Chebyshev Differential Flows for Shape Matching and Interpolation with Endpoint Guidance

Jiin Im
Hanyang University
2oij11nn@gmail.com

Sisung Liu
Hanyang University
liuss98@hanyang.ac.kr

Je Hyeong Hong[†]
Hanyang University
jhh37@hanyang.ac.kr

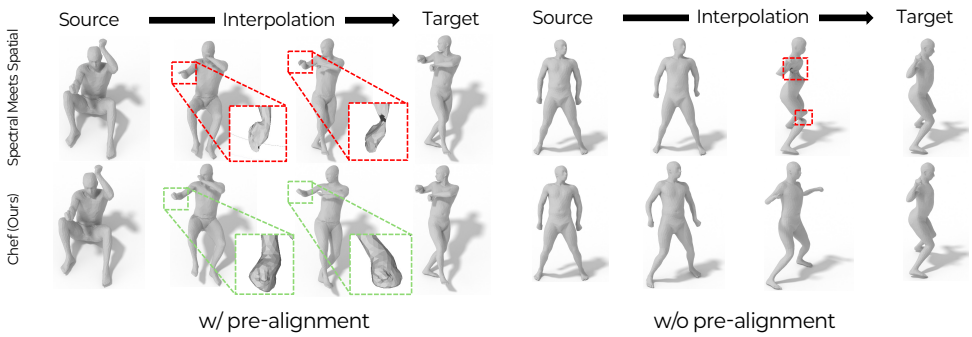


Figure 1: **Joint shape matching and interpolation from endpoint shapes.** CheF estimates an endpoint correspondence and a continuous deformation trajectory from two 3D shapes. Compared with Spectral Meets Spatial under rigidly pre-aligned and raw unaligned inputs, ours qualitatively reduces local interpolation artifacts. Red boxes highlight artifacts in Spectral Meets Spatial, while green boxes show the corresponding regions in ours.

Abstract

Estimating relationships between 3D shapes requires both dense endpoint correspondence and a deformation path connecting the source to the target. These tasks are coupled, as correspondence constrains the interpolation endpoint and a plausible trajectory can help disambiguate the map. However, when only endpoint shapes are observed, the interior trajectory is under-constrained. Existing joint frameworks typically predict sampled vertex displacements and use local rigidity priors that can restrict non-isometric stretch and shear. To address these limitations, we propose *Chebyshev Differential Flows* (CheF), an unsupervised framework that models motion through continuous trajectories of per-face Jacobians. Each Jacobian trajectory is represented by a low-order shifted Chebyshev expansion, yielding a compact deformation model for arbitrary times. The target endpoint depends only on the sum of the odd-order coefficients, while even-order modes shape the interior trajectory without changing the endpoint. A spectral functional-map branch estimates endpoint correspondence, and a differentiable Poisson solver integrates the Jacobian field into globally consistent intermediate shapes. We further introduce a trend regularizer that guides local stretch and shear over time without imposing per-snapshot As-Rigid-As-Possible (ARAP) rigidity. Experiments on near-isometric and non-isometric benchmarks show competitive correspondence accuracy, improved interpolation accuracy and geometric quality, and better stability when rigid pre-alignment is omitted. CheF is publicly available at https://github.com/SpatialAILab/chef

[†]Corresponding author.

1 Introduction

Estimating relationships between 3D shapes is a fundamental problem in computer vision and graphics, with applications in shape modeling [31], animation [26], editing [25] and deformation-aware generation [35]. Joint shape matching and interpolation aim to find both a dense endpoint correspondence and a deformation path describing how the source evolves into the target. These tasks are coupled, as correspondence constrains the deformation endpoint and a plausible trajectory can help disambiguate the map. However, in many practical settings only endpoint shapes are available during training, without ground-truth correspondences or intermediate states.

Recent unsupervised methods [20, 11] address this setting by learning correspondence and deformation jointly. In particular, Spectral Meets Spatial (SMS) combines spectral correspondence estimation with a time-conditioned spatial deformation model to handle both pose and shape changes. These frameworks typically represent motion through per-vertex displacements and use geometric regularization to guide deformation during training. Their interpolation behavior therefore depends on both the deformation representation and the priors used to learn the connecting trajectory.

Even in this joint setting, endpoint correspondence specifies where source points should arrive but leaves the evolution of local deformation under-constrained. Geometric priors help guide this evolution by favoring particular deformation paths. Local rigidity provides a useful prior for near-isometric motion, where deformation is largely explained by changes in pose. However, for non-isometric shape pairs, the connecting motion may require local stretch and shear that rigidity priors penalize. This motivates a representation that makes local deformation explicit and supports temporal guidance while allowing these non-isometric changes.

To address this problem, we propose *Chebyshev Differential Flows* (CheF), an unsupervised framework for joint dense shape matching and continuous shape interpolation from endpoint shape pairs. Our key idea is to represent motion through time-varying per-face Jacobians that explicitly encode local rotation, stretch, and shear. We parameterize each Jacobian trajectory with a low-order shifted Chebyshev expansion. A single coefficient set defines the entire trajectory, coupling its evolution across time and supporting arbitrary-time evaluation. Endpoint correspondence provides local deformation cues for predicting the trajectory, and the resulting Jacobian field is integrated to reconstruct intermediate shapes.

To guide this trajectory, we introduce a trend regularizer on the temporal evolution of local stretch and shear. The regularizer uses an endpoint-conditioned reference path to guide intermediate deformation without imposing per-snapshot As-Rigid-As-Possible (ARAP) rigidity. Experiments on near-isometric and non-isometric benchmarks demonstrate competitive correspondence accuracy and improvements in interpolation accuracy and geometric quality. CheF also shows greater stability when the standard rigid pre-alignment step is omitted.

Our contributions are summarized as follows: (i) an endpoint-conditioned Jacobian-flow formulation for joint shape matching and interpolation in local differential space (ii) a compact shifted Chebyshev trajectory parameterization that supports arbitrary-time evaluation and (iii) a trend regularizer that guides the temporal evolution of local stretch and shear without per-snapshot ARAP rigidity.

2 Related work

2.1 Shape matching

Dense shape matching seeks pointwise correspondences between 3D shapes. Classical methods enforce geometric consistency through combinatorial optimization or non-rigid registration [46, 5], often requiring iterative optimization over spatial correspondences or deformations. Functional maps [37] provide a compact spectral representation by encoding correspondence as a linear operator in Laplace–Beltrami eigenbases. This formulation supports efficient map refinement [38, 34] and learning with geometric structural constraints [14, 10]. Recent approaches improve matching by enriching spectral representations or incorporating spatial and deformation consistency [4, 12, 48]. A complementary direction recovers correspondence by deforming a source shape toward a target [22]. These approaches primarily focus on endpoint correspondence, whereas our work jointly estimates correspondence and a continuous deformation trajectory between the shapes.

2.2 Shape deformation and interpolation

Shape interpolation constructs a deformation path between endpoint shapes whose intermediate states are not uniquely determined. Classical formulations select a path through geometric distortion or shape-space objectives [9, 24, 47]. Local rigidity provides another common prior for interpolation and surface deformation [2, 44, 8], while constrained deformation flows incorporate geometric or physical structure into the motion [13, 16, 17].

Neural Jacobian Fields [1] represent deformation through per-face Jacobians and recover vertex positions through Poisson reconstruction. This representation supports correspondence refinement [45], semantic-guided registration [30], and localized deformation control [6]. Other neural approaches model continuous surface evolution through space-time implicit fields [36] or implicit surfaces coupled with velocity fields [40]. 4Deform [41] extends neural interpolation to partial and non-isometric shapes. These works model deformation through differential representations or continuous surface evolution, whereas our work introduces an endpoint-conditioned metric trend regularizer to guide how local stretch and shear evolve along continuous Jacobian trajectories.

2.3 Joint correspondence and interpolation

Joint approaches exploit the mutual dependence between correspondence and deformation. NeuroMorph [20] jointly learns correspondence and vertex deformation without supervision using geodesic-distance preservation and ARAP regularization. Spectral Meets Spatial [11] combines functional-map regularization with a time-conditioned vertex displacement model to learn spectral correspondence and spatial interpolation. ARC-Flow [23] jointly addresses matching and interpolation through smooth time-varying flow fields with geometric and physical constraints. It additionally uses a source skeleton to guide articulated deformation and resolve correspondence ambiguities.

However, endpoint alignment alone leaves intermediate motion under-constrained, while local rigidity priors can restrict the stretch and shear required for non-isometric interpolation. CheF addresses these challenges through continuous Jacobian trajectories with an endpoint-conditioned metric trend regularizer that guides local stretch and shear over time.

3 Method

Given two triangle meshes $\mathcal{X} = (V_\mathcal{X}, F_\mathcal{X})$ and $\mathcal{Y} = (V_\mathcal{Y}, F_\mathcal{Y})$, where $V_\mathcal{X} \in \mathbb{R}^{n_\mathcal{X} \times 3}$ and $V_\mathcal{Y} \in \mathbb{R}^{n_\mathcal{Y} \times 3}$ denote the coordinates of $n_\mathcal{X}$ and $n_\mathcal{Y}$ vertices, respectively and $F_\mathcal{X}, F_\mathcal{Y}$ contain triangular faces as vertex index triples. We normalize time to the interval $[0, 1]$, designating $\mathcal{X}$ as the source at $t = 0$ and $\mathcal{Y}$ as the target at $t = 1$. Our objective is to estimate both a dense endpoint correspondence and a continuous trajectory $V(t) \in \mathbb{R}^{n_\mathcal{X} \times 3}$ defined on the source topology for arbitrary $t \in [0, 1]$, ensuring $V(0) = V_\mathcal{X}$ and aligning $V(1)$ to $\mathcal{Y}$ through the predicted endpoint map.

The key challenge is that the endpoint map constrains only the final destination of the source mesh, without defining the temporal evolution of local surface elements. To address this, CheF transforms endpoint correspondence evidence into a continuous trajectory of local deformation operators. As shown in Fig. 2, a spectral branch first estimates a soft endpoint map where each row represents a probability distribution over target vertices for a given source vertex. Multiplying this map with target coordinates yields a soft target embedding on the source topology, from which we compute a deformation operator for each face. A coefficient generator then predicts Chebyshev coefficients that parameterize a continuous time field of face Jacobians. Finally, a differentiable Poisson solver integrates these Jacobians into globally consistent vertex positions for any queried time.

