

SGMatch: Semantic-Guided Non-Rigid Shape Matching with Flow Regularization

Tianwei Ye , Xiaoguang Mei , Yifan Xia ,
Fan Fan , Jun Huang , and Jiayi Ma  *

Wuhan University, Wuhan 430072, China
twye2001@gmail.com, jyma2010@gmail.com

Abstract. Establishing accurate point-to-point correspondences between non-rigid 3D shapes remains a critical challenge, particularly under non-isometric deformations and topological noise. Existing functional map pipelines suffer from ambiguities that geometric descriptors alone cannot resolve, and spatial inconsistencies inherent in the projection of truncated spectral bases to dense pointwise correspondences. In this paper, we introduce SGMatch, a learning-based framework that couples 3D-lifted semantic cues with trajectory-level feature transport regularization. Specifically, we design a Semantic-Guided Local Cross-Attention module that integrates semantic features from vision foundation models into geometric descriptors while preserving local structural continuity. Furthermore, we adapt conditional flow matching as a time-conditioned feature transport regularizer that promotes spatially coherent point-wise recovery. Experimental results on multiple benchmarks demonstrate that SGMatch achieves competitive performance across near-isometric settings and consistent improvements under non-isometric deformations and topological noise.

Keywords: Shape Matching · Functional Maps · Flow Matching

1 Introduction

Establishing point-to-point correspondences between non-rigid 3D shapes is a long-standing challenge in computer vision and graphics, serving as a prerequisite for a wide range of downstream applications such as texture transfer [16], pose transfer [62, 63], and statistical shape analysis [20, 36, 42]. While recent deep functional map frameworks [6, 13, 24, 72] have shown significant progress, they still struggle with strong non-isometric deformations and topological noise. Moreover, relying solely on geometric descriptors can lead to ambiguities in symmetric regions. More critically, pointwise correspondences reconstructed from truncated spectral bases are prone to local inconsistencies [12, 68], even when the alignment of the spectra appears globally plausible.

Since correct mappings should align meaningful parts across different shapes, correspondence is inherently semantic. Recent studies [1, 19, 47, 70] suggest that

* Corresponding author.




Fig. 1: (Left): Colormap transfer on the SHREC’19 dataset demonstrates that incorporating semantic features resolves ambiguity and yields globally consistent correspondences. **(Right):** Vertex transfer on the SMAL dataset shows that the proposed conditional flow matching regularization promotes spatially smooth correspondences.

semantic cues from foundation models can be lifted onto 3D surfaces to produce robust zero-shot descriptors. However, integrating these signals into functional map pipelines is non-trivial; naive use of global semantic information tends to override local geometric structure, potentially disrupting rather than stabilizing the recovery of pointwise correspondences. We instead treat semantic cues as structure-aware anchors and constrain them to respect the manifold locality. This allows global semantics to disambiguate symmetries while preserving geometric continuity. Despite these semantic enhancements, the spatial inconsistency inherent in spectral truncation calls for a complementary treatment: we introduce continuous feature transport across manifolds as a regularization process to encourage smoothness in the recovered correspondences.

To this end, we propose SGMATCH, a unified functional map framework that addresses the two sources of failure above. For descriptor ambiguity, the proposed Semantic-Guided Local Cross-Attention (SGLCA) module uses lifted semantic features to modulate geometric descriptors and restricts semantic interaction to mesh neighborhoods, thereby preserving local structure while improving part-level discrimination. For local inconsistency in dense recovery, we adapt conditional flow matching (CFM) [38, 39] to regularize feature transport between shapes, encouraging recovered correspondences to follow coherent trajectories in feature space. As shown in Fig. 1, SGMATCH improves semantic consistency in ambiguous regions and spatial coherence under challenging deformations. Extensive experiments on diverse benchmarks demonstrate consistent gains over state-of-the-art approaches under non-isometric deformations and topological noise, while maintaining competitive performance in near-isometric cases.

Our contributions are summarized as follows:

- We propose a semantic-guided functional map framework that couples 3D-lifted semantic cues with trajectory-level feature transport regularization.
- We design the Semantic-Guided Local Cross-Attention module to inject semantic context into geometric descriptors via channel-wise modulation and mesh-local attention, reducing ambiguity while preserving local structure.
- We adapt conditional flow matching as an amortized feature-transport regularizer that promotes spatially coherent point-wise recovery.

- SGMatch achieves consistent improvements over state-of-the-art methods, particularly under non-isometric deformations and topological noise.

2 Related Work

In this section, we focus on the research directions most relevant to the current study, rather than providing an exhaustive survey of the field [59, 67].

2.1 Functional Maps

The functional maps framework [50] is a foundational paradigm in non-rigid shape matching by shifting the problem from a search for discrete pointwise assignments to the estimation of a linear operator between functional spaces. By projecting functions onto the Laplace–Beltrami eigenbasis [53], the pipeline facilitates the representation of dense correspondences as compact, low-dimensional spectral matrices [50]. This formulation inherently enforces a low-frequency structural prior, which ensures stability under near-isometric deformations while significantly reducing the search space for optimization.

Subsequent work has improved functional maps via spectral refinement [25, 46], structured regularization [56], and extensions to non-isometric [48, 55, 57], partial [41, 58], and multi-shape matching [26, 30, 31, 69]. Learning-based pipelines further replace hand-crafted descriptors [3, 9] with task-adaptive features, progressing from supervised [18, 27, 40, 65] to unsupervised settings [6, 10, 13, 23, 64], with recent generative priors also explored in the spectral domain [52, 75]. Nevertheless, functional maps remain sensitive to descriptor discriminability, and point-wise correspondences recovered from truncated spectral bases can still exhibit local inconsistencies under strong non-isometric deformation and topological noise. Recent regularization methods therefore revisit smoother point-wise recovery through diffusion-based consistency constraints [12] or coupled spectral-spatial interpolation [11].

2.2 Semantic Features in Shape Matching

The emergence of large-scale vision foundation models has significantly enhanced the semantic richness of learned representations. Models from the DINO family [14, 49, 61] leverage self-supervised distillation to capture fine-grained semantic structures, while CLIP [54] aligns visual and textual modalities to endow image features with high-level conceptual understanding. The strong cross-instance consistency exhibited by these features across 2D tasks—ranging from segmentation [37, 73] and retrieval [4, 49] to dense matching [32, 71, 74]—has motivated the adoption of these features as semantic priors for the geometric analysis of 3D shapes [1].

Traditionally, 3D shape matching has relied on intrinsic descriptors such as HKS [9] and WKS [3], which can be ambiguous across symmetric or cross-category parts. Recent works lift 2D foundation-model features onto 3D surfaces:

Diff3F [19] extracts multi-view semantic descriptors, while Surface-Aware Distillation [66] improves the feature extraction stage. EchoMatch [70] uses semantic cues for partial-to-partial matching, while DV-Matcher [15] operates on point clouds with deformation-aware extrinsic alignment.

3 Preliminaries

3.1 Deep Functional Maps Pipeline

Given a pair of 3D shapes $\mathcal{X}$ and $\mathcal{Y}$, represented as triangular meshes with $n_{\mathcal{X}}$ and $n_{\mathcal{Y}}$ vertices, the deep functional maps pipeline proceeds as follows:

- **Spectral decomposition.** The discrete Laplace–Beltrami operators $\mathbf{L}_{\mathcal{X}}$ and $\mathbf{L}_{\mathcal{Y}}$ are computed for each shape [53]. From these operators, the first k eigenfunctions $\Phi_{\mathcal{X}} \in \mathbb{R}^{n_{\mathcal{X}} \times k}$ and $\Phi_{\mathcal{Y}} \in \mathbb{R}^{n_{\mathcal{Y}} \times k}$ are obtained, along with their corresponding eigenvalue matrices $\Lambda_{\mathcal{X}}, \Lambda_{\mathcal{Y}} \in \mathbb{R}^{k \times k}$.
- **Descriptor extraction.** Vertex-wise descriptors $\mathbf{F}_{\mathcal{X}}$ and $\mathbf{F}_{\mathcal{Y}}$ are computed through the use of a learnable network $\mathcal{F}_{\theta}$ [60].
- **Functional map estimation.** The functional map $\mathbf{C}_{\mathcal{X}\mathcal{Y}} \in \mathbb{R}^{k \times k}$ is estimated by solving the following optimization problem:

$$\mathbf{C}_{\mathcal{X}\mathcal{Y}} = \arg \min_{\mathbf{C}} E_{\text{data}}(\mathbf{C}) + \lambda E_{\text{reg}}(\mathbf{C}), \quad (1)$$

where the data term $E_{\text{data}}(\mathbf{C}) = \left\| \mathbf{C} \Phi_{\mathcal{X}}^{\dagger} \mathbf{F}_{\mathcal{X}} - \Phi_{\mathcal{Y}}^{\dagger} \mathbf{F}_{\mathcal{Y}} \right\|_F^2$ enforces the preservation of descriptors within the spectral domain, and E_{reg} imposes structural constraints such as the commutativity of the Laplacian [50].

