NHO: A NEURAL HAMILTONIAN OPERATOR FOR ANCHOR-BASED REGION LOCALIZATION AND DENSE CORRESPONDANCE

Anonymous authors
Paper under double-blind review

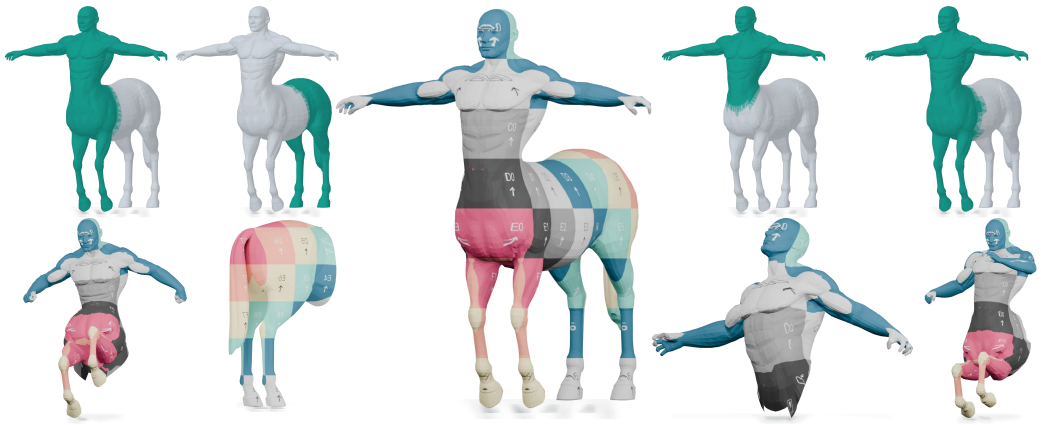

Figure 1: **NHO (Neural Hamiltonian Operator)** models a neural eigenspace using spatial cues, supporting both region localization and dense correspondence estimation for non-rigid partial-to-full shape matching.

ABSTRACT

Non-rigid partial-to-full shape correspondence from sparse anchors requires identifying the corresponding region on the full surface and recovering dense correspondences between the partial shape and that region. We present NHO, which combines sparse anchors with the intrinsic geometry of the partial shape to learn a neural Hamiltonian operator whose localized eigenspace encodes both the region support and intrinsic coordinates for dense correspondence. NHO parameterizes the Hamiltonian potential as an intrinsic neural field and optimizes it using anchor evidence together with spectral and geometric constraints. To resolve the spatial ambiguity left by sparse anchors, we introduce reciprocal refinement between operator estimation and correspondence recovery. At each round, the current eigenspace provides spectral coordinates and restricts matching to its induced support, while geometrically reliable correspondences provide additional evidence for updating the potential. After refinement, aggregated eigenfunction energy yields the final localization, and the recovered map initializes dense correspondence refinement. Experiments demonstrate competitive accuracy on both tasks and robustness to uniform scaling and rotation.

1 INTRODUCTION

Establishing dense correspondence between non-rigid partial and full shapes is a fundamental problem in 3D shape analysis, with broad applications in shape retrieval, animation, and geometric learning. Compared with full-to-full matching, partial-to-full correspondence is inherently more ambiguous: only an unknown subset of the full shape admits valid correspondences, while the remaining geometry introduces structured distractors. A small number of reliable correspondences can often

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

be obtained from manual landmarks or confident descriptor matches, providing a practical starting point for addressing this problem (Rakprayoon et al., 2021; Bensaïd et al., 2023).

However, sparse correspondences provide spatial constraints at only a few locations. The extent of the corresponding region and the correspondences for the remaining points are still unknown. These two estimation problems are coupled. Identifying the region excludes irrelevant geometry and reduces correspondence ambiguity, while additional correspondences provide spatial evidence for determining the region. With limited spatial cues, uncertainty in either estimate can propagate to the other.

Existing methods differ in which aspect of this coupled inference they emphasize. Correspondence-oriented methods (Rakprayoon et al., 2021; Attaiki et al., 2021; Cao et al., 2023) recover dense pointwise correspondences either by expanding sparse initial matches or by estimating functional maps from learned vertex-wise features. Localization-oriented methods (Rampini et al., 2019; Bensaïd et al., 2022; 2023) focus on the other unknown: they represent partiality as a region indicator or a spatially varying potential and recover the corresponding support through neural interpolation or spectral alignment. Joint spectral formulations (Rodolà et al., 2017; Litany et al., 2017; Postolache et al., 2020; Wu et al., 2020) explicitly connect the two by jointly estimating partial support and a functional map, or by constructing spectral bases localized in the latent corresponding region. When sparse correspondences are available, existing methods primarily use them to initialize correspondence refinement or to fit a partiality indicator. How to lift such sparse evidence into a local spectral representation that supports both region localization and dense correspondence recovery remains underexplored. App. A provides a comparison with representative methods.

In this work, we introduce NHO, a method for learning a **neural Hamiltonian operator** on the full surface by parameterizing its potential as an intrinsic neural field. Sparse correspondences serve as anchors that provide spatial evidence about the corresponding region, while spectral agreement with the partial shape and geometric regularization constrain the Hamiltonian potential beyond the observed locations. To further reduce the spatial ambiguity left by the sparse anchors, we introduce a reciprocal refinement mechanism that alternates between operator estimation and correspondence refinement. At each round, the eigenfunctions of the current operator guide correspondence estimation, while the induced support restricts the search to the candidate region. The refined correspondences, in turn, provide additional spatial evidence for updating the potential. This reciprocal process yields two complementary outputs (Fig. 1): the aggregated eigenfunction energy of the learned operator localizes the corresponding region, while the recovered map provides a reliable initialization for subsequent dense correspondence refinement. Our key contributions are as follows:

- We propose an intrinsic neural-field parameterization of the Hamiltonian potential for learning a local spectral representation from sparse anchors and the intrinsic geometry of the partial shape.
- We introduce reciprocal operator–correspondence refinement to exploit the coupling between region localization and dense correspondence recovery.
- Extensive experiments demonstrate that the learned Hamiltonian operator supports both region localization and dense correspondence recovery from sparse anchors, while remaining robust to uniform scaling and rotation.

2 Related Work

Partial Shape Correspondence. Functional maps (Ovsjanikov et al., 2012) provide a compact spectral representation for shape correspondence. For partial matching, Rodolà et al. (2017) jointly estimate the corresponding region and a functional map, while FSPM (Litany et al., 2017) constructs localized quasi-harmonic bases through joint approximate diagonalization. Learning-based methods use learned per-vertex features within functional-map frameworks, as in DPFM (Attaiki et al., 2021) and ULRSSM (Cao et al., 2023), or directly estimate pointwise correspondences from feature similarities (Bracha et al., 2024a;b). Another line of work (Ehm et al., 2024b; Roetzer & Bernard, 2025) uses discrete optimization to enforce geometric consistency in the recovered matchings. Our approach instead uses reliable sparse correspondences to guide the learning of a Hamiltonian potential, yielding a localized spectral representation for both region localization and dense correspondence recovery.

Hamiltonian Spectral Localization. Hamiltonian operators (Choukroun et al., 2018) augment the Laplace–Beltrami operator (LBO) with a spatially varying potential, allowing low-energy eigenfunctions to concentrate in selected regions. Rampini et al. (2019) exploit this property for correspondence-free region localization through Hamiltonian spectrum alignment. Postolache et al. (2020) further analyze the Hamiltonian–Dirichlet connection and exploit it for functional-map-based partial shape matching. Subsequent work (Bensaïd et al., 2022; Bensaïd & Kimmel, 2024) extends spectral region localization by aligning operator spectra under multiple intrinsic metrics. Building on this foundation, our method parameterizes the Hamiltonian potential as an intrinsic neural field and iteratively updates it using sparse anchors and refined correspondences.

Correspondence Refinement. Spectral correspondence refinement typically improves an initial map by alternating between spectral and pointwise representations. The original functional-map framework (Ovsjanikov et al., 2012) introduces ICP-style refinement in spectral embedding space, while BCICP (Ren et al., 2018) further promotes bijectivity, continuity, and coverage during refinement. ZoomOut (Melzi et al., 2019) refines coarse or noisy initial maps by progressively increasing the spectral resolution. For partial matching, Wu et al. (2020) combine Hamiltonian spectral alignment with ZoomOut-based iterative upsampling. More recently, NAM (Viganò et al., 2025) generalizes functional-map refinement through a nonlinear neural representation and Neural ZoomOut. We adopt this approach to refine the correspondences obtained by NHO.

3 PRELIMINARIES

Laplace–Beltrami Operator. We model each shape as a compact, connected, 2-manifold $\mathcal{X}$, possibly with a smooth boundary $\partial\mathcal{X}$. We denote its interior by $\mathrm{int}(\mathcal{X})$. The positive semi-definite LBO $\Delta_{\mathcal{X}}$ generalizes the basic differential operator from Euclidean analysis to Riemannian manifolds. It admits an eigendecomposition

$$\Delta_{\mathcal{X}}\phi_i(x) = \lambda_i\phi_i(x) \qquad x \in \mathrm{int}(\mathcal{X}) \tag{1}$$

$$\phi_i(x) = 0 \qquad x \in \partial\mathcal{X}, \tag{2}$$

with homogeneous Dirichlet boundary conditions (2), where $\{0 \le \lambda_1 \le \lambda_2 \le \cdots\}$ are the eigenvalues and ϕ_i are the corresponding eigenfunctions.

Hamiltonian Operator. As a theoretical foundation for our method, the Hamiltonian operator augments the LBO with a scalar potential function defined over the manifold (Choukroun et al., 2018). Given a non-negative potential $v : \mathcal{X} \to \mathbb{R}_+$, the Hamiltonian is defined as $H_{\mathcal{X}} = \Delta_{\mathcal{X}} + v$, acting on scalar functions as

$$H_{\mathcal{X}} f = \Delta_{\mathcal{X}} f + vf, \tag{3}$$

where vf denotes pointwise multiplication. For $v \equiv 0$, $H_{\mathcal{X}}$ reduces to $\Delta_{\mathcal{X}}$. On a compact manifold, $H_{\mathcal{X}}$ is self-adjoint on the same domain as $\Delta_{\mathcal{X}}$ and has a discrete spectrum. Its eigenpairs satisfy

$$H_{\mathcal{X}}\psi_i(x) = \mu_i\psi_i(x), \tag{4}$$

where, as before, the eigenvalues μ_i are listed in nondecreasing order, and the eigenfunctions ψ_i are chosen to form an orthonormal basis of $L^2(\mathcal{X})$.