3.1 Soft endpoint map and local deformation operator

Following [11], we predict a row-stochastic soft endpoint map $\Pi_{\mathcal{XY}}$ from source to target. Each row represents a probability distribution over target vertices for a source vertex. Given descriptors $D_\mathcal{X}$ and $D_\mathcal{Y}$ extracted by a Siamese DiffusionNet $\mathcal{F}_\theta$ [43], we compute $\Pi_{\mathcal{XY}} = \mathrm{softmax}_{\mathrm{row}}(D_\mathcal{X} D_\mathcal{Y}^\top / \tau)$, where τ is the temperature. The same descriptors are also used to compute bidirectional functional maps. We train this branch with structural and coupling losses. The functional-map solver additionally uses descriptor preservation and Laplacian commutativity. We denote the spectral matching objective by $\mathcal{L}_{\mathrm{spectral}}$ and provide details in Appendix A.9.

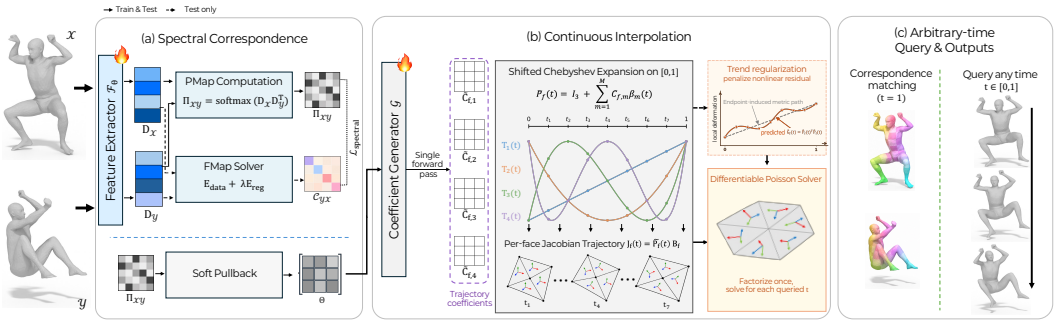


Figure 2: **Overview of CheF.** A spectral branch predicts a soft endpoint map, which induces per-face endpoint deformation cues. A coefficient generator maps these cues to Chebyshev coefficients defining a continuous per-face Jacobian trajectory. A Poisson solver reconstructs the mesh at any queried time, while endpoint and trend losses supervise correspondence and interpolation.

The soft map induces a target embedding on the source topology, $V_{\mathcal{XY}} = \Pi_{\mathcal{XY}} V_{\mathcal{Y}}$. Its i-th row is the expected target position of source vertex $\mathbf{x}_i$. For each source face $f = (i, j, k)$, we construct a translation-invariant local frame $E_f = [\, \mathbf{v}_j - \mathbf{v}_i,\ \mathbf{v}_k - \mathbf{v}_i,\ \mathbf{n}_f \,] \in \mathbb{R}^{3\times 3}$, where $\mathbf{n}_f$ is the face normal. Let $E_{\mathcal{X},f}$ and $E_{\mathcal{XY},f}$ denote the frames computed from $V_{\mathcal{X}}$ and $V_{\mathcal{XY}}$, respectively. We define the local endpoint deformation operator as $\Theta_f = E_{\mathcal{XY},f} E_{\mathcal{X},f}^{-1}$. It satisfies $\Theta_f E_{\mathcal{X},f} = E_{\mathcal{XY},f}$ and provides the input to the coefficient generator.

3.2 Chebyshev Jacobian trajectory

Inspired by Neural Jacobian Field (NJF) [1], we represent deformation through local face differentials. Existing registration frameworks [1, 45, 30, 21] reconstruct a source shape aligned with a target using predicted Jacobians. We model the evolution of these differentials through a time-varying Jacobian field $J_f(t)$. For each face, a single set of coefficients defines the entire trajectory through a low-order shifted Chebyshev expansion.

The coefficient generator $\mathcal{G}$ maps the endpoint deformation cue Θ_f to M ambient coefficient matrices, $\{\tilde{C}_{f,m}\}_{m=1}^{M} = \mathcal{G}(\Theta_f)$. Each coefficient $\tilde{C}_{f,m} \in \mathbb{R}^{3\times 3}$ is shared across all queried times. We define the ambient deformation trajectory as

$$\tilde{P}_f(t) = I_3 + \sum_{m=1}^{M} \tilde{C}_{f,m}\, \beta_m(t), \qquad \beta_m(t) = T_m(2t-1) - (-1)^m, \tag{1}$$

where T_m is the m-th Chebyshev polynomial of the first kind. The predicted coefficients allow this trajectory to be evaluated at any $t \in [0, 1]$. Restricting the ambient transformation to the source tangent basis gives $J_f(t) = \tilde{P}_f(t) B_f \in \mathbb{R}^{3\times 2}$, where $B_f \in \mathbb{R}^{3\times 2}$ is the source face tangent basis. Since $\beta_m(0) = 0$ for every m, the representation satisfies $\tilde{P}_f(0) = I_3$ and $J_f(0) = B_f$.

Poisson integration. Given the predicted Jacobian field at a queried time t, we recover vertex positions via Poisson reconstruction [1, 45, 30, 21]. Since $J_f(t)$ specifies local differentials independently per face, an arbitrary face-level Jacobian field may not correspond to a globally consistent vertex embedding. We therefore recover the embedding whose induced face differentials best match the predicted Jacobians,

$$V(t) = \arg\min_{Z} \sum_{f\in F_{\mathcal{X}}} A_f \left\| Z\nabla_f^\top - J_f(t) \right\|_F^2, \qquad V(t) = L^{-1} A \nabla^\top J(t), \quad L = \nabla^\top A \nabla. \tag{2}$$

A_f is the face area, ∇_f is the discrete face gradient operator, and L is the cotangent Laplacian. Since L depends only on the source mesh, its factorization is reused across all queried time steps, with temporal dependence entering exclusively through $J(t)$. Beyond yielding a continuous trajectory, the shifted Chebyshev parameterization in Eq. 1 also enables a structured separation between endpoint deformation and intermediate trajectory control, which we exploit in the following formulations.

3.3 Endpoint–interior mode separation and trend regularization

Building on the shifted Chebyshev expansion of $\tilde{P}_f(t)$, we now analyze how the endpoint values of the basis determine which coefficients are constrained by endpoint losses and which exclusively govern the intermediate trajectory.

Proposition 1 (Odd–even endpoint separation)**.** *At the target endpoint $t = 1$, only odd Chebyshev modes contribute to the ambient deformation, $\tilde{P}_f(1) = I_3 + 2\sum_{k=1}^{\lceil M/2 \rceil} \tilde{C}_{f,2k-1}$. All even modes vanish at both endpoints.*

This follows from the endpoint values of the shifted basis. At $t = 0$, all modes vanish since $\beta_m(0) = T_m(-1) - (-1)^m = 0$. At $t = 1$, $\beta_m(1) = T_m(1) - (-1)^m = 1 - (-1)^m$, which equals 2 for odd m and 0 for even m. Hence odd coefficients dictate the endpoint deformation and are directly penalized by endpoint matching losses, while even coefficients form latent degrees of freedom that shape the intermediate geometry without altering the established endpoints. Fig. 3 confirms this behavior, omitting even modes preserves the endpoints but significantly degrades the intermediate geometric transitions.

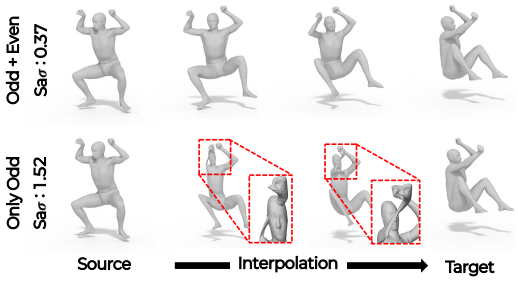


Figure 3: **Odd–even mode separation.** Removing even Chebyshev modes (*Only Odd*) preserves endpoints but degrades interior geometry.

While advantageous for endpoint alignment, this decoupling leaves the intermediate trajectory under-constrained when only endpoint losses are used, since such losses cannot regularize how the latent degrees of freedom behave inside the interval. We therefore introduce a trend loss that guides intermediate evolution in the local deformation space. Rather than regularizing $\tilde{P}_f(t)$ directly, we regularize its induced ambient metric $\tilde{G}_f(t) = \tilde{P}_f(t)^\top \tilde{P}_f(t)$, which factors out the rotational component and guides the evolution of local stretch and shear.

For each face, we define an endpoint-conditioned reference path and its residual,

$$L_f(t) = (1-t)I_3 + t\tilde{G}_f(1), \qquad R_f(t) = \tilde{G}_f(t) - L_f(t). \tag{3}$$

$L_f(t)$ is not intended to impose linear vertex trajectories, but to provide a low-frequency guide for local scale and shear evolution. The trend loss is

$$\mathcal{L}_{\text{trend}} = \frac{1}{\sum_f A_f} \sum_{f \in F_{\mathcal{X}}} A_f \int_0^1 \|R_f(t)\|_F^2 \, dt. \tag{4}$$

The residual vanishes at both endpoints by construction. The loss guides the intermediate evolution of local stretch and shear and is optimized jointly with the endpoint losses. The polynomial representation also permits exact evaluation of this integral using finitely many quadrature nodes.

Proposition 2 (Exact integration of trend energy)**.** *Assume each entry of $\tilde{P}_f(t)$ is a polynomial in t of degree at most M. Then each entry of $\tilde{G}_f(t) = \tilde{P}_f(t)^\top \tilde{P}_f(t)$ is a polynomial in t of degree at most $2M$. Consequently, the trend integrand $h_f(t) = \left\| \tilde{G}_f(t) - \big((1-t)I_3 + t\tilde{G}_f(1)\big) \right\|_F^2$ is a polynomial in t of degree at most $4M$. Gauss–Legendre quadrature with Q nodes therefore evaluates $\int_0^1 h_f(t)\, dt$ exactly whenever $Q \geq 2M + 1$.*

Unlike uniform grid sampling, the trend loss is thus evaluated as a continuous-time polynomial integral over the full interval, providing endpoint-conditioned path guidance for the local strain evolution without invoking a local rigidity prior. The residual form also links endpoint smoothness to the interior metric trajectory. Since

$$\tilde{G}_f(t) = (1-t)I_3 + t\tilde{G}_f(1) + R_f(t), \tag{5}$$

endpoint smoothness of the metric controls the low-frequency spatial variation, while $\mathcal{L}_{\text{trend}}$ controls nonlinear interior deviations in stretch and shear, as formalized below.