- **Point-wise map recovery.** The point-wise correspondence matrix $\Pi_{\mathcal{Y}\mathcal{X}}$ is recovered from $\mathbf{C}_{\mathcal{X}\mathcal{Y}}$ using the relationship $\mathbf{C}_{\mathcal{X}\mathcal{Y}} = \Phi_{\mathcal{Y}}^{\dagger} \Pi_{\mathcal{Y}\mathcal{X}} \Phi_{\mathcal{X}}$. This recovery is typically achieved via a nearest-neighbor search or other post-processing strategies [46, 68].

3.2 Flow Matching

Flow Matching [28, 39] is a framework for learning continuous-time probabilistic paths $(p_t)_{t \in [0,1]}$ that transform a source distribution p_0 into a target distribution p_1 . The central idea is to learn a time-dependent velocity field $u_t(x)$ parameterized by a neural network, which governs the dynamics of the probability flow. For a given spatial point x , the velocity field induces a flow $\psi_t(x)$ defined by the ordinary differential equation (ODE) $\frac{d}{dt} \psi_t(x) = u_t(\psi_t(x))$, $\psi_0(x) = x$, which describes the trajectory of samples evolving from the initial distribution.

Since directly modeling the marginal velocity field u_t is in general intractable, Lipman et al. [38] introduced Conditional Flow Matching (CFM), which replaces $u_t(x)$ with a conditional velocity field $u_t(x_t | x_1)$ while preserving the same direction of the training gradient. The model is trained by minimizing

$$\mathcal{L}_{\text{CFM}}(\theta) = \mathbb{E} \left[\left\| v_t^{\theta}(x_t) - u_t(x_t | x_1) \right\|^2 \right], \quad (2)$$




Fig. 2: Overview of SGMatch. Given a pair of shapes $\mathcal{X}$ and $\mathcal{Y}$, we extract geometric features $\mathbf{F}_{\mathcal{X}}^{\text{geo}}, \mathbf{F}_{\mathcal{Y}}^{\text{geo}}$ and semantic features $\mathbf{F}_{\mathcal{X}}^{\text{sem}}, \mathbf{F}_{\mathcal{Y}}^{\text{sem}}$, which are subsequently fused via the proposed SGLCA module. The resulting fused representations $\mathbf{F}_{\mathcal{X}}^{\text{fuse}}$ and $\mathbf{F}_{\mathcal{Y}}^{\text{fuse}}$ are then used to estimate functional maps $\mathbf{C}_{\mathcal{X}\mathcal{Y}}$ and to recover dense point-wise correspondences $\Pi_{\mathcal{Y}\mathcal{X}}$. In parallel, spectral heat diffusion followed by conditional flow matching regularization constrains feature transport, thereby suppressing local mismatches and promoting locally smooth correspondences.

where v_t^θ denotes the velocity predicted by the neural network. In practice, a simple and effective choice is to adopt a linear interpolation path $x_t = (1-t)x_0 + tx_1$, under which the corresponding conditional velocity reduces to the constant displacement $u_t = x_1 - x_0$. This formulation eliminates the need for numerical ODE integration during training, yielding improved optimization efficiency and numerical stability in comparison to simulation-based generative models.

4 Our SGMatch Method

An overview of the proposed framework is illustrated in Fig. 2. We begin by detailing the feature extraction process and the subsequent feature fusion module, followed by a description of the functional map solver for estimating correspondences in the spectral domain. Finally, we introduce the conditional flow matching regularization and the overall training objective.

4.1 Feature Extractor

Geometric Features Following [6, 13], we adopt DiffusionNet [60] to compute per-vertex descriptors that are robust to variations in mesh resolution and sampling density. Given shapes $\mathcal{X}$ and $\mathcal{Y}$, the extracted geometric features are denoted as $\mathbf{F}_{\mathcal{X}}^{\text{geo}} \in \mathbb{R}^{n_{\mathcal{X}} \times D^g}$ and $\mathbf{F}_{\mathcal{Y}}^{\text{geo}} \in \mathbb{R}^{n_{\mathcal{Y}} \times D^g}$, respectively, where D^g is the geometric feature dimensionality.

Semantic Features Following Diff3F [19], we employ a multi-view feature distillation strategy to transfer 2D visual semantics onto 3D surfaces, thereby constructing vertex-level descriptors enriched with high-level semantic awareness.

For a given shape $\mathcal{X}$, we render it from multiple viewpoints $\{\xi_j\}_{j=1}^n$ to obtain depth- and normal-guided images $P(\cdot \mid \xi_j) : \mathcal{X} \rightarrow I_j^{\mathcal{X}} \in \mathbb{R}^{H \times W}$, where

$P(\cdot)$ denotes the rendering operator. A pre-trained DINOv2 encoder [49] with frozen parameters then extracts dense semantic features $\mathbf{S}_j = \psi(I_j^{\mathcal{X}})$, where $\psi(\cdot)$ denotes the visual backbone. The 2D semantic features are subsequently back-projected onto the 3D surface using known camera parameters:

$$\mathbf{S}_j^{3D} = P^{-1}(\mathbf{S}_j), \quad (3)$$

thereby assigning viewpoint-specific semantic descriptors to surface vertices. Features from all viewpoints are then aggregated via averaging to obtain the vertex-level semantic representation:

$$\mathbf{F}_{\mathcal{X},i}^{\text{sem}} = \frac{1}{n} \sum_{j=1}^n \mathbf{S}_j^{3D}(i), \quad (4)$$

where $\mathbf{F}_{\mathcal{X}}^{\text{sem}} \in \mathbb{R}^{n_{\mathcal{X}} \times D^s}$ denotes the semantic feature matrix defined on $\mathcal{X}$, and D^s denotes the semantic feature dimensionality. The semantic features for $\mathcal{Y}$, denoted as $\mathbf{F}_{\mathcal{Y}}^{\text{sem}}$, are obtained in the same manner.

4.2 Semantic-Guided Local Cross-Attention Module

Given geometric features $\mathbf{F}^{\text{geo}}$ and semantic features $\mathbf{F}^{\text{sem}}$, we aim to inject semantic context into geometric representations while preserving intrinsic structural properties and local spatial coherence. To this end, we design the Semantic-Guided Local Cross-Attention (SGLCA) module, comprising a semantic-guided gating mechanism and a neighborhood-restricted cross-attention.

Semantic-Guided Gating We first project semantic features to the same dimensionality as geometric features via $\tilde{\mathbf{F}}^{\text{sem}} = \phi(\mathbf{F}^{\text{sem}})$, where $\phi(\cdot)$ is a linear projection. A lightweight MLP then generates channel-wise gating weights:

$$\mathbf{G} = \sigma(\text{MLP}_{\text{gate}}(\tilde{\mathbf{F}}^{\text{sem}})), \quad (5)$$

where $\sigma(\cdot)$ denotes the sigmoid function. The geometric features are then modulated as

$$\tilde{\mathbf{F}}^{\text{geo}} = \mathbf{F}^{\text{geo}} \odot (1 + \alpha \mathbf{G}), \quad (6)$$

where $\odot$ denotes element-wise multiplication and α is a learnable scaling parameter. This gating mechanism enables semantic information to adaptively amplify or attenuate geometric feature channels.

Local Cross-Attention To avoid introducing irrelevant global interactions, attention is restricted to local neighborhoods in the mesh. For each vertex i with neighborhood $\mathcal{N}(i)$, we compute $\mathbf{Q}_i = \mathbf{W}_{\mathbf{Q}} \tilde{\mathbf{F}}_i^{\text{geo}}$, $\mathbf{K}_j = \mathbf{W}_{\mathbf{K}} \tilde{\mathbf{F}}_j^{\text{sem}}$, and $\mathbf{V}_j = \mathbf{W}_{\mathbf{V}} \tilde{\mathbf{F}}_j^{\text{sem}}$, where $\mathbf{W}_{\mathbf{Q}}$, $\mathbf{W}_{\mathbf{K}}$, and $\mathbf{W}_{\mathbf{V}}$ are learnable projection matrices. The attention weights are defined as

$$\omega_{ij} = \text{Softmax}_{j \in \mathcal{N}(i)} \left((\mathbf{Q}_i \mathbf{K}_j^{\top}) / \sqrt{d} \right), \quad (7)$$

where d denotes the attention embedding dimension. The fused feature is then obtained via neighborhood aggregation:

$$\mathbf{F}_i^{\text{fuse}} = \tilde{\mathbf{F}}_i^{\text{geo}} + \text{LN} \left(\sum_{j \in \mathcal{N}(i)} \omega_{ij} \mathbf{V}_j \right), \quad (8)$$

where $\text{LN}(\cdot)$ denotes layer normalization. This design allows geometric features to selectively incorporate semantically relevant information from local spatial neighborhoods, achieving complementary integration of geometry and semantics while preserving structural coherence and avoiding spurious global interactions.

