

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

Supplementary Material

In the supplementary materials we first provide additional mathematical background in Sec. 1. Sec. 2 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 Sec. 3. We then detail datasets and splits used for evaluation, followed by further experimental details in Sec. 4. Additional experiment results is provided at Sec. 5. Ablation studies concerning our design choices are provided in Sec. 6. Finally, we present runtime analysis in Sec. 7, additional visualization of the basis embedding in Sec. 8 and additional qualitative results in Sec. 9.

1. Mathematical Background

1.1. Elastic Energy

We defer to Hartwig et al. [6] for a complete definition of the previously described elastic energy [1, 6, 7]. For all our experiments, we use the same elastic energy hyperparameters as Hartwig et al. [6], 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 .

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

1.2. 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.

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 [6, Sec 3.4]:

Table 1. Summary of notations used in this work.

Symbol	Description
$\mathcal{S}_1, \mathcal{S}_2$	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

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

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 [6]:

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

2. Generalization to the Hilbert-Schmidt Norm

2.1. 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 Sec. 1.2, 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.

□

As previously established [3], 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. (1), 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 $x := \text{vec}(C)$:

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

□

2.2. Solving the Combined Optimization Problem

To solve $E(C)$ for a vectorized functional map C , E_{data} must be expanded similarly. Using Sec. 2.1 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:

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

3. Hybrid Formulation

3.1. Basis Non-orthogonality




Figure 1. 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. 1. The matrix exhibits block form and is defined as the mass matrix in the reduced hybrid basis:

$$M_{k_{\text{total}}} = \begin{bmatrix} I & M^{12} \\ M^{21} & M_{k_{\text{Elas}}} \end{bmatrix} \quad (2)$$

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 $M_{k_{\text{Elas}}}$ for the elastic basis [6]. Deviation from the identity matrix is expected as the elastic basis is non-orthogonal. Inter-basis 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.




Figure 2. 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.

3.2. 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. 2. 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.

3.3. Optimization Block Matrix Formulation

As demonstrate in Fig. 2, we observe that inter-basis matchings in a hybrid map do not improve performance. In addition, such a map is harder to regularize (Sec. 6.1) 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 3.1. *Let the off-diagonal blocks in the hybrid functional map Eq. (4) and the mass matrix in the reduced hybrid basis Eq. (2) 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. 3.1. 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 [6] 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:

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

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.

$$\begin{aligned} C_{*}^{11} &= \arg \min_{C^{11}} E_{LB}(C^{11}) \\ C_{*}^{22} &= \arg \min_{C^{22}} E_{\text{Elas}}(C^{22}) \end{aligned}$$

□

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 Sec. 6.

3.4. 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 [6]:

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

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. 3. Note that the block structure assumption can still be taken advantage of for efficient matrix multiplications prior to nearest neighbor search.

4. Implementation Details

4.1. Datasets

We evaluate our method across near-isometric, non-isometric, and topologically noisy settings. Splits are chosen based on standard practices in the recent literature [2, 4].

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 [2, 4, 8]. 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. [4], resulting in a train/test split of 29/20 shapes. We further evaluate the large animation dataset DeformingThings4D (DT4D-H) [9], using the same inter- and intra-category splits as Donati et al. [4].

Topological Noise: The TOPKIDS dataset [10] 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 [2, 4, 5].

4.2. 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 [3] and ULRSSM [2]) are trained with PyTorch, using DiffusionNet as the feature extractor




Figure 3. 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.

and WKS descriptors as input features, except for the SMAL dataset where we use XYZ signal with augmented random rotation as in recent methods [2, 8]. 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. [8]. 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 [3]. 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 adapt the elastic loss as follows (note here C refers to $\mathcal{F}(S_1) \rightarrow \mathcal{F}(S_2)$ as the original work [3]):

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

Hybrid ULRSSM. The ULRSSM baseline [2] 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 [12] 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 [6]:

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

$$= \|C_{21}^* C_{21} - I\|_F^2 + \|C_{12}^* C_{12} - I\|_F^2$$

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:

$$\mathcal{L}_{\text{couple}} = \|C_{12} - \Psi_2^\dagger \Pi_{21} \Psi_1\|_{\text{HS}}^2 + \|C_{21} - \Psi_1^\dagger \Pi_{12} \Psi_2\|_{\text{HS}}^2$$

$$= \left\| \sqrt{M_{k,2}}(C_{12} - \Psi_2^\dagger \Pi_{21} \Psi_1) \sqrt{M_{k,1}^{-1}} \right\|_F^2$$

$$+ \left\| \sqrt{M_{k,1}}(C_{21} - \Psi_1^\dagger \Pi_{12} \Psi_2) \sqrt{M_{k,2}^{-1}} \right\|_F^2$$

For the definition of the point-to-point maps Π_{21} and Π_{12} we refer readers to the original method [2].