Proposition 3 (Endpoint smoothness propagates through trend residuals)**.** *Let $R_f(t)$ be defined as in Eq. 3. For adjacent faces (f, g), define $\mathcal{E}_f = \int_0^1 \|R_f(t)\|_F^2 \, dt$. Then*

$$\int_0^1 \|\tilde{G}_f(t) - \tilde{G}_g(t)\|_F^2 \, dt \leq \frac{2}{3}\|\tilde{G}_f(1) - \tilde{G}_g(1)\|_F^2 + 4\mathcal{E}_f + 4\mathcal{E}_g. \tag{6}$$

Endpoint smoothness and the trend loss therefore jointly control the spacetime smoothness of the induced metrics in a continuous-time sense. In our objective, $\mathcal{L}_{\text{smooth}}$ is applied to the endpoint deformation matrices $\tilde{P}_f(1)$, which also controls endpoint metric variation when these matrices remain bounded, since

$$\|\tilde{G}_f(1) - \tilde{G}_g(1)\|_F \leq \big(\|\tilde{P}_f(1)\|_2 + \|\tilde{P}_g(1)\|_2\big)\|\tilde{P}_f(1) - \tilde{P}_g(1)\|_F. \tag{7}$$

The determinant and endpoint smoothness terms thus complement the trend loss by discouraging collapse and excessive endpoint variation.

3.4 Training objective and inference

Endpoint losses. The spectral branch is trained with $\mathcal{L}_{\text{spectral}}$. We group the deformation losses into $\mathcal{L}_{\text{interp}}$. Endpoint fitting uses the same soft pullback $V_{\mathcal{XY}}$ used to form the deformation cue. For Jacobian fitting, we use the local endpoint deformation operator Θ_f as a pseudo-target. Omitting face indices for brevity, the endpoint losses are $\mathcal{L}_{\text{jac}} = \|\tilde{P}(1) - \Theta\|_F^2$ and $\mathcal{L}_{\text{vert}} = \|V(1) - V_{\mathcal{XY}}\|_2^2$. The former matches endpoint deformation matrices, while the latter aligns reconstructed endpoint vertices with the soft pullback. We compute them as weighted averages, using triangle areas for $\mathcal{L}_{\text{jac}}$ and vertex masses for $\mathcal{L}_{\text{vert}}$.

Endpoint regularization. We also regularize the endpoint deformation using smoothness and determinant terms. These are $\mathcal{L}_{\text{smooth}} = \|\tilde{P}(1) - H\tilde{P}(1)\|_F^2$ and $\mathcal{L}_{\text{det}} = \|\det(\tilde{P}(1)) - 1\|_2^2$, where H is the row-normalized neighboring-face averaging operator. The smoothness term pulls each endpoint deformation matrix toward the average of its adjacent faces. The determinant term softly penalizes deviations from unit determinant. Both terms are computed as weighted averages using source triangle areas.

Full objective. The interpolation loss combines endpoint fitting, endpoint regularization, and trend regularization, $\mathcal{L}_{\text{interp}} = \lambda_{\text{jac}}\mathcal{L}_{\text{jac}} + \lambda_{\text{vert}}\mathcal{L}_{\text{vert}} + \lambda_{\text{smooth}}\mathcal{L}_{\text{smooth}} + \lambda_{\text{det}}\mathcal{L}_{\text{det}} + \lambda_{\text{trend}}\mathcal{L}_{\text{trend}}$. The full objective is $\mathcal{L} = \mathcal{L}_{\text{spectral}} + \mathcal{L}_{\text{interp}}$. The endpoint and trend terms are optimized jointly through the predicted Chebyshev coefficients. Together, they guide endpoint alignment and the evolution of intermediate deformation.

Inference. At inference time, the network predicts the soft endpoint map and the Chebyshev coefficients. We reconstruct intermediate shapes from the predicted Jacobian field through Poisson integration. We center the reconstructed shapes to resolve the global translation ambiguity. For final correspondence recovery, we follow the inference procedure of Spectral Meets Spatial [11]. Detailed inference settings are provided in Appendix A.10.

4 Experimental results

Datasets. We evaluate matching on remeshed FAUST and SCAPE [7, 3, 14, 38], SHREC'19 [33], SMAL [49], and DT4D-H [32]. Splits follow [10], with FAUST+SCAPE training for SHREC'19 and disjoint animal categories for SMAL. Interpolation uses disjoint training and test animations of `elephantLN0` [28] in (a) and (b), and 114 short-term pairs from 19 fixed-connectivity sequences in (c). Exact splits and pair construction are provided in Appendix A.7.

Preprocessing. Meshes are centered and independently normalized to unit surface area. We obtain ULRSSM correspondences [10], as in [41], and use them for pair-specific orthogonal alignment. Tab. 2 omits this alignment. Test-time parameter refinement is separate and detailed in Appendix A.10.

Table 1: Shape matching results under near-isometric and non-isometric settings. Training uses each benchmark's dataset, except FAUST+SCAPE for SHREC'19. Published baselines retain their inference protocols. † denotes official-code reproductions. Bold/underline indicate the best/second-best distinct results, and "–" unavailable results.

		Near-isometric			Non-isometric		
	Train Dataset	FAUST	SCAPE	FAUST + SCAPE	SMAL	DT4D-H	
Category	Method	FAUST	SCAPE	SHREC'19	SMAL	DT4D-H intra	DT4D-H inter
Axiomatic	BCICP [38]	6.1	11.0	–	–	–	–
	ZoomOut [34]	6.1	7.5	–	38.4	4.0	29.0
	Smooth Shells [18]	2.5	4.7	–	36.1	1.1	6.3
	DiscreteOp [39]	5.6	13.1	–	38.1	3.6	27.6
Supervised	FMNet [29]	11.0	17.0	–	42.0	9.6	38.0
	GeomFMaps [14]	2.6	3.0	7.9	8.4	2.1	4.1
	Shape matching only						
Unsupervised	WSupFMNet [42]	3.8	4.4	–	7.6	3.3	22.6
	Deep Shells [19]	1.7	2.5	21.1	29.3	3.4	31.1
	DUO-FMNet [15]	2.5	2.6	6.4	6.7	2.6	15.8
	AttnFMaps [27]	1.9	2.2	5.8	5.4	1.7	11.6
	ULRSSM [10]	1.6	<u>1.9</u>	4.6	3.9	**0.9**	4.1
	HybridFMap [4]	<u>1.5</u>	**1.8**	**3.6**	**3.3**	<u>1.0</u>	<u>3.5</u>
	Joint shape matching and interpolation						
	NeuroMorph [20]	2.3	5.6	8.5	5.9	14.4	25.3
	Spectral Meets Spatial† [11]	**1.4**	2.0	<u>3.8</u>	<u>3.7</u>	2.1	3.7
	CheF (Ours)	**1.4**	**1.8**	**3.6**	**3.3**	**0.9**	**3.4**

Interpolation protocols. Tab. 3 compares (a) joint correspondence and interpolation with predicted maps and (b) interpolation trained with ground-truth (GT) endpoint maps and no spectral branch. In (c), a randomly initialized interpolation network is fitted to each pair using GT endpoint maps. GT intermediate geometry is used only for evaluation. For these near-isometric experiments, polar rotations are additionally guided along an endpoint-conditioned reference path on $\mathrm{SO}(3)$.

4.1 Shape matching results

Tab. 1 evaluates correspondence accuracy while CheF also learns a continuous deformation trajectory. Compared with the reproduced Spectral Meets Spatial baseline, CheF matches or improves correspondence accuracy across all evaluated benchmarks. It also remains competitive with dedicated unsupervised matchers, matching the best reported values in the table on SCAPE, SHREC'19, SMAL, and DT4D-H intra. On DT4D-H inter, it achieves the lowest error among the compared methods. Fig. 4 visualizes the predicted correspondences through texture transfer.



Figure 4: **Qualitative correspondence results.** Texture transfer visualizes the predicted dense correspondences across near-isometric and non-isometric benchmarks.

Table 3: Interpolation accuracy for (a) joint matching/interpolation and (b,c) ground-truth endpoint correspondences, with (c) fitted per short-term pair. Spectral Meets Spatial and CheF use matched test-time refinement in (a). Ground-truth intermediate shapes are used only for evaluation, with results averaged within and then across pairs.

Method	Local J ($\times 10$) ↓	J-velocity ($\times 10^2$) ↓	nCD ($\times 10^2$) ↓	nHD ($\times 10^2$) ↓
(a) Joint correspondence and interpolation				
NeuroMorph [20]	2.77	7.08	2.78	12.56
Spectral Meets Spatial [11]	1.71	5.41	2.04	7.86
CheF (Ours)	**1.63**	**5.05**	**1.99**	**7.41**
(b) Interpolation with ground-truth correspondences				
Linear vertices / Jacobians + Poisson	1.32	4.83	1.99	7.99
NeuroMorph / Spectral Meets Spatial	1.35	5.13	2.09	7.88
CheF (Ours)	**1.27**	**4.54**	**1.97**	**7.19**
(c) Per-pair interpolation with ground-truth correspondences				
4Deform [41]	4.52	24.86	1.28	7.18
NeuroMorph / Spectral Meets Spatial	1.11	4.35	1.01	3.73
CheF (Ours)	**0.94**	**3.93**	**0.97**	**3.62**

Metrics. We report mean normalized geodesic error for matching and endpoint accuracy. Local J measures source area-weighted Frobenius RMSE of deformation matrices formed from face edges and unit normals against GT geometry. J-velocity evaluates the corresponding consecutive-frame differences without time normalization. Surface accuracy uses nCD and nHD from 4,096 area-weighted samples per mesh. We evaluate intermediate frames and their internal differences, averaging errors temporally within each pair and then equally across pairs. J-velocity uses temporal RMS. Geometric ablations report surface-area standard deviation (SA_σ) [41], deformation distortion (Strain), mean minimum determinant (Det.), and the percentage of pairs with any face flip (Flip). Surface distance definitions and normalization details are provided in the appendix.

