

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

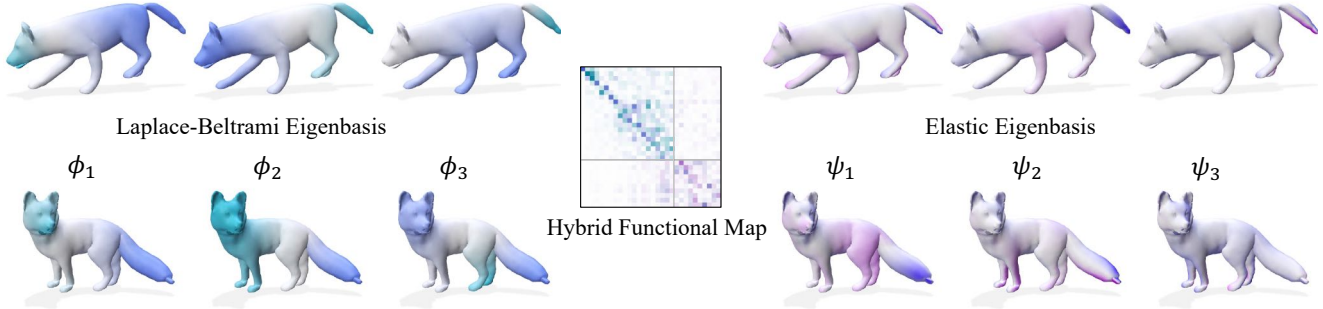


Figure 1. We propose a novel approach of hybridizing the eigenbases originating from different operators for mapping between function spaces in deformable shape correspondence. While the Laplace-Beltrami operator (LBO) eigenbasis is robust to coarse isometric deformations, it fails to encapsulate extrinsic 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 [26, 41], 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 [25] or heat kernel

[53] 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. [38], 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, 29] and unsupervised settings [11, 12, 27, 46, 52]; and while other basis choices have been proposed [21, 37, 40], 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 [45] and coarse-to-fine optimization [18, 36]. Other basis sets have been studied and shown to be effective in specific cases [13, 40], 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 (see Fig. 3). 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, 55]. 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, 47]. 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 [48], including surface orientation information [16, 42], anisotropic information [2], or spatial smoothness of the point map [52]. 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 [38] 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 [51], WKS [6] or SHOT [48].

Follow-up work has been proposed to improve the correspondence quality [36, 39], extend it to more general settings [24, 45], and learn to generate optimal descriptors [20, 29, 49]. 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, 27, 46], 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, 32].

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 [37], the landmark-adapted basis [40], a basis derived from Gaussians [13], or localized manifold harmonics [34]. 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 [23, 33, 50], 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, 30, 54], 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 [38] 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 [36] 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, 38]:

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

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

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

where Λ_1 a diagonal matrix of the eigenvalues of \mathcal{T} [15] or the resolvent [43]. 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 [49] 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, 46]. 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. [55] 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 [36] framework for functional map refinement. For a thorough treatment of these fundamental definitions, we refer the reader to the relevant literature [21, 55]. 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 [38] 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




Figure 3. Correspondence quality visualized through the transfer of normals (left) and vertex positions (right) from the source to the target shape. We compare the results from LB, Elastic, and Hybrid basis functions by encoding and recovering **ground truth** correspondences through functional map representations at a spectral resolution of $k = 60$. Additional results at a spectral resolution of $k = 200$ are provided in the supplementary.

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 sur-

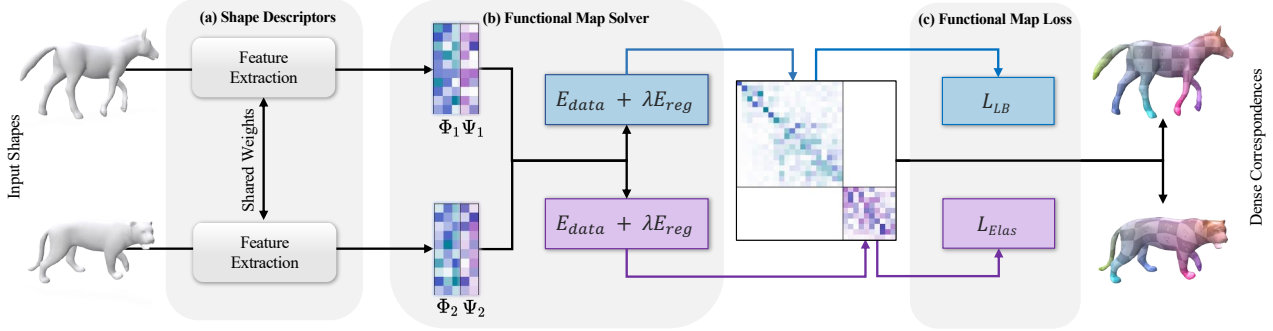


Figure 4. **Hybrid Functional Maps** in a typical pipeline. Features are first extracted from a pair of shapes with a Siamese network (a). They are then projected onto eigenbasis sets from different linear operators (b). We then solve for a block diagonal functional map spanning the constructed hybrid function space (b). Additional regularization can be used to impose structure on parts of the hybrid functional map (c).

face 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, 38] or its resolvent [43]. 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 $(C\Lambda_1 - \Lambda_2C) : \mathcal{F}(S_1) \rightarrow \mathcal{F}(S_2)$, 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:*

$$\|C\Lambda_1 - \Lambda_2C\|_{HS} = \|\sqrt{M_{k,2}}(C\Lambda_1 - \Lambda_2C)\sqrt{M_{k,1}^{-1}}\|_F$$

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, 43]. 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 re-

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, 27].

	Geodesic Error ($\times 100$)	FAUST	SCAPE	SHREC’19†	SMAL	DT4D-H		TOPKIDS
						intra-class	inter-class	
Axiomatic	ZoomOut [36]	6.1	7.5	-	38.4	4.0	29.0	33.7
	DiscreteOp [44]	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 [29]	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 [27]	1.9	2.1	6.3	5.8	1.2	14.6	28.5
	AttentiveFMaps [27]	1.9	2.2	5.8	5.4	1.7	11.6	23.4
	SSCDFM [52]	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

sulting 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 [36].

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 $C^{12} = \mathbf{0}$ and $C^{21} = \mathbf{0}$). Eq. (4) then separates into two optimization problems (one for each basis type), resulting in a block diagonal functional map (see Fig. 4). 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 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 [35], DeformingThings4D intra- [28]), *non-isometric* (SMAL [56], DeformingThings4D inter- [28]), and *topologically noisy* (TOPKIDS [31]) 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




Figure 5. **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 LB 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.

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. 5, Fig. 6 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 [29]. Notably, the proposed hybrid basis outperforms LBO GeomFMaps in most settings, spanning near-isometric and non-isometric shape matching, where a particular benefit




Figure 6. **Qualitative Results on FAUST.** Comparison of ULRSSM in the LBO, (orthogonalized) elastic, and proposed hybrid basis. Both normal (colored) and vertex transfer are visualized. The proposed hybrid basis yields accurate mappings, particularly on fine details such as the hands.

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,




Figure 7. **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).

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. 7).

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 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 pro-

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.

posed 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 implementation 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, 23].

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.