To illustrate the localization effect, consider a subdomain $\mathcal{R} \subset \mathcal{X}$ with smooth boundary and a finite step potential

$$v_\tau(x) = \begin{cases} 0 & x \in \mathcal{R} \\ \tau & x \in \mathcal{X} \setminus \mathcal{R} \end{cases} \qquad \tau > 0. \tag{5}$$

Here, τ denotes the potential barrier height outside $\mathcal{R}$. For sufficiently high barriers, low-energy Hamiltonian eigenfunctions concentrate in $\mathcal{R}$, although they generally do not vanish outside it (Choukroun et al., 2018). In this regime, the low-lying Hamiltonian eigenvalues approximate the corresponding Dirichlet eigenvalues of $\mathcal{R}$ (Postolache et al., 2020). This Hamiltonian–Dirichlet connection has been exploited for spectral region localization (Rampini et al., 2019) and motivates our use of the learned Hamiltonian eigenfunctions for region localization and alignment with the Dirichlet eigenbasis of the partial shape.

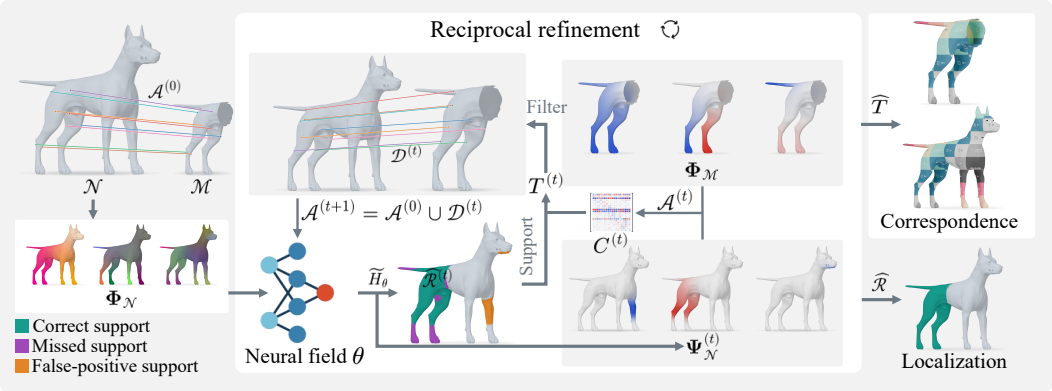

Figure 2: Overview of the NHO framework. Given a partial shape $\mathcal{M}$, a full shape $\mathcal{N}$, and sparse anchors $\mathcal{A}^{(0)}$, a neural field parameterizes the Hamiltonian potential on $\mathcal{N}$ using $\mathbf{\Phi}_{\mathcal{N}}$ as an intrinsic positional encoding. At round t, the Hamiltonian eigenspace induces the support $\widehat{\mathcal{R}}^{(t)}$ through aggregated energy and, together with $\mathbf{\Phi}_{\mathcal{M}}$ and $\mathcal{A}^{(t)}$, determines the alignment $C^{(t)}$. The support restricts the recovery of $T^{(t)}$, while geometrically reliable matches update the anchors and operator. The final support localizes the corresponding region, while the final-round map $\widehat{T}$ initializes dense correspondence refinement.

4 Methodology

Fig. 2 provides an overview of our framework. We first introduce the problem and notation. Sec. 4.1 then presents the intrinsic neural Hamiltonian field (INHF), followed by neural Hamiltonian optimization, reciprocal refinement, and dense correspondence recovery in Sec. 4.2.

Problem Formulation and Notation. Let $\mathcal{M}$ and $\mathcal{N}$ be connected partial and full triangle meshes with $n_{\mathcal{M}}$ and $n_{\mathcal{N}}$ vertices, respectively. We assume that $\mathcal{M}$ corresponds under a non-rigid deformation to an unknown region $\mathcal{R} \subseteq \mathcal{N}$. Given a sparse set of reliable anchors $\mathcal{A}^{(0)} = \{(p_\ell, q_\ell)\}_{\ell=1}^{m}$, where $p_\ell \in \mathcal{M}$ and $q_\ell \in \mathcal{N}$, our goal is to localize $\mathcal{R}$ and recover a dense map $T : \mathcal{M} \to \mathcal{N}$.

For vertices $p \in \mathcal{M}$ and $q \in \mathcal{N}$, let $a_p^{\mathcal{M}}$ and $a_q^{\mathcal{N}}$ denote the lumped vertex-area weights consistent with the Laplacian discretization. We denote the corresponding diagonal mass matrices by $A_{\mathcal{M}}$ and $A_{\mathcal{N}}$, and the resulting discrete surface areas by $\mathrm{Area}(\mathcal{M})$ and $\mathrm{Area}(\mathcal{N})$.

Let $\mathbf{\Phi}_{\mathcal{N}}^{j}$ denote the first j mass-orthonormal LBO eigenfunctions of $\mathcal{N}$, with $\Lambda_{\mathcal{N}}^{j}$ the diagonal matrix of the corresponding eigenvalues. On $\mathcal{M}$, we impose Dirichlet boundary conditions and denote its first j mass-orthonormal eigenfunctions by $\mathbf{\Phi}_{\mathcal{M}}^{j}$, with positive eigenvalues $\{\lambda_i^{\mathcal{M}}\}_{i=1}^{j}$.

4.1 Intrinsic Neural Hamiltonian Field

Building on the Hamiltonian–Dirichlet connection, we formulate region localization as spectral fitting. When $\mathcal{M}$ and $\mathcal{R}$ are approximately isometric and share a common scale, aligning the low-frequency Hamiltonian spectrum on $\mathcal{N}$ with the Dirichlet spectrum of $\mathcal{M}$ encourages the low-potential region to recover $\mathcal{R}$. However, spectrum-based potential estimation is non-convex and does not uniquely determine the location of the support, making it sensitive to initialization (Rampini et al., 2019). In our setting, sparse anchors provide additional spatial evidence, while the area of $\mathcal{M}$ serves as a soft prior on the size of $\mathcal{R}$. These observations motivate a neural parameterization that allows spatial and spectral constraints to jointly guide operator estimation.

Following Bensaïd et al. (2023), we use $\mathbf{\Phi}_{\mathcal{N}}^{r}$ as an intrinsic positional encoding. A single neural network $g_\theta : \mathbb{R}^r \to \mathbb{R}$ is applied row-wise to this encoding, with parameters shared across vertices, yielding an unconstrained scalar field $w_\theta \in \mathbb{R}^{n_{\mathcal{N}}}$:

$$w_\theta = g_\theta\left(\mathbf{\Phi}_{\mathcal{N}}^{r}\right). \tag{6}$$

To obtain a differentiable representation of the candidate region, we transform w_θ into a soft region indicator:

$$s_\theta = \frac{\mathbf{1} - \tanh(w_\theta)}{2}, \tag{7}$$

where $\mathbf{1}$ is the all-ones vector and $\tanh$ is applied elementwise. We then define the bounded potential (Rampini et al., 2019) and neural Hamiltonian:

$$v_\theta = \tau(\mathbf{1} - s_\theta), \tag{8}$$

$$H_{\mathcal{N},\theta} = \Delta_\mathcal{N} + \operatorname{diag}(v_\theta). \tag{9}$$

High region scores correspond to low potential, with $s_\theta \in (0,1)^{n_\mathcal{N}}$ and $v_\theta \in (0,\tau)^{n_\mathcal{N}}$. As s_θ approaches a binary indicator, v_θ approaches the step potential in Eq. 5. We denote the first k Hamiltonian eigenpairs by $\{(\mu_i, \psi_i)\}_{i=1}^k$ and collect the eigenfunctions in $\mathbf{\Psi}_\mathcal{N}^k$, leaving their dependence on θ implicit. The field thus links the candidate region to a learned Hamiltonian spectral representation.

4.2 Reciprocal Localization and Correspondence

With the neural Hamiltonian representation defined above, we first estimate its parameters using sparse spatial evidence and the intrinsic geometry of $\mathcal{M}$.

Neural Hamiltonian optimization. Each anchor identifies a vertex $q_\ell \in \mathcal{N}$, so we encourage high region scores at these locations:

$$\mathcal{L}_{\text{anc}} = -\frac{1}{m}\sum_{\ell=1}^{m} \log s_\theta(q_\ell). \tag{10}$$

Anchors constrain location but not region size. We therefore use $\operatorname{Area}(\mathcal{M})$ as a soft area prior. The soft area induced by s_θ is $\sum_{q\in\mathcal{N}} a_q^\mathcal{N} s_\theta(q)$, leading to

$$\mathcal{L}_{\text{area}} = \left(\frac{\sum_{q\in\mathcal{N}} a_q^\mathcal{N} s_\theta(q) - \operatorname{Area}(\mathcal{M})}{\operatorname{Area}(\mathcal{N})}\right)^2. \tag{11}$$

Area alone does not enforce connectivity. We regularize the superlevel filtration of s_θ using zero-dimensional persistent homology (Edelsbrunner et al., 2002; Hu et al., 2019):

$$\mathcal{L}_{\text{topo}} = \sum_{(b,d)\in\mathcal{P}_0(s_\theta)} \big(s_\theta(b) - s_\theta(d)\big)^2, \tag{12}$$

where $\mathcal{P}_0(s_\theta)$ contains finite persistence pairs indexed by vertices (b,d) attaining their birth and death values. Penalizing these pairs, excluding the unique essential component, encourages a single connected region.

These spatial terms alone do not ensure intrinsic compatibility with $\mathcal{M}$. Motivated by the Hamiltonian–Dirichlet connection (Postolache et al., 2020), we encourage agreement with its first k Dirichlet eigenvalues. We set the barrier above the target spectral range (Rampini et al., 2019):

$$\tau = \beta\lambda_k^\mathcal{M}, \qquad \beta > 1. \tag{13}$$

The spectral loss is

$$\mathcal{L}_{\text{spec}} = \frac{1}{k}\sum_{i=1}^{k} \log\left(1 + \left(\frac{\mu_i - \lambda_i^\mathcal{M}}{\lambda_i^\mathcal{M}}\right)^2\right). \tag{14}$$

Relative normalization balances errors across eigenvalue magnitudes, while the logarithm reduces the influence of large residuals.

For efficiency, we approximate the Hamiltonian eigensystem in the subspace spanned by the first K LBO eigenfunctions of $\mathcal{N}$, where $k \le K \ll n_\mathcal{N}$. Using the same mass matrix, we form the reduced operator:

$$\widetilde{H}_\theta = \Lambda_\mathcal{N}^K + (\mathbf{\Phi}_\mathcal{N}^K)^\top A_\mathcal{N} \operatorname{diag}(v_\theta)\mathbf{\Phi}_\mathcal{N}^K. \tag{15}$$

We solve the resulting $K \times K$ differentiable eigenproblem and lift the first k eigenvectors back to the mesh through $\mathbf{\Phi}_{\mathcal{N}}^{K}$.

Finally, we minimize

$$\mathcal{L}_{\mathrm{op}}(\theta; \mathcal{A}) = \alpha_1 \mathcal{L}_{\mathrm{anc}}(\mathcal{A}) + \alpha_2 \mathcal{L}_{\mathrm{area}} + \alpha_3 \mathcal{L}_{\mathrm{topo}} + \alpha_4 \mathcal{L}_{\mathrm{spec}}, \tag{16}$$

where $\alpha_1, \ldots, \alpha_4 \geq 0$ weight the four terms. Optimization with $\mathcal{A}^{(0)}$ yields the initial parameters $\theta^{(0)}$.