Tab. 2 examines dependence on the pair-specific orthogonal alignment used in preprocessing. When this step is removed, CheF has lower geodesic error than the reproduced Spectral Meets Spatial baseline on both SCAPE and SMAL. Together with the pre-aligned results, this comparison suggests lower sensitivity to the alignment step on the evaluated benchmarks.

Table 2: Robustness without rigid pre-alignment on SCAPE and SMAL. We report normalized geodesic error ($\times 10^2$). Lower is better. † denotes results reproduced using the official code.

Method	SCAPE *(near-iso)*	SMAL *(non-iso)*
Spectral Meets Spatial† [11]	10.7	43.8
CheF (Ours)	1.9	5.1

4.2 Shape interpolation results

In setting (a) of Tab. 3, CheF achieves lower errors across all four metrics than NeuroMorph and Spectral Meets Spatial. The comparison with Spectral Meets Spatial uses matched test-time refinement.

Lower Local J and J-velocity errors indicate closer agreement with the intermediate deformation and its consecutive-frame changes. The reductions in nCD and nHD also indicate closer agreement with the ground-truth intermediate surfaces. These results evaluate the complete joint framework with predicted endpoint correspondences.

Setting (b) examines interpolation after removing endpoint correspondence prediction errors. With ground-truth endpoint maps, CheF improves over both NeuroMorph / Spectral Meets Spatial and the linear baseline across all four metrics. For fixed reconstructed endpoints, linearly interpolating their Jacobians followed by the same source

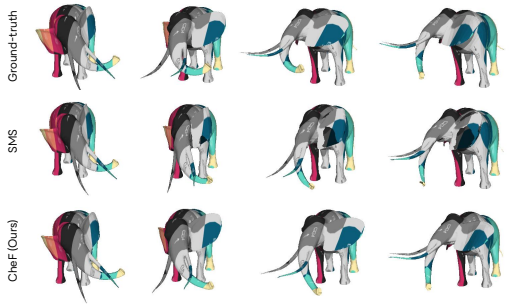


Figure 5: Joint correspondence and interpolation results on DeformingThings4D under setting (a).

Table 4: Ablation on representation design.

Variant	End. error ($\times 10^2$) ↓	SA_σ ($\times 10$) ↓
Full	**4.5**	**0.39**
Vertex coefficients	Deg.	Deg.
Sampled Jacobians	10.7	1.07
Endpoint Jacobian	4.8	0.60

Full uses Chebyshev coefficients, per-face Jacobians, and Poisson integration. *Deg.* denotes degenerate interpolation where intermediate states remain close to the source identity and do not form a meaningful trajectory.

Table 5: Ablation on trend regularization.

Variant	End. error ($\times 10^2$) ↓	SA_σ ($\times 10$) ↓
Jacobian trend	**4.5**	**0.39**
No trend	Deg.	Deg.
Vertex trend	**4.5**	1.28
ARAP trend	**4.5**	0.56

Jacobian trend denotes our trend loss in local deformation space. *Vertex trend* applies the trend loss directly to vertex trajectories. *Deg.* denotes a degenerate trajectory that remains near the source, for which SA_σ is not informative.

Poisson reconstruction yields the same trajectory as linear vertex interpolation when the translation convention is shared. The comparison with this baseline therefore tests whether learning the trajectory helps beyond a linear transition between the endpoints. The improvement supports the proposed interpolation formulation even when endpoint correspondence is given.

Setting (c) further examines per-pair fitting without pretrained interpolation weights. CheF achieves lower errors than 4Deform and NeuroMorph / Spectral Meets Spatial across all four metrics. The benefit thus persists when each pair is optimized independently, complementing the learned interpolation results in settings (a) and (b). Fig. 5 illustrates intermediate shapes reconstructed along a DT4D-H trajectory.

4.3 Ablation studies

We examine the trajectory representation and regularization on DT4D-H inter, where large shape variation challenges both endpoint alignment and intermediate deformation. The main ablations are evaluated on held-out pairs without test-time refinement. Endpoint error is computed from the interpolator's reconstructed endpoint rather than the initial spectral correspondence prediction. We read this error together with SA_σ, since low surface-area variation alone does not establish interpolation quality. Variants that remain near the source without producing meaningful intermediate motion are marked as degenerate.

Representation design. Tab. 4 examines the deformation domain and temporal structure of the trajectory representation. The vertex-coefficient variant predicts displacement coefficients directly without Poisson integration and produces a degenerate trajectory under the evaluated training setup. Predicting sampled Jacobian states independently avoids this collapse but yields higher endpoint error and SA_σ than the shared coefficient representation. The endpoint-only variant linearly interpolates the local deformation from the source to the predicted endpoint. It improves over sampled states but still trails the full model on both metrics. These comparisons support learning a temporally coupled Jacobian trajectory beyond a linear endpoint transition in the evaluated setting.

Trend regularization. Tab. 5 examines how temporal guidance affects the learned trajectory. Without the trend loss, training produces a degenerate trajectory that remains near the source. Endpoint fitting alone therefore does not yield useful interpolation in this setting. The Jacobian, vertex, and ARAP variants have the same reported endpoint error but differ in SA_σ. Our trend achieves lower surface-area variation, suggesting that the choice of regularization affects intermediate deformation beyond endpoint accuracy. This supports guiding local stretch and shear through the proposed reference path.

Temporal basis. Tab. 6 examines the choice of temporal basis under the same coefficient budget, generator, and training objective. Chebyshev achieves lower endpoint error and surface-area variation than the monomial and Legendre alternatives. Bernstein yields the lowest SA_σ but substantially higher endpoint error. Lower surface-area variation alone therefore does not indicate better interpolation. These results support Chebyshev as a suitable parameterization for balancing endpoint fitting and surface-area variation in this setting.

Table 6: Effect of temporal basis choice.

Variant	End. error ($\times 10^2$) ↓	SA_σ ($\times 10$) ↓
Chebyshev	**4.5**	0.39
Monomial	4.7	0.79
Legendre	4.6	0.49
Bernstein	12.0	**0.31**

All temporal bases use the same coefficient budget and generator. The losses and optimization settings are also shared. Chebyshev denotes our shifted low-order basis.

Table 7: Efficiency comparison.

Metric	Spectral Meets Spatial [11]	CheF (Ours)
Training memory	20.37 GB	2.55 GB
Training time	1429 ms	551 ms
Inference time	100 ms	88 ms

All measurements use the same preprocessing, batch size, hardware setting, and seven queried interpolation time steps. Inference time includes trajectory evaluation and reconstruction.

Table 8: Effect of even-order Chebyshev terms.

Variant	End. error ($\times 10^2$) ↓	SA_σ ($\times 10$) ↓
Full (proposed)	**4.5**	**0.39**
Odd inference	**4.5**	1.12
Odd training	**4.5**	0.71

Full uses both odd and even Chebyshev modes. *Odd inference* removes even-order terms from the trained trajectory only at inference. *Odd training* removes even modes during training.

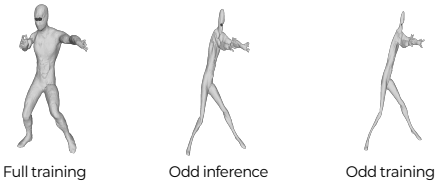


Figure 6: **Effect of even-order terms.**

Even-order terms. Tab. 8 examines whether the learned even-order terms contribute to intermediate deformation. Removing these terms only at inference increases SA_σ from 0.39 to 1.12 while leaving the reported endpoint error unchanged. This comparison shows that the trained model uses these terms to shape the interior trajectory. Retraining with odd terms alone reduces SA_σ to 0.71 but does not recover the full model's area stability. The qualitative comparison in Fig. 6 complements this analysis by showing the resulting intermediate shapes.

4.4 Efficiency

Tab. 7 compares the computational cost of CheF and Spectral Meets Spatial under seven queried time steps on an NVIDIA RTX A6000 GPU. The coefficient generator and the SMS interpolator use the same ResnetECPos architecture with approximately 764k and 755k parameters, respectively. Their similar parameter counts help separate the comparison from a large increase in deformation-network capacity. Under this setting, CheF uses less training memory and time, with slightly lower interpolation inference time. It predicts temporal coefficients once and reuses the source Poisson factorization across time queries. The measured results show that the proposed trajectory formulation provides interpolation benefits without increasing computational cost in this setting.

5 Conclusion

We introduced *Chebyshev Differential Flows* (CheF), an unsupervised framework for joint dense endpoint correspondence and continuous shape interpolation from endpoint pairs. Shifted Chebyshev polynomials parameterize per-face Jacobian trajectories with guidance from a spectral functional-map branch. A differentiable Poisson solver reconstructs globally consistent intermediate shapes. The endpoint depends on the sum of odd coefficients, while even modes shape the interior trajectory. A trend loss guides stretch and shear over the full interval without per-snapshot ARAP rigidity. Across near-isometric and non-isometric benchmarks, CheF achieves competitive correspondence accuracy and improves interpolation accuracy and geometric quality over the reproduced Spectral Meets Spatial baseline. Ablations support the benefits of Jacobian-space trajectories, even modes, and trend regularization. Experiments without rigid pre-alignment further suggest reduced sensitivity to ambient-frame alignment.

Limitations. CheF requires reliable endpoint correspondences and a source mesh suitable for stable differential reconstruction. Its fixed source connectivity prevents explicit topology changes such as cuts or merges. Partial inputs require reliable overlap estimation, while severely noisy or non-manifold meshes may need preprocessing.

Acknowledgments This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2026-25486241, 50%; No. RS-2022-NR070832, 25%), the Institute of Information & communications Technology Planning & Evaluation (IITP) under the artificial intelligence semiconductor support program to nurture the best talents (IITP-(2025)-RS-2023-00253914, 25%) grant funded by the Korea government (MSIT), and Hanyang University (No. HY202500000003991).