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, 9
- [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, 12
- [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, 9
- [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, 16, 17, 18, 19
- [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](#), [5](#), [6](#), [7](#), [13](#), [16](#), [17](#)
- [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](#), [16](#)
- [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](#), [8](#), [16](#), [18](#)
- [19] Marvin Eisenberger, Aysim Toket, 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 Proceedings*, pages 1–11, Los Angeles CA USA, 2023. ACM. [1](#), [2](#), [3](#), [4](#), [5](#), [12](#), [13](#), [14](#), [15](#), [18](#), [19](#)
- [22] B. Heeren, M. Rumpf, P. Schröder, M. Wardetzky, and B. Wirth. Exploring the Geometry of the Space of Shells. *Computer Graphics Forum*, 33(5):247–256, 2014. [12](#)
- [23] 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](#), [9](#)
- [24] 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](#)
- [25] S. Kezurer, Z. Kovalsky, R. Basri, and Y. Lipman. Tight relaxation of quadratic matching. *Computer Graphics Forum (CGF)*, 34, 2015. [1](#)
- [26] 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](#)
- [27] 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](#), [16](#)
- [28] 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](#), [16](#)
- [29] 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](#), [7](#)
- [30] 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](#)
- [31] 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](#), [16](#)
- [32] 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](#)
- [33] 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](#)
- [34] Simone Melzi, Emanuele Rodolà, Umberto Castellani, and Michael M. Bronstein. Localized manifold harmonics for spectral shape analysis. *Computer Graphics Forum (CGF)*, 2017. [3](#)
- [35] 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](#)
- [36] 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](#), [6](#), [18](#), [19](#)
- [37] 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](#)
- [38] 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](#), [5](#)
- [39] 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](#)
- [40] 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](#)

- [41] Francois Pomerleau, Francis Colas, Roland Siegwart, and Stéphane Magnenat. Comparing icp variants on real-world data sets. *Autonomous robots*, 34, 2013. 1
- [42] 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
- [43] Jing Ren, Mikhail Panine, Peter Wonka, and Maks Ovsjanikov. Structured Regularization of Functional Map Computations. *Computer Graphics Forum*, 38(5):39–53, 2019. 3, 5, 17
- [44] Jing Ren, Simone Melzi, Peter Wonka, and Maks Ovsjanikov. Discrete Optimization for Shape Matching. *Computer Graphics Forum*, 40(5):81–96, 2021. 6
- [45] 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
- [46] Jean-Michel Roufosse, 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
- [47] Yusuf Sahillioglu. Recent advances in shape correspondence. *The Visual Computer*, 36(8):1705–1721, 2020. 2
- [48] 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
- [49] 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
- [50] 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
- [51] 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
- [52] 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
- [53] 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
- [54] 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
- [55] 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
- [56] 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

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

Supplementary Material

In the supplementary materials we first provide additional mathematical background in Appendix A. Appendix B contains complete derivations for the generalization to the HS norm (Lemma 4.1, Thm. 4.2). Detailed analysis for the hybrid formulation is included in Appendix C. We then detail datasets and splits used for evaluation, followed by further experimental details in Appendix D. Additional experiment results is provided at Appendix E. Ablation studies concerning our design choices are provided in Appendix F. Finally, we present runtime analysis in Appendix G, additional visualization of the basis embedding in Appendix H and additional qualitative results in Appendix I.

A. Mathematical Background

A.1. Elastic Energy

We defer to Hartwig et al. [21] for a complete definition of the previously described elastic energy [6, 21, 22]. For all our experiments, we use the same elastic energy hyperparameters as Hartwig et al. [21], including a bending weight of 10^{-2} . We can then solve the generalized eigenvalue problem for the Hessian of the elastic energy at the identity to obtain the basis functions Ψ .

$$\text{Hess } \mathcal{W}_S[\text{Id}]v_\lambda = \lambda M v_\lambda$$

A.2. Problem Setting

In the non-rigid correspondence literature, descriptors D_i are commonly characterized as functions over the shapes S_i . In the discretized setting, many operations reduce to matrix-vector products. However, to derive the proper operations weighted by the non-uniform weight matrices M in the regularization of the functional map, we utilize the more general Hilbert space setting.

Symbol	Description
S_1, S_2	3D shapes (triangle mesh) with $n_{1,2}$ verts
M_i	mass matrix on shape i
D_i	vertex-wise descriptors for shape i
Δ_i	Laplacian operator applied to shape S_i
$\mathcal{W}_S[\cdot]$	Elastic energy associated with S_i
Φ_i	eigenbasis of Laplacian matrix Δ_i
Ψ_i	eigenbasis of Elastic Hessian $\text{Hess}\mathcal{W}_S[I]$
C_{ij}	functional map between shapes S_i and S_j
P_{ij}	point-wise map between shapes S_i and S_j
k	the total spectral resolution
$\ \cdot\ _{\{2,F,HS\}}$	the L2, Frobenius, and HS norms

Table 3. Summary of notations used in this work.

We assume that all functions on the spaces $\mathcal{F}(S_i)$ are L^2 integrable:

$$L^2(\Omega) := \left\{ f : \Omega \rightarrow \mathbb{R} \mid \int_{\Omega} |f(x)|^2 dx < \infty \right\}$$

Then the inner product on each space $\mathcal{F}(S_i)$ is given by $\langle \cdot, \cdot \rangle_M$ in the space induced by M :

$$\langle x, y \rangle_M = \int_{\Omega} x(t)y(t) dM(t) \stackrel{?}{=} x^T M y$$

where in the last step we emphasize that the discretized operations reduce to matrix-vector multiplications in the finite-dimensional setting.

We use the definition of the Hilbert-Schmidt norm for a general operator A between (unweighted) Hilbert spaces [21, Sec 3.4]:

$$\|A\|_{HS} := \sqrt{\text{tr}(A^*A)}$$

When $A : \mathcal{F}(S_1) \rightarrow \mathcal{F}(S_2)$ (the spaces under the anisotropic metric), we have the following equivalence with the Frobenius norm [21]:

$$\begin{aligned} \|A\|_{HS}^2 &:= \text{tr}(M_{k,1}^{-1}A^T M_{k,2}A) \\ &= \text{tr}(\sqrt{M_{k,1}^{-1}}A^T \sqrt{M_{k,2}}\sqrt{M_{k,2}}A\sqrt{M_{k,1}^{-1}}) \\ &= \left\| \sqrt{M_{k,2}}A\sqrt{M_{k,1}^{-1}} \right\|_F^2 \end{aligned} \quad (6)$$

B. Generalization to the Hilbert-Schmidt Norm

B.1. Derivation of Eq. (2) for General Hilbert Spaces.

Proof of Thm. 4.1. The data term can be interpreted as the difference of the descriptor functions $D_1, D_2 \in \mathcal{F}(S_2)$ after D_1 was transferred to $\mathcal{F}(S_2)$ via the functional map C . We denote the first k eigenfunctions by $\Psi_{k,i}$ and the coefficients of D_i projected into the basis spanned by these eigenfunctions by $D_{\Psi_i} := \Psi_{k,i}^\dagger D_i$. We then have the following:

$$\begin{aligned} &\|CD_{\Psi_1} - D_{\Psi_2}\|_{M_{k,2}} \\ &= \sqrt{\langle CD_{\Psi_1} - D_{\Psi_2}, CD_{\Psi_1} - D_{\Psi_2} \rangle_{M_{k,2}}} \\ &= \sqrt{\text{tr}((CD_{\Psi_1} - D_{\Psi_2})^T M_{k,2} (CD_{\Psi_1} - D_{\Psi_2}))} \end{aligned}$$

where we use the definition of the inner product Appendix A.2, the cyclicity of the trace. The identity then follows by splitting $M_{k,2} = \sqrt{M_{k,2}}\sqrt{M_{k,2}}$ and applying the definition of the Frobenius norm again, and using that $M_{k,2}$ is symmetric. \square

As previously established [15], the energy in Eq. (2) can be solved for C in closed form by solving k different $k \times k$ linear systems (for each row of C). In our case, the mass matrices M prohibit this, requiring an expansion to a $k^2 \times k^2$ system. This expansion is detailed below.

