DOI: 10.1111/cgf.14913
Eurographics Symposium on Geometry Processing 2023
P. Memari and J. Solomon
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 42 (2023), Number 5
e14913 *pp. 1 - 10*

Partial Matching of Nonrigid Shapes by Learning Piecewise Smooth Functions

David Bensaïd Noam Rotstein Nelson Goldenstein Ron Kimmel
Technion - Israel Institute of Technology

Abstract

Learning functions defined on non-flat domains, such as outer surfaces of non-rigid shapes, is a central task in computer vision and geometry processing. Recent studies have explored the use of neural fields to represent functions like light reflections in volumetric domains and textures on curved surfaces by operating in the embedding space. Here, we choose a different line of thought and introduce a novel formulation of partial shape matching by learning a piecewise smooth function on a surface. Our method begins with pairing sparse landmarks defined on a full shape and its part, using feature similarity. Next, a neural representation is optimized to fit these landmarks, efficiently interpolating between the matched features that act as anchors. This process results in a function that accurately captures the partiality. Unlike previous methods, the proposed neural model of functions is intrinsically defined on the given curved surface, rather than the classical embedding Euclidean space. This representation is shown to be particularly well-suited for representing piecewise smooth functions. We further extend the proposed framework to the more challenging part-to-part setting, where both shapes exhibit missing parts. Comprehensive experiments highlight that the proposed method effectively addresses partiality in shape matching and significantly outperforms leading state-of-the-art methods in challenging benchmarks. Code is available at `https://github.com/davidgip74/Learning-Partiality-with-Implicit-Intrinsic-Functions`

1. Introduction

Non-rigid shape matching is an essential task in computer vision, with applications spanning augmented reality, 3D medical modeling, and computer-aided design. Real-world scenarios frequently involve handling 3D data captured with partial views and occlusions, necessitating the development of methods capable of addressing shape matching *under partiality*. In these situations, only certain regions of the shapes are expected to exhibit similarity due to missing geometry. In this paper, we present a novel formulation of partiality in non-rigid shape matching. Partiality is represented as a piecewise smooth function defined on a target shape and is modeled as an indicator function that identifies the effective support of a part within a full shape. This formulation leads to an effective framework for partial shape matching consisting of two principal stages: feature matching and partiality function optimization. During the feature matching stage, sparse anchors are generated between shapes according to the similarity of features extracted using a pretrained neural network. The function representing partiality is then encoded by a neural network optimized to fit these anchors.

Representing functions with neural networks, also known as neural fields, aligns with recent studies that parameterize the physical properties of scenes using coordinate-based Multilayer Perceptrons (MLPs). Here, the neural representation serves as a regularization technique that encourages piecewise smoothness of the partiality function. This approach enables efficient interpolation of the sparse anchors, leading to a function that accurately captures partiality. In contrast to previous studies that employ neural fields for 3D objects operating on Cartesian coordinates, we adopt an intrinsic neural representation for functions defined on curved surfaces. This representation circumvents dependence on the geometry of the Euclidean embedding space and is based on the eigenvectors of the Laplace-Beltrami Operator (LBO). Our choice is motivated by their optimality as a truncated basis for approximating smooth functions on Riemannian manifolds [ABK15].

Initially, we address the partial shape matching setting, considering a full and a partial shape. Our framework is then extended to the more challenging partial-to-partial configuration. Extensive experiments demonstrate that the proposed approach achieves state-of-the-art results while maintaining robustness to mesh discretization and computational efficiency.

The main contributions of this paper include,

- A novel perspective on partiality in shape analysis that is effectively modeled as a piecewise smooth function, leading to a general, robust, and flexible framework for addressing partiality in non-rigid shape matching.
- An experimental analysis of the approximation of piecewise smooth functions defined on 2-Riemannian manifolds, showcasing the advantages of the proposed intrinsic approach over previous extrinsic ones.
- A novel framework for both partial and partial-to-partial shape matching, demonstrated through extensive experiments on stan-

dard benchmarks, such as SHREC'16 CUTS [CRB*16] and CP2P [APO21], outperforming state-of-the-art methods while maintaining computational efficiency.

Figure 1: *The first eight eigenvectors of the Laplace-Beltrami Operator (LBO) of quasi-isometric shapes. The eigenvectors of the LBO are invariant to isometry up to a sign change.*

2. Related Efforts

We provide an overview of recent advances related to the proposed approach. Our review begins with methods addressing shape matching under partiality, followed by an examination of neural representations of functions on meshes.

2.1. Non-Rigid Partial Shape Matching

The task of non-rigid shape matching under partiality was initially addressed by Bronstein *et al.* in 2009 [BBBK09], demonstrating superior performance compared to early non-rigid ICP algorithms [CR03]. Later, Partial Functional Maps (PFM) [RCB*17] extended the functional maps framework [OBS*12] to handle partial shapes by alternately optimizing partial shape localization and correspondence. Notable follow-up papers include a joint diagonalization method that aligns the spectral bases of a complete and a partial shape in the spectral domain [LRBB17] and an iterative refinement procedure [MRR*19] that alternates between coarse matching estimated with PFM and an up-sampling strategy. Recently, Deep Partial Functional Map (DPFM) [APO21] was introduced as the first learning method specifically dedicated to partial non-rigid shape correspondence. DPFM uses a neural architecture based on Diffusion Net [APO21] to extract features and generate a functional map, which is trained using a loss function inspired by the partial functional map framework [OBS*12]. Additionally, Rampini *et al.* [RTO*19] and Bensaïd *et al.* [BBK23] solve the problem of partial shape similarity using spectra alignment procedures.

2.2. Partial-to-Partial Shape Matching

Partial-to-partial shape matching involves identifying similar overlapping regions between two incomplete shapes. It generalizes the partial shape matching configuration, where one shape is assumed full. Early methods for partial-to-partial shape matching relied on rotation-invariant features, which were unsuitable for non-rigid objects [LSP08]. The Fully Spectral Partial (FSP) shape matching framework [LRBB17] marked a notable attempt to address partial-to-partial non-rigid shape matching. Litany et al. [LRB*18] build upon FSP [LRB*18] and PFM [RCB*17] by proposing a solution for non-rigid puzzles that combines part-to-full matching with clutter handling. In 2021, Attaiki *et al.* demonstrated that DPFM [APO21], the current state-of-the-art learning method for partial shape matching, can be adapted to the partial-to-partial configuration.

2.3. Neural Fields

Representing implicit functions with Multilayer Perceptrons (MLPs) is a widely used technique in computer vision and graphics, known as neural fields. MLPs operating in the Euclidean embedding space have been probed to represent visual signals such as images [WPG10; Sta07] and 3D scenes [MON*19; PFS*19]. Tancik *et al.* [TSM*20] proposed to replace coordinate based MLPs with MLPs operating on a Random Fourier Features (RFF) embedding [RR07] of the Cartesian coordinates $p \in \mathbb{R}^3$,

$$f_\theta([\sin(2\pi b_j^{\mathrm{T}} p), \cos(2\pi b_j^{\mathrm{T}} p)]_{j=1}^{k}), \tag{1}$$

where $\{b_j | b_j \in \mathbb{R}^3\}_{j=1}^{k}$ is a set of normal (Gaussian) random vectors and $f_\theta : \mathbb{R}^{2k} \to \mathbb{R}^m$ is a MLP. The injection of high frequency harmonics in Eq. 1 improves the representation capacities of neural fields and leads to state-of-the-art results in numerous applications [TSM*20]. This idea paved the way for learned representations parameterized with MLPs, such as Neural Radial Field (NeRF) [MST*21] and its extensions [BMV*22; MESK22; YYTK21].