References

[1] Noam Aigerman, Kunal Gupta, Vladimir G. Kim, Siddhartha Chaudhuri, Jun Saito, and Thibault Groueix. Neural Jacobian Fields: Learning intrinsic mappings of arbitrary meshes. *ACM TOG*, 41(4):109:1–109:17, 2022. doi: 10.1145/3528223.3530141.

[2] Marc Alexa, Daniel Cohen-Or, and David Levin. As-rigid-as-possible shape interpolation. In *SIGGRAPH*, pages 157–164, 2000. doi: 10.1145/344779.344859.

[3] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. SCAPE: Shape completion and animation of people. *ACM TOG*, 24(3):408–416, 2005. doi: 10.1145/1073204.1073207.

[4] Lennart Bastian, Yizheng Xie, Nassir Navab, and Zorah Lähner. Hybrid functional maps for crease-aware non-isometric shape matching. In *CVPR*, pages 3313–3323. IEEE, 2024.

[5] Florian Bernard, Zeeshan Khan Suri, and Christian Theobalt. MINA: Convex mixed-integer programming for non-rigid shape alignment. In *CVPR*, 2020.

[6] Thomas Besnier, Emery Pierson, Sylvain Arguillère, Maks Ovsjanikov, and Mohamed Daoudi. Pandas: Learnable shape interpolation modeling with localized control. In *CVPR*, 2026.

[7] Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. FAUST: Dataset and evaluation for 3D mesh registration. In *CVPR*, pages 3794–3801, 2014. doi: 10.1109/CVPR.2014.491.

[8] Mario Botsch, Mark Pauly, Markus Gross, and Leif Kobbelt. PriMo: Coupled prisms for intuitive surface modeling. In *SGP*, pages 11–20, 2006.

[9] Christopher Brandt, Christoph von Tycowicz, and Klaus Hildebrandt. Geometric flows of curves in shape space for processing motion of deformable objects. *Comput. Graph. Forum*, 35(2):295–305, 2016. doi: 10.1111/cgf.12832.

[10] Dongliang Cao, Paul Roetzer, and Florian Bernard. Unsupervised learning of robust spectral shape matching. *ACM TOG*, 42(4):132:1–132:15, 2023. doi: 10.1145/3592107.

[11] Dongliang Cao, Marvin Eisenberger, Nafie El Amrani, Daniel Cremers, and Florian Bernard. Spectral Meets Spatial: Harmonising 3D shape matching and interpolation. In *CVPR*, pages 3658–3668, 2024.

[12] Dongliang Cao, Zorah Lähner, and Florian Bernard. Synchronous Diffusion for unsupervised smooth non-rigid 3D shape matching. In *ECCV*, pages 262–281, 2024. doi: 10.1007/978-3-031-72652-1_16. URL https://doi.org/10.1007/978-3-031-72652-1_16.

[13] Guillaume Charpiat, Pierre Maurel, Jean-Philippe Pons, Renaud Keriven, and Olivier Faugeras. Generalized gradients: Priors on minimization flows. *IJCV*, 73:325–344, 2007. doi: 10.1007/s11263-006-9966-2.

[14] Nicolas Donati, Abhishek Sharma, and Maks Ovsjanikov. Deep geometric functional maps: Robust feature learning for shape correspondence. In *CVPR*, pages 8589–8598, 2020. doi: 10.1109/CVPR42600.2020.00862.

[15] Nicolas Donati, Etienne Corman, and Maks Ovsjanikov. Deep orientation-aware functional maps: Tackling symmetry issues in shape matching. In *CVPR*, pages 732–741, 2022.

[16] Ilya Eckstein, Jean-Philippe Pons, Yiying Tong, C.-C. Jay Kuo, and Mathieu Desbrun. Generalized surface flows for mesh processing. In *SGP*, pages 183–192, 2007.

[17] Marvin Eisenberger and Daniel Cremers. Hamiltonian dynamics for real-world shape interpolation. In *ECCV*, pages 179–196, 2020. doi: 10.1007/978-3-030-58548-8_11.

[18] Marvin Eisenberger, Zorah Lähner, and Daniel Cremers. Smooth Shells: Multi-scale shape registration with functional maps. In *CVPR*, pages 12265–12274, 2020.

[19] Marvin Eisenberger, Aysim Toker, Laura Leal-Taixé, and Daniel Cremers. Deep Shells: Unsupervised shape correspondence with optimal transport. In *NeurIPS*, volume 33, pages 10491–10502, 2020.

[20] Marvin Eisenberger, David Novotny, Gael Kerchenbaum, Patrick Labatut, Natalia Neverova, Daniel Cremers, and Andrea Vedaldi. NeuroMorph: Unsupervised shape interpolation and correspondence in one go. In *CVPR*, pages 7473–7483, 2021.

[21] William Gao, Noam Aigerman, Thibault Groueix, Vladimir G. Kim, and Rana Hanocka. TextDeformer: Geometry manipulation using text guidance. In *SIGGRAPH*, pages 1–11, 2023. doi: 10.1145/3588432.3591552.

[22] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan C. Russell, and Mathieu Aubry. 3D-CODED: 3D correspondences by deep deformation. In *ECCV*, pages 230–246, 2018.

[23] Adam Hartshorne, Allen Paul, Tony Shardlow, and Neill D. F. Campbell. ARC-Flow: Articulated, resolution-agnostic, correspondence-free matching and interpolation of 3D shapes under flow fields. In *3DV*, pages 11–33, 2025. doi: 10.1109/3DV66043.2025.00009. URL https://doi.org/10.1109/3DV66043.2025.00009.

[24] Behrend Heeren, Martin Rumpf, Max Wardetzky, and Benedikt Wirth. Time-discrete geodesics in the space of shells. *Comput. Graph. Forum*, 31(5):1755–1764, 2012. doi: 10.1111/j.1467-8659.2012.03180.x.

[25] Jin Huang, Yiying Tong, Kun Zhou, Hujun Bao, and Mathieu Desbrun. Interactive shape interpolation through controllable dynamic deformation. *IEEE TVCG*, 17(7):983–992, 2011. doi: 10.1109/TVCG.2010.109.

[26] J. P. Lewis, Matt Cordner, and Nickson Fong. Pose space deformation: A unified approach to shape interpolation and skeleton-driven deformation. In *SIGGRAPH*, pages 165–172, 2000. doi: 10.1145/344779.344862.

[27] Lei Li, Nicolas Donati, and Maks Ovsjanikov. Learning multi-resolution functional maps with spectral attention for robust shape matching. In *NeurIPS*, volume 35, pages 29336–29349, 2022.

[28] Yang Li, Hikari Takehara, Takafumi Taketomi, Bo Zheng, and Matthias Nießner. 4DComplete: Non-rigid motion estimation beyond the observable surface. In *ICCV*, pages 12686–12696, 2021. doi: 10.1109/ICCV48922.2021.01247.

[29] Or Litany, Tal Remez, Emanuele Rodolà, Alex M. Bronstein, and Michael M. Bronstein. Deep functional maps: Structured prediction for dense shape correspondence. In *ICCV*, 2017.

[30] Haolin Liu, Xiaohang Zhan, Zizheng Yan, Zhongjin Luo, Yuxin Wen, and Xiaoguang Han. Stable-SCore: A stable registration-based framework for 3D shape correspondence. In *CVPR*, pages 917–928, 2025.

[31] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM TOG*, 34(6):248:1–248:16, 2015. doi: 10.1145/2816795.2818013.

[32] Robin Magnet, Jing Ren, Olga Sorkine-Hornung, and Maks Ovsjanikov. Smooth non-rigid shape matching via effective Dirichlet energy optimization. In *3DV*, 2022.

[33] Simone Melzi, Riccardo Marin, Emanuele Rodolà, Umberto Castellani, Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. SHREC 2019: Matching humans with different connectivity. In *3DOR*, 2019.

[34] Simone Melzi, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, and Maks Ovsjanikov. ZoomOut: Spectral upsampling for efficient shape correspondence. *ACM TOG*, 38(6):155:1–155:14, 2019. doi: 10.1145/3355089.3356524.

[35] Sanjeev Muralikrishnan, Siddhartha Chaudhuri, Noam Aigerman, Vladimir G. Kim, Matthew Fisher, and Niloy J. Mitra. GLASS: Geometric latent augmentation for shape spaces. In *CVPR*, 2022.

[36] Tiago Novello, Vinicius Da Silva, Guilherme Schardong, Luiz Schirmer, Helio Lopes, and Luiz Velho. Neural implicit surface evolution. In *ICCV*, pages 14233–14243. IEEE, 2023.

[37] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: A flexible representation of maps between shapes. *ACM TOG*, 31(4):30:1–30:11, 2012. doi: 10.1145/2185520.2185526.

[38] Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. Continuous and orientation-preserving correspondences via functional maps. *ACM TOG*, 37(6):248:1–248:16, 2018. doi: 10.1145/3272127.3275040.

[39] Jing Ren, Simone Melzi, Peter Wonka, and Maks Ovsjanikov. Discrete optimization for shape matching. *Comput. Graph. Forum*, 40(5):81–96, 2021. doi: 10.1111/cgf.14359.

[40] Lu Sang, Zehranaz Canfes, Dongliang Cao, Florian Bernard, and Daniel Cremers. Implicit neural surface deformation with explicit velocity fields. In *ICLR*, volume 2025, pages 12213–12237, 2025.

[41] Lu Sang, Zehranaz Canfes, Dongliang Cao, Riccardo Marin, Florian Bernard, and Daniel Cremers. 4Deform: Neural surface deformation for robust shape interpolation. In *CVPR*, pages 6542–6551, 2025.

[42] Abhishek Sharma and Maks Ovsjanikov. Weakly supervised deep functional maps for shape matching. In *NeurIPS*, volume 33, pages 19264–19275, 2020.

[43] Nicholas Sharp, Souhaib Attaiki, Keenan Crane, and Maks Ovsjanikov. DiffusionNet: Discretization agnostic learning on surfaces. *ACM TOG*, 41(3):27:1–27:16, 2022. doi: 10.1145/3507905.

[44] Olga Sorkine and Marc Alexa. As-rigid-as-possible surface modeling. In *SGP*, pages 109–116, 2007.

[45] Ramana Sundararaman, Nicolas Donati, Simone Melzi, Etienne Corman, and Maks Ovsjanikov. Deformation recovery: Localized learning for detail-preserving deformations. *ACM TOG*, 43(6):1–16, 2024. doi: 10.1145/3687968.

