



Multi-part shape matching by simultaneous partial functional correspondence

Yan Wu^{1,2} · Jun Yang¹

Accepted: 22 October 2021 / Published online: 8 January 2022

© The Author(s), under exclusive licence to Springer-Verlag GmbH Germany, part of Springer Nature 2021

Abstract

Non-rigid multi-part shape matching has proven to be essential and challenging in many applications. This paper analyzes the aforementioned problem and proposes a novel multi-part shape matching method to simultaneously compute correspondences between a full shape and its multiple parts undergoing a non-rigid deformation. The main idea is to simultaneously integrate the Hamiltonian eigenvalue equivalence strategy as a part regularization term being fully spectral with the partial functional map. Moreover, we introduce a new upsampling refinement approach based upon ZoomOut in conjunction with the regularized point-wise map recovery algorithm to obtain high-quality partial matches. Our method naturally handles various challenges and noise that commonly occur in real scans, like non-rigid deformations, strong partiality, topological noise, and symmetric ambiguity. Finally, we demonstrate superior qualitative and quantitative results on several datasets. We show that our method produces more accurate, smoother results than other competing methods in realistic scenarios.

Keywords Multi-part shape matching · Partial functional maps · Fully spectral · Hamiltonian eigenvalue equivalence · Upsampling refinement

1 Introduction

Shape correspondence between 3D shapes is a fundamental problem in computer graphics and computer vision. Different flavors of this problem arise in a variety of applications ranging from animation [1], texture transfer [2], statistical shape analysis [3, 4], and shape retrieval [5], to robotic vision [6] and computational archaeology [7, 8].

Finding dense matches among multi-part shapes is a research topic that has received much attention in recent years [9, 10]. Particularly, challenging settings of this problem include: non-rigid correspondence, where the near-isometric shape undergoes certain degrees of elastic deformations [11]; partial matching, where a subset of the shape has to be matched to another full shape; and multiple part shape correspondence, where a reference shape is given and multiple

parts partially matching the reference. Each of these parts may contain overlapping parts and have additional clutter.

In general, such multi-part shape matching problems, which can be described as solving multiple pairwise partial matching problems simultaneously, are more computationally difficult than pairwise matching issues. Practical scenarios combine the above settings and clutter and topological noise, which are notoriously hard. Most existing methods for computing correspondence within shape collections can only be suitable for full shapes and do not consider these practical scenarios.

While working in the same general setting, we proposed a novel method for computing a detailed, high-quality correspondence between non-rigid multiple parts and full shape in the collection. Our work is motivated by the fact that this formulation can be used as a first step toward automatic reconstruction of the deformable shapes, in which one tries to match multiple scans to a near-isometric general model. Examples of applications include computational archaeology (assembly of fractured objects). Additional uses are assisting orthopedic surgeons in putting fragmented bones back together using a healthy bone 3D model. We have conducted extensive evaluations on various benchmark datasets. Experimental results show that the proposed

✉ Jun Yang
yangj@mail.lzjtu.cn

¹ School of Electronic and Information Engineering, Lanzhou Jiaotong University, Lanzhou, China

² School of Electronic and Information Engineering, Fujian Polytechnic Normal University, Fuqing, China

algorithms significantly outperform existing state-of-the-art partial matching techniques of shape collections.

The main contributions of our work can be summarized as follows:

- (1) We propose a new part regularization term with a purely spectral nature to promote the localizing accuracy in the simultaneous partial functional maps framework, based on the fully spectral partial matching (FSPM) [12] technology.
- (2) We introduce a new refinement scheme for map recovery that combines the regularized point-wise map recovery (RPMR) [13] algorithm and the iterative upsampling technique dubbed ZoomOut [14] while leading to a significant improvement in accuracy in several challenging cases.

2 Related work

In its full generality, shape matching is an extremely well-studied area of computer graphics and computer vision [15]. Its complete overview is beyond the scope of our paper. Therefore, below we review the methods most closely related to ours, concentrating on techniques to produce correspondences within non-rigid partial shapes and shape collections. We refer the interested readers to recent surveys, including [1, 16] for more in-depth discussions.

2.1 Functional correspondence

Our method fits within the functional map representation, initially introduced in [17], modeling correspondences as linear operators between spaces of functions on manifolds. This framework's key innovation can express maps as small matrices encoded on a reduced basis, significantly simplifying the associated optimization problems. However, because the basis consists of the first few eigenfunctions of the Laplace–Beltrami operator, functional maps can only transfer sufficiently smooth functions, leaving out high-frequency details. It severely limits the applicability of the functional map.