The method proposed by Tancik *et al.* [TSM*20] relies on the notion of positional encoding. The design of positional encoding for representing the position of an element in an organized set has received significant attention in various fields such as natural language processing and graph learning. Early methods include sinusoidal functions [VSP*17; XRK*19], while more recent approaches consider learned positional embeddings [LSL*21]. The use of the eigenvectors of the Laplacian to encode locations has recently been explored in the context of graphs [DLL*21] and for texture reconstruction of meshes from multiple views [KGM*22].

3. Representing Piecewise Smooth Functions in non-Euclidean Domains

In geometry processing, a shape in 3D is usually modeled by its boundary as a two dimensional Riemannian manifold $\mathcal{M} = (S, g)$, where S is a smooth two-dimensional surface embedded in $\mathbb{R}^3$ and g is a metric tensor, also referred to as the *first fundamental form*.

3.1. The Laplace-Beltrami Operator

The Laplace-Beltrami operator (LBO) Δ_g is an ubiquitous operator in shape analysis [BBL*17]. It generalizes the Laplacian operator to Riemannian manifolds and can intuitively be interpreted as a local measure of the deviation of a function a from its value around a point [BBL*17]. Formally,

$$\Delta_g f = \frac{1}{\sqrt{|g|}} \mathrm{div}(\sqrt{|g|} g^{-1} \nabla a), \; a \in \mathcal{L}^2(\mathcal{M}),$$

where $|g|$ is the determinant of the metric tensor g and $\mathcal{L}^2(\mathcal{M})$ is the Hilbert space of square-integrable scalar functions defined on $\mathcal{M}$.

The LBO is a semi-negative definite operator and thus admits a spectral decomposition under homogeneous Dirichlet boundary conditions [Tau95],

$$\begin{aligned} -\Delta_g \phi_i(x) &= \lambda_i \phi_i(x), & x &\in \mathcal{M} \setminus \partial\mathcal{M} \\ \phi_i(x) &= 0, & x &\in \partial\mathcal{M} \end{aligned} \tag{2}$$

where $\partial\mathcal{M}$ denotes the boundary of the manifold $\mathcal{M}$. The basis $\{\phi_i\}_{i\geq 0}$ is invariant to isometric transformations, as illustrated in Fig. 1. Since Eq. 2 also holds for the Fourier basis on Euclidean domains, the basis defined by the LBO spectral decomposition is commonly interpreted as its generalization on Riemannian manifolds [Tau95]. Truncating this basis has been proven to be unique and optimal in approximating smooth functions on a given manifold [ABK15].

3.2. Intrinsic Neural Representation

Neural architectures for non-Euclidean domains have received significant attention in recent years. The success of convolutional architecture in Euclidean domains has motivated the development of numerous generalizations of convolutional layers for graphs and meshes. One prominent approach is the spectral convolutional layer [BZSL13; BBL*17], which is grounded in a fundamental theorem of signal processing stating that convolution in the time domain is equivalent to multiplication in the Fourier domain. It is defined as

$$\xi((\hat{g} \odot \Phi^{\mathrm{T}} a)\Phi),$$

where a is the signal processed, $\hat{g}$ is a learned filter, $\odot$ represents the Hadamard product, ξ is an activation function such as ReLU, and $\Phi \in \mathbb{R}^{n\times k}$ is the truncated basis, determined by the LBO, with k being its cardinality and n the number of vertices in the mesh considered. We propose a novel neural representation that operates in the spectral domain, building on a similar idea. In our architecture, the i-th layer of the neural network is defined as follows:

$$l_i = \begin{cases} \xi(\mathbf{W}_1 \Phi^{\mathrm{T}} + \mathbf{b}_1) & \text{if } i = 1 \\ \xi(\mathbf{W}_i l_{i-1} + \mathbf{b}_i) & \text{otherwise,} \end{cases}$$

where $\{\mathbf{W}_i, \mathbf{b}_i\}$ are the parameters of the i-th layer l_i. The proposed neural representation can be regarded as an extension of [TSM*20], wherein the RFF embedding utilized in Eq. 1 is substituted with an embedding based on the LBO eigenvectors.

Fig. 2 compares the eigenvectors of the LBO with the RFF as positional encoding. By accounting for connectivity, the LBO basis notably allows for distinct values on separate body parts, which may have close Euclidean embedding in certain poses. This property is advantageous in the context of partiality, as the functions representing partiality should isolate different body parts. Henceforth, we denote the proposed intrinsic neural representation as $f_\theta : \mathbb{R}^n \to [0,1]$.

4. Partial Shape Matching

4.1. Overview

We propose a novel framework for localizing the effective support of a partial shape within a full shape. Initially, anchors across shapes

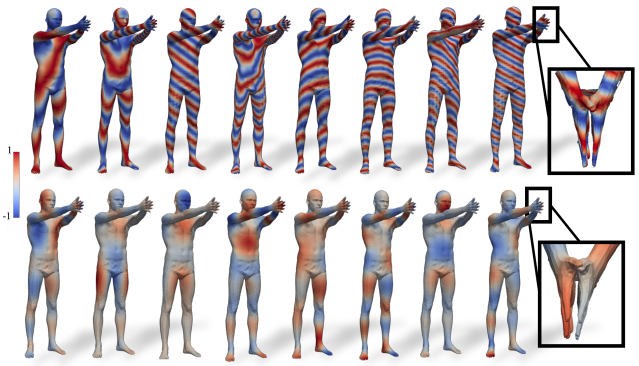

Figure 2: *Comparison of Random Fourier Features (RFF) (top) and LBO Eigenvectors (bottom) for a FAUST subject. The intrinsic nature of the LBO eigenvectors, which account for mesh connectivity, makes them well-suited for representing indicator functions capable of isolating semantically meaningful body parts, such as a single hand, even when both hands are close in the Euclidean domain.*

are identified through a feature matching procedure. Specifically, we employ a pre-trained Siamese neural network to extract features from every vertex in both the complete and partial shapes. These features are subsequently matched using our proposed coarse feature matching stage, which generates landmarks across the shapes. The neural representation, detailed in Section 3.2, is then iteratively optimized to fit the anchors on the target shape. This representation is leveraged to effectively interpolate between the sparse landmarks, enabling the generation of a function that accurately captures the partiality. Fig. 3 provides an overview of the proposed framework and its various components.

Problem formulation. We formulate the problem of partial shape similarity as the task of finding a suitable binary function defined on a full shape. Specifically, given a full shape $\mathcal{M}_{\mathrm{f}}$ and a partial shape $\mathcal{M}_{\mathrm{p}}$, the proposed framework outputs an indicator function $o : \mathcal{M}_{\mathrm{f}} \to \{0,1\}$, which indicates the effective support of the partial shape within the full shape. The indicator function should satisfy

$$\mathcal{M}_o \sim \mathcal{M}_{\mathrm{p}},$$

where $\mathcal{M}_o = \{x \in \mathcal{M}_{\mathrm{f}} | \, o(x) = 1\}$ denotes the subset of the full shape that is effectively supported by the partial shape, and $\sim$ stands for an isometry relation.