4.3 Functional Maps Module

Functional Maps Computation We adopt the regularized functional map formulation [56] to compute the bidirectional functional maps $\mathbf{C}_{\mathcal{X}\mathcal{Y}}$ and $\mathbf{C}_{\mathcal{Y}\mathcal{X}}$, as described in Sec. 3.1. Unlike classical approaches that rely solely on geometric descriptors, the functional maps are constructed using the fused vertex-level features $\mathbf{F}_{\mathcal{X}}^{\text{fuse}}$ and $\mathbf{F}_{\mathcal{Y}}^{\text{fuse}}$ obtained from Sec. 4.2.

During training, we minimize a structural regularization loss defined as

$$\mathcal{L}_{\text{struct}} = \lambda_{\text{bij}} \mathcal{L}_{\text{bij}} + \lambda_{\text{orth}} \mathcal{L}_{\text{orth}}, \quad (9)$$

where $\mathcal{L}_{\text{bij}} = \|\mathbf{C}_{\mathcal{X}\mathcal{Y}}\mathbf{C}_{\mathcal{Y}\mathcal{X}} - \mathbf{I}\|_F^2 + \|\mathbf{C}_{\mathcal{Y}\mathcal{X}}\mathbf{C}_{\mathcal{X}\mathcal{Y}} - \mathbf{I}\|_F^2$ encourages cycle-consistency (bijectivity), and $\mathcal{L}_{\text{orth}} = \|\mathbf{C}_{\mathcal{X}\mathcal{Y}}^\top \mathbf{C}_{\mathcal{X}\mathcal{Y}} - \mathbf{I}\|_F^2 + \|\mathbf{C}_{\mathcal{Y}\mathcal{X}}^\top \mathbf{C}_{\mathcal{Y}\mathcal{X}} - \mathbf{I}\|_F^2$ promotes local area preservation [56].

Point-wise Map Computation In theory, a point-wise map $\mathbf{\Pi}_{\mathcal{X}\mathcal{Y}}$ is a (partial) permutation matrix satisfying

$$\{\mathbf{\Pi} \in \{0, 1\}^{n_{\mathcal{X}} \times n_{\mathcal{Y}}} : \mathbf{\Pi} \mathbf{1}_{n_{\mathcal{Y}}} = \mathbf{1}_{n_{\mathcal{X}}}, \mathbf{1}_{n_{\mathcal{X}}}^\top \mathbf{\Pi} \leq \mathbf{1}_{n_{\mathcal{Y}}}^\top\}, \quad (10)$$

where $\mathbf{\Pi}_{\mathcal{X}\mathcal{Y}}(i, j)$ indicates that the i -th vertex of $\mathcal{X}$ corresponds to the j -th vertex of $\mathcal{Y}$. Following prior methods [13, 23], we compute a soft correspondence matrix using the fused features:

$$\mathbf{\Pi}_{\mathcal{X}\mathcal{Y}} = \text{Softmax} \left((\mathbf{F}_{\mathcal{X}}^{\text{fuse}} \mathbf{F}_{\mathcal{Y}}^{\text{fuse}\top}) / \tau_{\text{T}} \right), \quad (11)$$

where τ_{T} is a temperature parameter controlling the sharpness of the correspondence distribution. The reverse map $\mathbf{\Pi}_{\mathcal{Y}\mathcal{X}}$ is computed analogously.

4.4 Conditional Flow Matching Regularization

Building on the CFM formulation in Sec. 3.2, we design a task-specific regularization to encourage smooth feature evolution under soft correspondences. Unlike generative settings operating on distributions, the objective here regularizes vertex-level feature transport between corresponding shapes.

Spectral Heat Diffusion To improve robustness to local noise and stabilize feature transport, we first apply spectral heat diffusion to the fused features. Given the Laplace–Beltrami eigenbasis $\Phi \in \mathbb{R}^{n \times k}$ with eigenvalues $\Lambda \in \mathbb{R}^{k \times k}$ and mass matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$, the diffused features are computed as [7, 60]

$$\mathbf{Z} = \Phi \exp(-\tau \Lambda) \Phi^\top \mathbf{M} \mathbf{F}^{\text{fuse}}, \quad (12)$$

where $\tau > 0$ controls the diffusion scale and $\exp(-\tau \Lambda) = \text{diag}(e^{-\tau \lambda_1}, \dots, e^{-\tau \lambda_k})$. This spectral smoothing preserves intrinsic geometry while promoting spatial smoothness in the feature domain.

Flow Path and Velocity Modeling Let the diffused source feature be $\mathbf{z}_0 = \mathbf{Z}_\mathcal{X}$ and the transported target feature be $\mathbf{z}_1 = \Pi_{\mathcal{X}\mathcal{Y}} \mathbf{Z}_\mathcal{Y}$. Following the commonly adopted linear interpolation strategy, we define

$$\mathbf{z}_t = (1 - t)\mathbf{z}_0 + t\mathbf{z}_1, \quad t \sim \mathcal{U}(0, 1), \quad (13)$$

with the target velocity

$$\mathbf{v}_{\text{target}} = \mathbf{z}_1 - \mathbf{z}_0. \quad (14)$$

Intuitively, the interpolation path $\mathbf{z}_t$ defines a continuous feature-space trajectory from each source vertex to its transported target feature. Although the CFM loss is evaluated on sampled vertices, it is not an independent post-processing step or a formal substitute for explicit pairwise smoothness. Its spatial bias comes from two sources: the endpoints are heat-smoothed on their respective manifolds, and all vertices share the same time-conditioned velocity network. As a result, locally inconsistent correspondences create high-frequency and mutually incompatible displacement targets that are harder to fit with the shared finite-capacity model. CFM therefore acts as an implicit trajectory-level transport prior that promotes spatially coherent point-wise recovery, rather than as a Laplacian smoothing loss.

We parameterize a learnable velocity field $\mathbf{v}_\theta(\mathbf{z}_t, t)$ using an MLP. To effectively encode temporal information, the scalar time variable t is first mapped into a high-dimensional representation via sinusoidal embeddings, enabling the network to model non-linear temporal dependencies across the interpolation trajectory. The time-conditioned features are then injected into the MLP through Feature-wise Linear Modulation (FiLM) [51].

Importance-Weighted Objective Since the soft correspondence matrix $\Pi_{\mathcal{X}\mathcal{Y}}$ may contain uncertain matches, we introduce similarity-based importance sampling. For each vertex i , we compute a confidence weight

$$w_i = \exp(\alpha \cos(\mathbf{z}_{0,i}, \mathbf{z}_{1,i})), \quad (15)$$

where α is a scaling factor controlling the concentration of confidence weights. A subset $\mathcal{S}$ of vertices is then sampled according to $\{w_i\}$, and the conditional

flow matching objective is optimized using the Charbonnier loss in place of the standard MSE loss:

$$\mathcal{L}_{\text{cfm}} = \mathbb{E}_{t, i \in \mathcal{S}} \left[\sqrt{\|\mathbf{v}_\theta(\mathbf{z}_{t,i}, t) - \mathbf{v}_{\text{target},i}\|^2 + \varepsilon^2} \right], \quad (16)$$

where ε is a small constant. Compared to the MSE loss, the Charbonnier loss is less sensitive to outliers from inaccurate soft correspondences in early training, thereby stabilizing optimization and improving robustness under non-isometric deformations.

4.5 Loss Functions

The overall training objective consists of a spectral loss and a conditional flow matching (CFM) regularization term (see Eq. (16)). The spectral loss comprises structural regularization (Eq. (9)) and a coupling term that enforces consistency between functional and point-wise maps:

$$\mathcal{L}_{\text{couple}} = \left\| \mathbf{C}_{\mathcal{X}\mathcal{Y}} - \Phi_{\mathcal{Y}}^\dagger \Pi_{\mathcal{Y}\mathcal{X}} \Phi_{\mathcal{X}} \right\|_F^2 + \left\| \mathbf{C}_{\mathcal{Y}\mathcal{X}} - \Phi_{\mathcal{X}}^\dagger \Pi_{\mathcal{X}\mathcal{Y}} \Phi_{\mathcal{Y}} \right\|_F^2. \quad (17)$$

The full spectral loss is then given by

$$\mathcal{L}_{\text{spectral}} = \mathcal{L}_{\text{struct}} + \lambda_{\text{couple}} \mathcal{L}_{\text{couple}}. \quad (18)$$

The final training objective is therefore

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{spectral}} + \lambda_{\text{cfm}} \mathcal{L}_{\text{cfm}}. \quad (19)$$

5 Experiments

In this section, we conduct a comprehensive evaluation of our method against state-of-the-art approaches across multiple established benchmarks.

5.1 Near-isometric Shape Matching

Datasets We evaluate on three near-isometric datasets: FAUST [8], SCAPE [2], and SHREC’19 [45], using their remeshed versions [18, 57]. FAUST contains 100 human meshes (10 subjects, 10 poses) with an 80/20 train-test split. SCAPE includes 71 meshes of one subject in various poses, split into 51 training and 20 testing samples. SHREC’19 contains 44 human shapes with diverse body types and articulations, used solely for evaluation, excluding shape 40 due to its non-closed geometry.