Several authors have proposed extensions of the functional map framework. Pokrass et al. [18] improved the framework by employing sparsity-based priors and extended it to solve the unknown permutation of the unordered inputs. Gasparetto et al. [19] defined a novel approach to compute a continuous bijective map between two surfaces moving from the low rank spectral representation to a sparse spatial representation. In [20], the authors presented a spectral formulation for the generalized multi-dimensional scaling method. Further, Yang et al. [21] calculated shape correspondences using functional maps by calibrating the basis matrix

between 3D geometric shapes. Nevertheless, the above methods' common deficiency is that they always fail to compute correspondence under moderate non-isometries.

Recently, some approaches have extended this framework to computing non-isometric matching. Eynard et al. [22] exploited functional maps in different directions to make their method more robust to non-isometric deformations. This approach enforces orthonormality as a hard constraint, which is unlikely to hold on a reduced basis and leads to a non-convex problem. Ezuz et al. [23] formulate a non-isometric correspondence optimization problem, aiming to minimize the energy that tries to preserve a given landmark correspondence or functional correspondence while penalizing the Dirichlet energies of the forward and backward maps. In follow-up work [24], the authors used similar reversibility energy to obtain more accurate non-isometric maps. This energy is composed of nonlinear membrane energy that favors isometry and bending energy that is rotation invariant and promotes feature alignment. However, this method needs to deal with the bijective matching problem to get the membrane and elastic energies, leading to more massive computation than [23].

As another possible remedy, several recent works considered alternatives to the Laplace–Beltrami basis. Melzi et al. [25] introduced a new framework for local spectral shape analysis. The key idea is to construct localized bases by spectral decomposition of a modified Laplacian operator, crafted primarily to provide eigenfunctions with local support. This approach's limitation is that the region information may not be available in specific unsupervised applications to carry out the localized spectral analysis. Closely related to [25] is the recent approach of Wang et al. [26]. Using the Steklov operator as an extrinsic alternative to the Laplace–Beltrami operator, they introduced a practical and mathematically justified spectral system to extrinsic geometry for shape analysis. Nevertheless, the most prominent drawback is that the primary boundary element method may fail when triangle meshes self-overlap.

Indeed, despite these methods' success, recovering a point-to-point map from its functional representation is still considered a common requirement in many practical applications. In [27], the authors showed that by representing descriptors as linear operators acting on functions through point-wise multiplication, it is possible to improve the quality of the recovered functional map significantly. Unlike [27], which has always restricted the functional subspaces to linear combinations of basis functions, Nogneng et al. [28] showed that they could construct a much richer space by exploiting point-wise products of basis functions. However, this method suffers from the restrictive assumptions about the inputs and only uses low-order product pairs of basis functions. Rodolà et al. [13] introduced a simple probabilistic model for map recovery and refinement. Nevertheless, it can-

not achieve high-quality transfer accuracy while dealing with partial correspondence.

2.2 Partial correspondence

Iterative closest point (ICP)-like approaches [29, 30] have tackled early rigid partial correspondence problems arising from fusion or completion of multiple 3D scans. In recent years, people attempt to extend these ideas to the non-rigid case in single nonlinear optimization [11]. Unfortunately, it requires an excellent initial input to guarantee that the final map is smooth or bijective in most cases.

Most recent works addressing the partial dense correspondence problem can cope with non-rigid shapes. Bronstein et al. [31] used generalized multi-dimensional scaling to embed one mesh in another for partial matching. In follow-up work, Bronstein et al. [32] combined metric distortion minimization proposed by [31] with optimization over regular matching parts, showing an algorithm that maximizes the regularity of corresponding parts in the given shapes. Being a superior mathematical formulation, the methods mentioned above rely upon repeatedly computing geodesic distances between arbitrary pairs of points on a mesh, and so it is expensive to calculate. Windheuser et al. [33] studied a framework for non-rigid three-dimensional shape matching, which is geometrically consistent because it uses discrete graph surfaces to induce continuous and orientation-preserving matchings. It shows that although this method allows dealing with partiality, it depends on the assumption of mesh water tightness, which is hugely complex to compute. Sahillioğlu et al. [34] proposed a voting-based formulation to match shape extremities, which are assumed to be preserved by the partiality transformation. However, conventional ways of computing geodesic distances, which the techniques rely on, become invalid on noisy surfaces with holes and gaps.

