference on 3D Vision (3DV)*, 2017. [2](#)
- [53] Yu Wang, Mirela Ben-Chen, Iosif Polterovich, and Justin Solomon. Steklov spectral geometry for extrinsic shape analysis. *ACM Transactions on Graphics (TOG)*, 38(1):1–21, 2018. [3](#)
- [54] Benedikt Wirth, Leah Bar, Martin Rumpf, and Guillermo Sapiro. A Continuum Mechanical Approach to Geodesics in

Shape Space. *International Journal of Computer Vision*, 93 (3):293–318, 2011. [2](#), [3](#), [4](#)

- [55] Silvia Zuffi, Angjoo Kanazawa, David W. Jacobs, and Michael J. Black. 3D Menagerie: Modeling the 3D Shape and Pose of Animals. pages 6365–6373, 2017. [6](#)