Proof of Thm. 4.2. Let S_1 and S_2 be Hilbert spaces defined on two shapes associated with mass matrices $M_{k,1}$ and $M_{k,2}$, respectively, which induce the inner product on each space. Let Λ_1 and Λ_2 be the diagonal matrices of eigenvalues of the respective linear operator on S_1 and S_2 , and let $C: \mathcal{F}(S_1) \rightarrow \mathcal{F}(S_2)$ be a linear map between the function spaces. The weighted linear operator commutativity regularization term can be expressed using the Hilbert-Schmidt norm as follows:

$$\begin{aligned} & \| (C\Lambda_1 - \Lambda_2 C) \|_{HS}^2 \\ &= \text{tr}(M_{k,1}^{-1}(C\Lambda_1 - \Lambda_2 C)^T M_{k,2}(C\Lambda_1 - \Lambda_2 C)) \\ &= \left\| \sqrt{M_{k,2}}(C\Lambda_1 - \Lambda_2 C) \sqrt{M_{k,1}^{-1}} \right\|_F^2 \\ &= \left\| \sqrt{M_{k,2}}C\Lambda_1 \sqrt{M_{k,1}^{-1}} - \sqrt{M_{k,2}}\Lambda_2 C \sqrt{M_{k,1}^{-1}} \right\|_F^2 \end{aligned}$$

where we apply the definition of the HS-norm Eq. (6), the definition of the Frobenius norm, and multiply out the terms.

Now, we can use the definition of the Kronecker product for matrices E, F, G :

$$\text{vec}(EFG) = (G^T \otimes E) \text{vec}(F)$$

to expand and rearrange this into the form $\|\zeta x\|_F$ for a matrix ζ and vector $x := \text{vec}(C)$:

$$\begin{aligned} & \left\| \sqrt{M_{k,2}}(C\Lambda_1 - \Lambda_2 C) \sqrt{M_{k,1}^{-1}} \right\|_F^2 \\ &= \left\| ((\Lambda_1 \sqrt{M_{k,1}^{-1}}) \otimes \sqrt{M_{k,2}} - \sqrt{M_{k,1}^{-1}} \otimes (\sqrt{M_{k,2}}\Lambda_2)) \text{vec}(C) \right\|_F^2 \end{aligned}$$

\square

B.2. Solving the Combined Optimization Problem

To solve $E(C)$ for a vectorized functional map C , E_{data} must be expanded similarly. Using Appendix B.1 we have:

$$\begin{aligned} & \left\| \sqrt{M_{k,2}}(CD_{\Psi_1} - D_{\Psi_2}) \right\|_F \\ &= \left\| \text{vec}(\sqrt{M_{k,2}}CD_{\Psi_1}) - \text{vec}(\sqrt{M_{k,2}}D_{\Psi_2}) \right\|_2 \\ &= \left\| ((\sqrt{M_{k,2}}D_{\Psi_1})^T \otimes I) \text{vec}(C) - \text{vec}(\sqrt{M_{k,2}}D_{\Psi_2}) \right\|_2 \end{aligned}$$

Where we use the fact that the Frobenius norm of a matrix is just the L_2 norm of its stacked column vectors and the definition of the Kronecker product. Combining the expanded forms of E_{data} and E_{reg} and observing the first variation of $E(C)$ yields a $k^2 \times k^2$ linear system which can be solved for C :

$$(A^T A + \lambda \zeta^T \zeta) \text{vec}(C) - A^T \text{vec}(B) = 0$$

Here, we made the following substitutions for readability:

$$\begin{aligned} A &= (\sqrt{M_{k,2}}D_{\Psi_1})^T \otimes I \\ B &= \sqrt{M_{k,2}}D_{\Psi_2} \\ \zeta &= (\Lambda_1 \sqrt{M_{k,1}^{-1}}) \otimes \sqrt{M_{k,2}} - \sqrt{M_{k,1}^{-1}} \otimes (\sqrt{M_{k,2}}\Lambda_2) \end{aligned}$$

C. Hybrid Formulation

C.1. Basis Non-orthogonality




Figure 8. Matrix of inner product between hybrid basis $M_{k_{\text{total}}} = [\Phi\Psi]^T M [\Phi\Psi]$, with $k_{\text{total}} = 30 + 30$. We show the heatmap of the resulting matrix from three animal shapes of the SMAL dataset.

To better understand the non-orthogonality between the bases, we study the inner product matrix of the hybrid basis induced by the mass M on the shape in Fig. 8. The matrix exhibits block form and is defined as the mass matrix in the reduced hybrid basis:

$$M_{k_{\text{total}}} = \begin{bmatrix} I & M^{12} \\ M^{21} & M_{k_{\text{Elast}}} \end{bmatrix} \quad (7)$$

The top-left block clearly depicts an identity matrix as we know the LBO eigenfunctions are orthogonal. The bottom-right block corresponds to the spectral mass matrix $M_{k_{\text{Elast}}}$ for the elastic basis [21]. Deviation from the identity matrix is expected as the elastic basis is non-orthogonal. Inter-basis

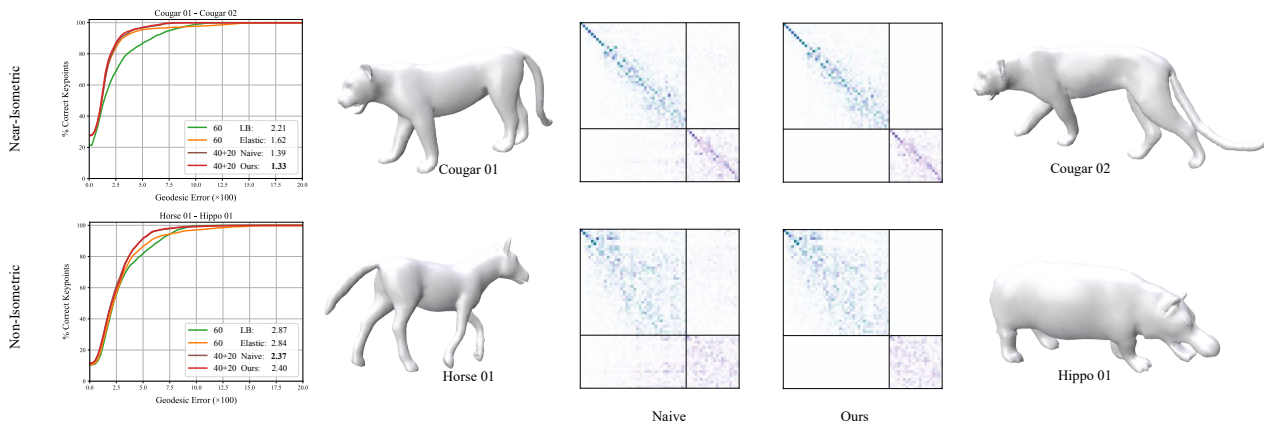


Figure 9. Point-to-point Map recovery accuracy from ground truth hybrid functional maps without any assumptions (Naive) and with zero cross region assumptions (Ours) between two representative pairs of shapes from the SMAL dataset, showcasing both near-isometric and non-isometric scenarios. We use $k = 40 + 20$ as the spectral resolution.

regions (M^{12} and M^{21}) have non-zero entries, indicating e.g. the first 30 LB bases and the first 30 Elastic bases are mutually non-orthogonal. However, we observe that the cross region blocks are sparse. Approximately 90% of the cross region blocks have values below 0.1.