Several works tried to employ machine learning methods to compute partial correspondence. Convolutional neural networks (CNNs) on non-Euclidean domains were first considered by Masci et al. [35] who introduced the geodesic CNN model, a deep learning framework for computing dense correspondences between deformable and deformed shapes. Nevertheless, a vital drawback of this method lies in its emphasis on learning a descriptor. The learning process remains agnostic to compute the final correspondence, and costly post-processing steps are often necessary to obtain accurate solutions from the learned descriptors. Wei et al. [36] employed a deep CNN architecture that finds exact and dense correspondences between clothed human body shapes with partial input data. Although successful in the presence of sufficient training data, this approach can lead to artifacts, such as outliers, requiring post-processing to achieve high-quality maps.

Most of the partial correspondence methods, as mentioned above, are point-wise. Follow-up work [37] showed that functional maps could also deal with specific settings with missing parts. Rodolà et al. [38] extended the functional correspondence framework to deal with partial correspondence, dubbed partial functional maps (PFMs). Nevertheless, the primary deficiency of PFM is its explicit model of the part, requiring a somewhat cumbersome solver alternating between optimization in the spatial domain and the spectral domain. To address this challenge, Litany et al. [12] proposed an efficient method called fully spectral partial matching (FSPM), computing a dense intrinsic correspondence between non-rigid partial and full shapes in the entire spectral domain. This approach suffers from the requirement of a proper initialization in the spectral domain or will lead to an inaccurate part-to-full correspondence. More recently, Wu et al. [39] introduced a novel approach using fully spectral eigenvalue alignment (FSEA) and upsampling refinement (UR), dubbed FSEAU, for partial computing correspondence, demonstrating noticeable improvement upon the classical partial matching method FSPM.

2.3 Multiple shape correspondence

While most rigid shape correspondence techniques concentrate on the pairwise setting, several methods have been proposed to find correspondences in the context of shape collections. The authors of [8] presented an algorithm for assembling fractured surfaces. They can generate a collection of cycle-consistent maps from maps associated with a spanning tree in the model graph. However, this approach only optimizes the spanning trees greedily and locally, often resulting in sub-optimal solutions. Torsello et al. [40] cast the multi-view registration problem into the diffusion of rigid transformations over the graph of adjacent scans. It only transfers the registration errors between coordinate frames but does not update the correspondences through the registration process. Litany et al. [41] proposed an approach to address simultaneous matching and segmentation of multiple shapes. The authors gave a reference shape and used the approach to match multiple parts to the reference. Litany et al. [42] extended PFM [38] to the multiple partial shape correspondence, which can be considered as a non-rigid generalization of the problem of the rigid puzzles solved in [41] dubbed as non-rigid puzzles (NRP). While coming with significant advantages, the restriction of this work is still the complicated alternating optimization. At the same time, Cosmo et al. [43] proposed a method for dense matching of deformable objects in cluttered 3D scenes, which is the first attempt at solving this problem in a fully deformable setting. This method's main limitation is that it relies on local features to initialize the pipeline, and the presence of clutter and missing parts result in distorted distance values.

Sahillioğlu et al. [44] explicitly minimize the average isometric distortion over all possible pairs of collection shapes via dynamic programming efficiently without any initial set of maps. Still, it is limited to provide only a sparse correspondence. In [45], the authors constructed a consistent functional map framework that discovers shared structures within heterogeneous shape collections. Nevertheless, this work cannot deal with partial correspondence. Yang et al. [46] presented a method to obtain correspondence in the non-rigid model cluster based on the functional map framework and cycle consistency constraints. While showing impressive quality in challenging settings, the work needs to adjust the parameters manually during the calculation. Cohen et al. [47] proposed a method to automatically match two shape collections with a similar shape space structure and compute the collections’ inter-maps. However, this approach requires the parallel variations within each cluster, and the non-symmetric shape space embeddings, which is slow. Further, Huang et al. [48] formulated a novel multi-scale synchronization approach and extended the basic functional map synchronization. This work suffers at the quality of the initial maps and cannot provide theoretical guarantees on processing collections of partial shapes.