Train	FAUST	SCAPE	FAUST + SCAPE
Test	FAUST	SCAPE	SHREC'19
<i>Axiomatic Methods</i>			
ZoomOut [46]	6.1	7.5	-
Smooth Shells [22]	2.5	4.2	-
DiscreteOp [55]	5.6	13.1	-
<i>Supervised Methods</i>			
FMNet [40]	11.0	33.0	-
GeomFMaps [18]	2.6	3.0	7.9
<i>Unsupervised Methods</i>			
Deep Shell [24]	1.7	2.5	21.1
DUO-FMNet [17]	2.5	4.2	6.4
AttnFMaps [35]	1.9	2.2	5.8
ULRSSM [13]	1.6	1.9	4.6
HybridFMap [6]	1.5	1.8	3.6
DenoisFMap [75]	1.7	2.1	3.6
DiffuMatch [52]	1.9	4.4	3.9
DeepFAFM [43]	1.5	1.9	3.6
Ours	1.4	1.8	3.3




Fig. 3: Left: Near-isometric matching and cross-dataset generalisation on FAUST, SCAPE, and SHREC'19. **Best** results are highlighted. **Right:** Qualitative results on the challenging SHREC'19 dataset.

Results We adopt the mean geodesic error as the evaluation metric [33]. Quantitative and qualitative results in Fig. 3 demonstrate that our method achieves competitive performance under the near-isometric setting, with marginal differences from leading baselines on FAUST and SCAPE where geometric invariance already provides strong discriminative power. Notably, on the SHREC'19 dataset, which evaluates cross-dataset generalization to shapes with diverse body types and articulations, our approach attains the best results. This suggests that the SGLCA module contributes most prominently when generalizing beyond the training distribution, where semantic priors provide complementary cues that purely geometric descriptors tend to underutilize.


5.2 Non-isometric Shape Matching

Datasets We conduct non-isometric shape matching experiments on two widely adopted benchmarks: SMAL [76] and DT4D-H [44]. The SMAL dataset comprises 49 animal shapes spanning 8 distinct species, of which 5 species are used for training and the remaining 3 are reserved for testing, resulting in a 29/20 train-test split. The DT4D-H dataset is designed for non-isometric human shape matching across multiple categories. We use 9 categories for evaluation, with 198 shapes allocated for training and 95 for testing.

Results As shown in Tab. 1, our method achieves strong performance on non-isometric benchmarks, outperforming most existing approaches including supervised ones. On the SMAL dataset, we achieve the best overall performance, sur-

Table 1: Quantitative comparison on non-isometric shape matching.

Geo.Err (×100)	SMAL	DT4D-H	
		intra	inter
<i>Axiomatic Methods</i>			
ZoomOut [46]	38.4	4.0	29.0
Smooth Shells [22]	30.0	1.2	6.4
DiscreteOp [55]	38.1	3.6	27.6
<i>Supervised Methods</i>			
FMNet [40]	42.0	9.6	38.0
GeomFMaps [18]	8.4	1.9	4.2
<i>Unsupervised Methods</i>			
Deep Shell [24]	29.3	3.4	31.1
DUO-FMNet [17]	6.7	2.6	15.8
AttnFMaps [35]	5.4	1.7	11.6
ULRSSM [13]	3.9	0.9	4.1
HybridFMap [6]	3.3	1.0	3.5
DenoisFMap [75]	4.3	5.8	16.9
DiffuMatch [52]	10.1	1.8	8.6
DeepFAFM [43]	3.8	0.9	3.9
Ours	2.5	1.0	3.4

**Fig. 4:** PCK curves and AUC values. **Left:** Non-isometric matching on SMAL and DT4D-H. **Right:** Matching with topological noise on TOPKIDS. Our approach achieves strong performance in both settings, outperforming existing methods.

passing the previous state-of-the-art by 24%. On DT4D-H, our approach attains the best intra-class result and competitive performance in the more challenging inter-class setting, where cross-category semantic consistency is harder to maintain through geometric descriptors alone. This highlights the limitations of relying solely on intrinsic geometric descriptors under large non-isometric deformations. In contrast, incorporating semantic features yields more stable cross-category correspondences. The PCK curves in Fig. 4 (left) and qualitative comparisons in Fig. 5 further confirm our method’s robustness.

5.3 Matching with Topological Noise

Datasets Real-world scans often exhibit self-intersections and local geometric artifacts, leading to degraded mesh topology. Such topological noise distorts




Fig. 5: Qualitative Results on SMAL and DT4D-H. Comparison of our method against DeepFAFM and HybridFMap, via texture transfer.

		Source	DeepFAFM	HybridFMap	Ours
Geo.Err	TOPKIDS				
<i>Axiomatic Methods</i>					
ZoomOut [46]	33.7				
Smooth Shells [22]	10.8				
DiscreteOp [55]	35.5				
<i>Unsupervised Methods</i>					
Deep Shell [24]	13.7				
AttnFMaps [35]	23.4				
ULRSSM [13]	9.2				
HybridFMap [6]	5.0				
DeepFAFM [43]	6.2				
Ours	2.9				

Fig. 6: Left: Quantitative comparison with state-of-the-art methods on TOPKIDS under topological noise. **Right:** Qualitative comparison on TOPKIDS; our method produces more accurate, coherent correspondences than existing approaches.

the intrinsic geometric structure and poses substantial challenges to functional map-based methods, which rely on stable Laplace-Beltrami operators. To assess robustness under severe topological perturbations, we conduct experiments on the TOPKIDS dataset [34]. Owing to its limited training set (26 shapes), the comparison is restricted to axiomatic and unsupervised methods.

Table 2: Map smoothness comparison with HybridFMap. We report conformal distortion, surface coverage, Dirichlet energy, and bijectivity to evaluate map smoothness beyond point-wise geodesic accuracy.

Metric	Method	FAUST	SCAPE	SHREC'19	SMAL	DT4D-inter	TOPKIDS
Conf. ↓	HybridFMap	0.654	0.779	1.030	2.473	1.866	3.201
	Ours	0.639	0.764	0.787	1.956	1.563	2.517
Cov. (%) ↑	HybridFMap	83.3	82.5	75.9	65.0	67.1	63.3
	Ours	83.5	81.9	76.1	71.2	66.3	63.9
Dir.E ↓	HybridFMap	2.96	3.11	12.30	19.29	8.44	98.30
	Ours	2.95	3.27	6.30	10.20	8.70	48.33
Bij. ↓	HybridFMap	0.0049	0.0060	–	0.0285	0.0130	–
	Ours	0.0052	0.0064	–	0.0175	0.0166	–

Results As shown in Fig. 6, our method achieves the best overall performance, yielding a 42% improvement in mean geodesic error. The PCK curves in Fig. 4 (right) further illustrate this advantage. This can be attributed to semantic features, which provide complementary high-level structural cues when intrinsic geometric descriptors become unreliable.

5.4 Map Smoothness

While average geodesic error quantifies overall correspondence accuracy, it does not explicitly capture the local smoothness of the point-wise maps. We therefore evaluate map smoothness using complementary metrics: conformal distortion [29], surface coverage, bijectivity [57], and Dirichlet energy [44].

As shown in Tab. 2, SGMatch improves most smoothness metrics, with the clearest gains on challenging non-isometric and topology-noisy datasets. The conformal-distortion PCK curves in Fig. 7 further show that these gains hold across distortion thresholds. On SMAL and TOPKIDS, where intrinsic geometry is less reliable, our method substantially reduces conformal distortion and Dirichlet energy while improving coverage. These results indicate that SGMatch improves not only point-wise accuracy but also local map smoothness. Some near-isometric metrics on SCAPE and DT4D-inter remain comparable rather than strictly better, suggesting that the main benefit appears under stronger ambiguity, deformation, or topological noise.

5.5 Comparison with Alternative Smoothness Regularizers

We further isolate the role of CFM by comparing it with alternative regularization objectives under the same SGMatch backbone. This comparison tests whether trajectory-level feature transport supervision is redundant with end-point matching, cycle-consistency-style constraints [56], or synchronous-diffusion-style constraints [12].

Endpoint-only removes the velocity MLP and trajectory sampling, and instead directly matches the heat-diffused source features with the transported




Fig. 7: Conformal-distortion-based smoothness comparison. Our method achieves competitive or lower average distortion, indicating smoother correspondences under non-isometric and topological-noise settings.

Table 3: Comparison with alternative smoothness regularizers under the same SG-Match backbone.