4.2. Finding Anchors

Feature extractor. We employ a feature extraction module based on Diffusion-Net [SACO22], a geometric neural network that relies on intrinsic operations. This architecture has been shown robust to various shape discretizations and invariant to isometries. We apply the feature extractor in a Siamese manner, using the same network and identical weights for both the full and partial shapes. The extracted features are denoted as $\mathcal{F}_{\mathrm{f}} \in \mathbb{R}^{n_{\mathrm{f}}\times k}$ and $\mathcal{F}_{\mathrm{p}} \in \mathbb{R}^{n_{\mathrm{p}}\times k}$, where k represents the number of extracted features, and n_{f}, n_{p} correspond to the number of vertices in the full and partial shape, respectively.

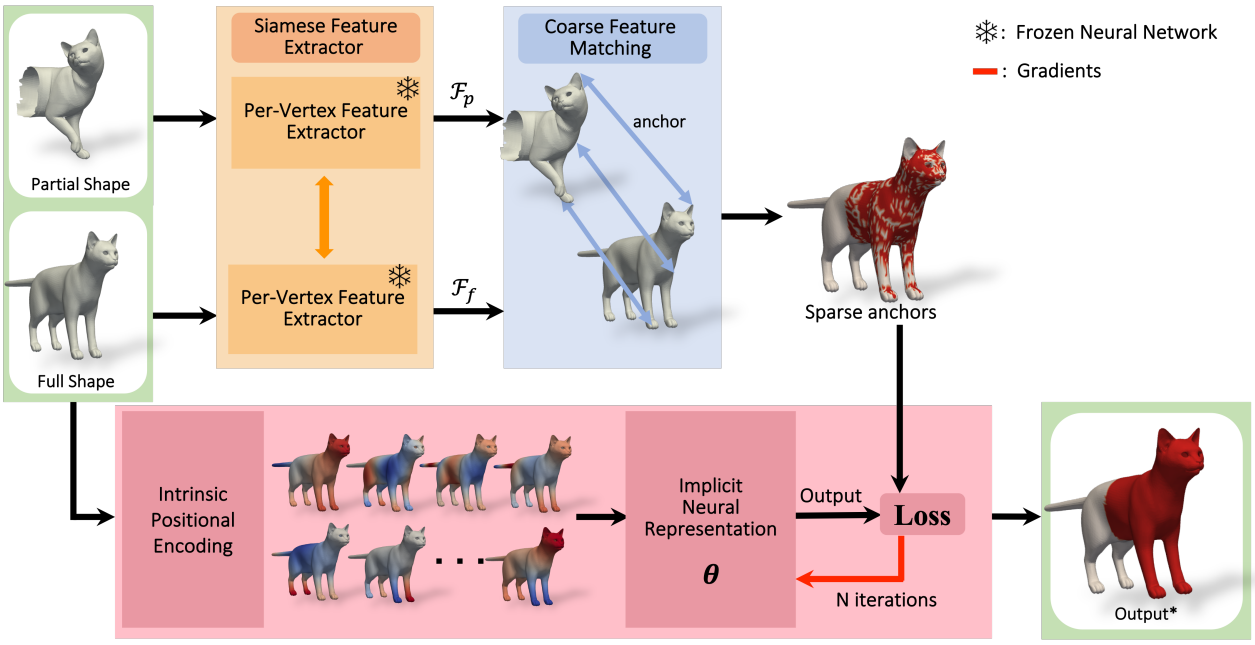

Figure 3: ***Overview of the proposed framework for partial shape matching.*** *Features of the partial and full shapes are extracted using a previously trained Siamese Feature Extractor. These features undergo a Feature Matching stage to identify corresponding sparse anchors across the shapes. An intrinsic neural representation is then iteratively optimized to fit these landmarks, enabling efficient interpolation between them and accurately capturing the partiality within the full shape.*

The feature extractor is pretrained independently from the remaining components of the framework. To achieve this, we utilize a regularized functional map module [DSO20] to predict a functional map $\mathbf{C}$ between shapes. The functional map $\mathbf{C}$ is obtained by solving the following least squares problem in a differentiable manner [LRR*17],

$$\mathbf{C} = \underset{C'}{\operatorname{argmin}} \|\Phi^{\mathrm{T}}\mathcal{F}_{\mathrm{f}} - C'\Psi^{\mathrm{T}}\mathcal{F}_{\mathrm{p}}\|_F, \tag{3}$$

where $\|\cdot\|_F$ denotes the Froebenius norm, Ψ is the truncated basis for the partial shape, and Φ is the truncated basis for the full shape. Following [DSO20], we incorporate to Eq. 3 additional regularization terms on $\mathbf{C}$. The feature extractor is trained using a spectral loss

$$\mathbf{L}_{\text{spectral}} = \|\mathbf{C} - \mathbf{C}_{gt}\|_F^2,$$

which compares $\mathbf{C}$ to the ground truth functional map $\mathbf{C}_{gt}$.

It is important to emphasize that, during inference, functional maps are not employed; rather, we only rely on the features produced by the Siamese neural network. The proposed approach harnesses the functional map framework [OBS*12] to efficiently train the feature extractor by solving function-based matching problem within a low-dimensional space rather than a point-based one. Functional maps also offer the substantial advantage of being adaptable to unsupervised setting [HLR*19; RSO19]. Nevertheless, the transformation of a functional map into a point-wise correspondence map presents a non-trivial challenge and is an active research field [CRB23; PRM*21; VRBC17; EB17]. It can be computationally expensive and typically results in a many-to-one mapping [CB22]. To circumvent these challenges, at inference, we resort to a direct feature similarity approach to derive anchors. This strategy aligns with recent studies [LAO22; CB22] which also rely on the similarity of features learned via the functional map framework to determine point-to-point correspondences during inference.

Matching scheme. Following feature extraction, the goal is to identify *landmarks* in the full shape that correspond to points of the partial shape. This task is achieved by constructing a permutation matrix $\mathbf{P} \in \mathbb{R}^{n_{\mathrm{f}} \times n_{\mathrm{p}}}$, which accounts for partiality using the generated features. The matrix is defined as

$$\mathbf{P} = \mathrm{Softmax}(\mathcal{F}_{\mathrm{f}}\mathcal{F}_{\mathrm{p}}^{\mathrm{T}}),$$

where $\mathrm{Softmax}(\cdot)$ denotes a softmax operation applied to each column. Each row in $\mathbf{P}$ corresponds to a vertex in the full shape and the associated column indicates the probability of matching the vertex in the partial shape corresponding to that particular column.

Subsequently, a mask $\mathbf{m} \in \mathbb{R}^{n_f}$ indicating the anchors is inferred from $\mathbf{P}$ by considering the corresponding points of the partial shape predicted with *high confidence*,

$$\mathbf{m}_i = \begin{cases} 0 & \text{if } (\mathbf{P1})_i < c \\ 1 & \text{if } (\mathbf{P1})_i \geq c \end{cases},$$

where $\mathbf{1}$ is a vector of ones, and $c \in \mathbb{R}$. $(\mathbf{P1})_i$ corresponds to the estimated probability of vertex i belonging to the partial shape. The threshold c determines the required probability for a vertex to be considered as an anchor.