3 Background

3.1 Manifolds

We model shapes as smooth two-dimensional Riemannian manifolds S embedded into \mathbb{R}^3 . We further equip the manifold with area elements ds induced by the standard metric. Given the intrinsic gradient ∇_S , we consider the positive semidefinite Laplace–Beltrami operator Δ_S as the divergence of the gradient

$$\Delta_S = \text{div}_S(\nabla_S), \tag{1}$$

which provides us with all the tools of Fourier analysis on the manifold.

We define two scalar functions $h_1, h_2 : S \rightarrow \mathbb{R}$, and the standard inner product is $\langle h_1, h_2 \rangle_S = \int_S h_1 h_2 ds$. At each point x , the Laplace–Beltrami operator admits an orthonormal eigendecomposition on the compact manifold S

$$\Delta_S \phi_i(x) = \lambda_i \phi_i(x) \quad x \in \text{int}(S), \tag{2}$$

$$\phi_i(x) = 0 \quad x \in \partial S, \tag{3}$$

with homogeneous Dirichlet boundary conditions (3), where $\text{int}(S)$ and ∂S denote the interior and boundary of S , respectively. Here, $\{0 = \lambda_1 \leq \lambda_2 \leq \dots\}$ are eigenvalues and

ϕ_1, ϕ_2, \dots are the corresponding eigenfunctions, forming the orthonormal basis of square-integrable functions space $L^2(S) = \{h : S \rightarrow \mathbb{R} \mid \langle h, h \rangle_S < \infty\}$ on S . We can represent any function $h \in L^2(S)$ via the Fourier series expansion

$$h(x) = \sum_{i \geq 1} \langle h, \phi_i \rangle_S \phi_i(x). \tag{4}$$

3.2 Functional map

We build our work upon the functional map representation [17] and its estimation pipeline. Suppose a pair of shapes X and Y consists of n_1 and n_2 points, respectively. Let $T : X \rightarrow Y$ be a bijective mapping. The functional map’s essential view is to identify correspondences between shapes by a linear operator $T_F : L^2(X) \rightarrow L^2(Y)$, mapping functions on X to functions on Y by the composition. T_F is the functional representation of the mapping T . The basic pipeline of the functional map for computing correspondences between X and Y consists of the following general steps:

- (1) Construct a set of basis functions on both X and Y . The most common basis choice is the Laplace–Beltrami eigenfunctions, which satisfy the critical requirements of choosing a basis for functional maps: compactness and stability. Compute sets $k_1 \ll n_1, k_2 \ll n_2$ of basis functions $\{\phi_i\}_{i \geq 1}, \{\psi_j\}_{j \geq 1}$ on each shape by taking the first few eigenfunctions of the respective Laplace–Beltrami operators, and store them as columns of matrices $\Phi = (\phi_1, \dots, \phi_{k_1})$ and $\Psi = (\psi_1, \dots, \psi_{k_2})$, of sizes $n_1 \times k_1$ and $n_2 \times k_2$, respectively.
- (2) Assume given a set of q corresponding descriptor (also called scalar) functions $\{f_1, \dots, f_q\} \subseteq L^2(X)$ and $\{g_1, \dots, g_q\} \subseteq L^2(Y)$ on the shapes. The descriptor functions (e.g., heat kernel signature (HKS) [49], wave kernel signature (WKS) [50], and average mixing kernel signature (AMKS) [51]) identify each point uniquely and are approximately preserved by the mapping. Express them in the given (Laplace–Beltrami) basis and store their coefficients as columns of matrices $\mathbf{A} = (\langle f, \phi_i \rangle_X)$ and $\mathbf{B} = (\langle g, \psi_j \rangle_Y)$, of sizes $k_1 \times q$ and $k_2 \times q$, respectively.
- (3) The correspondence between X and Y can be written simply as $\mathbf{C}\mathbf{A} = \mathbf{B}$, where \mathbf{C} is the functional representation of the map. Compute the optimal functional map $k_1 \times k_1$ matrix \mathbf{C} by solving the following optimization problem:

$$\begin{aligned} \mathbf{C} &= \arg \min_{\mathbf{C}_{XY}} E_{desc}(\mathbf{C}_{XY}) + E_{reg}(\mathbf{C}_{XY}), \\ &= \arg \min_{\mathbf{C}_{XY}} \|\mathbf{C}_{XY}\mathbf{A} - \mathbf{B}\|^2 + \zeta \|\Lambda_Y \mathbf{C}_{XY} - \mathbf{C}_{XY} \Lambda_X\|^2, \end{aligned} \tag{5}$$