[46] Thomas Windheuser, Ulrich Schlickewei, Frank R. Schmidt, and Daniel Cremers. Geometrically consistent elastic matching of 3D shapes: A linear programming solution. In *ICCV*, 2011. doi: 10.1109/ICCV.2011.6126489.

[47] Benedikt Wirth, Leah Bar, Martin Rumpf, and Guillermo Sapiro. A continuum mechanical approach to geodesics in shape space. *IJCV*, 93:293–318, 2011. doi: 10.1007/s11263-010-0416-9.

[48] Yifan Xia, Tianwei Ye, Jun Huang, Xiaoguang Mei, and Jiayi Ma. Probabilistic deformation consistency for unsupervised shape matching. In *AAAI*, volume 40, pages 10951–10959, 2026. doi: 10.1609/aaai.v40i13.38072. URL https://doi.org/10.1609/aaai.v40i13.38072.

[49] Silvia Zuffi, Angjoo Kanazawa, David W. Jacobs, and Michael J. Black. 3D menagerie: Modeling the 3D shape and pose of animals. In *CVPR*, pages 6365–6373, 2017.

A Appendix

We provide proofs for the structural properties stated in Section 3. We recall the per-face ambient deformation trajectory

$$\tilde{P}_f(t) = I_3 + \sum_{m=1}^{M} \tilde{C}_{f,m}\, \beta_m(t), \qquad \beta_m(t) = T_m(2t-1) - (-1)^m, \tag{8}$$

where T_m is the m-th Chebyshev polynomial of the first kind. The restricted face Jacobian is

$$J_f(t) = \tilde{P}_f(t) B_f, \tag{9}$$

where $B_f \in \mathbb{R}^{3\times 2}$ is a tangent basis of the source face. For the trend loss, we use the induced ambient metric

$$\tilde{G}_f(t) = \tilde{P}_f(t)^\top \tilde{P}_f(t), \tag{10}$$

and the metric residual

$$R_f(t) = \tilde{G}_f(t) - \big((1-t)I_3 + t\tilde{G}_f(1)\big). \tag{11}$$

A.1 Proof of Proposition 1

Proof. At $t = 1$, the shifted time coordinate is $2t - 1 = 1$. The Chebyshev polynomial satisfies

$$T_m(1) = 1. \tag{12}$$

Therefore,

$$\beta_m(1) = T_m(1) - (-1)^m = 1 - (-1)^m = \begin{cases} 0, & m \text{ even}, \\ 2, & m \text{ odd}. \end{cases} \tag{13}$$

Substituting this into Eq. 8 gives

$$\begin{aligned} \tilde{P}_f(1) &= I_3 + \sum_{m=1}^{M} \tilde{C}_{f,m}\, \beta_m(1) \\ &= I_3 + 2 \sum_{k=1}^{\lceil M/2 \rceil} \tilde{C}_{f,2k-1}. \end{aligned} \tag{14}$$

Thus only odd-order coefficients affect the target endpoint. Moreover, since $\beta_m(0) = T_m(-1) - (-1)^m = 0$ for all m, even-order coefficients vanish at both endpoints. Therefore, even modes are endpoint-invisible degrees of freedom for the ambient deformation trajectory. □

A.2 Proof of Proposition 2

Proof. Each Chebyshev polynomial $T_m(2t-1)$ is a polynomial in t of degree m. Therefore, each shifted basis function

$$\beta_m(t) = T_m(2t-1) - (-1)^m \tag{15}$$

is a polynomial in t of degree at most m. By Eq. 8, each entry of $\tilde{P}_f(t)$ is therefore a polynomial in t of degree at most M.

Since

$$\tilde{G}_f(t) = \tilde{P}_f(t)^\top \tilde{P}_f(t), \tag{16}$$

each entry of $\tilde{G}_f(t)$ is a sum of products of two degree-M polynomials. Thus each entry of $\tilde{G}_f(t)$ is a polynomial in t of degree at most $2M$. The endpoint metric $\tilde{G}_f(1)$ is constant with respect to t, so

$$(1-t)I_3 + t\tilde{G}_f(1) \tag{17}$$

is polynomial in t of degree at most 1. Hence the metric residual

$$R_f(t) = \tilde{G}_f(t) - \big((1-t)I_3 + t\tilde{G}_f(1)\big) \tag{18}$$

has entries that are polynomials in t of degree at most $2M$. Consequently, the trend integrand

$$h_f(t) = \|R_f(t)\|_F^2 \tag{19}$$

is a scalar polynomial in t of degree at most $4M$.

Gauss–Legendre quadrature with Q nodes integrates all polynomials of degree at most $2Q-1$ exactly. Therefore, $\int_0^1 h_f(t)\,dt$ is evaluated exactly whenever

$$2Q-1 \geq 4M, \qquad \text{equivalently} \qquad Q \geq 2M+1. \tag{20}$$

This proves the claim. □

A.3 Proof of Proposition 3

Proof. From the definition of the metric residual,

$$\tilde{G}_f(t) = (1-t)I_3 + t\tilde{G}_f(1) + R_f(t). \tag{21}$$

For adjacent faces f and g, subtracting the corresponding decompositions gives

$$\tilde{G}_f(t) - \tilde{G}_g(t) = t\big(\tilde{G}_f(1) - \tilde{G}_g(1)\big) + R_f(t) - R_g(t). \tag{22}$$

The identity terms cancel because all faces share the same source metric I_3.

Using the inequality

$$\|a+b\|_F^2 \leq 2\|a\|_F^2 + 2\|b\|_F^2, \tag{23}$$

with

$$a = t\big(\tilde{G}_f(1) - \tilde{G}_g(1)\big), \qquad b = R_f(t) - R_g(t), \tag{24}$$

we obtain

$$\|\tilde{G}_f(t) - \tilde{G}_g(t)\|_F^2 \leq 2t^2\|\tilde{G}_f(1) - \tilde{G}_g(1)\|_F^2 + 2\|R_f(t) - R_g(t)\|_F^2. \tag{25}$$

Applying the same inequality to the residual difference gives

$$\|R_f(t) - R_g(t)\|_F^2 \leq 2\|R_f(t)\|_F^2 + 2\|R_g(t)\|_F^2. \tag{26}$$

Combining Eqs. 25 and 26, we get

$$\|\tilde{G}_f(t) - \tilde{G}_g(t)\|_F^2 \leq 2t^2\|\tilde{G}_f(1) - \tilde{G}_g(1)\|_F^2 + 4\|R_f(t)\|_F^2 + 4\|R_g(t)\|_F^2. \tag{27}$$

Integrating over $t \in [0,1]$ yields

$$\begin{aligned}\int_0^1 \|\tilde{G}_f(t) - \tilde{G}_g(t)\|_F^2\,dt \leq 2\left(\int_0^1 t^2\,dt\right)\|\tilde{G}_f(1) - \tilde{G}_g(1)\|_F^2 \\ + 4\int_0^1 \|R_f(t)\|_F^2\,dt + 4\int_0^1 \|R_g(t)\|_F^2\,dt.\end{aligned} \tag{28}$$

Since

$$\int_0^1 t^2\,dt = \frac{1}{3}, \tag{29}$$

and since

$$\mathcal{E}_f = \int_0^1 \|R_f(t)\|_F^2\,dt, \qquad \mathcal{E}_g = \int_0^1 \|R_g(t)\|_F^2\,dt, \tag{30}$$

we obtain

$$\int_0^1 \|\tilde{G}_f(t) - \tilde{G}_g(t)\|_F^2\,dt \leq \frac{2}{3}\|\tilde{G}_f(1) - \tilde{G}_g(1)\|_F^2 + 4\mathcal{E}_f + 4\mathcal{E}_g. \tag{31}$$

This proves the claim. □

A.4 Endpoint deformation smoothness controls endpoint metric smoothness

The trend smoothness bound in Proposition 3 involves endpoint metric variation. Our objective applies endpoint smoothness to the deformation matrices $\tilde{P}_f(1)$. The following bound connects the two quantities.

Lemma A.1 (Endpoint matrix smoothness controls endpoint metric smoothness)**.** *Let*

$$\tilde{G}_f(1) = \tilde{P}_f(1)^\top \tilde{P}_f(1), \qquad \tilde{G}_g(1) = \tilde{P}_g(1)^\top \tilde{P}_g(1). \tag{32}$$

Then

$$\|\tilde{G}_f(1) - \tilde{G}_g(1)\|_F \le \big(\|\tilde{P}_f(1)\|_2 + \|\tilde{P}_g(1)\|_2\big)\|\tilde{P}_f(1) - \tilde{P}_g(1)\|_F. \tag{33}$$

Proof. Let

$$P_f = \tilde{P}_f(1), \qquad P_g = \tilde{P}_g(1). \tag{34}$$

Then

$$\begin{aligned} P_f^\top P_f - P_g^\top P_g &= P_f^\top P_f - P_f^\top P_g + P_f^\top P_g - P_g^\top P_g \\ &= P_f^\top (P_f - P_g) + (P_f - P_g)^\top P_g. \end{aligned} \tag{35}$$

Using the triangle inequality and submultiplicativity of the Frobenius norm with the spectral norm, we obtain

$$\begin{aligned} \|P_f^\top P_f - P_g^\top P_g\|_F &\le \|P_f^\top (P_f - P_g)\|_F + \|(P_f - P_g)^\top P_g\|_F \\ &\le \|P_f\|_2 \|P_f - P_g\|_F + \|P_g\|_2 \|P_f - P_g\|_F \\ &= \big(\|P_f\|_2 + \|P_g\|_2\big)\|P_f - P_g\|_F. \end{aligned} \tag{36}$$

Substituting $P_f = \tilde{P}_f(1)$ and $P_g = \tilde{P}_g(1)$ proves the bound. □

A.5 Auxiliary result: ambient and tangent metrics

The trend loss is defined using the ambient metric $\tilde{G}_f(t) = \tilde{P}_f(t)^\top \tilde{P}_f(t)$. The actual face Jacobian is restricted to the tangent basis as $J_f(t) = \tilde{P}_f(t) B_f$. The tangent metric induced by this restricted Jacobian is

$$J_f(t)^\top J_f(t) = B_f^\top \tilde{P}_f(t)^\top \tilde{P}_f(t) B_f = B_f^\top \tilde{G}_f(t) B_f. \tag{37}$$

Thus the ambient metric controls the metric observed on the source tangent plane after restriction. If the same orthonormal tangent basis B is used for two matrices, then

$$\|B^\top (\tilde{G}_f(t) - \tilde{G}_g(t)) B\|_F \le \|\tilde{G}_f(t) - \tilde{G}_g(t)\|_F. \tag{38}$$

For neighboring faces, the tangent bases can differ; this introduces an additional tangent-frame mismatch term, analogous to the restricted Jacobian decomposition below.