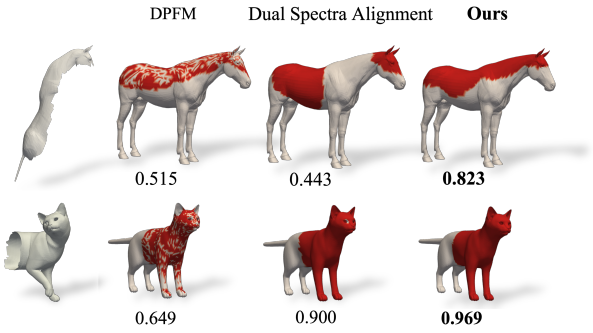

Figure 4: *Region localization. Comparison of the proposed approach for partial shape similarity with state-of-the-art learning [APO21] and geometric [BBK23] methods applied to SHREC'16. The red regions indicate the estimated match of the query parts (left column); the geometrical IoU is shown below each mask. The proposed method is effective in handling challenging cuts (top) and produces precise results thanks to the regularization on the queried area (bottom).*

4.3. Learning Partiality

The intrinsic neural representation $f_\theta : \mathbb{R}^{n_f} \to [0,1]$ introduced in Section 3.2 is optimized via the following objective,

$$\mathbf{L} = \mathbf{L}_{\text{anchor}} + \lambda \mathbf{L}_{\text{area}}\,,$$

where $\lambda \in \mathbb{R}^+$ is a regularization parameter. $\mathbf{L}_{\text{anchor}}$ is a binary cross entropy loss promoting the inclusion of the produced anchors,

$$\begin{aligned}\mathbf{L}_{\text{anchor}} &= \text{BCE}(f_\theta, \mathbf{m}) \\ &= \frac{1}{n_f}\sum_{i=1}^{n_f} \mathbf{m}_i \log(f_\theta(p_i)) + (1-\mathbf{m}_i)\log(1-f_\theta(p_i))\end{aligned}$$

where p_i is the embedding of the vertex i. $\mathbf{L}_{\text{area}}$ regularizes the area of the mask predicted by f_θ to ensure it corresponds to the partial shape's area,

$$\mathbf{L}_{\text{area}} = \|\text{trace}(\mathbf{A}_{\text{f}}\mathbf{h}(f_\theta)) - \text{trace}(\mathbf{A}_{\text{p}})\|_2^2,$$

where $\mathbf{A}_{\text{f}}$ and $\mathbf{A}_{\text{p}}$ are the mass element matrices [Tau95] of the full and partial shapes, respectively and $\mathbf{h} : \mathbb{R}^{n_f} \to [0,1]^{n_f}$ is a function applying onto each entry $h(x) = \frac{1}{2}(1+\tanh(\alpha x - \frac{1}{2}))$, with $\alpha \in \mathbb{R}$.

The proposed intrinsic neural representation can be regarded as a regularization technique that controls the smoothness of the partiality function over the manifold considered. As explained in Section 3.1, the LBO generalizes the Fourier basis, and the degree of smoothness is controlled by the number of LBO eigenvectors considered.

4.4. Generalization to Partial-to-Partial Shape Matching

Here, the framework introduced for partial shape matching is extended to address the partial-to-partial matching task. Specifically, building upon the formulation proposed in Subsection 4.1, we formulate the partial-to-partial shape matching problem as determining two piecewise smooth functions. These functions represent the maximal overlapping region between the two partial shapes under consideration. Formally, let $\mathcal{M}_{\text{p}_1}$ and $\mathcal{M}_{\text{p}_2}$ be two partial shapes and $o_1 : \mathcal{M}_{\text{p}_1} \to \{1,0\}$, $o_2 : \mathcal{M}_{\text{p}_2} \to \{1,0\}$ two indicator functions. o_1 and o_2 should satisfy

$$\{o_1, o_2\} = \underset{\{o_1', o_2'\}}{\text{argmax}}\ \mathbf{A}_{\text{p}_1} o_1' + \mathbf{A}_{\text{p}_2} o_2' \quad \text{s.t. } \mathcal{M}_{o_1} \sim \mathcal{M}_{o_2},$$

where $\mathbf{A}_{\text{p}_1}$ and $\mathbf{A}_{\text{p}_2}$ are the mass element matrices [Tau95] of $\mathcal{M}_{\text{p}_1}$ and $\mathcal{M}_{\text{p}_2}$, respectively, and $\mathcal{M}_{o_i} = \{x \in \mathcal{M}_{\text{p}_i} \mid o_i(x) = 1\}$ for $i = 1,2$.

In partial shape matching, each point of the partial shape has a corresponding point in the full shape. This assumption does not hold in the partial-to-partial setting, and the permutation matrix $\mathbf{P}$ is thus defined as

$$\mathbf{P} = \mathcal{F}_{\text{p}_1}\mathcal{F}_{\text{p}_2}^{\text{T}},$$

where $\mathcal{F}_{\text{p}_1}$ and $\mathcal{F}_{\text{p}_2}$ are normalized features (each channel has a L_2 norm of 1) extracted with the Siamese network introduced in Subsection 4.2. Two masks, denoted $\mathbf{m}^1 \in \mathbb{R}^{n_1}$ and $\mathbf{m}^2 \in \mathbb{R}^{n_2}$, are subsequently inferred from $\mathbf{P}$.

Next, *two* intrinsic neural representations, $f_\theta^1 : \mathcal{M}_{\text{p}_1} \to \mathbb{R}$ and $f_\theta^2 : \mathcal{M}_{\text{p}_2} \to \mathbb{R}$, are concurrently optimized to represent o_1 and o_2, respectively. We consider the loss function

$$\mathbf{L} = \mathbf{L}_{\text{anchor}_1} + \mathbf{L}_{\text{anchor}_2} + \lambda_1 \mathbf{L}_{\text{area}_1} + \lambda_2 \mathbf{L}_{\text{area}_2}\,,$$

with $\lambda_1, \lambda_2 \in \mathbb{R}^+$.

$\mathbf{L}_{\text{anchor}_i}$, with $i = 1,2$, is a binary cross entropy loss enforcing the inclusion of the anchors on $\mathcal{M}_{\text{p}_i}$,

$$\begin{aligned}\mathbf{L}_{\text{anchor}_i} &= \text{BCE}(f_\theta^i, \mathbf{m}) \\ &= \frac{1}{n_{p_i}}\sum_{j=1}^{n_{p_i}} \mathbf{m}_j^i \log(f_\theta^1(p_j^i)) + (1-\mathbf{m}_j^i)\log(1-f_\theta^1(p_j^i))\end{aligned} \tag{4}$$

$\mathbf{L}_{\text{area}_1}$ ensures that the overlapping regions defined by f_θ^1 and f_θ^2 have equal areas

$$\mathbf{L}_{\text{area}_1} = \|\text{trace}(\mathbf{A}_{\text{p}_1}\mathbf{h}(f_\theta^{\text{p}_1})) - \text{trace}(\mathbf{A}_{\text{p}_2}\mathbf{h}(f_\theta^{\text{p}_2}))\|_2^2,$$

Finally, $\mathbf{L}_{\text{area}_2}$ favors the largest possible overlapping region

$$\mathbf{L}_{\text{area}_2} = -\sum_{j=1}^{2} \text{trace}(\mathbf{A}_{p_j}\mathbf{h}(f_\theta^{\text{p}_j})).$$

4.5. Implementation Considerations

The intrinsic neural representation f_θ is implemented with three hidden layers of width 128 and we consider 15 eigenvectors of the LBO as input. The discrete LBO is computed using the cotangent scheme (first order finite elements approximation). For partial shape matching, we set $\lambda = 0.1$ and $\alpha = 5$. For the partial-to-partial setting we set $\lambda_1 = 10$, $\lambda_2 = 0.1$ and $\alpha = 5$. The selection of c impacts the trade-off between precision and recall. Specifically, a smaller c value enhanced precision, while a larger c fosters higher recall. We tuned c on the training data to obtain a satisfactory balance between precision and recall and set $c = 0.02$ for all experiments. The optimization is performed for 150 steps using the Adam optimizer with a learning rate of 10^{-4}. We use an Intel i7 processor and an NVIDIA RTX 2080 Ti GPU card for all experiments.