C.2. Starting Assumptions

In Sec. 4.2, we make the assumption that the off-diagonal blocks of both the hybrid functional map and reduced hybrid mass matrix $M_{k_{\text{total}}}$ contain no off-diagonal blocks to separate the hybrid Fmap optimization. To see that such assumptions are plausible, we conduct an experiment on how well a functional map can represent an underlying ground truth correspondence by recovering the point-to-point map from a hybrid functional map via standard nearest neighbor search. This experiment is carried out both without any assumptions (Naive) and with the assumptions of zero off-diagonal blocks (Ours). Results can be seen in Fig. 9. As can be observed, even though the optimal naive map exhibits "leakage", with mild assumptions and less parameters one can solve for a block-diagonal map representing the same level of accuracy, outperforming the baseline LB and Elastic Basis in overall detail alignment and coarse global alignments.

C.3. Optimization Block Matrix Formulation

As demonstrate in Fig. 9, we observe that inter-basis matchings in a hybrid map do not improve performance. In addition, such a map is harder to regularize (Appendix F.1) and more expensive to compute. We, therefore, impose the constraint $C^{21} = C^{12} = 0$ and show that this is equivalent to solving the optimization problems in Eq. (2) separately.

Theorem C.1. *Let the off-diagonal blocks in the hybrid functional map Eq. (4) and the mass matrix in the reduced hybrid basis Eq. (7) be zero, i.e., there are no inter-basis matchings, and the bases are considered as if mutually orthogonal. In hybrid function space, the energy in Eq. (2) can*

then be equivalently formulated as two separate optimization problems:

$$C_*^{11} = \arg \min_C E_{LB}(C) \quad C_*^{22} = \arg \min_C E_{Elas}(C)$$

$$C_* = \begin{pmatrix} C_*^{11} & \mathbf{0} \\ \mathbf{0} & C_*^{22} \end{pmatrix}$$

Proof of Thm. C.1. In the following we use k_{total} as the total basis size, k_{LB} as LB basis size, k_{Elas} as the non-orthogonal elastic basis size. The block matrix representation of the functional map C in the hybrid vector space is given by

$$C = \begin{pmatrix} C^{11} & \mathbf{0} \\ \mathbf{0} & C^{22} \end{pmatrix},$$

where C^{11} and C^{22} represent the functional maps within the same basis types, and off diagonal blocks are fixed to zero as per our starting assumptions.

By our second assumption, the spectral mass matrix in the reduced hybrid basis assumes zero off diagonal blocks and is similarly given by

$$M_{k_{\text{total}}} = \begin{bmatrix} I & \mathbf{0} \\ \mathbf{0} & M_{k_{Elas}} \end{bmatrix}$$

We denote the hybrid basis as $\Theta_i := [\Phi_i \Psi_i]$. The orthogonal projector operator [21] of the hybrid basis, is given by $\Theta_i^\dagger := M_{k_{\text{total},i}}^{-1} \Theta_i^T M_i$. Therefore the orthogonal projector can be written to block form:

$$\Theta_i^\dagger = \begin{bmatrix} \Phi_i^T M_i \\ M_{k_{Elas,i}}^{-1} \Psi_i^T M_i \end{bmatrix} = \begin{bmatrix} \Phi_i^\dagger \\ \Psi_i^\dagger \end{bmatrix}$$

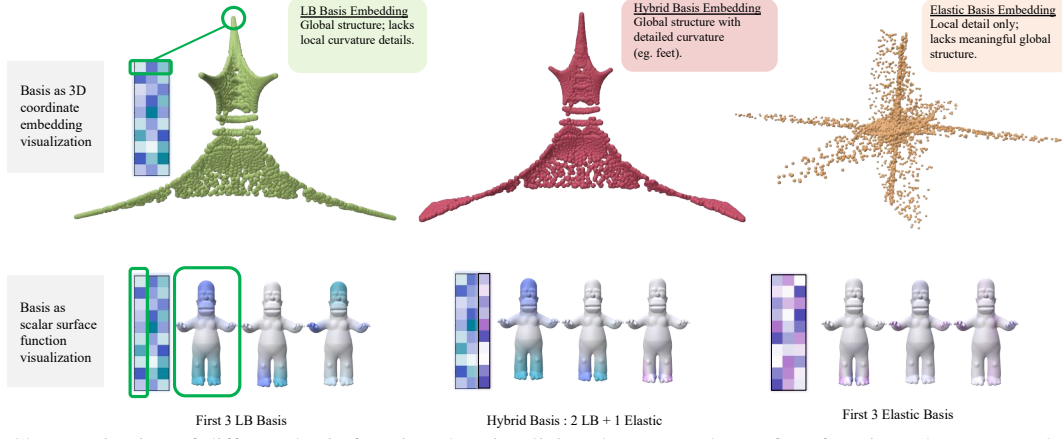


Figure 10. Characterization of different basis functions by visualizing them as scalar surface functions (bottom), and as 3d coordinate embeddings (top). While the elastic basis lacks global structure, it naturally extends the smooth approximation of the LBO basis to incorporate geometric details (creases) when hybridized.

Then we have for the descriptor projecting to the hybrid basis as:

$$D_{\Theta_i} := \Theta_i^\dagger D_i = \begin{bmatrix} \Phi_i^\dagger D_i \\ \Psi_i^\dagger D_i \end{bmatrix} = \begin{bmatrix} D_{\Phi_i} \\ D_{\Psi_i} \end{bmatrix}$$

We let $\Sigma_i := \text{diag}(\lambda_1, \dots, \lambda_{k_{LB}}, \gamma_1, \dots, \gamma_k)$ the diagonal matrix of combined eigenvalues from Δ_i and $\text{Hess}\mathcal{W}_S[\text{Id}]$, respectively. Then, both the data and regularization terms in Eq. (2) can be expanded:

$$E(C) = \|CD_{\Theta_1} - D_{\Theta_2}\|_{M_{k_{\text{total},2}}}^2 + \lambda \|C\Sigma_1 - \Sigma_2 C\|_{HS}^2$$

We can express the data term in the block matrix format:

$$\begin{aligned} E_{\text{data}}(C) &= \|CD_{\Theta_1} - D_{\Theta_2}\|_{M_{k_{\text{total},2}}}^2 \\ &= \left\| \begin{bmatrix} C^{11} D_{\Phi_1} - D_{\Phi_2} \\ \sqrt{M_{k_{\text{Elas},2}}} (C^{22} D_{\Psi_1} - D_{\Psi_2}) \end{bmatrix} \right\|_F^2 \end{aligned}$$

Due to the additivity of the Frobenius norm, the two terms can be minimized separately. A similar condition holds for the regularization term as:

$$\begin{aligned} E_{\text{reg}}(C) &= \|C\Sigma_1 - \Sigma_2 C\|_{HS}^2 \\ &= \left\| \begin{bmatrix} C^{11} \Lambda_1 - \Lambda_2 C^{11} \\ \sqrt{M_{k_{\text{Elas},2}}} (C^{22} \Gamma_1 - \Gamma_2 C^{22}) \sqrt{M_{k_{\text{Elas},1}}}^{-1} \end{bmatrix} \right\|_F^2 \end{aligned}$$

Now, the optimization problem decouples into two separate problems, one for each basis type. These can be solved independently to obtain the optimal functional maps C_{11}^* and C_{22}^* within the LBO and elastic bases, respectively.

$$C_{*}^{11} = \arg \min_{C^{11}} E_{\text{LB}}(C^{11})$$

$$C_{*}^{22} = \arg \min_{C^{22}} E_{\text{Elas}}(C^{22})$$

□

Estimating a hybrid functional map in this manner is an effective regularization which aids computational efficiency, however the two bases are not separated everywhere. The notion of "hybrid" is conceptually important and essential for the rest of optimization stages, including point map conversion and the training of neural network feature extractors as seen in Appendix F.