A.6 Auxiliary result: ambient smoothness and restricted Jacobians

The main text regularizes endpoint smoothness at the level of the ambient deformation matrices $\tilde{P}_f(1)$. Here we record a simple relation between ambient matrix variation and restricted face-Jacobian variation.

Lemma A.2 (Restricted Jacobian gap decomposition)**.** *For adjacent faces f and g, let*

$$J_f(t) = \tilde{P}_f(t) B_f, \qquad J_g(t) = \tilde{P}_g(t) B_g. \tag{39}$$

Then

$$\|J_f(t) - J_g(t)\|_F \le \|(\tilde{P}_f(t) - \tilde{P}_g(t)) B_f\|_F + \|\tilde{P}_g(t)(B_f - B_g)\|_F. \tag{40}$$

If B_f has orthonormal columns, then

$$\|(\tilde{P}_f(t) - \tilde{P}_g(t)) B_f\|_F \le \|\tilde{P}_f(t) - \tilde{P}_g(t)\|_F. \tag{41}$$

Proof. We add and subtract $\tilde{P}_g(t) B_f$:

$$\begin{aligned} J_f(t) - J_g(t) &= \tilde{P}_f(t) B_f - \tilde{P}_g(t) B_g \\ &= \big(\tilde{P}_f(t) - \tilde{P}_g(t)\big) B_f + \tilde{P}_g(t)(B_f - B_g). \end{aligned} \tag{42}$$

Applying the triangle inequality gives Eq. 40. If B_f has orthonormal columns, then multiplication by B_f cannot increase the Frobenius norm:

$$\|A B_f\|_F^2 = \operatorname{tr}(B_f^\top A^\top A B_f) \le \operatorname{tr}(A^\top A) = \|A\|_F^2. \tag{43}$$

Taking the square root proves Eq. 41. □

Table 9: Ablation on Chebyshev order.

Order	End. error ($\times 10^2$) ↓	SA_σ ($\times 10$) ↓
$M = 1$	4.6	0.97
$M = 2$	**4.5**	0.72
$M = 4$	**4.5**	0.39
$M = 6$	4.6	**0.17**

All variants use the same architecture and losses, changing only the Chebyshev order M. We use $M = 4$ as the default setting in the main experiments.

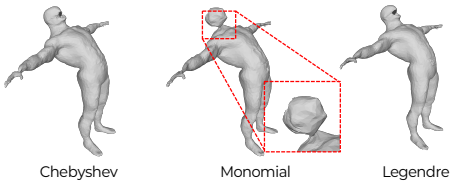


Figure 7: **Temporal basis comparison.**

Table 10: Effect of endpoint regularization. Endpoint accuracy and surface-area variation are evaluated on DT4D-H inter. Deformation validity is evaluated on short-term pairs under interpolation setting (c) in Tab. 3. Det. denotes the mean minimum determinant. Flip denotes the percentage of pairs with any face flip.

	DT4D-H inter		Interpolation setting (c)		
Variant	End. error ($\times 10^2$) ↓	SA_σ ($\times 10$) ↓	Strain ($\times 10^2$) ↓	Det. ↑	Flip (%) ↓
Full	**4.5**	0.39	1.43	**0.91**	**0.0**
w/o $\mathcal{L}_{\text{det}}$	8.2	1.29	2.18	0.62	5.3
w/o $\mathcal{L}_{\text{smooth}}$	6.9	**0.26**	**1.29**	0.82	4.4
w/o $\mathcal{L}_{\text{det}}$ & $\mathcal{L}_{\text{smooth}}$	8.5	1.19	3.98	−2.53	7.0

A.7 Datasets and evaluation protocols

Shape matching. We follow the benchmark splits of [10], using independently remeshed FAUST and SCAPE meshes. SHREC'19 uses models trained on FAUST+SCAPE, and SMAL uses disjoint training and test animal categories.

Learned interpolation: settings (a) and (b). We use `elephantLN0`, training on `act2,3,4,5,8` and testing on `act1,6,7`. Each training animation contributes 500 endpoint pairs with frame gaps of 7–11. Evaluation uses adjacent pairs from keyframes 0, 10, 18, 28, 39, 47, and 54. Setting (a) uses predicted endpoint maps, whereas setting (b) uses ground-truth endpoint correspondences with the spectral branch removed.

Per-pair interpolation and ablations. Setting (c) uses 19 fixed-connectivity animal sequences, each contributing the same six adjacent keyframe pairs, for 114 pairs in total. For CheF and NeuroMorph / Spectral Meets Spatial, a new interpolation network is optimized from random initialization for 1,000 steps per pair using identity ground-truth endpoint maps. The 4Deform comparison uses its Pairs S/S configuration on these pairs. The validity ablation follows this per-pair protocol, while the learned DT4D-H inter ablations use held-out pairs without parameter refinement. Ground-truth intermediate shapes are used only for evaluation in all settings.

A.8 Additional ablations

We provide additional ablations for the trend regularizer and the Chebyshev order. All experiments are conducted under the same setting as the main ablation study. We report endpoint error and SA_σ. Degenerate cases are marked as *Deg.*, where the learned trajectory collapses to an identity-like interpolation and therefore does not provide meaningful intermediate shapes.

Temporal basis visualization. Fig. 7 provides a qualitative comparison of the temporal bases evaluated in Tab. 6. All variants use the same coefficient budget, generator, losses, and optimization settings.

Chebyshev order. Tab. 9 studies the effect of the Chebyshev order M. The lowest order $M = 1$ yields higher temporal surface-area variation. Increasing the order reduces SA_σ, with $M = 6$

Table 11: Comparison of spectral branches with and without parameter refinement. Branch only denotes the standalone spectral matcher. SMS denotes Spectral Meets Spatial [11]. Refinement consists of 15 network parameter updates per test pair using each method's training objective. Bold indicates the lowest error within each spectral branch and refinement setting.

Spectral branch	Method	w/o refine		w/ refine	
		SMAL	DT4D-H inter	SMAL	DT4D-H inter
ULRSSM [10]	Branch only	7.1	4.7	3.9	4.1
	+ SMS	4.4	4.0	3.7	3.7
	+ CheF	**4.2**	**3.8**	**3.3**	**3.4**
HybridFMap [4]	Branch only	5.2	3.9	3.3	3.5
	+ SMS	4.8	3.7	3.4	3.5
	+ CheF	**4.5**	**3.6**	**3.2**	**3.4**

achieving the lowest value. However, its endpoint error is slightly higher than those of $M = 2$ and $M = 4$. We therefore use $M = 4$ as the default setting. It matches the lowest endpoint error while providing lower surface-area variation than $M = 2$.

Endpoint regularization. Tab. 10 evaluates the individual roles of $\mathcal{L}_{\mathrm{det}}$ and $\mathcal{L}_{\mathrm{smooth}}$ under two protocols. On DT4D-H inter, removing either term increases endpoint error. Removing $\mathcal{L}_{\mathrm{det}}$ also increases surface-area variation. Without $\mathcal{L}_{\mathrm{smooth}}$, SA_σ decreases despite worse endpoint accuracy. This shows that lower surface-area variation alone does not imply better interpolation quality.

We further examine deformation validity through short-term per-pair optimization with ground-truth endpoint correspondences. The spectral branch is removed and a new interpolation network is randomly initialized for each pair. Each network is optimized for 1,000 steps without intermediate-shape supervision. Removing either regularizer introduces face flips in some pairs. Removing both terms also produces a negative mean minimum determinant. Although removing $\mathcal{L}_{\mathrm{smooth}}$ yields slightly lower strain, the full objective produces no face flips on the evaluated pairs. Together, these results support the complementary roles of the two regularizers.

Spectral branches and refinement. Tab. 11 compares standalone spectral matchers and their joint extensions with and without parameter refinement. On SMAL and DT4D-H inter, CheF achieves lower matching error than Spectral Meets Spatial with either spectral branch in both refinement settings. However, we observed instability during joint training with the Hybrid Functional Maps branch. In particular, starting joint optimization without preliminary training of the spectral branch sometimes led to degenerate solutions. For this branch, we therefore first train the spectral component before jointly optimizing the spectral and interpolation branches. We use ULRSSM for the main results and report the Hybrid branch results here as an additional comparison.

Matching result on TOPKIDS. Tab. 12 reports matching results on TOPKIDS. For this dataset, we use a staged training procedure: we first train the spectral branch, then train the interpolation branch, and finally optimize both branches jointly. While CheF improves over the standalone spectral matchers, its advantage over Spectral Meets Spatial depends on the spectral branch. With ULRSSM, the error decreases from 6.9 to 5.7. With HybridFMap, both joint models achieve an error of 4.8, compared with 5.0 for the standalone matcher. The limited gain with HybridFMap and the need for staged training highlight remaining challenges for our framework on TOPKIDS.

Table 12: Shape matching on TOPKIDS.

Spectral branch	Method	Geo. error ↓
ULRSSM [10]	Branch only	9.2
	+ SMS	6.9
	+ CheF	**5.7**
HybridFMap [4]	Branch only	5.0
	+ SMS	**4.8**
	+ CheF	**4.8**

A.9 Implementation details: spectral regularization

We follow the deep functional map formulation used in prior unsupervised shape matching methods [11]. Given descriptors $D_{\mathcal{X}}$ and $D_{\mathcal{Y}}$ extracted by the Siamese feature network, we compute

bidirectional functional maps $C_{\mathcal{XY}}$ and $C_{\mathcal{YX}}$ in the Laplace–Beltrami eigenbases of the two shapes. Each functional map is obtained as the solution to a regularized least-squares problem,

$$C_{\mathcal{XY}} = \arg\min_{C} E_{\text{data}}(C) + \lambda E_{\text{reg}}(C), \tag{44}$$

where E_{data} enforces descriptor preservation and E_{reg} imposes Laplacian commutativity. These terms are part of the functional map solver, not the training objective.