Reciprocal Refinement. The initial estimate may remain spatially ambiguous, as illustrated in App. B. We therefore alternate operator updates with support-restricted correspondence estimation, using geometrically consistent matches as additional spatial evidence.

A superscript (t) denotes quantities at round t. We extract the current region from aggregated eigenfunction energy:

$$e^{(t)}(q) = \frac{1}{k} \|\mathbf{\Psi}_{\mathcal{N}}^{(t)}(q, :)\|_2^2, \tag{17}$$

$$\widehat{\mathcal{R}}^{(t)} = \left\{ q \in \mathcal{N} \mid e^{(t)}(q) > \eta \mu_1^{(t)} \right\}, \tag{18}$$

where $\eta > 0$ is a fixed empirical threshold factor. The estimated region restricts the target domain for matching. Using $\mathcal{A}^{(t)}$, we align the spectral coordinates:

$$C^{(t)} \in \underset{C \in \mathbb{R}^{k \times k}}{\arg\min} \sum_{(p,q) \in \mathcal{A}^{(t)}} \left\| \mathbf{\Psi}_{\mathcal{N}}^{(t)}(q, :) C - \mathbf{\Phi}_{\mathcal{M}}^{k}(p, :) \right\|_2^2. \tag{19}$$

We select the minimum-Frobenius-norm least-squares solution using a pseudoinverse. Restricted nearest-neighbor search then gives an intermediate map:

$$T^{(t)}(p) = \underset{q \in \widehat{\mathcal{R}}^{(t)}}{\arg\min} \left\| \mathbf{\Phi}_{\mathcal{M}}^{k}(p, :) - \mathbf{\Psi}_{\mathcal{N}}^{(t)}(q, :) C^{(t)} \right\|_2^2. \tag{20}$$

To limit unreliable feedback, we retain matches whose local mapping distortion (LMD) falls below a fixed threshold (Xiang et al., 2021). We merge them with the fixed original anchors, resolving conflicts among additional matches in favor of lower-LMD pairs. Denoting the retained additional pairs by $\mathcal{D}^{(t)}$, we update

$$\mathcal{A}^{(t+1)} = \mathcal{A}^{(0)} \cup \mathcal{D}^{(t)}, \tag{21}$$

and minimize $\mathcal{L}_{\mathrm{op}}(\theta; \mathcal{A}^{(t+1)})$, warm-starting from $\theta^{(t)}$. The additional pairs are reselected at each round to adapt the spatial constraints on the operator.

Region Localization and Correspondence Recovery. After reciprocal refinement, the low-frequency Hamiltonian eigenfunctions exhibit stronger spatial concentration within the corresponding region, as visualized in App. B. We then freeze the operator and obtain the final region estimate $\widehat{\mathcal{R}}$ using Eq. 18. For dense correspondence recovery, we denote the final-round map obtained by the restricted nearest-neighbor search in Eq. 20 as $\widehat{T}$. We then refine $\widehat{T}$ using NAM-based Neural ZoomOut (Viganò et al., 2025) to obtain the final correspondence T.

5 Experiments

Implementation Details. We optimize NHO independently for each shape pair. By default, we use $m = 25$ anchors selected by farthest-point sampling over the annotated partial vertices. Ours (DPFM) uses DPFM feature matches as initial anchors (App. C); complete settings are provided in App. D.

Comparison Methods. We compare PFM (Rodolà et al., 2017), FSPM (Litany et al., 2017), DPFM (Attaiki et al., 2021), and EchoMatch (Xie et al., 2025) for both region localization and dense correspondence recovery. For region localization, we additionally include the original correspondence-free Hamiltonian spectrum alignment method of Rampini et al. (2019) and the piecewise-smooth localization method of Bensaïd et al. (2023). For dense correspondence recovery, we further evaluate DPFM refined with ZoomOut (Melzi et al., 2019), ULRSSM (Cao et al., 2023), and Wormhole (Bracha et al., 2024b).

Table 1: Quantitative comparison of region localization on CUTS'24. All metrics are surface-area-weighted percentages; higher is better. Best in bold; second-best underlined.

	Default								Normalized							
	Unrotated				Rotated				Unrotated				Rotated			
Method	IoU	Precision	Recall	F1-score	IoU	Precision	Recall	F1-score	IoU	Precision	Recall	F1-score	IoU	Precision	Recall	F1-score
PFM	71.88	78.40	85.27	81.60	72.22	78.60	85.50	81.82	50.58	59.88	67.61	63.28	49.56	58.88	66.42	62.22
FSPM	51.59	57.17	80.02	66.28	51.56	57.21	79.84	66.25	50.52	56.38	79.32	65.47	50.39	56.35	79.30	65.46
Hamiltonian	50.73	60.34	62.37	61.02	50.38	60.02	61.58	60.48	52.43	62.04	64.40	62.85	52.96	62.69	65.37	63.63
DPFM	17.03	53.48	20.06	26.16	13.33	42.37	15.82	20.73	80.45	85.63	93.24	88.86	37.65	62.76	46.01	51.12
Piecewise Smooth	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	73.88	83.77	84.74	82.84	49.84	64.43	63.15	61.59
EchoMatch	43.15	50.25	76.34	58.19	41.30	49.58	75.21	56.43	**82.14**	**86.95**	93.80	**89.94**	**76.69**	**86.18**	87.65	**86.21**
Ours	**75.24**	**79.43**	**94.11**	**85.37**	**75.23**	**79.38**	**93.95**	**85.36**	75.15	80.31	92.18	85.13	75.06	80.16	**92.08**	85.02
Ours (DPFM)	44.39	53.67	64.66	58.06	22.51	31.80	36.97	33.65	80.92	84.92	**93.91**	89.11	53.01	61.46	68.08	64.45

Table 2: Mean geodesic correspondence error ($\times 100$, lower is better) on four partial shape matching benchmarks under four input settings. Avg. denotes the unweighted average over the four input settings within each dataset. Best in bold; second-best underlined.

	CUTS'24					PFAUST-M					PFAUST-H					PFARM				
	Default		Normalized		Avg.	Default		Normalized		Avg.	Default		Normalized		Avg.	Default		Normalized		Avg.
Method	Unrotated	Rotated	Unrotated	Rotated		Unrotated	Rotated	Unrotated	Rotated		Unrotated	Rotated	Unrotated	Rotated		Unrotated	Rotated	Unrotated	Rotated	
PFM	**5.03**	**4.97**	24.27	26.15	15.10	50.07	50.41	50.07	49.93	50.12	44.84	45.48	45.43	44.93	45.17	59.38	59.12	60.06	59.80	59.59
FSPM	19.33	19.22	19.97	19.99	19.63	49.93	50.09	50.04	50.06	50.03	45.89	45.76	45.78	45.65	45.77	63.20	63.38	62.79	62.69	63.02
DPFM	67.67	72.84	2.38	24.94	41.96	34.46	36.64	37.88	38.93	36.98	38.60	38.68	39.62	41.71	39.65	49.53	53.14	51.59	45.82	50.02
DPFM + ZoomOut	52.59	61.23	**2.17**	24.22	35.05	36.30	37.16	40.34	40.52	38.58	37.95	39.70	41.10	42.52	40.32	47.38	53.11	49.93	45.56	49.00
ULRSSM	39.92	54.93	2.85	29.11	31.70	35.33	36.57	35.35	36.83	36.02	42.01	42.34	42.63	43.25	42.56	52.92	48.93	54.12	48.02	51.00
EchoMatch	26.53	28.25	3.15	**4.95**	15.72	**9.21**	14.80	**10.02**	14.51	12.14	**16.91**	**21.53**	**17.80**	**22.64**	**19.72**	23.16	26.70	24.62	28.01	25.62
Wormhole	38.17	55.60	4.56	42.75	35.27	28.82	42.12	29.17	42.19	35.58	30.65	45.50	33.43	46.06	38.91	46.43	47.45	48.27	51.10	48.31
Ours	7.03	7.17	6.53	6.58	**6.83**	12.20	**11.85**	12.36	**11.63**	**12.01**	26.94	27.91	27.60	27.50	27.49	**10.87**	**10.30**	**9.12**	**10.13**	**10.11**
Ours (DPFM)	31.99	49.69	3.29	26.45	27.86	23.13	35.27	20.77	35.47	28.66	26.89	36.76	31.64	38.20	33.37	33.56	38.74	46.38	48.24	41.73

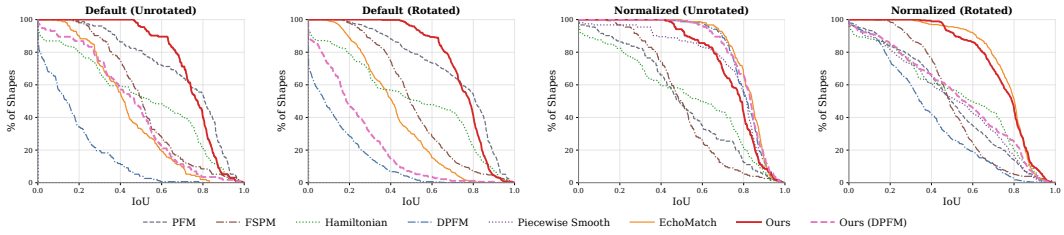

Figure 3: IoU success curves for region localization on CUTS'24 under four input settings. Each point reports the percentage of shapes whose surface-area-weighted IoU exceeds the corresponding threshold.

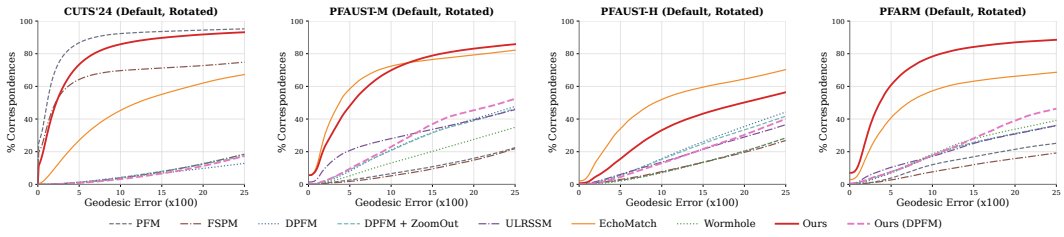

Figure 4: PCK curves under the default-scale, rotated setting.