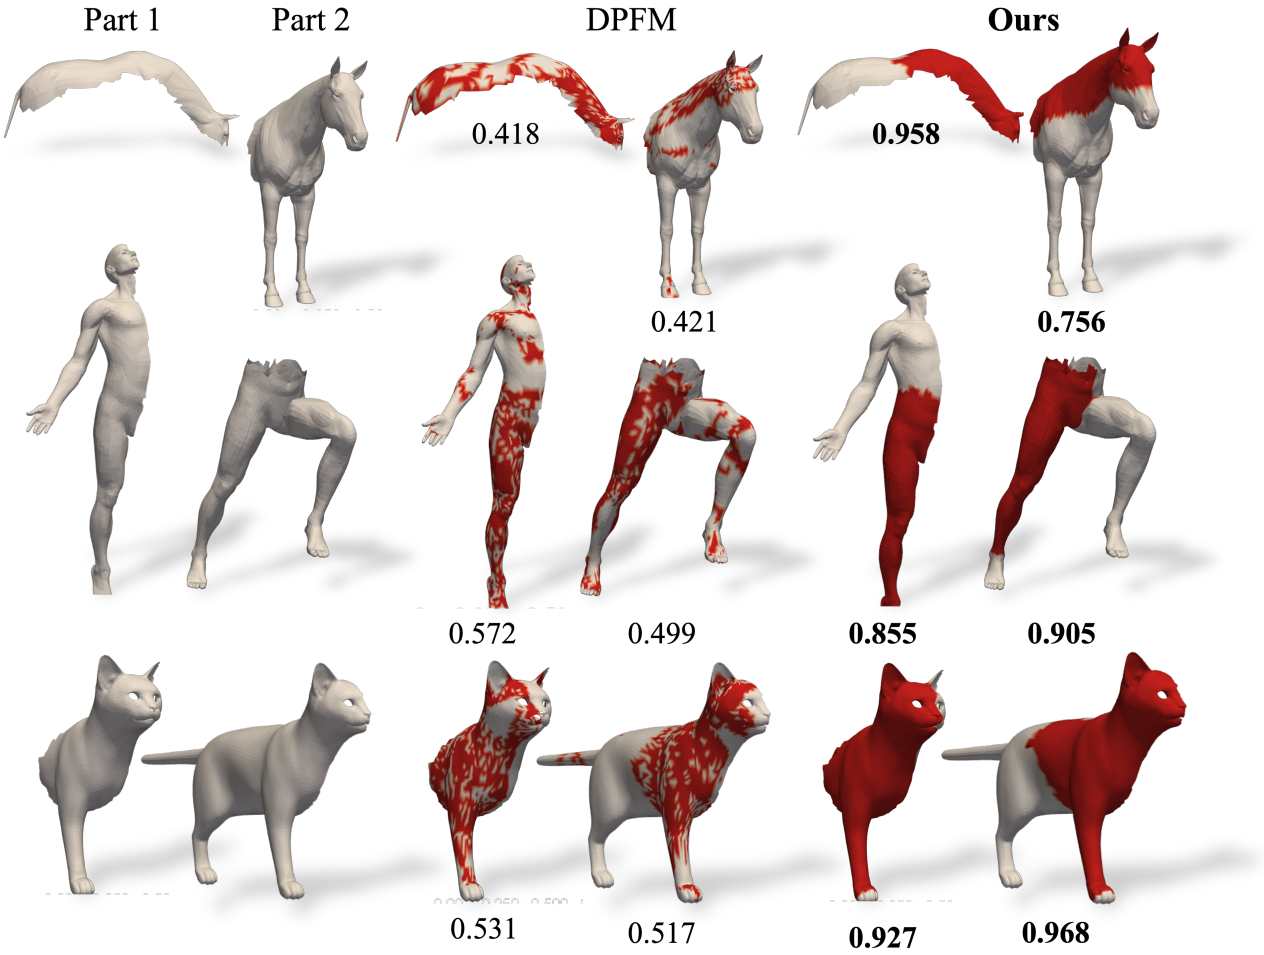

Figure 5: ***Partial-to-Partial Shape Matching***. *Comparison between the proposed framework and DPFM, the current state-of-the-art approach, [APO21] for partial-to-partial shape matching. The red regions represent the estimated match between the overlapping regions, and the geometrical Intersection over Union (IoU) is shown below each mask.*

5. Experiments and Results

The proposed framework is evaluated on established benchmarks in non-rigid partial and partial-to-partial shape matching. Our method is compared with state-of-the-art approaches for partial shape matching, including the spectrum alignment procedure proposed by Rampini *et al.* [RTO*19], the dual spectra alignment procedure proposed by Bensaïd *et al.* [BBK23], Partial Functional Correspondences (PFC) [RCB*17], and a bag-of-words aggregation [TCF09] of SHOT descriptors [STD14]. Additionally, we compare the proposed method to Deep Partial Functional Map (DPFM) [APO21], the state-of-the-art deep learning framework for partial and partial-to-partial shape matching. We employ the original codes provided by the authors and apply the best reported hyper-parameters.

5.1. Partial Shape Matching

We evaluate the proposed framework on SHREC'16 CUTS [CRB*16], the standard benchmark for non-rigid partial shape matching. The benchmark consists of 120 partial shapes from 8 classes, including quadrupeds such as dogs, horses, centaurs, wolves, and cats, as well as human subjects. Fig. 4 provides a qualitative comparison of frameworks for partial shape matching and demonstrates the efficacy of our approach in handling challenging cuts and produces precise results. For quantitative evaluation, we consider standard segmentation metrics, namely geometric Intersection over Union (IoU), precision, and recall. Fig. 6 (Left) compares the cumulative IoU of our framework and alternative techniques, while Table 1 summarizes their performances. The proposed method outperforms competing methods and achieves a significant improvement of 28% in IoU over the current state of the art in SHREC'16 CUTS [CRB*16]. Table 1 also presents the running time of each method in identical conditions. Our method executes in less than 5 seconds when applied to high-resolution meshes of SHREC'16, which comprise over 20,000 vertices, using the hardware configuration described in Subsection 4.5. Specifically, the proposed approach offers a significant speed-up advantage over competing iterative methods,