where Λ_X and Λ_Y are diagonal matrices of Laplace–Beltrami eigenvalues on the two shapes and ζ is a small scalar weighting parameter. Unless stated otherwise, it is common to use the Frobenius norm $\|\cdot\|$ to compute these matrices’ distance. The first term $E_{desc}(\mathbf{C}_{XY}) = \|\mathbf{C}_{XY}\mathbf{A} - \mathbf{B}\|^2$ is the general form of functional correspondence problem. The second term $E_{reg}(\mathbf{C}_{XY}) = \zeta \|\Lambda_Y \mathbf{C}_{XY} - \mathbf{C}_{XY} \Lambda_X\|^2$ regularizes the map by promoting its overall correctness, which is associated with the standard assumption that the sought map should be approximately intrinsically isometric [52–54].

- (4) Recovering point-wise map from the functional map \mathbf{C} . This post-processing step (i.e., the refinement process) can be performed via map reconstruction approaches, such as ICP or the deblurring and denoising approach [55].

3.3 Fully spectral partial matching

The functional map framework can handle uncertainty and incomplete information more gracefully but is not designed to address partial correspondences in principle and require more information (e.g., point landmarks) to obtain accurate results. FSPM is an efficient spectral domain method for calculating partial dense correspondence proposed by Litany et al. [12], build on PFM [38] and joint approximate diagonalization (JAD) of Laplacians [56].

Assume given a partial shape X and a full shape Y , nearly isometric to sub-part $X' \subseteq Y$. The manifolds X and Y are discretized as triangular meshes with n_1 and n_2 vertices, respectively. We aim at calculating partial functional map $T_F : L^2(X) \rightarrow L^2(Y)$ mapping functions supported on X to functions in the region X' . Truncating the Fourier series at the first k coefficients, we give the first k Laplace–Beltrami eigenvalues of the partial shape X and the full shape Y , respectively. In other words, the first k terms are a low-pass approximation of the full map. Since X and \bar{X} are disjoint parts, the $n_2 \times n_2$ Laplacian matrix Δ_Y is composed of $n_1 \times n_1$ Δ_X and $(n_2 - n_1) \times (n_2 - n_1)$ $\Delta_{\bar{X}}$. The eigenvalues of Δ_Y form a mixed sequence composed of the eigenvalues from Δ_X and $\Delta_{\bar{X}}$, so only the first $r < k$ eigenvalues of Δ_X will exist in the first k eigenvalues of Δ_Y (see Fig. 1).

Note that different from the full-to-full setting, the eigenfunctions $\{\phi_i\}_{i \geq 1}$ of partial shape X correspond to the “full” eigenfunctions $\{\psi_j\}_{j \geq 1}$ of Y with $j \geq i$. Under this result, it can derive that the matrix $\mathbf{C} = (\langle T_F(\phi_i), \psi_j \rangle)$ of the partial functional map has a slanted–diagonal structure, whose diagonal slope is obtained using the following approximate form:

$$\theta = \frac{r}{k} \tag{6}$$

Let $\mathbf{A} = (\langle f, \phi_i \rangle_X)$ and $\mathbf{B} = (\langle g, \psi_j \rangle_Y)$ be the $k \times q$ matrices of the respective Fourier coefficients. The partial functional correspondence has the form

$$\mathbf{C}\mathbf{A} = \mathbf{B}(v), \tag{7}$$

where the soft indicator function v is defined as

$$v(x) = \begin{cases} 1 & x \in X' \\ 0 & x \notin X' \end{cases} \tag{8}$$

$\mathbf{B}(v) = (\langle v \cdot g, \psi_j \rangle_Y)$ is the $r \times q$ matrix of Fourier coefficients weighted by v , which acts as a mask restricted to the area X' . From Eq. (7), we obtain that it gives rise to a $k \times k$ matrix \mathbf{C} encoding the partial functional correspondence.

In light of the previous analysis, it can conclude that the computation of \mathbf{C} by the functional map framework relies on a pair of coefficient matrices on shapes. On the contrary, FSPM tacitly assumes that \mathbf{C} is an identity matrix and then calculates the transformation coefficients matrix, as we elucidate in the following.