C.4. Point-to-Point Map Conversion

In order to obtain a point-to-point map P from a general functional map C between non-orthogonal bases, the following minimization objective is considered [21]:

$$\begin{aligned} \min_{P \in \{0,1\}^{m \times n}} & \|CM_{k,1}^{-1} \Psi_1^T - M_{k,2}^{-1} \Psi_2^T P^T\|_{M_{k,2}}^2 \\ \text{s.t. } & P^T \mathbf{1} = \mathbf{1} \end{aligned}$$

This formulation can be solved efficiently by considering the nearest neighbor in $\sqrt{M_{k,2}} CM_{k,1}^{-1} \Psi_1^T$ for every point in $\sqrt{M_{k,2}^{-1}} \Psi_2^T$. For a hybrid functional map, the above objective cannot be decoupled, and the search yields a single point-to-point map based on the embeddings in the hybrid space, as illustrated in Fig. 10. Note that the block structure assumption can still be taken advantage of for efficient matrix multiplications prior to nearest neighbor search.

D. Implementation Details

D.1. Datasets

We evaluate our method across near-isometric, non-isometric, and topologically noisy settings. Splits are chosen




Figure 11. Comparison of correspondence visualizations by transferring vertex positions from the source (horse) to the target shape (hippo). We examine the results from LB, Elastic, and Hybrid basis functions by encoding and recovering **ground truth** point-to-point correspondences through functional map representations at a spectral resolution of $k = 60$ and $k = 200$. Notably, the hybrid functional map representation can encode reasonable correspondences at a resolution of 60, while the other two methods struggle even at a resolution of 200.

based on standard practices in the recent literature [12, 17].

Near-isometric: The FAUST, SCAPE, and SHREC’19 datasets represent near-isometric deformations of humans, with 100, 71, and 44 subjects, respectively. We follow the standard train/test splits for FAUST and SCAPE: : 80/20 for FAUST and 51/20 for SCAPE. Evaluation of our method on SHREC’19 is conducted with a model trained on a combination of FAUST and SCAPE inline with recent methods [12, 17, 27]. We use the more challenging re-meshed versions as in recent works.

Non-isometric: The SMAL dataset features non-isometric deformations between 49 four-legged animal shapes from eight classes. The dataset is split 5/3 by animal category as in Donati et al. [17], resulting in a train/test split of 29/20 shapes. We further evaluate the large animation dataset DeformingThings4D (DT4D-H) [28], using the same inter- and intra-category splits as Donati et al. [17].

Topological Noise: The TOPKIDS dataset [31] consists of shapes of children featuring significant topological variations and poses a significant challenge for unsupervised functional map-based works. Considering its limited size of 26 shapes, we restrict our comparisons to axiomatic and unsupervised methods and use shape 0 as a reference for matching with the other 25 shapes, following recent methods [12, 17, 18].

D.2. Experimental Details

In this section we provide additional details regarding the evaluation of our proposed hybrid basis from Sec. 5, including the axiomatic, supervised, and unsupervised settings. Unless otherwise mentioned, implementations and parameters are left unaltered for the hybrid adaptation.

We first provide general details regarding learning and then the individual adaptations for each method. Learned methods (GeomFMaps [15] and ULRSSM [12]) are trained with PyTorch, using DiffusionNet as the feature extractor and WKS descriptors as input features, except for the SMAL dataset where we use XYZ signal with augmented random rotation as in recent methods [12, 27]. The dimension of the output features is fixed at 256 for all experiments.

For unsupervised learned methods, we propose the following linear annealing scheme for learning in a hybrid basis, as mentioned in Sec. 4.3.

$$\mathcal{L}_{\text{total}} = \alpha \mathcal{L}_{\text{LB}} + \mu \beta \mathcal{L}_{\text{Elas}}$$

$$\alpha = \frac{1}{2} \cdot \frac{k^2}{(k_{\text{LB}})^2} \quad \beta = \frac{1}{2} \cdot \frac{k^2}{(k_{\text{Elas}})^2}$$

Where k is the total spectral resolution. The parameters α and β ensure the losses are normalized w.r.t. the number of entries in the functional map similar to the approach of Li et al. [27]. We increase μ over the first 2000 iterations so that the less-robust elastic basis functions do not adversely affect feature initialization.

In the following, C represents the block-functional map’s elastic part for clarity. Without loss of generality we let $C_{12} : \mathcal{F}(S_1) \rightarrow \mathcal{F}(S_2)$.

Hybrid GeomFMaps. We use 30 total eigenfunctions as in the original work. For the hybrid adaptation, 20 LBO and 10 Elastic basis functions are used as the spectral resolution. To compute the functional map, we use the standard regularized functional map solver and set $\lambda = 1 \times 10^{-3}$ as in the original work [15]. For the hybrid adaptation, we empirically set $\lambda = 5 \times 10^{-4}$ for the elastic solver.

GeomFMaps is supervised using a functional map constructed from the ground-truth correspondences. We thus