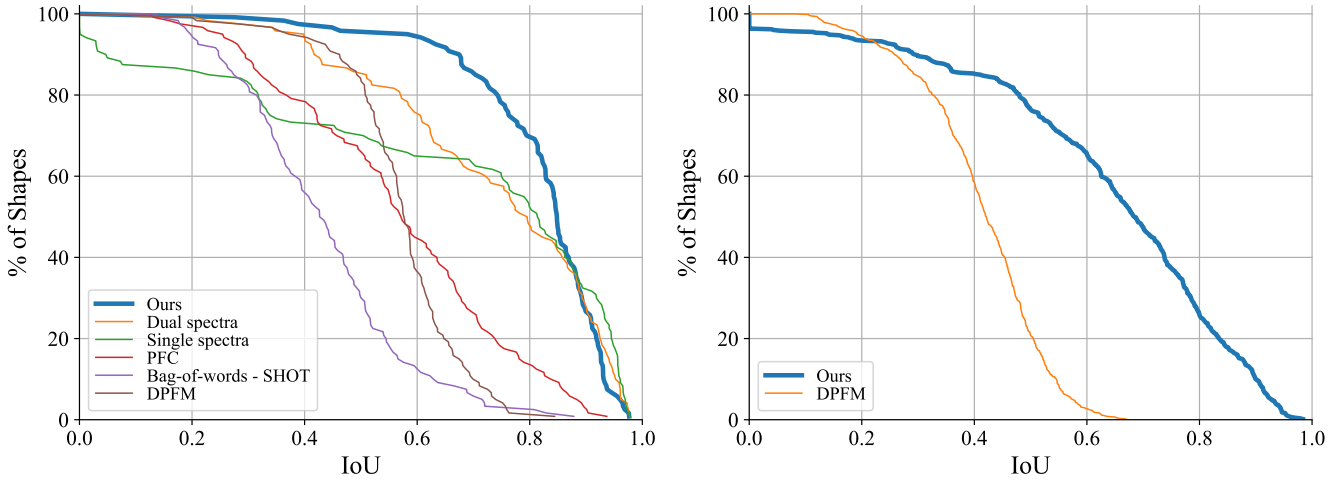

Figure 6: *(Left)* ***Partial Shape Matching.*** *Cumulative Intersection over Union (IoU) comparison of various methods on SHREC'16 CUTS dataset. The areas under the curves are the mean IoUs reported in Table 1. (Right)* ***Partial-to-partial.*** *Cumulative IoU comparison of the proposed method and DPFM on CP2P [APO21]. The areas under the curves correspond to the mean IoUs reported in Table 2.*

Method	Precision	Recall	IoU	Running Time
Bag-of-words of SHOT descriptors [STD14]	0.653	189	0.430	$\sim 5s$
PFC [RCB*17]	0.938	0.573	0.564	$\sim 30s$
Single spectra alignment [RTO*19]	0.775	0.738	0.668	$\sim 800s$
DPFM [APO21]	**0.975**	0.576	0.569	$\sim 1s$
Dual Spectra Alignment [BBK23]	0.859	0.838	0.751	$\sim 850s$
Proposed method	0.847	**0.949**	**0.818**	$\sim 4s$

Table 1: *Quantitative comparison of partial shape matching approaches on SHREC'16 CUTS [CRB*16]. The proposed framework achieves state of the art results.*

such as [BBK23; RTO*19], which require several minutes to complete. To ensure a fair evaluation, we conducted a random three-fold split of SHREC'16. Models were trained on two folds and evaluated on the remaining distinct fold. This procedure was repeated for each of the splits. In contrast to the evaluation methodology for partial shape matching suggested in [APO21], we evaluate models solely on pairs of shapes that were not included in the training set.

5.2. Partial-to-Partial Shape Matching

Method	Precision	Recall	IoU
DPFM [APO21]	0.653	0.589	0.430
Proposed method	**0.835**	**0.721**	**0.643**

Table 2: ***Partial-to-Partial Shape Matching.*** *Quantitative comparison of the proposed framework with DPFM, the current state-of-the-art approach, on CP2P [APO21]. Our method outperforms DPFM in all the metrics considered, showcasing its superiority over the current state-of-the-art approach.*

Figure 7: *Piecewise smooth functions from our newly introduced dataset. The considered functions are indicators on body parts such as the head, single or couple of hands, and feet.*

In Section 4.4, we extend our framework for partial shape matching to the more challenging partial-to-partial setting. In this context, we evaluate the proposed approach using the CP2P benchmark

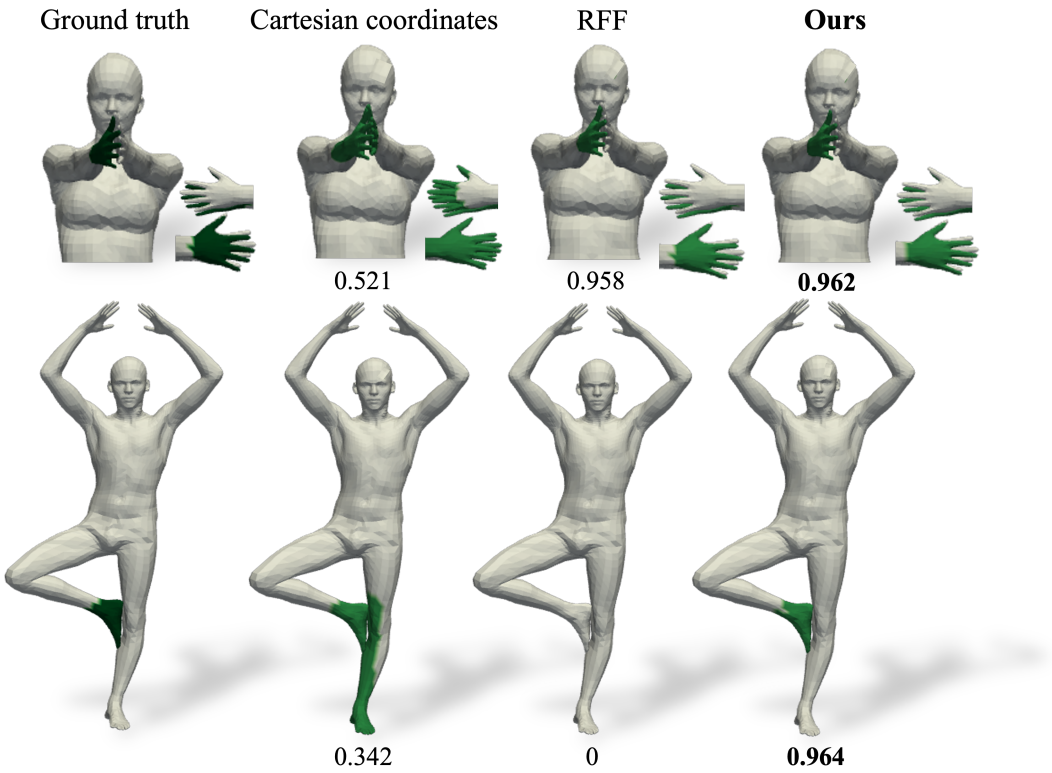

Figure 8: *Neural representations of ground-truth binary functions (first column). We consider three types of neural representations: a Multilayer Perceptron (MLP) operating on Cartesian coordinates, Random Fourier Features (RFF), and our proposed approach operating on a Laplacian Beltrami Operator (LBO) embedding. The Intersection over Union (IoU) value obtained by each neural representation is indicated below the corresponding shape.*

[APO21]. CP2P comprises 300 pairs of partial shapes with overlaps ranging from 10% to 90% of the total area of each shape. In Table 2 and Fig. 6 (Right) we present a quantitative comparison of our method with DPFM [APO21], the current state-of-the-art approach for partial-to-partial shape matching. A qualitative comparison is also shown in Fig. 5. The proposed method outperforms DPFM in all quantitative measures considered and achieves a significant improvement over the existing state-of-the-art by effectively addressing challenging overlaps and generating piecewise smooth functions instead of fractioned masks produced by DPFM [APO21].