Equation (7) can be formulated as

$$\mathbf{C}^T \mathbf{C}\mathbf{A} = \mathbf{C}^T \mathbf{B}(v). \tag{9}$$

Since the matrix \mathbf{C} would manifest a slanted diagonal structure with most off-diagonal elements close to zero, $\mathbf{W} = \mathbf{C}^T \mathbf{C} = \begin{bmatrix} \mathbf{I}_r & 0 \\ 0 & 0 \end{bmatrix}_{k \times k}$, where \mathbf{I}_r is an identity matrix of size $r \times r$. To make $\mathbf{W}\mathbf{A}$ containing the first r rows of \mathbf{A} , \mathbf{W} is cut to a $r \times k$ matrix $\mathbf{W}_{r \times k} = [\mathbf{I}_r \ 0]_{r \times k}$. The key novelty of FSPM is to search for a new basis locating at the potential part of the full shape instead of computing the partial functional map between shapes. To this end, they replace the matrix \mathbf{C} by \mathbf{Q} and rewrite Eq. (9) as

$$\mathbf{W}_{r \times k} \mathbf{A} = \mathbf{Q}^T \mathbf{B}(v), \tag{10}$$

where the $r \times k$ matrix \mathbf{Q} is not considered as the functional representation of the map, but a transformation coefficients




Fig. 1 The eigenvalues and eigenfunctions of Δ_Y consist of those of the blocks Δ_X and $\Delta_{\bar{X}}$

matrix for the basis [12]. Further, the new basis is constructed as $\widehat{\psi}_i = \sum_{j=1}^k q_{ji} \psi_j$, where q_{ji} is the elements of the matrix \mathbf{Q}^T . In a nutshell, $\mathbf{U}(v) = \mathbf{Q}^T \mathbf{B}(v)$ is the transformation coefficients of $\{g_i\}_{i=1}^q$ restricted to the area indicated by v in the new bases $\{\widehat{\psi}_j\}_{j \geq 1}$ and leads to the following equation:

$$\mathbf{U}(v) = \langle v \cdot g, \widehat{\psi}_j \rangle_Y. \tag{11}$$

Without loss of generality, the new basis functions $\{\widehat{\psi}_j\}_{j \geq 1}$ must be localized; thus, $\widehat{\psi}_j = v \cdot \widehat{\psi}_j$ for all j . Thus, we get to

$$\mathbf{U}(v) = \mathbf{U}, \tag{12}$$

and absorb the indicator function v into the new bases. The equality in (10) becomes

$$\mathbf{W}_{r \times k} \mathbf{A} = \mathbf{U} = \mathbf{Q}^T \mathbf{B}. \tag{13}$$

For the modified basis $\{\widehat{\psi}_j\}_{j \geq 1}$ to behave more like eigenfunctions for Δ_Y , it has to satisfy a constraint that the Dirichlet energy

$$\sum_i^k \langle \Delta_Y \widehat{\psi}_i, \widehat{\psi}_i \rangle_Y = \text{trace}(\mathbf{Q}^T \Delta_Y \mathbf{Q}) \tag{14}$$

should converge to a minimum, where $\text{trace}()$ is the trace of a square matrix denoting the sum of the matrix’s eigenvalues and $\Delta_Y = \text{diag}(\xi_1, \dots, \xi_i)$ is a diagonal matrix containing the first k eigenvalues of Δ_Y . Ideally, this behavior is explained by the fact that $\langle \Delta_Y \widehat{\psi}_i, \widehat{\psi}_j \rangle_Y = \xi_i \delta_{ij}$, where ξ_i are the eigenvalues of the full shape Y .

Replacing the trace term by an off-diagonal penalty $\text{off}(\mathbf{E}) = \sum_{i \neq j} |e_{ij}|^2$, where e_{ij} denotes the elements of matrix \mathbf{E} , we consider to solve the following optimization problem:

$$\arg \min_{\mathbf{Q} \in S(r,k)} \alpha \|\mathbf{W}_{r \times k} \mathbf{A} - \mathbf{Q}^T \mathbf{B}\|_{2,1} + \text{off}(\mathbf{Q}^T \Delta_Y \mathbf{Q}), \tag{15}$$

where optimization is done over the Stiefel manifolds $S(r, k) = \{\mathbf{P} \in \mathbb{R}^{r \times k} : \mathbf{P}^T \mathbf{P} = \mathbf{I}_k\}$ and α denotes a weight. The $L_{2,1}$ norm $\|\mathbf{E}\|_{2,1} = \sum_i \|e_i\|_2$, where e_i is the i th column of \mathbf{E} , affects the optimal value insensitive to outlier functional correspondences and promotes column-wise sparsity.