Evaluation Datasets and Metrics. We evaluate region localization and dense correspondence on the leakage-reduced CUTS'24 (Ehm et al., 2024a) split of SHREC'16 CUTS (Lähner et al., 2016), PFAUST-M and PFAUST-H (Bracha et al., 2024a), and PFARM (Attaiki et al., 2021). The main paper reports localization results on CUTS'24 and correspondence results across all four datasets. Additional localization results on PFAUST-M, PFAUST-H, and PFARM and additional correspondence results are provided in Apps. E.1 and E.2, respectively. To assess sensitivity to scale and orientation, we evaluate both default-scale and normalized inputs under unrotated and independently rotated $\mathrm{SO}(3)$ settings. Region localization is evaluated using surface-area-weighted intersection over union (IoU), precision, recall, and F1 score; IoU success curves report the fraction of test pairs exceeding each IoU threshold. Dense correspondence accuracy is reported using the mean geodesic

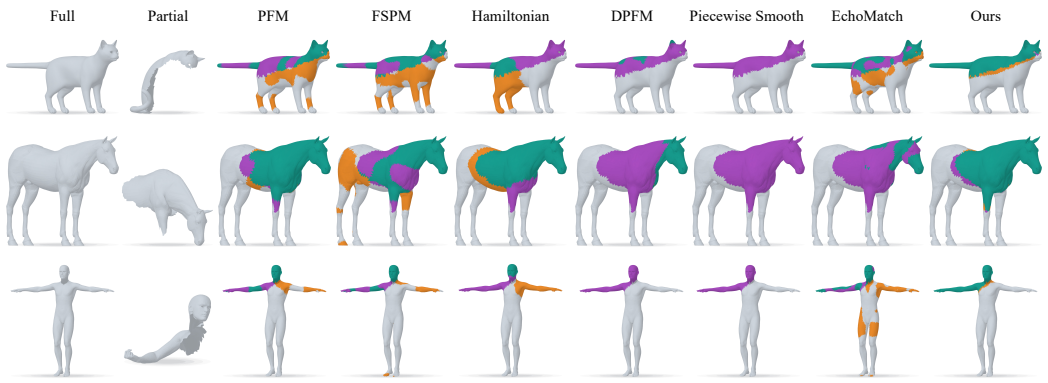

Figure 5: Qualitative comparison at the default scale under $\mathrm{SO}(3)$ rotations. Green, orange, and purple denote correct, false-positive, and missed support, respectively.

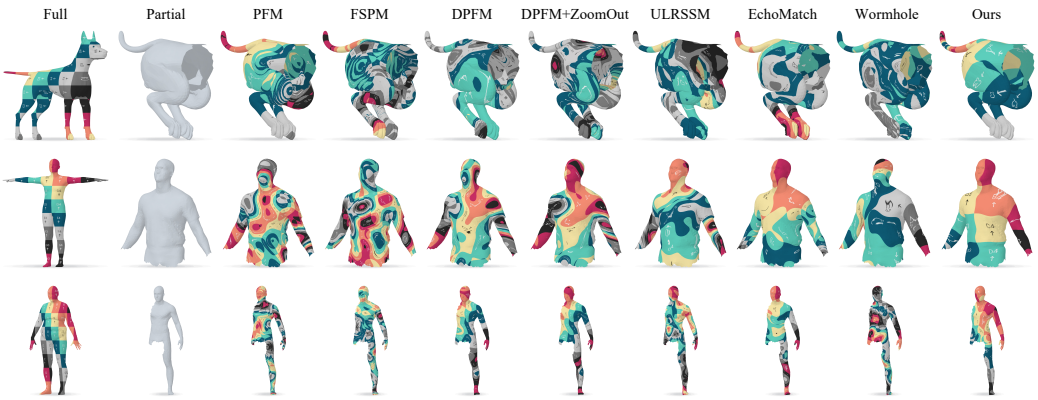

Figure 6: Qualitative comparison of partial-to-full correspondence. Our method yields more accurate correspondences than competing methods, even for challenging non-isometric shape pairs.

error ($\times 100$), with percentage of correct keypoints (PCK) curves illustrating performance across different error thresholds. Rotated results are averaged over three trials.

5.1 Comparative Studies

Region Localization. Tab. 1 compares region localization on CUTS'24 under four scale–rotation settings. Traditional optimization-based methods remain sensitive to scale, whereas feature-learning methods perform well in the normalized regime used for feature training and degrade at the default scale. In contrast, NHO maintains IoU between 75.06 and 75.24 and F1 between 85.02 and 85.37, corresponding to variations of only 0.18 and 0.35 percentage points, respectively. The rotation stability follows from the intrinsic nature of the LBO and Hamiltonian operator. Under uniform rescaling, the area ratio remains unchanged, while the eigenvalues and the potential height $\tau = \beta\lambda_k^{\mathcal{M}}$ scale by the same inverse-square factor, leaving the relative spectral residual invariant. These properties explain why the optimization remains balanced across the tested transformations. The success curves in Fig. 3 show that NHO is substantially more consistent across configurations.

Dense Correspondence. Tab. 2 reports dense correspondence accuracy across four benchmarks and four input settings. Across these settings, NHO achieves the lowest average error on CUTS'24, PFAUST-M, and PFARM, and the second-lowest average error on PFAUST-H. Its advantage lies primarily in consistency across transformations rather than peak performance in a single regime. For example, rotating normalized CUTS'24 increases the error of DPFM with ZoomOut from 2.17 to 24.22, whereas the error of NHO changes only from 6.53 to 6.58. This behavior is consistent with

the learned operator: its localized support removes irrelevant candidates from the search domain, while its eigenfunctions provide intrinsic coordinates adapted to that support. Correspondence recovery therefore depends less on transformation-sensitive feature similarities. The PCK curves in Fig. 4 show that this robustness extends across a range of error thresholds under the challenging default-scale, rotated setting. On normalized, unrotated CUTS'24, using reliable DPFM matches as anchors reduces the NHO error from 6.53 to 3.29, demonstrating that the operator can also incorporate stronger external correspondence evidence.

Qualitative Results. Fig. 5 shows that NHO produces fewer missed and false-positive regions, while Fig. 6 demonstrates accurate dense correspondence, including on challenging non-isometric shape pairs. Additional qualitative and quantitative results are provided in Apps. E.1 and E.2.

5.2 Ablation Study

Tab. 3 and Fig. 7 examine the components of our default sparse-anchor configuration. Removing the topology loss has only a modest effect on the aggregate localization metrics, but produces spurious disconnected regions in the qualitative results. This indicates that the topology term mainly regularizes the spatial structure of the estimated support rather than substantially changing its overall overlap.

Reciprocal refinement improves support selectivity and correspondence accuracy. It increases precision from 78.55 to 80.31 and IoU from 74.37 to 75.15. Correspondence error decreases from 9.23 to 8.37 before NAM. These changes are consistent with reliable correspondence feedback reducing residual spatial ambiguity in an already plausible support, improving the balance between region coverage and the exclusion of distracting geometry.

Replacing the learned Hamiltonian operator with the global full-shape LBO reduces IoU from 75.15 to 55.48 and F1-score from 85.13 to 70.94, while increasing the final geodesic error from 6.53 to 9.84. This replacement produces the largest degradation among the evaluated ablations, highlighting the importance of a spectral representation adapted to the partial support.

Additional analyses are provided in App. F.

Table 3: Component ablations on normalized, unrotated CUTS'24. Localization metrics are area-weighted percentages; correspondence is measured by mean geodesic error ($\times 100$) before and after external NAM refinement. The best result is shown in bold.

	Region localization				Dense correspondence	
Variant	IoU ↑	Precision ↑	Recall ↑	F1-score ↑	Geo. before NAM ↓	Geo. after NAM ↓
w/o topology loss	74.45	80.09	91.28	84.69	8.71	6.85
w/o reciprocal refinement	74.37	78.55	**93.26**	84.61	9.23	6.97
w/o Hamiltonian operator	55.48	68.08	74.20	70.94	12.69	9.84
Full method	**75.15**	**80.31**	92.18	**85.13**	**8.37**	**6.53**

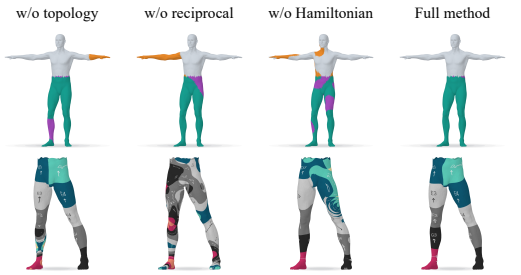

Figure 7: Qualitative ablation results for region localization and dense correspondence. Green, orange, and purple denote correct, false-positive, and missed support, respectively.

6 Conclusion

We presented NHO, which learns a localized Hamiltonian operator on a full surface from sparse anchors and the intrinsic geometry of a partial shape. The learned operator localizes the corresponding region and supports dense matching, while geometrically reliable correspondences feed back into operator estimation. Experiments across multiple benchmarks demonstrate competitive performance and robustness to uniform scaling and rotation, while ablations validate the contributions of the neural Hamiltonian representation, topology regularization, and reciprocal refinement. These results show that our method can use sparse spatial evidence to learn a local operator for coupled region localization and correspondence. They further support the value of localized spectral representations for partial-to-full shape matching.

AI USE STATEMENT

Generative AI tools were used to draft and revise parts of the manuscript, improve readability, and assist with LaTeX formatting. All AI-assisted text and formatting were reviewed by the authors. The authors take responsibility for the final content, including all claims and artifacts produced with the aid of generative AI.

REFERENCES

Souhaib Attaiki, Gautam Pai, and Maks Ovsjanikov. Dpfm: Deep partial functional maps. In *2021 International Conference on 3D Vision (3DV)*, pp. 175–185. IEEE, 2021.

David Bensaïd and Ron Kimmel. A multi-spectral geometric approach for shape analysis. *Journal of Mathematical Imaging and Vision*, 66(4):606–615, 2024.

David Bensaïd, Amit Bracha, and Ron Kimmel. Partial shape similarity via alignment of multi-metric hamiltonian spectra. *arXiv preprint arXiv:2207.03018*, 2022.

David Bensaïd, Noam Rotstein, Nelson Goldenstein, and Ron Kimmel. Partial matching of nonrigid shapes by learning piecewise smooth functions. In *Computer Graphics Forum*, volume 42, pp. e14913. Wiley Online Library, 2023.

Amit Bracha, Thomas Dagès, and Ron Kimmel. On unsupervised partial shape correspondence. In *Asian Conference on Computer Vision*, pp. 316–332. Springer, 2024a.

Amit Bracha, Thomas Dagès, and Ron Kimmel. Wormhole loss for partial shape matching. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024b. URL https://openreview.net/forum?id=gPhBvrPdEs.

Dongliang Cao and Florian Bernard. Hyper-network neural functional maps for unsupervised robust 3d shape matching. *arXiv preprint arXiv:2606.30131*, 2026.

Dongliang Cao, Paul Roetzer, and Florian Bernard. Unsupervised learning of robust spectral shape matching. *ACM Trans. Graph.*, 42(4), July 2023. ISSN 0730-0301. doi: 10.1145/3592107. URL https://doi.org/10.1145/3592107.

Dongliang Cao, Zorah Lähner, and Florian Bernard. Synchronous diffusion for unsupervised smooth non-rigid 3d shape matching. In *European conference on computer vision*, pp. 262–281. Springer, 2024a.