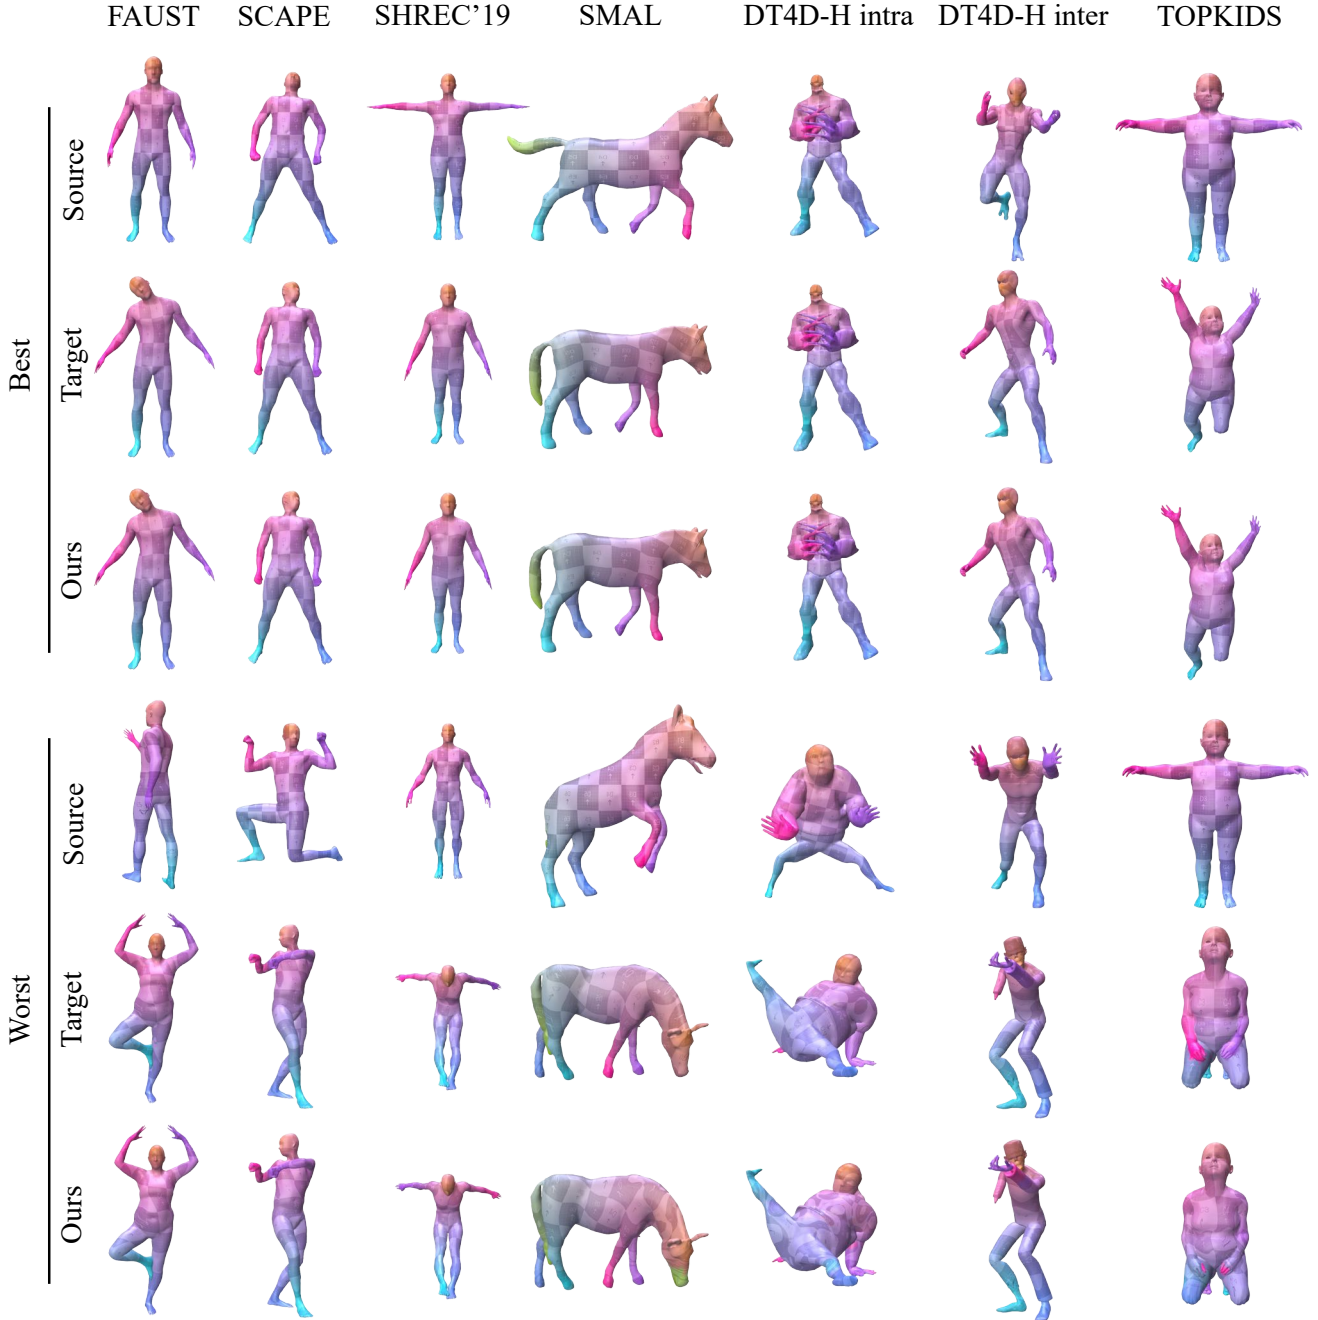


Figure 12. **Additional qualitative results** for the best and worst predictions of ULRSSM in the proposed hybrid basis (our method).

adapt the elastic loss as follows (note here C refers to $\mathcal{F}(S_1) \rightarrow \mathcal{F}(S_2)$ as the original work [15]):

$$\mathcal{L}_{\text{Elas}} = \|C - C_{\text{gt}}\|_{\text{HS}}^2 = \|\sqrt{M_{k,2}}(C - C_{\text{gt}})\sqrt{M_{k,1}^{-1}}\|_F^2$$

Hybrid ULRSSM. The ULRSSM baseline [12] uses a spectral resolution of $k = 200$. We keep the total spectral resolution fixed at $k = 200$, using 140 LBO and 60 Elastic eigenfunctions.

For the functional map computation, we use the Resolvent regularized functional map solver [43] for LB map block setting $\lambda = 100$ as in the original work. Our adapted

variant is weighted empirically with $\lambda = 50$ for the elastic block. ULRSSM regularizes the functional map obtained from Eq. (2) with 3 losses: bijectivity, orthogonality, and a coupling loss with the point-to-point map. The loss for the LBO functional map block \mathcal{L}_{LB} is kept the same as the baseline method while we adapt the bijectivity, orthogonality, and coupling terms for the elastic block in the HS-norm.

While the bijectivity loss is left unchanged, we adapt the $\mathcal{L}_{\text{orth}}$ term using the adjoint C^* as follows, similar to Hartwig et al. in their adapted ZoomOut [21]:

$$\begin{aligned}\mathcal{L}_{\text{orth}} &= \|C_{12}^* C_{12} - I\|_{\text{HS}}^2 + \|C_{21}^* C_{21} - I\|_{\text{HS}}^2 \\ &= \|C_{21}^* C_{21} - I\|_F^2 + \|C_{12}^* C_{12} - I\|_F^2\end{aligned}$$

We note that concerning the respective bijectivity and orthogonality losses, the operators $C_{12}C_{21} - I$ and $C_{21}^*C_{21} - I$ map to and from the same function space, thus the HS-norm is equivalent to the standard Frobenius norm and requires no non-uniform weighting.

The $\mathcal{L}_{\text{couple}}$ term is given by:

$$\begin{aligned}\mathcal{L}_{\text{couple}} &= \left\| C_{12} - \Psi_2^\dagger \Pi_{21} \Psi_1 \right\|_{\text{HS}}^2 + \left\| C_{21} - \Psi_1^\dagger \Pi_{12} \Psi_2 \right\|_{\text{HS}}^2 \\ &= \left\| \sqrt{M_{k,2}} (C_{12} - \Psi_2^\dagger \Pi_{21} \Psi_1) \sqrt{M_{k,1}^{-1}} \right\|_F^2 \\ &\quad + \left\| \sqrt{M_{k,1}} (C_{21} - \Psi_1^\dagger \Pi_{12} \Psi_2) \sqrt{M_{k,2}^{-1}} \right\|_F^2\end{aligned}$$

For the definition of the point-to-point maps Π_{21} and Π_{12} we refer readers to the original method [12].

Empirically, we set $\lambda_{\text{bij}} = \lambda_{\text{orth}} = \lambda_{\text{couple}} = 1.0$ for the LB part following the original work. We keep these parameters the same for the elastic block except setting $\lambda_{\text{orth}} = 0.0$ as we observed the orthogonality constraint adversely affects the method’s performance.

Hybrid SmoothShells. To demonstrate how the proposed hybrid basis can be used in an axiomatic method, we adapt the method SmoothShells [18] with the minimally needed changes.