The training objective consists of two groups of losses imposed on the resulting functional maps. The structural loss

$$\mathcal{L}_{\text{struct}} = \mathcal{L}_{\text{bij}} + \mathcal{L}_{\text{orth}} \tag{45}$$

combines a bijectivity term that encourages the two maps to be inverse-consistent and an orthogonality term that encourages near-orthogonal functional maps in the spectral domain,

$$\mathcal{L}_{\text{bij}} = \|C_{\mathcal{XY}}C_{\mathcal{YX}} - I\|_F^2 + \|C_{\mathcal{YX}}C_{\mathcal{XY}} - I\|_F^2, \tag{46}$$

$$\mathcal{L}_{\text{orth}} = \|C_{\mathcal{XY}}^\top C_{\mathcal{XY}} - I\|_F^2 + \|C_{\mathcal{YX}}^\top C_{\mathcal{YX}} - I\|_F^2. \tag{47}$$

The coupling loss connects the spectral map to the differentiable point-wise map,

$$\mathcal{L}_{\text{couple}} = \left\|C_{\mathcal{XY}} - \Phi_{\mathcal{Y}}^\dagger \Pi_{\mathcal{YX}} \Phi_{\mathcal{X}}\right\|_F^2 + \left\|C_{\mathcal{YX}} - \Phi_{\mathcal{X}}^\dagger \Pi_{\mathcal{XY}} \Phi_{\mathcal{Y}}\right\|_F^2. \tag{48}$$

The total spectral matching objective is

$$\mathcal{L}_{\text{spectral}} = \mathcal{L}_{\text{struct}} + \mathcal{L}_{\text{couple}}. \tag{49}$$

This objective is used to train the spectral correspondence branch, and the resulting soft point map is passed to the interpolation branch to construct the endpoint deformation cue.

A.10 Implementation details of inference

At inference time, a single forward pass produces the soft point map $\Pi_{\mathcal{XY}}$ and the Chebyshev coefficients defining $J(t)$. We reconstruct the vertex embedding at each queried time through Poisson integration in Eq. equation 2. The same reconstruction procedure is used for intermediate shapes and the target endpoint.

Since the Jacobian field does not determine global translation, the reconstructed embedding is defined only up to a translation. We resolve this ambiguity by centering each reconstructed shape,

$$V(t) = \widetilde{V}(t) - \mathbf{1}c(\widetilde{V}(t)), \qquad c(Z) = \frac{\mathbf{1}^\top M_{\mathcal{X}} Z}{\mathbf{1}^\top M_{\mathcal{X}} \mathbf{1}}, \tag{50}$$

where $\widetilde{V}(t)$ is the Poisson reconstruction and $M_{\mathcal{X}}$ is the source lumped mass matrix. The input shapes are also centered using the same convention. This operation fixes global translation while preserving the reconstructed local deformation.

When refinement is enabled, we update the pretrained network parameters for 15 iterations on each test pair. These updates use each method's training objective, following prior work [10, 4]. Results with and without parameter refinement are reported separately.

After parameter refinement, we reconstruct the endpoint $V(1)$ using the updated network parameters. We then perform spatial alignment following Spectral Meets Spatial [11]. An additional deformation field is optimized using a Chamfer alignment term and Dirichlet regularization. The network parameters remain fixed during this alignment. The final pointwise correspondence is obtained by nearest-neighbor search between the aligned endpoint and the target vertices.

A.11 Implementation details: Experimental setting

Training configuration. We use the AdamW optimizer with learning rate 1×10^{-4} and weight decay 1×10^{-2}, with batch size 1. Models are trained for 10 epochs. The random seed is fixed to 1234 for reproducibility.

Hardware and compute. All experiments are conducted on a single NVIDIA RTX A6000 GPU (48 GB) on a workstation equipped with two Intel Xeon Platinum 8568Y+ CPUs and 472 GB of system memory, running CUDA 12.4. Under the SMAL training configuration described above (batch size 1, $K = 128$, $M = 4$, $T = 7$), GPU memory usage is approximately 3GB. Training for 10 epochs takes approximately 20 hours (120 minutes per epoch). At inference, a forward pass on a mesh pair, including Jacobian interpolation across 7 time steps, takes on average 90ms. We did not track the total project compute precisely; exploratory experiments and failed runs during development consumed additional GPU time beyond the reported experiments.

A.12 Additional visualizations

We provide additional qualitative results of shape interpolation on three benchmarks: SCAPE (Fig. 8), SMAL (Fig. 9), and the inter-class split of DT4D-H (Fig. 10). For each source–target pair, we show the source shape, four uniformly sampled intermediate shapes along the predicted trajectory, and the target shape. Across all three benchmarks, the intermediate shapes transition smoothly between the endpoints without local distortion or degenerate geometry.

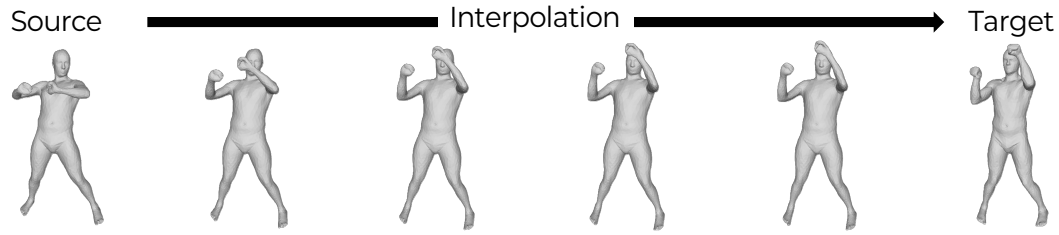


Figure 8: Additional visualization of SCAPE dataset

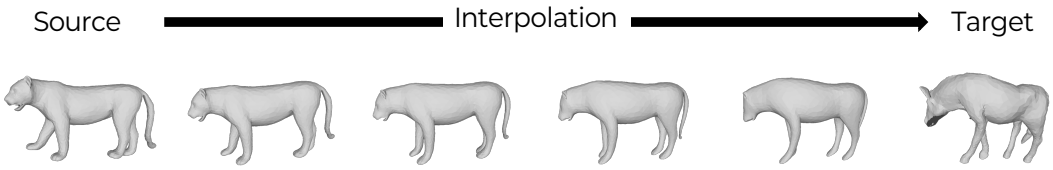


Figure 9: Additional visualization of SMAL dataset

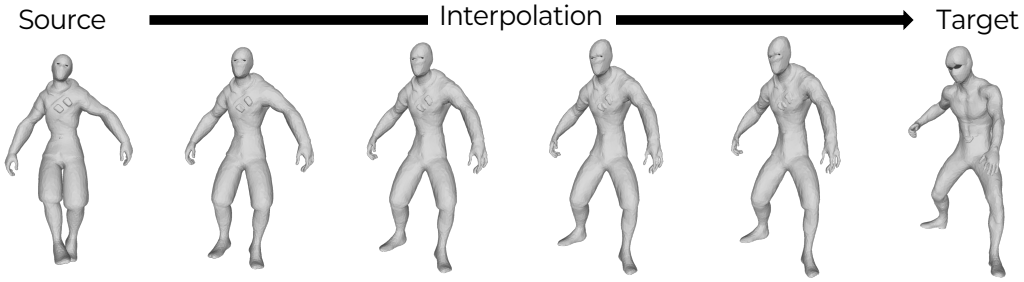


Figure 10: Additional visualization of DT4D-H inter dataset

NeurIPS paper checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [N/A].
- [N/A] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[N/A]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS paper checklist"**,
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers**.

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: Claims made in the abstract and introduction are supported by the method (Section 3), theoretical analyses (Propositions 1–3), and experiments (Section 4).

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

Justification: Limitations are discussed in the “Limitations” paragraph of Section 5.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren’t acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory Assumptions and Proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

   Justification: The paper includes three propositions stated in Section 3: Proposition 1 (endpoint behavior of odd/even Chebyshev modes), Proposition 2 (exactness of Gauss-Legendre quadrature for the trend loss), and Proposition 3 (endpoint smoothness bound for the trend regularization). All required assumptions are stated alongside each proposition, and complete proofs are provided in Appendix A.1, A.2, and A.3. Two supporting lemmas (Lemma A.1 and Lemma A.2) with full proofs are also provided in Appendix A.4 and A.6. All results are numbered and cross-referenced throughout.

   Guidelines:

   - The answer NA means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental Result Reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Details on training procedures and experiments are made clear in the appendix and main body. Code is also provided for exact replication.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: Code is provided to replicate the results.

   Guidelines:

   - The answer NA means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.

- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental Setting/Details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes]

   Justification: The full implementation details—including hyperparameters, optimizer configuration—are provided in Appendix A.11.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment Statistical Significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [No]

   Justification: Following common practice in the shape correspondence literature, we report results from a single training run on the standard benchmarks. Running multiple seeds across all datasets and ablations was not feasible within our compute budget. However, we follow identical training and evaluation protocols as prior work, ensuring fair comparison with reported numbers from baseline methods.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
   - For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
   - If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments Compute Resources**

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Compute resources used for the reported experiments, including GPU type, memory, and training time per run, are reported in Appendix A.11. We do not have an exact measurement of the total project compute, including failed experiments and time spent developing different experimental settings.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code Of Ethics**

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: The research conducted in this paper conforms with the NeurIPS Code of Ethics in every respect.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader Impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Broader impacts are discussed in the "Broader impacts" paragraph of Section 5.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.

- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: The paper does not release pretrained generative models or scraped datasets that pose a high risk for misuse. The proposed method addresses non-rigid shape correspondence/interpolation on standard benchmarks, which carries no such risk.

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: All sources of external data are mentioned.

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: The paper does not release new datasets or models. All experiments are conducted on publicly available benchmarks.

    Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing experiments or research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The paper does not involve research with human subjects, so IRB approval is not applicable.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: This work does not use LLMs as part of the proposed method. LLMs were not involved in the core methodology, experiments, or analysis.

    Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.