Dongliang Cao, Paul Roetzer, and Florian Bernard. Revisiting map relations for unsupervised non-rigid shape matching. In *2024 International Conference on 3D Vision (3DV)*, pp. 1371–1381. IEEE, 2024b.

Yoni Choukroun, Alon Shtern, Alex Bronstein, and Ron Kimmel. Hamiltonian operator for spectral shape analysis. *IEEE transactions on visualization and computer graphics*, 26(2):1320–1331, 2018.

Edelsbrunner, Letscher, and Zomorodian. Topological persistence and simplification. *Discrete & computational geometry*, 28(4):511–533, 2002.

Viktoria Ehm, Maolin Gao, Paul Roetzer, Marvin Eisenberger, Daniel Cremers, and Florian Bernard. Partial-to-partial shape matching with geometric consistency. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 27478–27487. IEEE, 2024a.

Viktoria Ehm, Paul Roetzer, Marvin Eisenberger, Maolin Gao, Florian Bernard, and Daniel Cremers. Geometrically consistent partial shape matching. In *2024 International Conference on 3D Vision (3DV)*, pp. 914–922. IEEE, 2024b.

Xiaoling Hu, Fuxin Li, Dimitris Samaras, and Chao Chen. Topology-preserving deep image segmentation. *Advances in neural information processing systems*, 32, 2019.

Zorah Lähner, Emanuele Rodola, Michael M Bronstein, Daniel Cremers, Oliver Burghard, Luca Cosmo, Alexander Dieckmann, Reinhard Klein, Y Sahillioğlu, et al. Shrec'16: Matching of deformable shapes with topological noise. In *Eurographics Workshop on 3D Object Retrieval, EG 3DOR*, pp. 55–60. Eurographics Association, 2016.

Or Litany, Emanuele Rodolà, Alexander M Bronstein, and Michael M Bronstein. Fully spectral partial shape matching. In *Computer Graphics Forum*, volume 36, pp. 247–258. Wiley Online Library, 2017.

Simone Melzi, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, and Maks Ovsjanikov. Zoomout: spectral upsampling for efficient shape correspondence. *ACM Trans. Graph.*, 38(6), November 2019. ISSN 0730-0301. doi: 10.1145/3355089.3356524. URL https://doi.org/10.1145/3355089.3356524.

Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (ToG)*, 31(4):1–11, 2012.

Emilian Postolache, Marco Fumero, Luca Cosmo, and Emanuele Rodola. A parametric analysis of discrete hamiltonian functional maps. In *Computer Graphics Forum*, volume 39, pp. 103–118. Wiley Online Library, 2020.

Panjawee Rakprayoon, Miti Ruchanurucks, Somying Thainimit, and Ikuhisa Mitsugami. Part-to-full shape matching of different human subjects. *Heliyon*, 7(10), 2021.

Arianna Rampini, Irene Tallini, Maks Ovsjanikov, Alex M Bronstein, and Emanuele Rodolà. Correspondence-free region localization for partial shape similarity via hamiltonian spectrum alignment. In *2019 International Conference on 3D Vision (3DV)*, pp. 37–46. IEEE, 2019.

Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics (ToG)*, 37(6): 1–16, 2018.

Emanuele Rodolà, Luca Cosmo, Michael M Bronstein, Andrea Torsello, and Daniel Cremers. Partial functional correspondence. In *Computer graphics forum*, volume 36, pp. 222–236. Wiley Online Library, 2017.

Paul Roetzer and Florian Bernard. Fast globally optimal and geometrically consistent 3d shape matching. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 912–922. IEEE, 2025.

Nicholas Sharp, Souhaib Attaiki, Keenan Crane, and Maks Ovsjanikov. Diffusionnet: Discretization agnostic learning on surfaces. *ACM Transactions on Graphics (ToG)*, 41(3):1–16, 2022.

Giulio Viganò, Maks Ovsjanikov, and Simone Melzi. Nam: Neural adjoint maps for refining shape correspondences. *ACM Transactions on Graphics (TOG)*, 44(4):1–15, 2025.

Yan Wu, Jun Yang, and Jinlong Zhao. Partial 3d shape functional correspondence via fully spectral eigenvalue alignment and upsampling refinement. *Computers & Graphics*, 92:99–113, 2020.

Rui Xiang, Rongjie Lai, and Hongkai Zhao. A dual iterative refinement method for non-rigid shape matching. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15925–15934. IEEE, 2021.

Yizheng Xie, Viktoria Ehm, Paul Roetzer, Nafie El Amrani, Maolin Gao, Florian Bernard, and Daniel Cremers. Echomatch: Partial-to-partial shape matching via correspondence reflection. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11665–11675. IEEE, 2025.

APPENDIX

A METHODOLOGICAL COMPARISON

Table 4: Comparison with representative partial-to-full correspondence methods. "Anchor-source agnostic" indicates that a method can accept sufficiently reliable sparse anchors without prescribing how they are obtained.

Method	Region-localizing operator	Explicit region localization	Dense correspondence	Uniform-scale and rotation robustness	Anchor-source agnostic
PFM(Rodolà et al., 2017)	✗	✓	✓	✗	✗
FSPM(Litany et al., 2017)	✗	✗	✓	✗	✗
Hamiltonian(Rampini et al., 2019)	✓	✓	✗	✗	✗
DPFM(Attaiki et al., 2021)	✗	✓	✓	✗	✗
Piecewise Smooth(Bensaïd et al., 2023)	✗	✓	✗	✗	✗
ULRSSM(Cao et al., 2023)	✗	✗	✓	✗	✗
RMR(Cao et al., 2024b)	✗	✗	✓	✗	✗
Synchronous Diffusion(Cao et al., 2024a)	✗	✗	✓	✗	✗
EchoMatch(Xie et al., 2025)	✗	✓	✓	✗	✗
NFM(Cao & Bernard, 2026)	✗	✗	✓	✗	✓
Ours	✓	✓	✓	✓	✓

Tab. 4 summarizes the methodological differences between NHO and representative partial-to-full correspondence approaches. NHO retains the standard Hamiltonian construction $H = \Delta + v$; its distinction lies in parameterizing the potential as an intrinsic neural field, estimating it from sparse anchors and partial-shape geometry, and using the resulting localized eigenspace for both support recovery and dense correspondence. Classical Hamiltonian spectrum alignment uses the operator primarily for region localization, whereas correspondence-oriented methods do not explicitly learn a region-localizing operator. NHO connects these roles through reciprocal operator–correspondence refinement.

B VISUALIZATION OF RECIPROCAL REFINEMENT

Fig. 8 visualizes the evolution of region localization and correspondence during reciprocal refinement. The initial estimate obtained from sparse anchors contains both false-positive and missed regions, accompanied by inaccurate correspondences. As geometrically reliable matches are fed back into operator estimation, the support becomes progressively more accurate and, in turn, provides a more reliable domain for subsequent correspondence recovery. The simultaneous reduction in support errors and improvement in correspondence consistency illustrate the reciprocal interaction between the two tasks.

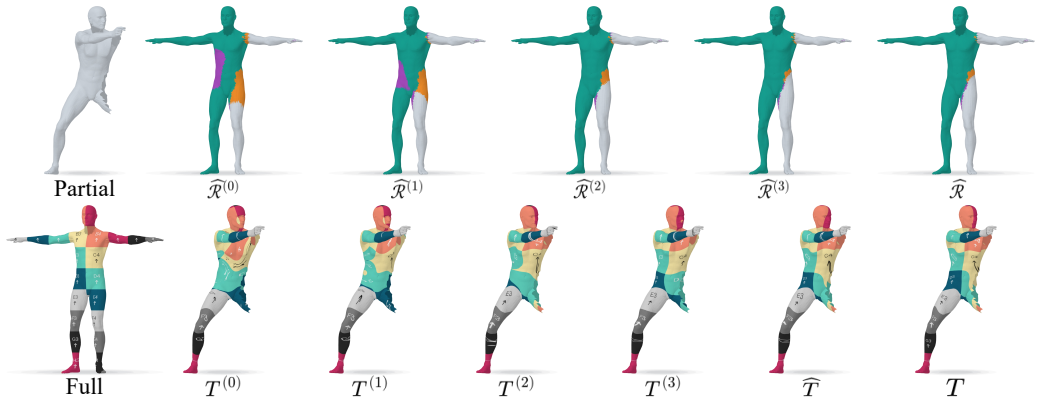

Figure 8: Visualization of reciprocal refinement. The estimated support and correspondence progressively improve across rounds. Green, orange, and purple denote correct, false-positive, and missed support, respectively.

Fig. 9 further shows how reciprocal refinement changes the learned spectral representation. For reference, the top row presents the LBO eigenfunctions of the full shape and the Dirichlet eigenfunctions of the partial shape. The bottom row compares the anchor-aligned Hamiltonian eigenfunctions

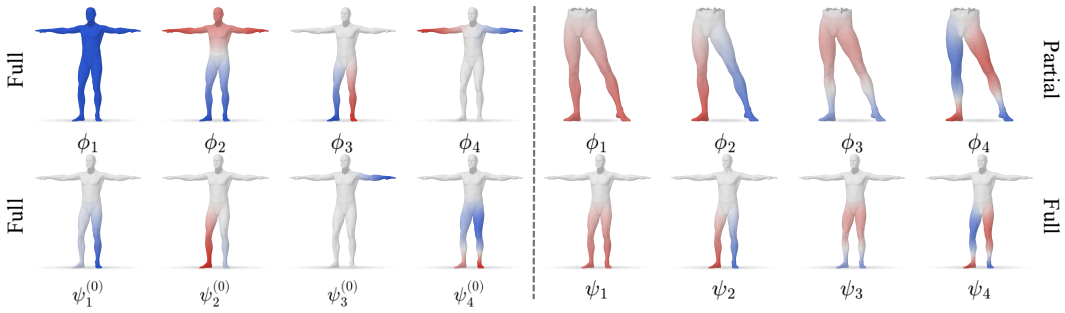

Figure 9: Evolution of the Hamiltonian eigenfunctions during reciprocal refinement. Top: the first four LBO eigenfunctions of the full shape (left) and Dirichlet eigenfunctions of the partial shape (right). Bottom: the anchor-aligned neural Hamiltonian eigenfunctions after initial operator estimation (left) and reciprocal refinement (right).

before and after reciprocal refinement. Initially, the Hamiltonian modes remain weakly localized and differ from the partial Dirichlet modes. After refinement, they concentrate on the corresponding region and exhibit more consistent spatial patterns, explaining the improved localization and correspondence.

C DPFM-Based Anchor Initialization

Ours (DPFM) replaces the annotated anchors with correspondences derived from features extracted by a pretrained Siamese DiffusionNet (Sharp et al., 2022; Attaiki et al., 2021). Let $\mathbf{F}_{\mathcal{M}} \in \mathbb{R}^{n_{\mathcal{M}} \times d_f}$ and $\mathbf{F}_{\mathcal{N}} \in \mathbb{R}^{n_{\mathcal{N}} \times d_f}$ denote the row-wise ℓ_2-normalized feature matrices of the partial and full shapes, respectively. We first compute a dense feature-based nearest-neighbor map

$$\widetilde{T}_{\mathcal{M}\to\mathcal{N}}(p) = \arg\max_{q\in\mathcal{N}} \langle \mathbf{F}_{\mathcal{M}}(p), \mathbf{F}_{\mathcal{N}}(q) \rangle . \tag{22}$$

Since the features are normalized, maximizing their inner product is equivalent to minimizing their Euclidean distance.