The initialization of SmoothShells consists of a low-frequency MCMC alignment. We keep this step as-is and do not replace the LBO smoothing with the hybrid eigenfunctions because the elastic eigenfunctions cannot achieve low-frequency smoothing by design. We fix the random seed and re-run the baseline Smooth Shells and the hybrid version with the same MCMC initialization to rule out noise.

The main idea of Smooth Shells [18] is to achieve a coarse-to-fine alignment by iteratively adding higher-frequency LBO eigenfunctions to the intrinsic-extrinsic embedding. Instead of only adding LBO eigenfunctions in a new iteration, we add a ratio of LBO and elastic eigenfunctions.

As in the other adaptations, we keep the total number of basis functions $k = 500$ fixed. We then empirically replace the highest 200 LBO eigenfunctions with elastic eigenfunctions, modifying the product embedding to be:

$$\begin{aligned}\mathbf{X}_k &:= \left(\Phi_{1,k}, \Psi_{1,k}, X_k, \mathbf{n}_k^{S_1} \right) \in \mathbb{R}^{n_1 \times (k+6)} \\ \mathbf{Y}_k &:= \left(\Phi_{2,k}, \Psi_{2,k}, Y_k, \mathbf{n}_k^{S_2} \right) \in \mathbb{R}^{n_2 \times (k+6)}\end{aligned}$$

where we use our notation of the LBO and elastic basis functions. $\mathbf{X}_k, \mathbf{Y}_k$ are the product embeddings for shape S_1 and S_2 with X_k, Y_k the respective smoothed cartesian coordinates, and \mathbf{n}_k the outer normals on each shape. The rest of the optimization follows directly from [18].

E. Additional Experiments

E.1. Near-isometric Cross-dataset Generalisation

We provide additional experimental results in near-isometric cross-dataset scenarios on FAUST, SCAPE, and SHREC’19 datasets using ULRSSM [12] and our hybrid adaptation in Tab. 4. Using the hybrid basis also improves performance under these settings.

E.2. Hybrid ZoomOut Refinement

We provide experimental results for ZoomOut (see Tab. 5) when refining a Hybrid GeomFMaps correspondence prediction to $k = 200$ in both near-isometric (FAUST) and non-isometric (SMAL) settings. It is worth noting that the performance of the hybrid ZoomOut algorithm relies on good initialization, similar to the original method [21, 36], and the performance can vary significantly given a poor initialization.

F. Further Ablation Studies

In this section, we present the results of several ablation studies, focusing on key design choices: separating the optimization problems from Thm. C.1, the ratio of the hybrid basis, and our training strategies.

F.1. Hybridization and Separation.

We argue that the proposed hybrid notion is conceptually important; it is essential for optimization and training stages, including point map conversion and training learned features. Preventing inter-basis matchings by setting the off-diagonal blocks of a hybrid functional map to 0 serves as a strong regularization, which is also computationally efficient (see Fig. 13). The proposed hybrid formulation generally raises questions on the perspective of *hybridization* and which components of a hybrid FM architecture could be optimized

Table 4. **Additional Experimental Results of Hybrid ULRSSM compared to ULRSSM [12] in cross-dataset near-isometric scenario on FAUST, SCAPE, and SHREC’19 datasets.** Our hybrid adaptation generalizes across datasets and also shows improved performance under such settings. ULRSSM results are reproduced from the official repository.

Train	FAUST			SCAPE			FAUST + SCAPE		
	FAUST	SCAPE	SHREC’19	FAUST	SCAPE	SHREC’19	FAUST	SCAPE	SHREC’19
ULRSSM [12]	1.6	2.2	6.7	1.6	1.9	5.7	1.6	2.1	4.6
Hybrid ULRSSM (Ours)	1.5	2.1	5.5	1.5	1.8	5.4	1.5	2.0	3.6

separately. To further support our design decisions, we provide ablation experiments for this regularization effect and analyze why training two entirely separate networks is not optimal.

Separating the Optimization As detailed in Thm. C.1, setting the off-diagonal blocks of the hybrid functional map to $\mathbf{0}$ is equivalent to solving the optimization problems separately. Here, we show this is both important computationally and in terms of regularization.

To demonstrate its regularization effect, we conduct an experiment in ULRSSM on the FAUST dataset comparing solving a hybrid functional map with our proposed method against naively via a single solve. This experiment is conducted with the orthogonalized elastic basis as a proof-of-concept, as solving a full-dimensional ($k = 200$) hybrid functional map from the $k^2 \times k^2$ system would be prohibitively expensive.

Results. The results of this ablation are depicted in Fig. 13. We observe that solving the two maps separately yields notably faster convergence compared to the naive approach with a marginal performance advantage. This suggests that a block-diagonal functional map is desirable; restricting inter-basis matches leads to faster convergence. Separately solving the optimization problems can be interpreted as a strong regularization of the off-diagonal blocks, reducing the search space.

Geo.err. (x100)		FAUST	SMAL
200	LB	2.1	7.6
200	Elastic	2.0	5.9
100+100	Hybrid	1.6	5.1

Table 5. Results of applying the proposed hybrid basis to ZoomOut [21, 36], compared to pure LB and pure elastic bases. We initialize each method with a prediction from Hybrid GeomFmaps (from Tab. 1) and carry out the spectral upsampling from $k = 30$ to $k = 200$, particularly, we upsample from $k = 20 + 10$ to $100 + 100$ for Hybrid ZoomOut. We demonstrate results on both near-isometric (FAUST) and non-isometric (SMAL) datasets.




Figure 13. Ablation study concerning separate optimization of the block matrix on FAUST with Hybrid ULRSSM. The y-axis depicts validation error, while the x-axis shows training steps. Separating the optimization problems primarily leads to faster convergence.

Separating the Networks. During the training of a deep hybrid FM, features are updated by gradients passed through both map solves. Tab. 6 includes an additional experiment where we instead train two entirely separated feature networks, only concatenating features during inference to obtain correspondences. The performance falls behind our approach, showing that we separate only where meaningful.

Geo.err.(x100)	Hybrid ULRSSM	Hybrid GeomFmaps
Separate Network	3.9	10.9
Ours	3.3	7.6

Table 6. Ablation study on whether separating neural networks for feature training is effective to our proposed framework. Instead of using one shared weight neural network as feature extractor, consider training two completely separate networks and concatenate the features only at test time. Note that under this scheme, our annealing training strategies in Appendix F.3 are no longer meaningful. Experiment conducted on SMAL dataset with our two adapted learning methods. Best result is reported.

F.2. Basis Ratio

To further validate the effectiveness of our choice of hybridizing between the LB (Laplace-Beltrami) and Elastic eigenfunctions, we conduct extensive ablation experiments showcasing the performance of different ratios of hybridized basis.

In all these experiments, we again fix the total number of basis functions used as k while replacing the highest frequency LB basis with the Elastic eigenfunctions corresponding to the smallest eigenvalues. This follows the intuition that the low-frequency LB basis functions enable coarse shape alignment while failing to capture fine details, while optimizing in the hybrid basis enables alignment to thin structures and high curvature details better than in the pure LBO basis. We conduct two ablations to demonstrate such a choice; both experiments were carried out on the SMAL dataset for its challenging non-isometry and practical relevance as a stress test.