5.3. Approximation of Piecewise Smooth Functions in 2-Riemmanian Manifolds

In this subsection, we conduct a comprehensive ablation study to examine the effectiveness of the proposed intrinsic neural representation for representing piecewise smooth functions. Specifically, we disentangle the feature matching stage and the intrinsic neural representation to accurately assess the latter. This analysis allows to isolate the impact of the intrinsic neural representation on the overall performance and gain a deeper understanding of its capabilities in approximating piecewise smooth functions.

5.3.1. Representation Capabilities of Intrinsic and Extrinsic Neural Representations

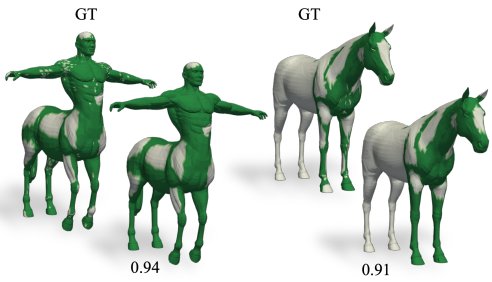

Figure 9: *Ground truth of complex indicators from SHREC HOLES [CRB*16] (left) and the approximation provided by the proposed neural representation (right). The representation of complex functions is notably enhanced by adjusting the number of eigenvectors considered. In this case, we consider the 50 first eigenvectors of the LBO to enhance higher frequencies. The IoU for each shape is indicated below it.*

Here, we evaluate the proposed intrinsic approach for representing piecewise smooth functions. We explore both intrinsic and extrin-

sic neural representations, focusing on their capacity to capture semantically meaningful parts of articulated shapes. To this end, we introduce a new dataset comprising 300 piecewise smooth functions defined on shapes from the FAUST dataset [BRLB14]. These functions serve as indicators of semantically meaningful regions, such as the head, foot, and hands, or combinations thereof in various shapes with different poses. The introduced dataset includes various individuals, poses, and challenging cases, such as close and in-contact body parts, that are not present in the random cuts from SHREC'16. Fig. 7 illustrates a subset of the considered functions.

All the neural representations are trained with a binary cross-entropy loss and with identical optimization schemes. Table 3 displays the Intersection over Union (IoU) values achieved by each neural representation. The proposed approach outperforms extrinsic alternatives, demonstrating its superior representation capabilities in approximating meaningful piecewise smooth functions in articulated shapes. Fig. 8 presents qualitative comparisons with neural representations operating on Cartesian embedding and Random Fourier Features (RFF). It highlights the advantage of our intrinsic approach, which, by accounting for connectivity, can isolate and separate different body parts, even when in contact. These results motivate the choice of the proposed neural representation to model functions representing parts of articulated shapes. To further assess the representation capacity of our intrinsic neural representation beyond the scope of the original dataset and the semantically meaningful dataset we introduced, we examine its performance on additional complex functions characterized by high frequencies. As demonstrated in Fig. 9, our approach successfully handles such cases by adjusting the cardinality of the LBO basis considered. This showcases its adaptability and versatility for handling a wide range of piecewise smooth functions.

Method	**IoU**
MLP - Cartesian coordinates	0.543
MLP - RFF [TSM*20]	0.839
Ours	**0.920**

Table 3: *Comparison of the proposed intrinsic approach with extrinsic alternatives on the dataset introduced in Subsection 5.3.1. Our framework significantly outperforms extrinsic alternatives.*

5.3.2. Interpolating Sparse Anchors

We evaluate the ability of the proposed intrinsic framework to represent the ground truth of SHREC CUTS using varying numbers of landmarks, as illustrated in Fig. 10. By utilizing ground truth matching anchors, this experiment helps to disentangle approximation errors that originate from the feature matching scheme used to determine the anchors and those stemming from the interpolation using the proposed intrinsic neural representation. Fig. 10 is obtained by sampling various proportions of the ground truth and training the proposed intrinsic neural representations on those sampled points using a binary cross-entropy loss. The results reveal that only a moderate performance drop is observed even when considering as few as 10% of the points, showcasing the ability of the method to represent functions with sparse sampling. This highlights the effectiveness of the proposed approach in accurately interpolating sparse anchors.

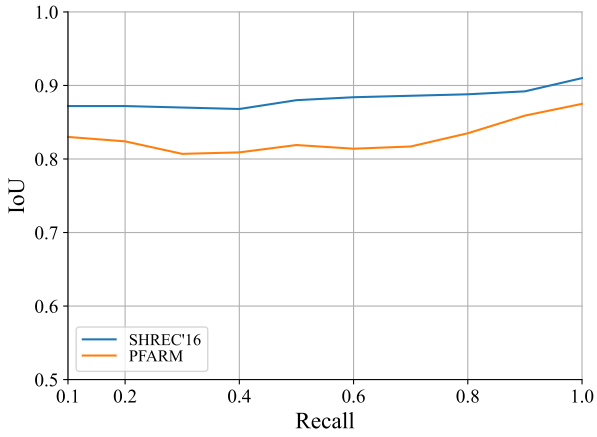

Figure 10: *IoU as a function of the recall on the ground-truth. We select a portion of the ground-truth, i.e. the recall, to serve as landmarks and applied the method explored for partial shape matching. The proposed method effectively represents the considered functions with sparse sampling, which motivates the use of this approach for accurately interpolating sparse anchors.*

5.3.3. Robustness to Mesh Discretization

To investigate the robustness of the proposed interpolation framework with respect to varying discretizations, we evaluate it on PFARM [APO21], a dataset that includes 25 test pairs of humans characterized by significantly distinct connectivity and vertex density. PFARM enables evaluation of the model in challenging scenarios that more closely resemble real-world applications. Fig. 10 demonstrates that the proposed framework for approximating piecewise smooth functions remains effective across different connectivity and vertex densities, exhibiting only a moderate performance drop compared to SHREC'16 [CRB*16].

6. Conclusion

In this paper, we introduce a novel perspective on partiality in nonrigid shape analysis by considering it as a piecewise smooth function defined on a 2-Riemannian manifold. This formulation leads to an effective framework for partial shape matching which employs an intrinsic neural representation for interpolating sparse anchors. The proposed approach significantly outperforms current state-of-the-art methods and is demonstrated to be efficiently extendable to the partial-to-partial configuration. In addition to achieving new best results for partial-to-partial shape matching, this showcases the versatility and broad applicability of the proposed approach for dealing with partiality in shape analysis. As a future research direction, we aim to enhance our framework to complex partiality scenarios, such as those found in SHREC HOLES [CRB*16], by exploring neural feature extractors capable of generating coherent features under these challenging conditions.

References

[ABK15] AFLALO, Y., BREZIS, H., and KIMMEL, R. "On the optimality of shape and data representation in the spectral domain". *SIAM Journal on Imaging Sciences* (2015) 1, 3.

[APO21] ATTAIKI, S., PAI, G., and OVSJANIKOV, M. "DPFM: Deep partial functional maps". (2021) 2, 5–9.

[BBBK09] BRONSTEIN, A., BRONSTEIN, M., BRUCKSTEIN, A., and KIMMEL, R. "Partial similarity of objects, or how to compare a centaur to a horse". *IJCV* (2009) 2.