Variant	SMAL					TOPKIDS				
	Geo. ↓	Conf. ↓	Dir.E ↓	Bij. ↓	Cov. ↑	Geo. ↓	Conf. ↓	Dir.E ↓	Cov. ↑	
w/o CFM	2.7	2.32	10.55	0.0210	66.8	3.4	2.92	45.64	63.5	
Endpoint-only	2.7	2.46	14.37	0.0225	65.0	3.4	2.76	50.92	63.8	
Cycle-consistency	2.8	2.29	10.53	0.0201	67.2	3.4	2.75	42.42	63.7	
Sync. Diff.	2.6	2.45	13.48	0.0221	64.7	3.5	2.79	50.14	63.8	
Ours	2.5	1.96	10.20	0.0175	71.2	2.9	2.52	48.33	63.9	

target features in both directions. This variant isolates whether the gain comes from trajectory-level supervision rather than endpoint feature matching alone. Cycle-consistency and Sync. Diff. provide time-free alternatives based on round-trip consistency and diffusion-style smoothness constraints. As shown in Tab. 3, CFM gives the best geodesic error and conformal distortion on both datasets, and improves most map regularity metrics on SMAL. This indicates that the trajectory-level transport objective is not redundant with endpoint matching or diffusion-based smoothing.

On TOPKIDS, cycle-consistency gives a lower Dirichlet energy than CFM, but it sacrifices geodesic accuracy, conformal distortion, and coverage. This suggests over-smoothing rather than a better correspondence, and supports using multiple regularity metrics rather than Dirichlet energy alone.

5.6 Ablation Study

After isolating the CFM regularizer above, we conduct ablation experiments on the SMAL dataset to evaluate the main components and the design of SGLCA.

Results in Tab. 4 show that both modalities are important: geometric features provide the structural basis, while semantic features alone are insufficient for accurate correspondence. The SGLCA variants further disentangle semantic gating from attention locality. Local attention provides the main gain by restricting semantic interactions to mesh neighborhoods, and semantic gating gives an additional improvement by allowing semantic cues to modulate geometric channels rather than replacing geometry. Global attention does not improve

Table 4: Quantitative results of the ablation experiments on the SMAL dataset.

Variant	Geo.	Sem.	Gating	Attention	Heat Diff.	CFM	Geo.Err ↓
Geo. only	✓	✗	–	–	✓	✓	3.2
Sem. only	✗	✓	–	–	✓	✓	21.2
Gating only	✓	✓	✓	none	✓	✓	3.6
Local attention only	✓	✓	✗	local	✓	✓	2.6
Gating + global attention	✓	✓	✓	global	✓	✓	2.6
w/o heat diffusion	✓	✓	✓	local	✗	✓	3.0
Ours	✓	✓	✓	local	✓	✓	2.5

over local attention, suggesting that unrestricted semantic interactions can introduce irrelevant long-range responses. Finally, removing spectral heat diffusion degrades performance, confirming that stabilized feature endpoints complement the trajectory-level transport regularization analyzed above.

6 Limitations

Despite its effectiveness, SGMatch has two main limitations. First, the current framework is designed for full-shape correspondence and does not explicitly handle partial matching scenarios, which remain an important and actively studied problem in non-rigid shape analysis [21, 70]. Second, its semantic features are derived from pretrained 2D visual models, so the performance of SGMatch depends on the quality and domain generalization capability of these external representations. Domains where such visual priors are weak or unavailable [5] may therefore require domain-specific encoders or task-specific semantic pretraining.

7 Conclusion

We presented SGMatch for robust non-rigid shape correspondence under challenging deformations and topology changes. The results show that semantic guidance is most beneficial when geometric descriptors become ambiguous, while flow-based transport regularization improves the spatial coherence of point-wise recovery from truncated spectral representations. Experiments across near-isometric, non-isometric, and topology-noisy benchmarks demonstrate improved correspondence accuracy and map regularity, with the clearest gains under large deformation and topological noise.

Acknowledgements

This work was supported by the National Key R&D Program of China under Grant 2024YFE0202500, and the National Natural Science Foundation of China under Grants 624B2107 and U23B2050.

References

1. Abdelreheem, A., Eldesokey, A., Ovsjanikov, M., Wonka, P.: Zero-shot 3d shape correspondence. In: SIGGRAPH Asia 2023 Conference Papers. pp. 1–11 (2023)
2. Anguelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., Davis, J.: Scape: shape completion and animation of people. In: ACM Siggraph 2005 Papers, pp. 408–416 (2005)
3. Aubry, M., Schlickewei, U., Cremers, D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In: ICCV. pp. 1626–1633. IEEE (2011)
4. Baldrati, A., Bertini, M., Uricchio, T., Del Bimbo, A.: Composed image retrieval using contrastive learning and task-oriented clip-based features. *ACM Transactions on Multimedia Computing, Communications and Applications* **20**(3), 1–24 (2023)
5. Bastian, L., Baumann, A., Hoppe, E., Bürgin, V., Kim, H.Y., Saleh, M., Busam, B., Navab, N.: S3m: scalable statistical shape modeling through unsupervised correspondences. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 459–469. Springer (2023)
6. Bastian, L., Xie, Y., Navab, N., Lähner, Z.: Hybrid functional maps for crease-aware non-isometric shape matching. In: CVPR. pp. 3313–3323 (2024)
7. Behmanesh, M., Krahn, M., Ovsjanikov, M.: Tide: Time derivative diffusion for deep learning on graphs. In: ICML. pp. 2015–2030. PMLR (2023)
8. Bogo, F., Romero, J., Loper, M., Black, M.J.: Faust: Dataset and evaluation for 3d mesh registration. In: CVPR. pp. 3794–3801 (2014)
9. Bronstein, M.M., Kokkinos, I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In: CVPR. pp. 1704–1711. IEEE (2010)
10. Cao, D., Bernard, F.: Unsupervised deep multi-shape matching. In: ECCV. pp. 55–71. Springer (2022)
11. Cao, D., Eisenberger, M., El Amrani, N., Cremers, D., Bernard, F.: Spectral meets spatial: Harmonising 3d shape matching and interpolation. In: CVPR. pp. 3658–3668 (2024)
12. Cao, D., Lähner, Z., Bernard, F.: Synchronous diffusion for unsupervised smooth non-rigid 3d shape matching. In: ECCV. pp. 262–281. Springer (2024)
13. Cao, D., Roetzer, P., Bernard, F.: Unsupervised learning of robust spectral shape matching. *ACM TOG* **42**(4), 1–15 (2023)
14. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: CVPR. pp. 9650–9660 (2021)
15. Chen, Z., Jiang, P., Huang, R.: Dv-matcher: Deformation-based non-rigid point cloud matching guided by pre-trained visual features. In: CVPR. pp. 27264–27274 (2025)
16. Dinh, H.Q., Yezzi, A., Turk, G.: Texture transfer during shape transformation. *ACM TOG* **24**(2), 289–310 (2005)
17. Donati, N., Corman, E., Ovsjanikov, M.: Deep orientation-aware functional maps: Tackling symmetry issues in shape matching. In: CVPR. pp. 742–751 (2022)
18. Donati, N., Sharma, A., Ovsjanikov, M.: Deep geometric functional maps: Robust feature learning for shape correspondence. In: CVPR. pp. 8592–8601 (2020)
19. Dutt, N.S., Muralikrishnan, S., Mitra, N.J.: Diffusion 3d features (diff3f): Decorating untextured shapes with distilled semantic features. In: CVPR. pp. 4494–4504 (2024)
20. Egger, B., Smith, W.A., Tewari, A., Wuhler, S., Zollhoefer, M., Beeler, T., Bernard, F., Bolkart, T., Kortylewski, A., Romdhani, S., et al.: 3d morphable face models: past, present, and future. *ACM TOG* **39**(5), 1–38 (2020)