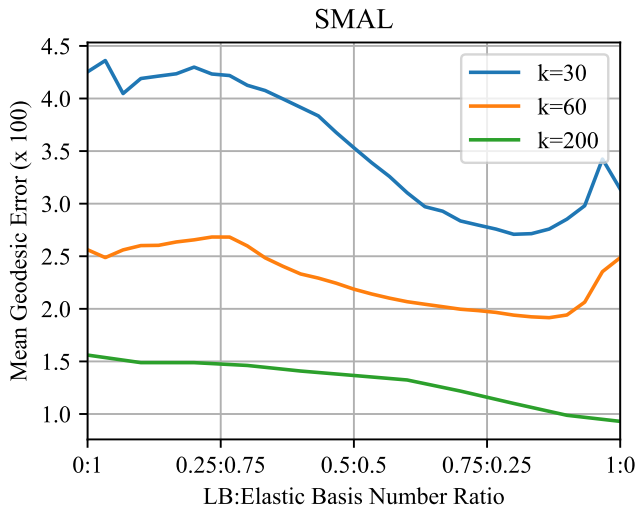


Figure 14. Point-to-point correspondence recovery in a hybrid basis, using a varying ratio of LBO and elastic eigenfunctions. We present several different total spectral resolutions k . Units are in geodesic distance ($\times 100$).

Point-to-Point Map Recovery For ground-truth map recovery, a functional map is obtained by projecting ground-truth point-to-point correspondences into the spectral domain. Subsequently, the point-to-point correspondences are reconstructed using a nearest neighbor search, upon which the discrepancy with the ground-truth point-to-point map can be measured by geodesic error. This simple experimental scenario enables a convenient way to measure the expressiveness of a functional map; In Fig. 14, we consider the hybrid basis composed with a varying ratio of LBO and

elastic basis functions and a different number of total basis functions: $k = 30, 60,$ and 200 . We measure the mean geodesic error between the ground-truth point-to-point map and the recovery from the hybrid functional maps.

Results. The hybridized basis can notably better represent the ground truth for $k = 30$ and 60 . We observe an optimum of around 80% LBO and 20% elastic eigenfunctions. This phenomenon diminishes at $k = 200$, suggesting the LB basis functions can indeed represent fine details with a sufficiently high number of basis functions. However, ground-truth recovery does not necessarily represent the setting where features are learned through backpropagation of the functional map loss. Our experiments indicate that learned pipelines cannot leverage the high-frequency LBO eigenfunctions to represent fine extrinsic details as effectively as the elastic basis functions, even with a large total number of basis functions. We, therefore, conduct a similar ablation in the learned setting.

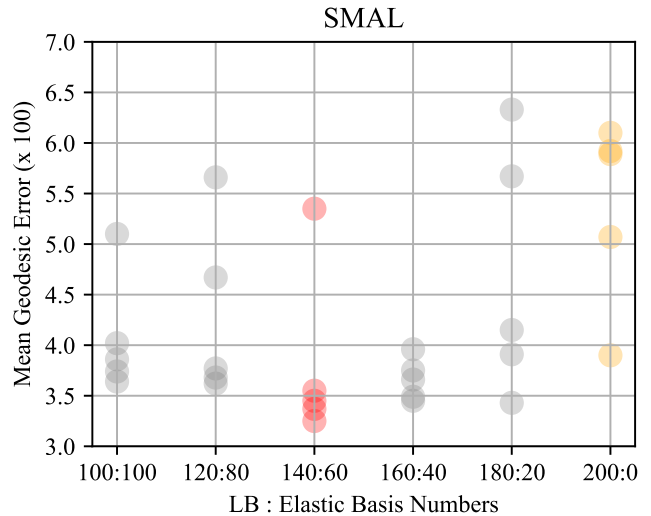


Figure 15. Hybrid ULRSSM is evaluated on SMAL using a varying ratio of LBO and elastic eigenfunctions. Our basis ratio can be observed in red, with the baseline pure LBO implementation in yellow.

Learned Setting. In Fig. 15, we consider the hybrid ULRSSM method with a fixing total basis number of $k = 200$ under different hybrid basis ratios with a step size of 20. Due to the high-order polynomial computational increase (see Appendix G) for the $k^2 \times k^2$ system, we limit the number of elastic basis functions to less than 100 (Point Map Recovery from GT also indicates inferior performance outside of this regime). We run each experiment 5 times to eliminate inherent noise and report all results.

Results. Here, we observe that a ratio of around 140:60 is optimal; the hybridized basis (red) shows consistent perfor-





Figure 16. The computation cost comparison of deep functional map solvers between classic, generalized, and hybrid solutions: runtime per solve (left), and VRAM consumption (right). Results were obtained in PyTorch on an NVIDIA A40.

mance improvements over the baseline (orange) and other basis ratios.

F.3. Training Strategies

Training a reliable shape correspondence estimation pipeline through hybrid functional maps involves several key modeling decisions. Both the linearly increasing scheduler for the elastic loss during training and normalizing factors for both Laplace-Beltrami (LB) and elastic losses play a large role in the obtained performance increases.

As mentioned in Sec. 4, we observed the elastic basis functions are not robust to uninitialized features. Easing in the elastic loss after feature initialization in the LBO basis mitigates convergence to undesirable local minima. Furthermore, the loss of each component in the hybrid functional map is normalized according to the number of matrix elements for this component, an important hyperparameter to balance the two blocks. During the ablation studies presented in Fig. 17, we selectively eliminate each one of these factors from our model and measure the mean geodesic error. We further demonstrate that fine-tuning from a pre-trained LBO checkpoint is ineffective, likely converging to local minima. **Results.** The results in Fig. 17 show that each component is indeed important for our final model. Fine-tuning from a checkpoint or training without normalization yields inferior results. Furthermore, except for a single outlier, our approach converges to a significantly lower minimum than learning without the linear-annealing strategy.

Method	Runtime
ULRSSM	610.70 ± 45.20 ms
Hybrid ULRSSM	623.19 ± 32.03 ms

Table 7. Comparison of per-iteration runtime on SMAL over 100 training iterations (mean ± st. dev).




Figure 17. Ablation study of optimization strategies using Hybrid ULRSSM on SMAL. Five random runs are shown for each training setting.

G. Runtime Analysis

We provide our runtime analysis for the Hybrid ULRSSM method in SMAL dataset under Table 7 and Fig. 16. Results are obtained on an NVIDIA A40. Our hybrid adaptation incurs minimal runtime overhead while yielding significant performance gains despite needing to solve an expanded $k^2 \times k^2$ system. This can be explained by analyzing the complexity. Assuming the complexity of solving a linear system for an $k \times k$ matrix is $\mathcal{O}(k^3)$, solving the combined optimization problem costs $\mathcal{O}(k^4)$ flops (since we solve k separate $k \times k$ systems). Solving the separate optimization problem costs $\mathcal{O}((k-l)^4 + l^6)$ flops, which for the total spectral resolution $k = 200$ and elastic eigenfunctions $l = 60$ is only one order of magnitude larger.

H. Additional Figures

We provide additional visualization of the basis embeddings in Fig. 10. The figure depicts 3D embeddings of the first three basis functions on the shape in different combinations. The visualization suggests that while the elastic basis lacks global structure, it naturally extends the smooth approximation of the LBO basis to incorporate geometric details (creases) when hybridized.

Additionally, we provide further comparison of correspondence visualizations by transferring vertex positions from the source (horse) to the target shape (hippo), the results of LB, Elastic, and Hybrid basis functions were compared by encoding and recovering ground truth correspondences through functional map representations. As illustrated in Fig. 11, the hybrid functional map representation notably encodes reasonable correspondences already at a spectral resolution of 60, whereas the other two methods still face challenges even at a resolution of 200.

I. Additional Qualitative Results

We provide additional qualitative results in Fig. 12 on each dataset. We evaluate and visualize the best and worst predictions of ULRSSM in the proposed hybrid basis.