The initial operator optimization and spectral alignment impose different requirements on these matches. The anchor loss in Eq. 10 uses their target vertices as positive spatial evidence; an incorrect match can therefore directly distort the estimated support, making precision the primary concern. In contrast, estimating the alignment matrix in Eq. 19 requires explicit correspondence pairs with sufficient coverage. We accordingly construct separate initial anchor sets for the two operations.

We measure the geometric consistency of $\widetilde{T}_{\mathcal{M}\to\mathcal{N}}$ using local mapping distortion (LMD) (Xiang et al., 2021), computed over the 30 nearest neighbors of each partial vertex under mesh-edge geodesic distance. The initial anchors for operator optimization are

$$\mathcal{A}_H^{(0)} = \left\{ \left(p, \widetilde{T}_{\mathcal{M}\to\mathcal{N}}(p)\right) \;\middle|\; \mathrm{LMD}_{\widetilde{T}}(p) < \delta_H \right\}, \tag{23}$$

where $\delta_H = 0.42$. We use $\mathcal{A}_H^{(0)}$ in $\mathcal{L}_{\mathrm{op}}(\theta; \mathcal{A}_H^{(0)})$ to obtain the initial operator.

For spectral alignment, we additionally require bidirectional feature consistency. The reverse nearest-neighbor map is

$$\widetilde{T}_{\mathcal{N}\to\mathcal{M}}(q) = \arg\max_{p\in\mathcal{M}} \langle \mathbf{F}_{\mathcal{N}}(q), \mathbf{F}_{\mathcal{M}}(p) \rangle , \tag{24}$$

from which we construct

$$\mathcal{A}_C^{(0)} = \left\{ (p,q) \;\middle|\; q = \widetilde{T}_{\mathcal{M}\to\mathcal{N}}(p),\ p = \widetilde{T}_{\mathcal{N}\to\mathcal{M}}(q),\ \mathrm{LMD}_{\widetilde{T}}(p) < \delta_C \right\}, \tag{25}$$

where $\delta_C = 1.5$. This set is used in Eq. 19 to estimate the initial alignment matrix $C^{(0)}$.

Table 5: Quantitative comparison of region localization on PFAUST-M, PFAUST-H, and PFARM under four input settings. All metrics are surface-area-weighted percentages; higher is better. Best in bold; second-best underlined.

	Default								Normalized							
	Unrotated				Rotated				Unrotated				Rotated			
Method	IoU	Precision	Recall	F1-score	IoU	Precision	Recall	F1-score	IoU	Precision	Recall	F1-score	IoU	Precision	Recall	F1-score
PFAUST-M																
PFM	44.93	55.32	71.39	61.84	46.45	55.93	73.97	63.23	51.32	57.44	83.14	67.50	50.14	56.87	81.50	66.49
FSPM	51.52	51.52	**85.00**	63.92	51.52	51.52	**85.00**	63.92	51.52	51.52	**85.00**	63.92	51.52	51.52	**85.00**	63.92
Hamiltonian	51.48	60.66	75.31	65.72	51.72	61.01	75.49	66.05	53.55	62.62	78.59	68.28	52.96	62.41	77.48	67.62
DPFM	13.95	68.74	15.01	24.02	19.20	61.35	22.05	31.77	19.75	66.73	21.94	32.69	17.82	57.22	20.54	29.94
Piecewise Smooth	36.77	57.72	50.01	53.17	50.39	60.36	75.03	65.96	48.37	64.30	67.35	64.46	51.76	61.33	77.32	67.37
EchoMatch	51.69	**81.42**	58.43	67.10	41.46	**77.42**	46.74	57.11	53.76	**79.44**	61.58	68.73	46.81	**76.44**	53.78	62.10
Ours	**55.89**	70.89	71.58	**70.77**	**56.39**	71.90	71.95	**71.53**	**56.28**	71.74	71.87	**71.40**	**55.71**	71.09	71.50	**70.90**
Ours (DPFM)	47.47	61.81	65.81	63.54	48.19	60.17	69.71	64.33	52.97	64.38	73.41	68.50	46.64	58.58	67.30	62.45
PFAUST-H																
PFM	46.36	52.37	80.53	**63.18**	46.09	52.30	80.05	**62.93**	**49.28**	53.37	**86.58**	**65.91**	**48.69**	53.09	**85.57**	**65.38**
FSPM	**46.75**	46.75	**85.00**	60.23	**46.75**	46.75	**85.00**	60.23	46.75	46.75	85.00	60.23	46.75	46.75	85.00	60.23
Hamiltonian	45.38	58.11	67.30	61.54	44.84	58.13	66.12	60.97	44.88	58.21	66.22	61.02	44.63	58.13	65.84	60.79
DPFM	13.62	61.75	14.87	23.70	18.52	54.37	22.00	30.80	19.25	58.67	22.13	32.03	17.62	51.84	20.98	29.66
Piecewise Smooth	32.31	53.45	45.06	48.17	43.49	55.02	67.91	59.83	36.30	56.96	50.61	52.15	43.09	55.52	67.24	59.48
EchoMatch	37.04	**77.44**	41.98	52.95	32.85	**69.71**	38.87	48.10	35.96	**77.56**	40.30	51.81	31.80	**69.72**	37.13	47.02
Ours	39.55	61.12	52.27	56.07	39.43	61.09	52.59	56.14	40.20	61.64	53.50	56.90	40.53	61.76	53.89	57.15
Ours (DPFM)	43.00	58.93	61.11	59.82	38.69	54.02	57.49	55.43	44.71	60.21	63.36	61.54	43.05	58.34	61.82	59.81
PFARM																
PFM	15.21	33.12	20.80	24.62	14.32	32.31	19.65	23.63	13.70	31.58	18.42	22.57	14.33	32.22	19.40	23.49
FSPM	20.04	20.04	36.00	25.35	19.03	20.38	34.36	24.75	22.64	25.20	39.50	29.50	21.98	24.67	37.95	28.39
Hamiltonian	37.18	46.27	46.71	45.64	37.64	47.89	45.24	45.77	30.59	40.25	38.85	38.86	33.10	42.91	39.91	40.74
DPFM	14.09	35.31	22.97	24.17	13.29	32.37	22.12	22.77	16.48	37.35	27.31	27.59	17.24	38.32	26.77	27.97
Piecewise Smooth	37.07	44.05	65.48	49.29	35.24	40.07	66.57	46.82	37.49	46.68	66.17	50.11	36.89	41.81	67.95	48.89
EchoMatch	48.76	**72.73**	59.33	62.71	45.30	**69.83**	56.15	59.21	47.32	**73.40**	53.46	60.64	42.68	**71.33**	49.40	56.56
Ours	**62.71**	68.92	**90.84**	**74.80**	**62.76**	68.83	**90.78**	**74.73**	**63.50**	69.36	**91.55**	**75.62**	**63.60**	69.34	**91.72**	**75.56**
Ours (DPFM)	36.36	46.51	58.29	51.37	42.60	52.98	61.03	56.32	26.81	37.79	42.89	39.98	37.78	50.29	55.76	52.64

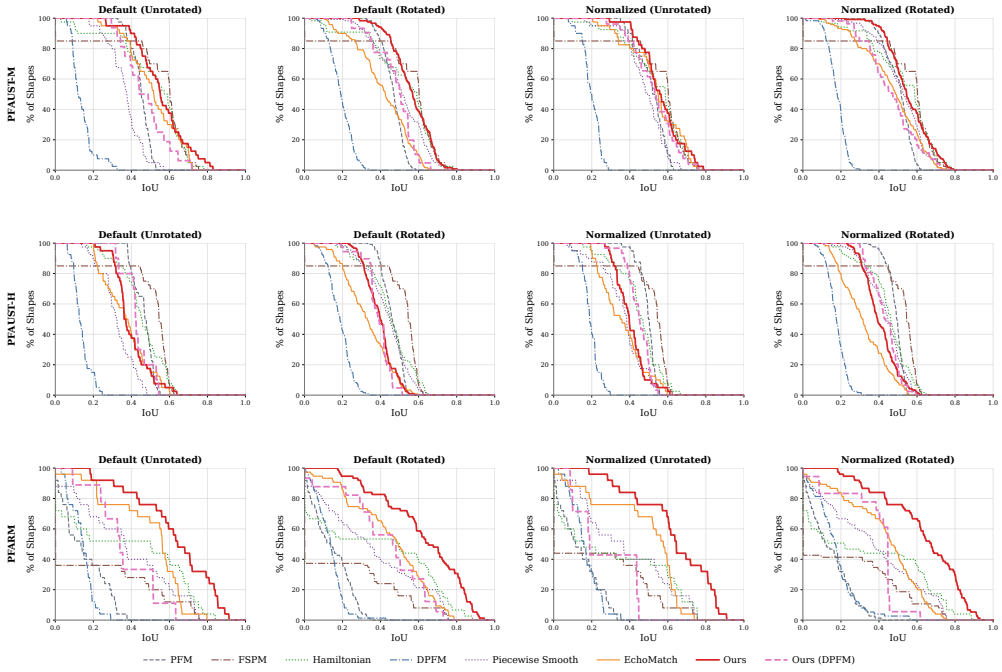

Figure 10: IoU success curves for region localization on PFAUST-M, PFAUST-H, and PFARM under four input settings.

Thus, $\mathcal{A}_H^{(0)}$ uses a strict LMD threshold to provide high-confidence spatial constraints for the anchor loss, whereas $\mathcal{A}_C^{(0)}$ combines mutual-nearest-neighbor consistency with a less restrictive threshold to retain sufficient pairs for spectral alignment. Subsequent operator–correspondence refinement



Figure 11: Additional qualitative comparisons of region localization. Numbers report the surface-area-weighted IoU for each prediction. Green, orange, and purple denote correctly localized, false-positive, and missed regions, respectively.

proceeds as described in Sec. 4.2. Neither initialization uses ground-truth region or correspondence annotations.