21. Ehm, V., Amrani, N.E., Xie, Y., Bastian, L., Gao, M., Wang, W., Sang, L., Cao, D., Weißberg, T., Löhner, Z., et al.: Beyond complete shapes: A benchmark for quantitative evaluation of 3d shape surface matching algorithms. In: *Computer Graphics Forum*. vol. 44, p. e70186. Wiley Online Library (2025)
22. Eisenberger, M., Lahner, Z., Cremers, D.: Smooth shells: Multi-scale shape registration with functional maps. In: *CVPR*. pp. 12265–12274 (2020)
23. Eisenberger, M., Novotny, D., Kerchenbaum, G., Labatut, P., Neverova, N., Cremers, D., Vedaldi, A.: Neuromorph: Unsupervised shape interpolation and correspondence in one go. In: *CVPR*. pp. 7473–7483 (2021)
24. Eisenberger, M., Toker, A., Leal-Taixé, L., Cremers, D.: Deep shells: Unsupervised shape correspondence with optimal transport. *NeurIPS* **33**, 10491–10502 (2020)
25. Eynard, D., Rodola, E., Glashoff, K., Bronstein, M.M.: Coupled functional maps. In: *2016 Fourth International Conference on 3D Vision (3DV)*. pp. 399–407. IEEE (2016)
26. Gao, M., Lahner, Z., Thunberg, J., Cremers, D., Bernard, F.: Isometric multi-shape matching. In: *CVPR*. pp. 14183–14193 (2021)
27. Groueix, T., Fisher, M., Kim, V.G., Russell, B.C., Aubry, M.: 3d-coded: 3d correspondences by deep deformation. In: *ECCV*. pp. 230–246 (2018)
28. Holderrieth, P., Erives, E.: An introduction to flow matching and diffusion models. *arXiv preprint arXiv:2506.02070* (2025)
29. Hormann, K., Greiner, G.: Mips: An efficient global parametrization method (2000)
30. Huang, Q., Wang, F., Guibas, L.: Functional map networks for analyzing and exploring large shape collections. *ACM TOG* **33**(4), 1–11 (2014)
31. Huang, R., Ren, J., Wonka, P., Ovsjanikov, M.: Consistent zoomout: Efficient spectral map synchronization. In: *Computer Graphics Forum*. vol. 39, pp. 265–278. Wiley Online Library (2020)
32. Jiang, H., Karpur, A., Cao, B., Huang, Q., Araujo, A.: Omniglue: Generalizable feature matching with foundation model guidance. In: *CVPR*. pp. 19865–19875 (2024)
33. Kim, V.G., Lipman, Y., Funkhouser, T.: Blended intrinsic maps. *ACM TOG* **30**(4), 1–12 (2011)
34. Löhner, Z., Rodolà, E., Bronstein, M.M., Cremers, D., Burghard, O., Cosmo, L., Dieckmann, A., Klein, R., Sahillioglu, Y., 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)
35. Li, L., Donati, N., Ovsjanikov, M.: Learning multi-resolution functional maps with spectral attention for robust shape matching. *NeurIPS* **35**, 29336–29349 (2022)
36. Li, T., Bolkart, T., Black, M.J., Li, H., Romero, J.: Learning a model of facial shape and expression from 4d scans. *ACM TOG* **36**(6), 194–1 (2017)
37. Liang, F., Wu, B., Dai, X., Li, K., Zhao, Y., Zhang, H., Zhang, P., Vajda, P., Marculescu, D.: Open-vocabulary semantic segmentation with mask-adapted clip. In: *CVPR*. pp. 7061–7070 (2023)
38. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747* (2022)
39. Lipman, Y., Havasi, M., Holderrieth, P., Shaul, N., Le, M., Karrer, B., Chen, R.T., Lopez-Paz, D., Ben-Hamu, H., Gat, I.: Flow matching guide and code. *arXiv preprint arXiv:2412.06264* (2024)
40. Litany, O., Remez, T., Rodola, E., Bronstein, A., Bronstein, M.: Deep functional maps: Structured prediction for dense shape correspondence. In: *ICCV*. pp. 5659–5667 (2017)

41. Litany, O., Rodolà, E., Bronstein, A.M., Bronstein, M.M.: Fully spectral partial shape matching. In: *Computer Graphics Forum*. vol. 36, pp. 247–258. Wiley Online Library (2017)
42. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: A skinned multi-person linear model. In: *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pp. 851–866 (2023)
43. Luo, F., Li, Q., Hu, L., Wang, H., Xu, H., Liu, X., Liu, S., Chen, H.: Deep frequency awareness functional maps for robust shape matching. *IEEE TVCG* (2025)
44. Magnet, R., Ren, J., Sorkine-Hornung, O., Ovsjanikov, M.: Smooth non-rigid shape matching via effective dirichlet energy optimization. In: *3DV*. pp. 495–504. IEEE (2022)
45. Melzi, S., Marin, R., Rodolà, E., Castellani, U., Ren, J., Poulenard, A., Ovsjanikov, P., et al.: Shrec’19: matching humans with different connectivity. In: *Eurographics Workshop on 3D Object Retrieval*. pp. 1–8. The Eurographics Association (2019)
46. Melzi, S., Ren, J., Rodola, E., Sharma, A., Wonka, P., Ovsjanikov, M.: Zoomout: Spectral upsampling for efficient shape correspondence. *arXiv preprint arXiv:1904.07865* (2019)
47. Morreale, L., Aigerman, N., Kim, V.G., Mitra, N.J.: Neural semantic surface maps. In: *Computer Graphics Forum*. vol. 43, p. e15005. Wiley Online Library (2024)
48. Nogneng, D., Ovsjanikov, M.: Informative descriptor preservation via commutativity for shape matching. In: *Computer Graphics Forum*. vol. 36, pp. 259–267. Wiley Online Library (2017)
49. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193* (2023)
50. Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L.: Functional maps: a flexible representation of maps between shapes. *ACM TOG* **31**(4), 1–11 (2012)
51. Perez, E., Strub, F., De Vries, H., Dumoulin, V., Courville, A.: Film: Visual reasoning with a general conditioning layer. In: *AAAI*. vol. 32 (2018)
52. Pierson, E., Li, L., Dai, A., Ovsjanikov, M.: Diffumatch: Category-agnostic spectral diffusion priors for robust non-rigid shape matching. In: *ICCV*. pp. 5745–5756 (2025)
53. Pinkall, U., Polthier, K.: Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics* **2**(1), 15–36 (1993)
54. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *ICML*. pp. 8748–8763. PmLR (2021)
55. Ren, J., Melzi, S., Wonka, P., Ovsjanikov, M.: Discrete optimization for shape matching. In: *Computer Graphics Forum*. vol. 40, pp. 81–96. Wiley Online Library (2021)
56. Ren, J., Panine, M., Wonka, P., Ovsjanikov, M.: Structured regularization of functional map computations. In: *Computer Graphics Forum*. vol. 38, pp. 39–53. Wiley Online Library (2019)
57. Ren, J., Poulenard, A., Wonka, P., Ovsjanikov, M.: Continuous and orientation-preserving correspondences via functional maps. *ACM TOG* **37**(6), 1–16 (2018)
58. Rodolà, E., Cosmo, L., Bronstein, M.M., Torsello, A., Cremers, D.: Partial functional correspondence. In: *Computer Graphics Forum*. vol. 36, pp. 222–236. Wiley Online Library (2017)
59. Sahillioglu, Y.: Recent advances in shape correspondence. *The Visual Computer* **36**(8), 1705–1721 (2020)

60. Sharp, N., Attaiki, S., Crane, K., Ovsjanikov, M.: Diffusionnet: Discretization agnostic learning on surfaces. *ACM TOG* **41**(3), 1–16 (2022)
61. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khaidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., et al.: Dinov3. *arXiv preprint arXiv:2508.10104* (2025)
62. Song, C., Wei, J., Li, R., Liu, F., Lin, G.: 3d pose transfer with correspondence learning and mesh refinement. *NeurIPS* **34**, 3108–3120 (2021)
63. Song, C., Wei, J., Li, R., Liu, F., Lin, G.: Unsupervised 3d pose transfer with cross consistency and dual reconstruction. *IEEE TPAMI* **45**(8), 10488–10499 (2023)
64. Sun, M., Mao, S., Jiang, P., Ovsjanikov, M., Huang, R.: Spatially and spectrally consistent deep functional maps. In: *CVPR*. pp. 14497–14507 (2023)
65. Trappolini, G., Cosmo, L., Moschella, L., Marin, R., Melzi, S., Rodolà, E.: Shape registration in the time of transformers. *NeurIPS* **34**, 5731–5744 (2021)
66. Uzolas, L., Eisemann, E., Kellnhofer, P.: Surface-aware distilled 3d semantic features. In: *Proceedings of the SIGGRAPH Asia 2025 Conference Papers*. pp. 1–12 (2025)
67. Van Kaick, O., Zhang, H., Hamarneh, G., Cohen-Or, D.: A survey on shape correspondence. In: *Computer Graphics Forum*. vol. 30, pp. 1681–1707. Wiley Online Library (2011)
68. Xia, Y., Lu, Y., Gao, Y., Ma, J.: Locality preserving refinement for shape matching with functional maps. In: *AAAI*. vol. 38, pp. 6207–6215 (2024)
69. Xia, Y., Ye, T., Zhou, H., Wang, Z., Ma, J.: Multi-shape matching with cycle consistency basis via functional maps. In: *AAAI*. vol. 39, pp. 8575–8583 (2025)
70. Xie, Y., Ehm, V., Roetzer, P., El Amrani, N., Gao, M., Bernard, F., Cremers, D.: Echomatch: Partial-to-partial shape matching via correspondence reflection. In: *CVPR*. pp. 11665–11675 (2025)
71. Yang, M., Fan, F., Li, Z., Deng, S., Ma, Y., Ma, J.: Distillmatch: Leveraging knowledge distillation from vision foundation model for multimodal image matching. *arXiv preprint arXiv:2509.16017* (2025)
72. Ye, T., Ma, Y., Mei, X.: Dcmatch: Unsupervised multi-shape matching with dual-level consistency. *arXiv preprint arXiv:2509.01204* (2025)
73. Yuan, M., Zhang, Y., Zhang, W., Ma, L., Gao, Y., Ying, J., Xin, Y.: Infoclip: Bridging vision-language pretraining and open-vocabulary semantic segmentation via information-theoretic alignment transfer. *arXiv preprint arXiv:2511.15967* (2025)
74. Zhang, Y., Zhao, X.: Mesa: Matching everything by segmenting anything. In: *CVPR*. pp. 20217–20226 (2024)
75. Zhuravlev, A., Löhner, Z., Golyanik, V.: Denoising functional maps: Diffusion models for shape correspondence. In: *CVPR*. pp. 26899–26909 (2025)
76. Zuffi, S., Kanazawa, A., Jacobs, D.W., Black, M.J.: 3d menagerie: Modeling the 3d shape and pose of animals. In: *CVPR*. pp. 6365–6373 (2017)

A Implementation Details

A.1 Experimental Settings

All learning-based methods are implemented using PyTorch 2.1.0 and CUDA 12.1, while axiomatic approaches are executed in MATLAB 2018a. All experiments are conducted on a single NVIDIA GeForce RTX 3090 GPU and an Intel Xeon(R) Platinum 8365A CPU (2.60GHz). Regarding geometric feature descriptors, we follow established practices and use WKS as geometric feature descriptors. For the SMAL dataset, we instead use raw XYZ coordinates of the 3D vertices and apply random rotation augmentation.