[BBK23] BENSAÏD, D., BRACHA, A., and KIMMEL, R. "Partial Shape Similarity via Alignment of Multi-Metric Hamiltonian Spectra". *SSVM* (2023) 2, 5–7.

[BBL*17] BRONSTEIN, M., BRUNA, J., LECUN, Y., et al. "Geometric deep learning: going beyond euclidean data". *IEEE Signal Processing Magazine* (2017) 2, 3.

[BMV*22] BARRON, J., MILDENHALL, B., VERBIN, D., et al. "Mip-nerf 360: Unbounded anti-aliased neural radiance fields". (2022) 2.

[BRLB14] BOGO, F., ROMERO, J., LOPER, M., and BLACK, M.J. "FAUST: Dataset and evaluation for 3D mesh registration". (2014) 9.

[BZSL13] BRUNA, J., ZAREMBA, W., SZLAM, A., and LECUN, Y. "Spectral networks and locally connected networks on graphs". *arXiv* (2013) 3.

[CB22] CAO, D. and BERNARD, F. "Unsupervised deep multi-shape matching". (2022) 4.

[CR03] CHUI, HAILI and RANGARAJAN, ANAND. "A new point matching algorithm for non-rigid registration". *CVIU* (2003) 2.

[CRB*16] COSMO, L., RODOLA, E., BRONSTEIN, M.M., et al. "SHREC'16: Partial matching of deformable shapes". (2016) 2, 6–9.

[CRB23] CAO, D., ROETZER, P., and BERNARD, F. "Unsupervised Learning of Robust Spectral Shape Matching". *arXiv* (2023) 4.

[DLL*21] DWIVEDI, V., LUU, A., LAURENT, T., et al. "Graph neural networks with learnable structural and positional representations". *arXiv* (2021) 2.

[DSO20] DONATI, N., SHARMA, A., and OVSJANIKOV, M. "Deep geometric functional maps: Robust feature learning for shape correspondence". (2020) 4.

[EB17] EZUZ, D. and BEN-CHEN, M. "Deblurring and denoising of maps between shapes". (2017) 4.

[HLR*19] HALIMI, O., LITANY, O., RODOLA, E., et al. "Unsupervised learning of dense shape correspondence". (2019) 4.

[KGM*22] KOESTLER, L., GRITTNER, D., MOELLER, M., et al. "Intrinsic neural fields: Learning functions on manifolds". (2022) 2.

[LAO22] LI, L., ATTAIKI, S., and OVSJANIKOV, M. "SRFeat: Learning locally accurate and globally consistent non-rigid shape correspondence". *arXiv* (2022) 4.

[LRB*18] LITANY, O., RODOLÀ, E., BRONSTEIN, A., et al. "Partial Single-and Multishape Dense Correspondence Using Functional Maps". *Handbook of Numerical Analysis* (2018) 2.

[LRBB17] LITANY, O., RODOLÀ, E., BRONSTEIN, A., and BRONSTEIN, M. "Fully spectral partial shape matching". (2017) 2.

[LRR*17] LITANY, O., REMEZ, T., RODOLA, E., et al. "Deep functional maps: Structured prediction for dense shape correspondence". (2017) 4.

[LSL*21] LI, Y., SI, S., LI, G., et al. "Learnable fourier features for multi-dimensional spatial positional encoding". *NeurIPS* (2021) 2.

[LSP08] LI, H., SUMNER, R., and PAULY, M. "Global correspondence optimization for non-rigid registration of depth scans". (2008) 2.

[MESK22] MÜLLER, T., EVANS, A., SCHIED, C., and KELLER, A. "Instant neural graphics primitives with a multiresolution hash encoding". *Transactions on Graphics* (2022) 2.

[MON*19] MESCHEDER, L., OECHSLE, M., NIEMEYER, M., et al. "Occupancy networks: Learning 3d reconstruction in function space". (2019) 2.

[MRR*19] MELZI, S., REN, J., RODOLA, E., et al. "Zoomout: Spectral upsampling for efficient shape correspondence". *ACM Trans. Graph.* (2019) 2.

[MST*21] MILDENHALL, B., SRINIVASAN, P., TANCIK, M., et al. "Nerf: Representing scenes as neural radiance fields for view synthesis". *Communications of the ACM* (2021) 2.

[OBS*12] OVSJANIKOV, M., BEN-CHEN, M., SOLOMON, J., et al. "Functional maps: a flexible representation of maps between shapes". *ACM ToG* (2012) 2, 4.

[PFS*19] PARK, J., FLORENCE, P., STRAUB, J., et al. "Deepsdf: Learning continuous signed distance functions for shape representation". (2019) 2.

[PRM*21] PAI, G., REN, J., MELZI, S., et al. "Fast sinkhorn filters: Using matrix scaling for non-rigid shape correspondence with functional maps". (2021) 4.

[RCB*17] RODOLÀ, E., COSMO, L., BRONSTEIN, M., et al. "Partial functional correspondence". *Comp. graphics forum* (2017) 2, 6, 7.

[RR07] RAHIMI, A. and RECHT, B. "Random features for large-scale kernel machines". *NeurIPS* (2007) 2.

[RSO19] ROUFOSSE, J., SHARMA, A., and OVSJANIKOV, M. "Unsupervised deep learning for structured shape matching". (2019) 4.

[RTO*19] RAMPINI, A., TALLINI, I., OVSJANIKOV, M., et al. "Correspondence-free region localization for partial shape similarity via hamiltonian spectrum alignment". (2019) 2, 6, 7.

[SACO22] SHARP, N., ATTAIKI, S., CRANE, K., and OVSJANIKOV, M. "Diffusionnet: Discretization agnostic learning on surfaces". *ACM ToG* (2022) 3.

[Sta07] STANLEY, K. "Compositional pattern producing networks: A novel abstraction of development". *Genetic programming and evolvable machines* (2007) 2.

[STD14] SALTI, S., TOMBARI, F., and DI STEFANO, L. "SHOT: Unique signatures of histograms for surface and texture description". *CVIU* (2014) 6, 7.

[Tau95] TAUBIN, G. "A signal processing approach to fair surface design". (1995) 3, 5.

[TCF09] TOLDO, R., CASTELLANI, U., and FUSIELLO, A. "A bag of words approach for 3d object categorization". (2009) 6.

[TSM*20] TANCIK, M., SRINIVASAN, P., MILDENHALL, B., et al. "Fourier features let networks learn high frequency functions in low dimensional domains". *NeurIPS* (2020) 2, 3, 9.

[VRBC17] VESTNER M.and Litman, R., RODOLA, E., BRONSTEIN, A., and CREMERS, D. "Product manifold filter: Non-rigid shape correspondence via kernel density estimation in the product space". (2017) 4.

[VSP*17] VASWANI, A., SHAZEER, N., PARMAR, N., et al. "Attention is all you need". *NeurIPS* (2017) 2.

[WPG10] WILLMORE, B., PRENGER, R., and GALLANT, J. "Neural representation of natural images in visual area V2". *Journal of Neuroscience* (2010) 2.

[XRK*19] XU, D., RUAN, C., KORPEOGLU, E., et al. "Self-attention with functional time representation learning". *NeurIPS* (2019) 2.

[YYTK21] YU, A., YE, V., TANCIK, M., and KANAZAWA, A. "pixelnerf: Neural radiance fields from one or few images". (2021) 2.