Figure 4. **Additional qualitative results** for the best and worst predictions of ULRSSM in the proposed hybrid basis (our method).

Empirically, we set $\lambda_{bij} = \lambda_{orth} = \lambda_{couple} = 1.0$ for the LB part following the original work. We keep these parameters the same for the elastic block except setting $\lambda_{orth} = 0.0$ 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 [5] 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 eigen-

functions 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 [5] 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:

$$\mathbf{X}_k := \left(\Phi_{1,k}, \Psi_{1,k}, X_k, \mathbf{n}_k^{S_1} \right) \in \mathbb{R}^{n_1 \times (k+6)}$$

$$\mathbf{Y}_k := \left(\Phi_{2,k}, \Psi_{2,k}, Y_k, \mathbf{n}_k^{S_2} \right) \in \mathbb{R}^{n_2 \times (k+6)}$$

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 [5].

5. Additional Experiments

5.1. Hybrid ZoomOut Refinement

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 2. Results of applying the proposed hybrid basis to ZoomOut [6, 11], 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.

We provide experimental results for ZoomOut (see Tab. 2) when refining a Hybrid GeomFMaps correspondence prediction to $k = 200$ in both near-isometric (FAUST) and non-isometric (SMAL) settings. It’s worth note that the performance of hybrid ZoomOut algorithm relies on good initialization, similar to the original method [6, 11] and the performance on poor initializations varies largely.

6. Further Ablation Studies

In this section, we present the results of several ablation studies, focusing on key design choices: separating the op-

timization problems from Thm. 3.1, the ratio of the hybrid basis, and our training strategies.

6.1. 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. 5). The proposed hybrid formulation generally raises questions on the perspective of *hybridization* and which components of a hybrid FM architecture could be optimized 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. 3.1, setting the off-diagonal blocks of the hybrid functional map to 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.




Figure 5. 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.

Results. The results of this ablation are depicted in Fig. 5.

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.

Separating the Networks. During the training of a deep hybrid FM, features are updated by gradients passed through both map solves. Tab. 3 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 3. 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 Sec. 6.3 are no longer meaningful. Experiment conducted on SMAL dataset with our two adapted learning methods. Best result is reported.

6.2. 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.

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




Figure 6. 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 ($\times 100$).

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. 6, 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.

Learned Setting. In Fig. 7, 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 Sec. 7) for the $k^2 \times k^2$ system, we limit the number of




Figure 7. 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.

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

6.3. 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. 8, 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. 8 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.




Figure 8. Ablation study of optimization strategies using Hybrid ULRSSM on SMAL. Five random runs are shown for each training setting.

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

Table 4. Comparison of per-iteration runtime on SMAL over 100 training iterations (mean ± st. dev).

7. Runtime Analysis

We provide our runtime analysis for the Hybrid ULRSSM method in SMAL dataset under Table 4. 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.

8. Additional Figure of Basis Embedding

We provide additional visualization of the basis embeddings in Fig. 3. 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.

9. Additional Qualitative Results

We provide additional qualitative results in Fig. 4 on each dataset. We evaluate and visualize the best and worst predictions of ULRSSM in the proposed hybrid basis.

References

- [1] 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. [1](#)
- [2] Dongliang Cao, Paul Roetzer, and Florian Bernard. Unsupervised learning of robust spectral shape matching. *ACM Transactions on Graphics (TOG)*, 42(4):1–15, 2023. [4](#), [5](#)
- [3] 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](#), [4](#), [5](#)
- [4] 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. [4](#)
- [5] 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. [4](#), [6](#), [7](#)
- [6] 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](#), [5](#), [7](#)
- [7] 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. [1](#)
- [8] 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. [4](#), [5](#)
- [9] Yang Li, Hikari Takehara, Takafumi Taketomi, Bo Zheng, and Matthias Nießner. 4DComplete: Non-Rigid Motion Estimation Beyond the Observable Surface. pages 12706–12716, 2021. [4](#)
- [10] Z Lähler, 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. [4](#)
- [11] 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. [7](#)
- [12] Jing Ren, Mikhail Panine, Peter Wonka, and Maks Ovsjanikov. Structured Regularization of Functional Map Computations. *Computer Graphics Forum*, 38(5):39–53, 2019. [5](#)