It is important to remark that, while the transformation coefficients matrix \mathbf{Q} is the solution to Eq. (15)’s optimization problem, this is not the only aim as it directly deduces matrix \mathbf{B} to obtain functional map \mathbf{C} . This way, it would be natural to calculate the partial correspondence entirely in the spectral domain.

3.4 Fully spectral eigenvalue alignment

The main drawback of FSPM is the lack of localized behavior in the spatial domain, leading to an inaccurate part-to-full correspondence. One of the main contributions of our work is a significant improvement allowing us to introduce an effective and straightforward part regularization term into the optimization problem, endowing our solutions with guarantees of high localizing accuracy. Our technique bears a resemblance to Hamiltonian spectrum alignment [57], dubbed fully spectral eigenvalue alignment (FSEA) [39]. Still, it has a significant difference that Rampini et al. [57] designed the spectral alignment algorithm to detect similar regions among deformable shapes, not to solve for partial correspondence. We aim to use a new part regularization to be fully spectral, viewed as an extension of the Hamiltonian spectrum alignment, which will promote localizing accuracy in the partial correspondence method.

3.4.1 Hamiltonian operator

The basis defined by the classical Hamiltonian operator from quantum mechanics is a new choice for shape analysis in many senses [58]. Furthermore, we can use the basis defined by the Hamiltonian to solve partial correspondences between non-rigid shapes. The scenario we consider concerns matching an approximately isometric deformation of part X to the full shape Y . One can obtain a Hamiltonian operator H_Y on the manifold Y by adding a potential function to the Laplacian. Then, the Hamiltonian operator H_Y has the elliptic operators form

$$H_Y = \Delta_Y + \beta, \tag{16}$$

with the Laplace–Beltrami operator Δ_Y and the scalar potential function $\beta : Y \rightarrow \mathbb{R}$. That is, β is also a Heaviside step function that places zero energy density on a subset X' of Y and $\tau > 0$ on the complement $\overline{X'}$. Formally,

$$\beta(x) = \begin{cases} 0 & x \in X' \\ \tau & x \in \overline{X'} \end{cases}. \tag{17}$$

Since Δ_Y and β are self-adjoint, the resulting Hamiltonian H_Y as a sum of them is also self-adjoint. We can almost directly derive its spectral theory from the regular Laplacian. Thus, the Hamiltonian H_Y admits a spectral decomposition:

$$H_Y \varepsilon_i(x) = \mu_i \varepsilon_i(x), \tag{18}$$

where the eigenvalues μ_i form a discrete spectrum $\{\mu_1, \mu_2, \dots\}$ and the eigenfunctions $\varepsilon_1, \varepsilon_2, \dots$ form an orthonormal basis.

From Eqs. (17) and (18), we note that the Hamiltonian eigenfunctions ε_i denote the particle’s energy, and the finite




Fig. 2 Discretization of the Hamiltonian operator on a triangular mesh

step potential function β divides the region Y into two constant potential spaces. In the null potential energy region where $\beta(x) = 0$, the particle is free to pass, and the approximate equality $H_Y \approx \Delta_Y$ holds for all $x \in X'$. On the other hand, for the high potential region where $\beta(x) = \tau \forall x \in \bar{X}'$, the particle cannot penetrate and its eigenenergy ε_i decays exponentially and vanishes when the corresponding eigenvalues $\mu_i < \tau$. As analyzed above, the scalar potential function $\beta(x)$ ensures that the eigenfunctions corresponding to eigenvalues $\mu_i < \tau$ are restricted to the null potential energy region, akin to the soft indicator function mentioned above.

3.4.2 Generalized eigenproblem

To compute the Hamiltonian eigenvalues and eigenvectors, we should solve a generalized eigenvalue problem. At the moment, one can do this job by analyzing the discrete operator for reconstructing triangle meshes [59, 60].

In the discrete setting, we approximate the shape Y by manifold triangle mesh sampled at vertices $\{p_i\}_{i=1}^{n_2}$. The discrete Hamiltonian operator of the shape Y assumes the form of a matrix

$$H_Y = \mathbf{S}_Y^{-1} \mathbf{M}_Y + \text{diag}(\mathbf{V}) \tag{19}$$

with an n_2 -dimensional vector \mathbf{V} containing the values of the potential β at each vertex, where \mathbf{S}_Y is a diagonal matrix of local area elements s_i and \mathbf{M}_Y is a stiffness matrix of cotangent weights, defined in terms of the discrete metric as

$$m_{ij} = \begin{cases} -(\cot \theta_{ij} + \cot \eta_{ij})/2 & \text{if } i \text{ and } j \text{ are adjacent} \\ \sum_{h \neq i} (\cot \theta_{ih} + \cot \eta_{ih})/2 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases} \tag{20}$$

Figure 2 depicts our notation. s_i is the area of the shaded region in the same figure.