D Implementation Details

The neural Hamiltonian field has three hidden layers of width 128. We optimize the initial operator and each of four reciprocal updates for 100 Adam steps with a learning rate of 10^{-2}. We use $r = 20$ LBO eigenfunctions for intrinsic positional encoding, $k = 20$ Hamiltonian eigenpairs for

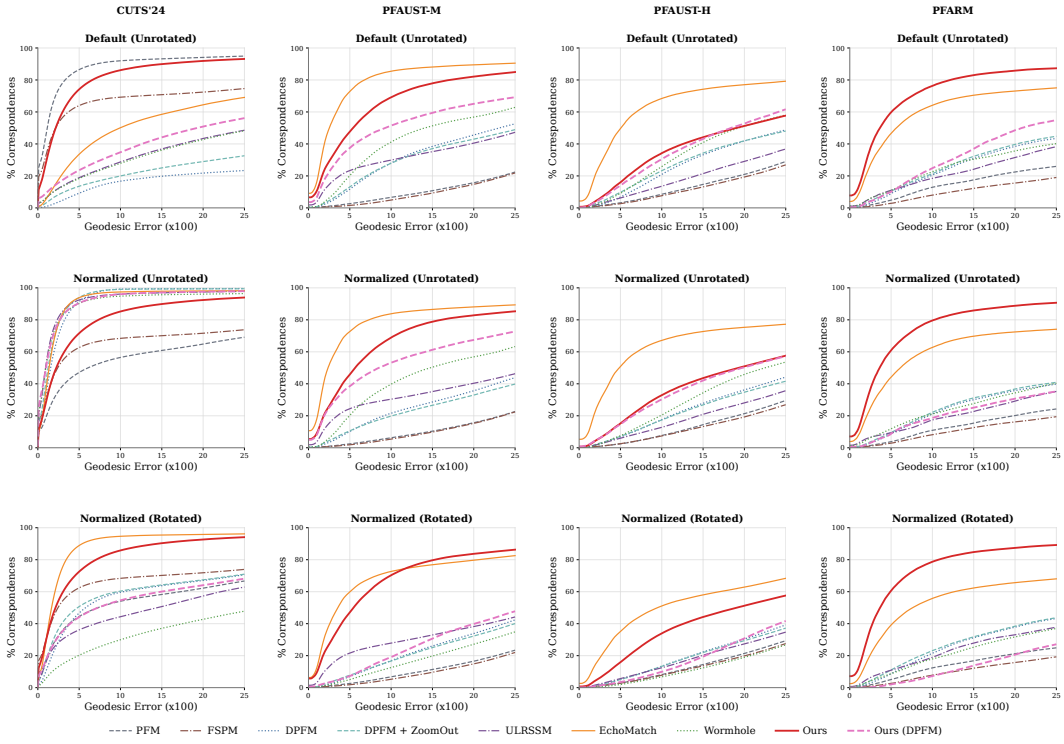

Figure 12: PCK curves on four benchmarks under the three input settings not shown in Fig. 4. Each point reports the percentage of correspondences whose geodesic error ($\times 100$) does not exceed the threshold.

spectral fitting and reciprocal alignment, and $K = 100$ dimensions for Galerkin projection. We fix $(\alpha_1, \alpha_2, \alpha_3, \alpha_4) = (0.1, 1.0, 0.3, 0.01)$, $\beta = 50$, and $\eta = 0.01$ across all datasets. LMD is computed over the 30 nearest neighbors under mesh-edge geodesic distance, and matches below 0.42 are retained.

After reciprocal refinement, we freeze the neural-field parameters and use the final-round map to initialize correspondence refinement. We independently apply per-coordinate RMS normalization to the Hamiltonian coordinates induced by the frozen operator and the partial Dirichlet coordinates. We then apply an adapted Neural ZoomOut (Viganò et al., 2025), increasing the spectral dimension from 20 to 100 in steps of 5.

E Additional Localization and Correspondence Results

We provide additional quantitative and qualitative results for region localization and dense correspondence, complementing the comparisons in Sec. 5.

E.1 Region Localization

Tab. 5 reports region localization results on PFAUST-M, PFAUST-H, and PFARM under four scale–rotation settings. On PFAUST-M, our default sparse-anchor configuration achieves the highest IoU and F1 score across all four settings. On PFARM, it also ranks first in both metrics. Its IoU ranges from 62.71% to 63.60%, and its F1 score ranges from 74.73% to 75.62%, while recall remains above 90% throughout the four settings.

PFAUST-H is substantially more challenging. It contains a larger number of small missing regions, producing multiple additional boundary components and more pronounced topological changes (Bracha et al., 2024a). Although the Hamiltonian–Dirichlet connection remains applicable to such regions, their fine boundary structure is difficult to recover from a truncated low-frequency

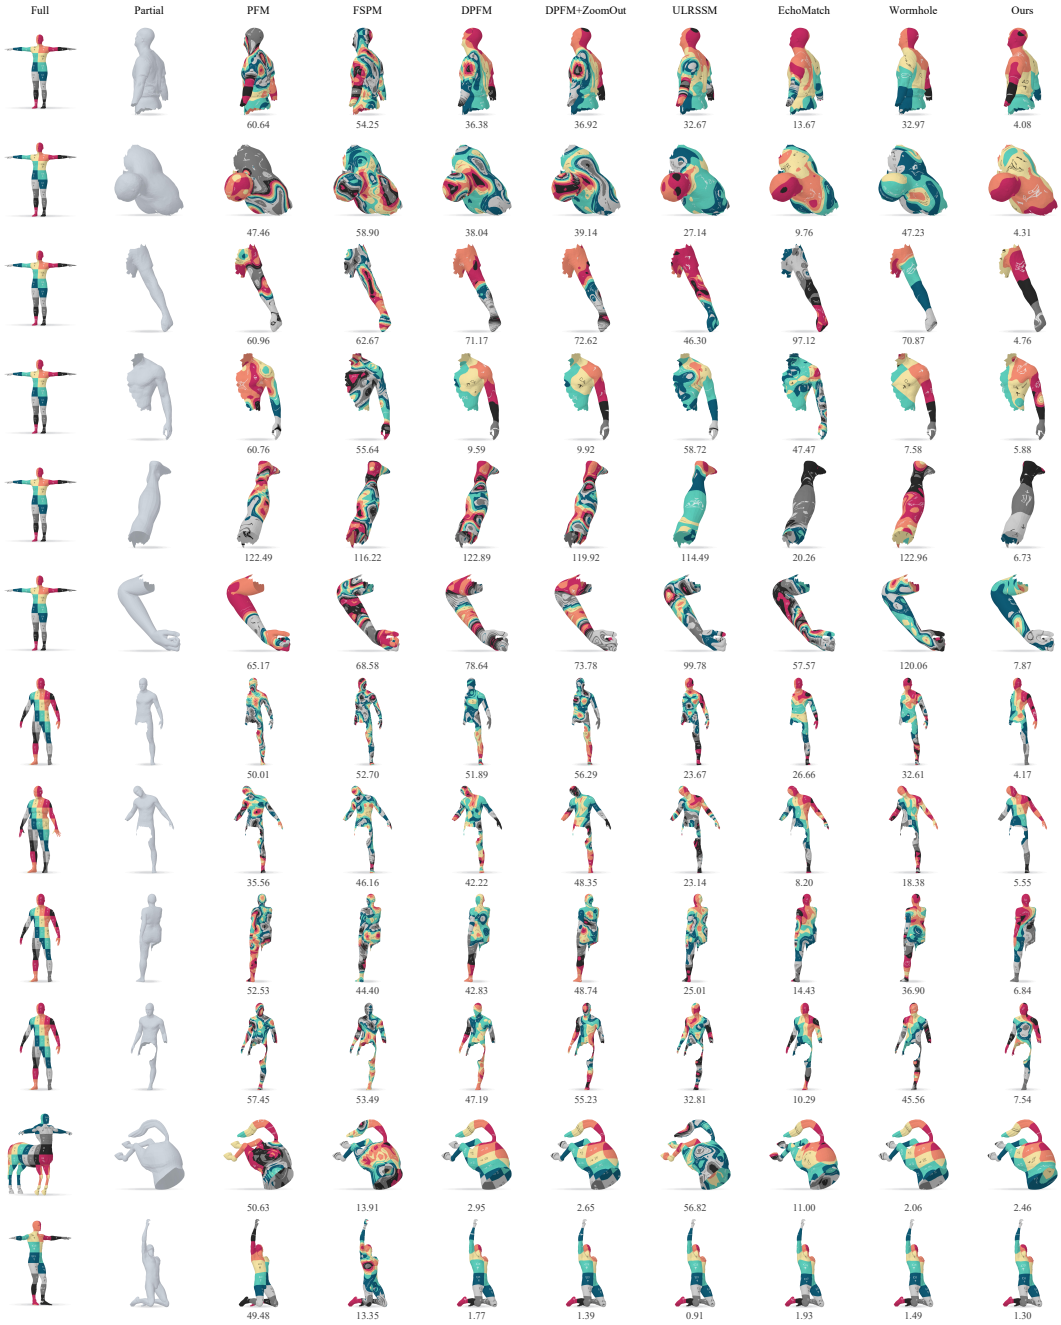

Figure 13: Additional qualitative comparisons of partial-to-full dense correspondence. Numbers report the mean geodesic error ($\times 100$) for each case; lower is better.

spectrum and a smoothly parameterized potential. Moreover, our zero-dimensional topological regularizer promotes connectivity but does not explicitly characterize the hole structure. The lower IoU and recall may therefore reflect the difficulty of resolving fine-scale hole boundaries using a truncated low-frequency representation.

Despite the variation in absolute accuracy across datasets, sparse-anchor NHO remains stable under changes in scale and orientation. Its IoU varies by only 0.68 percentage points on PFAUST-M, 1.10 on PFAUST-H, and 0.89 on PFARM across the four settings. The success curves in Fig. 10 provide a

Table 6: Anchor sampling on normalized, unrotated CUTS'24. Localization scores are area-weighted percentages; Geo. is the final mean geodesic error ($\times 100$). Best values are bold.

	Region localization				Dense correspondence
Sampling strategy	IoU ↑	Precision ↑	Recall ↑	F1-score ↑	Geo. ↓
Geodesic FPS	**75.15**	**80.31**	92.18	**85.13**	**6.53**
Euclidean FPS	75.08	79.97	**92.38**	85.05	6.81
Uniform random	72.38	78.67	89.73	83.21	7.92

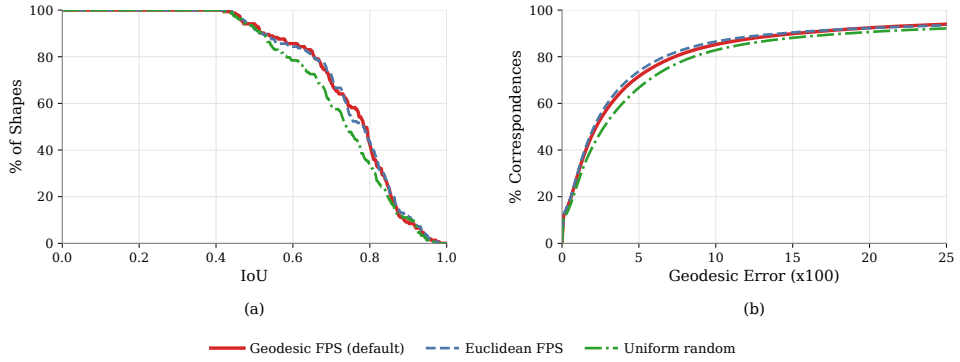

Figure 14: Anchor sampling: (a) IoU success and (b) correspondence PCK curves.