A.2 Technical Details

In the feature solver, the geometric feature dimension D^g is 256, and the semantic feature dimension D^s is 768. We adopt the pre-trained DINOv2-ViT-B/14 model. The number of viewpoints $\{\xi\}$ is set to 100, and the rendered image resolution is $H = 512$ and $W = 512$. We include normal maps to improve the representation of 3D shape details. For the SGLCA module, the attention embedding dimension d is 256. Neighborhoods use a 1-ring structure with self-loops to maintain a fixed size of 32. In the functional map module, the eigenbasis dimensions are hybridly set to $k_{LB} = 140$ and $k_{Elas} = 60$ (see [6]), and the temperature τ_T is set to 0.07. For CFM, the diffusion time τ in spectral heat diffusion varies by dataset. Following Cao et al. [12], we set τ to 10^{-2} for near-isometric cases, 10^{-4} for non-isometric datasets and 10^{-1} for topology-noisy datasets. Concurrently, the confidence sampling parameter α is set to 2, and the constant in the CFM loss is 10^{-3} .

The loss weights are $\lambda_{bij} = \lambda_{orth} = \lambda_{couple} = 1.0$ and $\lambda_{cfm} = 100$. We train the model end-to-end using the Adam optimizer with a learning rate of 0.001 and a cosine annealing schedule.

B Additional Ablation Experiments

B.1 SGLCA Module Design

We investigate alternative fusion strategies within the SGLCA module to validate architectural choices. Three variants are compared against our full design: (1) **element-wise addition**, where linear projection aligns feature dimensions before summation, treating modalities equally without channel modulation; (2) **raw concatenation**, which preserves cross-modal information but causes dimensional mismatch; and (3) **concatenation with MLP**, restoring original dimensions while lacking channel-wise selectivity. The main paper further disentangles semantic gating, local attention, and global attention in the unified ablation table.

Tab. 5 shows our full SGLCA design outperforms all alternatives. The gap between element-wise addition and our method indicates simple summation fails

Table 5: Ablation of feature fusion strategies and gating mechanism in the SGLCA module on the SMAL dataset.

Settings	(1) Add	(2) Raw concat	(3) Concat+MLP	Ours
Geo. Err ($\times 100$)	5.3	3.5	3.6	2.5

Table 6: Comparison of CFM regularization variants on the SMAL dataset.

Settings	(1) MSE	(2) Laplacian	Ours
Geo. Err ($\times 100$)	2.6	4.1	2.5

to capture geometry-semantic relationships, as it treats modalities uniformly regardless of local reliability. While concatenation variants improve upon addition, they remain inferior, suggesting information preservation is insufficient without fine-grained selectivity. Together with the main ablation table, these results show that semantic context is most effective when it modulates geometric channels through the proposed gated local attention design.

B.2 Conditional Flow Matching Regularization

We assess CFM regularization design choices by comparing two alternatives on SMAL: (1) **MSE loss**. Replacing the Charbonnier loss with a standard MSE objective:

$$\mathcal{L}_{\text{mse}} = \mathbb{E}_{t,i \in \mathcal{S}} \left[|\mathbf{v}_{\theta}(\mathbf{z}_{t,i}, t) - \mathbf{v}_{\text{target},i}|^2 \right]. \quad (20)$$

Although MSE and Charbonnier losses behave similarly under accurate correspondences, MSE is more sensitive to outliers from early-stage unreliable soft correspondences, potentially destabilizing optimization and hindering convergence.

(2) **Laplacian smoothing**. We replace CFM regularization with static Laplacian smoothing on transported target features:

$$\mathcal{L}_{\text{lap}} = \text{tr}(\mathbf{z}_1^{\top} \mathbf{L}_{\mathcal{X}} \mathbf{z}_1), \quad (21)$$

where $\mathbf{z}_1 = \mathbf{\Pi}_{\mathcal{X}\mathcal{Y}} \mathbf{Z}_{\mathcal{Y}}$ denotes transported target features and $\mathbf{L}_{\mathcal{X}}$ is the source shape’s Laplace–Beltrami operator. This approach enforces smoothness as a static pairwise constraint rather than through dynamic trajectories, encouraging nearby source vertices to match similar target features.

Tab. 6 shows both alternatives reduce performance. The small gap between MSE and Charbonnier losses suggests CFM’s primary benefit is trajectory-level regularization rather than the specific loss function, though Charbonnier loss adds noise robustness. The larger drop with Laplacian smoothing indicates static pairwise constraints cannot fully capture the dynamic feature transport required for non-isometric deformations. While Laplacian smoothing enforces final-state




Fig. 8: Feature similarity heatmaps on SMAL cross-species pairs. For each red-marked source query point, we visualize cosine similarity to all target vertices using geometric (middle) and semantic (right) features. Semantic features produce more localized and semantically consistent responses.

Table 7: Comparison of sampling strategies in the CFM objective on SMAL.

Settings	Uniform sampling	Confidence sampling (Ours)
Geo. Err ($\times 100$)	2.6	2.5

smoothness, it lacks transport process guidance, making optimization prone to locally inconsistent solutions a coherent velocity field prevents.

B.3 Importance Sampling Strategy

We evaluate the confidence-based sampling strategy used in the importance-weighted CFM objective (Eq. (16)), where vertices are sampled according to the confidence weights defined in Eq. (15). We compare it against a uniform sampling baseline ($w_i = 1 \forall i$) with equal subset size $|\mathcal{S}|$.

As shown in Tab. 7, confidence-based sampling yields small but steady gains over uniform sampling. This limited gap is expected as the Charbonnier loss inherently handles outliers, partially mitigating the lack of weighting in the baseline. Nonetheless, confidence weighting provides further benefits by filtering unreliable transport targets during sampling, particularly in early training.

C Semantic Feature Analysis

We investigate whether 2D-pretrained semantic features maintain cross-species discriminability on 3D surfaces without geometric supervision. We evaluate this by comparing per-vertex similarity between geometric and semantic descriptors using challenging SMAL shape pairs.

Table 8: Mean neighbor velocity deviation. Lower is better.

Dataset	w/o CFM	w/ CFM
SMAL	26.65	1.59
TOPKIDS	15.72	0.015

For a source query vertex, we visualize its cosine similarity to all target vertices as a heatmap. Two representative pairs are selected to evaluate feature behavior across matching difficulties: an easy case with low geodesic error (Horse $\rightarrow$ Cougar) and a hard case with high geodesic error (Hippo $\rightarrow$ Horse).

Fig. 8 shows that the two descriptors behave differently. In the easy pair (top row), geometric features produce a false high response near the target tail root, while semantic features correctly ignore this region and focus on the torso. In the hard pair (bottom row), geometric features yield diffuse responses under large deformations, whereas semantic features maintain a localized distribution, showing that DINOv2 captures stable cross-species part correspondences. These results suggest semantic features provide the discriminability geometric descriptors lack, particularly in resolving ambiguities from symmetry or self-similarity.

D CFM Regularization Analysis

We further analyze CFM from transport-field coherence and qualitative map smoothness.

D.1 Transport-Field Coherence

To directly assess the local coherence of the transport field, we compute the neighbor velocity deviation

$$d_i = \left\| \mathbf{u}_i - \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \mathbf{u}_j \right\|_2, \quad (22)$$

where $\mathcal{N}(i)$ denotes the mesh neighborhood. For SGMatch, $\mathbf{u}_i = \mathbf{v}_\theta(\mathbf{z}_{t,i}, t)$ is the learned velocity; for the variant without CFM, $\mathbf{u}_i = \mathbf{z}_{1,i} - \mathbf{z}_{0,i}$ is the endpoint displacement. As shown in Tab. 8, CFM substantially reduces the mean deviation on both SMAL and TOPKIDS, supporting its role as a smoothness-promoting transport prior.