Table 7: Initial-anchor budget on normalized, unrotated CUTS'24. Metrics follow Table 6; best values are bold.

	Region localization				Dense correspondence
m	IoU ↑	Precision ↑	Recall ↑	F1-score ↑	Geo. ↓
5	66.87	74.31	85.54	78.89	17.73
15	72.78	78.68	90.26	83.38	11.43
25	75.15	80.31	92.18	85.13	6.53
50	75.49	80.60	92.14	85.29	4.75
100	**77.47**	**81.73**	**93.81**	**86.67**	**3.95**

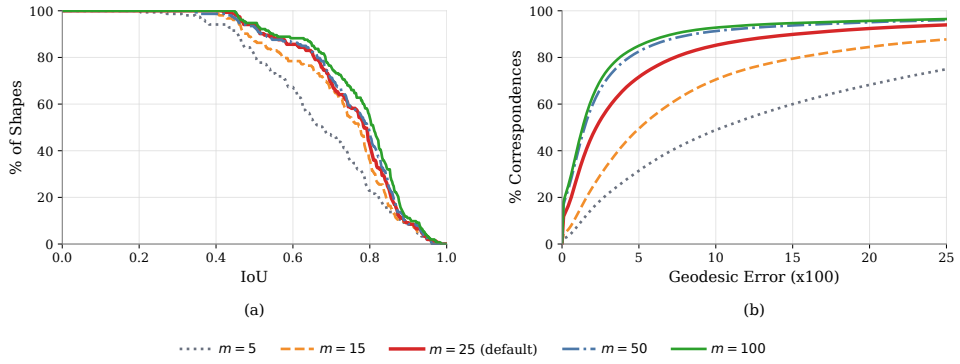

Figure 15: Initial-anchor budget: (a) IoU success and (b) correspondence PCK curves.

distribution-level view of this consistency, while Fig. 11 presents additional qualitative comparisons of region localization.

E.2 Dense Correspondence

Tab. 2 reports mean geodesic errors across four benchmarks and four input settings, while Figs. 4 and 12 provide the corresponding threshold-wise comparisons. The within-dataset error ranges of sparse-anchor NHO are 0.64, 0.73, 0.97, and 1.75 on the four benchmarks, respectively. This con-

Table 8: Robustness to mesh discretization on normalized, unrotated CUTS'24. Localization metrics are surface-area-weighted percentages; Geo. denotes the final mean geodesic error ($\times 100$). Best values are shown in bold.

	Region localization				Dense correspondence
Full faces retained	IoU ↑	Precision ↑	Recall ↑	F1-score ↑	Geo. ↓
50%	80.90	86.75	92.21	88.77	**6.74**
20%	82.24	**88.02**	92.64	89.67	7.20
10%	**82.45**	87.92	**92.93**	**89.74**	8.39
5%	80.44	86.16	92.17	88.52	10.56

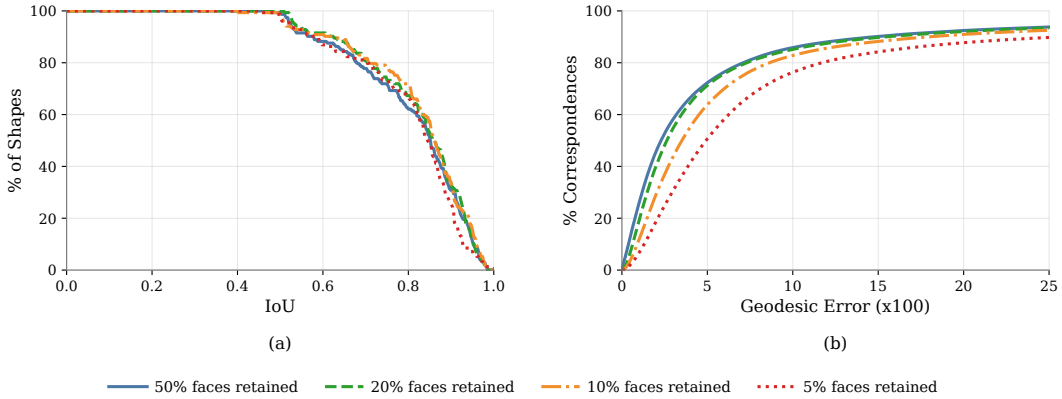

Figure 16: Robustness to mesh discretization on normalized, unrotated CUTS'24. (a) IoU success curves for region localization. (b) PCK curves for dense correspondence. Percentages denote the retained full-shape faces.

sistency indicates that the correspondence recovered from the learned operator is largely insensitive to the tested scale and orientation changes.

The PCK curves reveal additional behavior beyond the mean errors. On PFARM, NHO achieves higher PCK over most error thresholds in all four settings, showing that its advantage extends beyond the mean geodesic error to the overall distribution of correspondence accuracy. On PFAUST-M, EchoMatch is more accurate in the unrotated settings, whereas NHO becomes stronger after rotation: its error changes from 12.20 to 11.85 at the default scale and from 12.36 to 11.63 after normalization, while the corresponding EchoMatch errors increase from 9.21 to 14.80 and from 10.02 to 14.51. This comparison further highlights the rotation stability of the intrinsic operator representation.

Across the four benchmarks, sparse-anchor NHO records its lowest localization IoU and highest mean geodesic error on PFAUST-H. This shared degradation is consistent with the coupling between the two tasks: an inaccurate support may exclude valid target vertices or retain distracting regions, directly affecting the restricted correspondence search. Nevertheless, the correspondence error on PFAUST-H varies by only 0.97 across the four settings, suggesting that its lower absolute accuracy is primarily associated with the more challenging partiality pattern discussed in Sec. E.1, rather than sensitivity to scale or rotation. Fig. 13 provides additional qualitative comparisons of the recovered dense correspondences.

F Additional Analysis

Effect of Anchor Sampling. We fix the number of initial anchors at $m = 25$ and compare the default geodesic FPS with Euclidean FPS and uniform random sampling over the annotated vertices of the partial shape. All other settings remain unchanged, and random sampling uses a fixed draw for each shape pair.

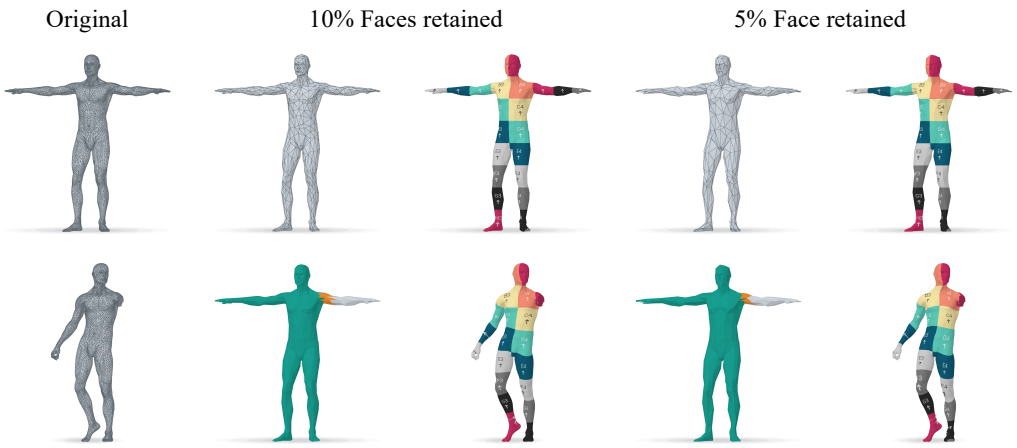

Figure 17: Qualitative results under severe discretization mismatch. The full meshes retain 10% and 5% of their original faces, while the partial meshes remain unchanged. Green, orange, and purple denote correct, false-positive, and missed support, respectively.

Tab. 6 and Fig. 14 show that geodesic and Euclidean FPS yield nearly identical performance, with a modest decline under uniform random sampling. These results support NHO's robustness across the evaluated sampling strategies, while the more uniform spatial coverage provided by FPS benefits both localization and correspondence.

Effect of the Anchor Budget. We vary the number of initial anchors $m \in \{5, 15, 25, 50, 100\}$ on normalized, unrotated CUTS'24 while keeping all other settings fixed. When $m < k = 20$, we set the spectral correspondence dimension to m to avoid an underdetermined anchor-based alignment.

Tab. 7 and Fig. 15 show that NHO remains effective with limited spatial evidence. Increasing m from 5 to the default budget of 25 improves IoU from 66.87 to 75.15 and reduces the final geodesic error from 17.73 to 6.53. Beyond 25 anchors, localization improves more gradually, whereas correspondence continues to benefit from the additional constraints, reaching an error of 3.95 with 100 anchors. We therefore use $m = 25$ uniformly across datasets as a sparse-anchor budget that balances supervision and accuracy, rather than as the best-performing anchor count.

Robustness to Mesh Discretization. We evaluate robustness to discretization mismatch by reducing the resolution of the full meshes to 50%, 20%, 10%, and 5% of their original faces while leaving the partial meshes unchanged. The geometric operators are recomputed at each resolution, with all other settings fixed.

Tab. 8 and Fig. 16 show that region localization remains stable across the tested resolutions. IoU varies only from 80.44% to 82.45%, while F1 ranges from 88.52% to 89.74%, and the corresponding success curves largely overlap. This stability is consistent with the low-frequency intrinsic formulation: when large-scale geometry is preserved, the low-frequency LBO and Hamiltonian subspaces remain comparatively stable under changes in mesh connectivity and sampling density. Localization through aggregated eigenfunction energy therefore does not rely on compatible discretizations of the partial and full shapes.

Dense correspondence is more sensitive to aggressive simplification. Reducing the retained full-shape faces from 50% to 5% increases the final mean geodesic error from 6.74 to 10.56, with a corresponding gradual decline in the PCK curves. Pointwise recovery requires resolving local geometry and selecting individual target vertices, both of which become less precise on coarser meshes. Moreover, the higher-frequency coordinates used during correspondence refinement are more sensitive to discretization than the low-frequency energy used for localization. Nevertheless, the gradual degradation shows that NHO retains useful correspondence accuracy even when the full and partial shapes have substantially different resolutions. Fig. 17 presents qualitative results at the two most aggressive simplification levels.

G COMPUTATIONAL COST

To reduce dependence on a dataset-specific training distribution, NHO optimizes each shape pair independently rather than using amortized feed-forward inference. This design improves adaptability across datasets and input transformations at the cost of additional test-time computation. The neural Hamiltonian field and the final 100-dimensional NAM stage contain 35.84K and 52.34K parameters, respectively, with forward costs of 0.709G and 1.040G FLOPs for $10{,}000$ vertices. Operator estimation requires 500 Adam steps, including the initial estimate and four reciprocal updates, while multi-resolution correspondence refinement requires $3{,}400$ steps. On an NVIDIA RTX 4090, the complete pipeline takes 143.15 seconds per shape pair on average, excluding cached spectral preprocessing.