D.2 Qualitative Smoothness Analysis

Fig. 9 displays texture transfer and per-face conformal distortion heatmaps on a hippo $\rightarrow$ cougar pair. In texture transfer (top row), omitting CFM causes local inconsistencies at tail and paw regions (red circles), where spatially adjacent




Fig. 9: Effect of CFM regularization on correspondence smoothness for a challenging SMAL cross-species pair. Top row shows texture transfer, and bottom row shows per-face conformal distortion. CFM suppresses local correspondence artifacts and reduces high-distortion regions.

vertices map to distant source regions, creating color discontinuities. With CFM, these artifacts are suppressed, and texture remains spatially coherent.

Conformal distortion heatmaps (bottom row) offer a quantitative perspective on these results. Without CFM, high-distortion regions appear across the surface, with severe distortions concentrated around the limbs and tail, matching the correspondence artifacts in the texture transfer. With CFM, the share of high-distortion regions is reduced, and the overall distortion distribution shifts toward lower values.

E Semantic Backbone and View Sensitivity

We evaluate how SGMatch depends on the quality of the lifted semantic features. We replace DINOv2 with weaker visual backbones while keeping the downstream matching architecture fixed. As shown in Tab. 9, stronger visual backbones improve correspondence quality under large non-isometric deformation.

Table 9: Semantic backbone ablation on SMAL.

Backbone	CLIP	DINOv1	DINOv2
Geo.Err ($\times 100$) $\downarrow$	5.8	4.1	2.5

We also vary the number of rendered DINOv2 views. Tab. 10 shows that 50 views preserve the final accuracy of the 100-view setting, while fewer views gradually degrade performance.

Table 10: View number ablation on SMAL using DINOv2 features.

Views	12	25	50	100
Geo.Err ($\times 100$) $\downarrow$	2.8	2.7	2.5	2.5

F Parameter Analysis

F.1 Neighborhood Size

We analyze SGLCA module sensitivity to neighborhood size k on SMAL. As shown in Tab. 11, performance is stable across k values, with $k = 16$ and $k = 32$ yielding lowest errors. Small sizes ($k = 8$) provide insufficient structural context for effective semantic-geometric fusion, degrading accuracy. Conversely, $k = 64$ extends the receptive field beyond local geometry. Since neighbor lists derive from direct adjacency, vertices with fewer than k neighbors are padded with self-loops, introducing redundant attention weights that dilute actual neighbor influence. We set $k = 32$ as the default to balance local context coverage with attention focus.

Table 11: Effect of neighborhood size k in the SGLCA module on the SMAL dataset.

Size	$k = 8$	$k = 16$	$k = 32$ (Ours)	$k = 64$
Geo. Err ($\times 100$)	2.7	2.5	2.5	2.6

F.2 Diffusion Time

Following [12], we adopt a diffusion time of $\tau = 10^{-2}$ for near-isometric matching, $\tau = 10^{-4}$ for non-isometric matching and $\tau = 10^{-1}$ for topology-noisy matching. We further examine the sensitivity of SGMatch to this parameter using the SMAL and TOPKIDS datasets.

Tab. 12 demonstrates that SMAL performance remains stable for τ between 10^{-5} and 10^{-3} , with degradation occurring only at $\tau = 10^{-2}$. For non-isometric deformations, excessive diffusion likely smooths out the local geometric details required for cross-category correspondence. Conversely, a moderate-to-small diffusion scale suppresses noise while preserving feature discriminability. The marginal fluctuation at $\tau = 10^{-5}$ is consistent with spectral heat diffusion approaching the identity operator as $\tau \rightarrow 0$: at extremely small scales, the smoothing effect becomes negligible, slightly reducing the stabilization benefit for CFM transport without the catastrophic detail loss seen at large τ . These results confirm the robustness of our method across a broad range of τ values.

On TOPKIDS, performance is consistent for $\tau \in \{10^{-2}, 10^{-1}\}$, while $\tau = 10^{-1}$ gives the best result, as shown in Tab. 13. We therefore use $\tau = 10^{-1}$

for TOPKIDS. The limited gap between 10^{-2} and 10^{-1} also indicates that the improvements from SGMATCH are robust and do not rely on intensive hyperparameter tuning.

Table 12: Effect of diffusion time τ on the SMAL dataset.

Diff. Time	10^{-2}	10^{-3}	10^{-4} (Ours)	10^{-5}	10^{-6}
Geo. Err ($\times 100$)	3.3	2.5	2.5	2.6	2.5

Table 13: Effect of diffusion time τ on the TOPKIDS dataset.

Diff. Time	1	10^{-1} (Ours)	10^{-2}	10^{-3}	10^{-4}
Geo. Err ($\times 100$)	3.3	2.9	3.3	3.5	3.5

F.3 Confidence Weight Parameter

In the importance-weighted CFM objective, α controls the concentration of the confidence weights defined in Eq. (15). Larger α values prioritize vertices with reliable soft correspondences, while smaller α approaches uniform sampling. We evaluate α sensitivity on SMAL.

As shown in Tab. 14, performance is stable for $\alpha \in \{0.5, 1.0, 2.0, 5.0\}$, indicating robustness to moderate variations. Excessively large α (e.g., 10.0) concentrates weights on a tiny vertex subset with highest similarity, reducing training diversity and degrading performance. The partial recovery at $\alpha = 20.0$ likely stems from the interaction between extreme concentration and the Charbonnier loss, which limits outlier influence. We set $\alpha = 2.0$ as default to balance effective weighting with sampling distribution diversity.

Table 14: Sensitivity analysis of the confidence weight parameter α in importance sampling on the SMAL dataset.

α	0.5	1.0	2.0 (Ours)	5.0	10.0	20.0
Geo.Err ($\times 100$)	2.6	2.6	2.5	2.5	3.6	2.7

G Statistical Analysis

We assess the robustness and stability of our approach by retraining on SHREC’19, SMAL, and TOPKIDS using three random seeds (42, 1234, and 3407). The

mean geodesic error and standard deviation are reported and compared with HybridFMap under identical settings.

Results in Tab. 15 show that SGMatch yields lower mean errors than HybridFMap across all tested datasets. SGMatch also maintains much lower standard deviation (e.g., 0.01 vs. 0.17 on SMAL; 0.06 vs. 1.02 on TOPKIDS), suggesting greater stability regarding random initialization. This performance gain likely stems from the SGLCA module’s semantic guidance and the regularization of CFM, which together lead to a more stable optimization process. In comparison, the higher variance of HybridFMap reflects a heavier reliance on initialization, potentially limiting its practical reliability.

Table 15: Mean geodesic error ($\times 100$) with standard deviation over three random seeds on SHREC’19, SMAL, and TOPKIDS.

Method	SHREC’19	SMAL	TOPKIDS
HybridFMap	4.17 ± 0.68	3.52 ± 0.17	5.38 ± 1.02
Ours	3.33 ± 0.07	2.51 ± 0.01	3.36 ± 0.06

H Runtime and Memory Analysis

We evaluate the computational efficiency of SGMatch by reporting the average inference time per pair and peak memory usage across several datasets, using HybridFMaps as a baseline. All measurements are conducted on the same hardware. Our model contains 2.8M parameters, while HybridFMaps contains 0.5M. Although our parameter count is higher due to the SGLCA module and CFM regularization, this trade-off yields significant performance improvements, particularly in non-isometric settings and those with topological noise.

Table 16: Inference time (seconds per shape pair) and peak GPU memory (GB) comparison on each benchmark.

Method	Metric	Dataset					
		FAUST	SCAPE	SHREC’19	SMAL	DT4D (inter)	TOPKIDS
HybridFMap	Time (s)	0.11	0.10	5.42	6.98	10.79	11.93
	Mem. (GB)	3.3	3.4	3.3	3.7	5.5	18.7
Ours	Time (s)	0.08	0.07	6.27	7.11	8.26	10.93
	Mem. (GB)	3.9	4.0	4.0	3.2	6.3	16.3

As shown in Tab. 16, our method achieves inference times similar to HybridFMap on near-isometric datasets (e.g., 0.08s vs. 0.11s for FAUST). We

Table 17: Semantic feature extraction time per shape.


Dataset	FAUST	SCAPE	SMAL	SHREC'19	DT4D	TOPKIDS
Time (s)	34.7	42.1	38.5	36.2	36.2	109.4

observe a slight increase in runtime on larger benchmarks such as SHREC'19 and DT4D, reflecting the added complexity of our modules. Regarding peak GPU memory, our method requires roughly 0.6 GB more than HybridFMap on near-isometric data but remains within a similar range on non-isometric and topological-noise datasets. These results indicate that the additional computational overhead is manageable and well-compensated by the accuracy gains under challenging deformations.

Semantic feature extraction is performed once per shape and can be cached for repeated pairwise evaluations. We report this one-time preprocessing cost in Tab. 17.

I More Qualitative Results

In the figures below, we provide additional qualitative results of our method on SHREC'19, SMAL, DT4D-H inter class, and TOPKIDS datasets.

**Fig. 10:** Additional qualitative results of our method on SHREC'19.




Fig. 11: Additional qualitative results of our method on SMAL.




Fig. 12: Additional qualitative results of our method on DT4D-H inter class.




Fig. 13: Additional qualitative results of our method on TOPKIDS.