In the case of the Laplacian, the generalized eigenproblem thus takes the form:

$$\mathbf{S}_Y^{-1} \mathbf{M}_Y \Phi = \Phi \text{diag}(\lambda), \tag{21}$$

where $\Phi = (\phi_1, \dots, \phi_k)$ is a matrix containing the first k Laplacian eigenfunctions as columns and $\lambda = \{\lambda_1, \dots, \lambda_k\}$ is a k -dimensional vector of the corresponding Laplacian eigenvalues.

In the case of the Hamiltonian, the generalized eigenproblem

$$(\mathbf{S}_Y^{-1} \mathbf{M}_Y + \text{diag}(\mathbf{V}))E = E \text{diag}(\mu) \tag{22}$$

is solved for computing the Hamiltonian eigenfunctions and eigenvalues. $E = (\varepsilon_1, \dots, \varepsilon_k)$ and $\mu = \{\mu_1, \dots, \mu_k\}$ are the first k eigenpairs. We readily obtain the Hamiltonian eigenvalues of Eq. (22) as $\mu(\mathbf{S}_Y^{-1} \mathbf{M}_Y + \text{diag}(\mathbf{V}))$.

3.4.3 Eigenvalue alignment

Let $\{(\mu_i, \varepsilon_i)\}_{i=1}^k$ and $\{(\lambda_i, \phi_i)\}_{i=1}^k$ be the first k Hamiltonian eigenpairs of the full shape Y and the first k Laplacian eigenpairs of the partial shape X with the homogeneous Dirichlet boundary conditions, respectively. We consider that the partial shape X corresponds to the region $X' \subseteq Y$ on the full shape Y up to isometry.

From (16) and (18), we can expand the eigendecomposition of H_Y as

$$\begin{aligned} H_Y \varepsilon_i(x) &= \Delta_Y \varepsilon_i(x) + \beta(x) \varepsilon_i(x) \\ &= \mu_i \varepsilon_i(x). \end{aligned} \tag{23}$$

In light of our previous analysis, if $x \in X'$, the approximate equality $H_Y \approx \Delta_{X'}$ holds point-wise. We can write Eq. (23) as

$$\Delta_{X'} \varepsilon_i(x) = \mu_i \varepsilon_i(x). \tag{24}$$

If we make the partial shape X an excellent localization to the corresponding latent region $X' \subseteq Y$, ϕ_i of Δ_X has the same values as ε_i of $\Delta_{X'}$ at corresponding points with the Dirichlet boundary (i.e., $\phi_i(x) = 0$ for $x \in \partial M$). Thus,

$$\Delta_{X'} \varepsilon_i(x) = \Delta_X \phi_i(x). \tag{25}$$

According to Eq. (24) and $\Delta_X \phi_i(x) = \lambda_i \phi_i(x)$, equality in (25) becomes

$$\mu_i \varepsilon_i(x) = \lambda_i \phi_i(x). \tag{26}$$

Therefore, we draw the following conclusion:

$$\mu_i = \lambda_i. \tag{27}$$

We conclude from the above analysis that under the assumption of providing an indicator for the correspondence region on a full shape, one can align the Hamiltonian eigenvalues on

Fig. 3 Results of eigenvalue alignment technology on four different classes from the SHREC' 16 Partial Correspondence benchmark. We show the partial shape and its sought localization on the full shape (in maroon). Below each result, we report the intersection over union (*IoU*) scores w.r.t. the ground-truth regions

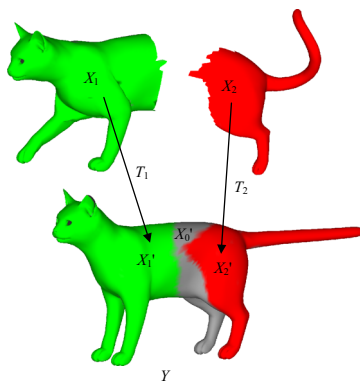
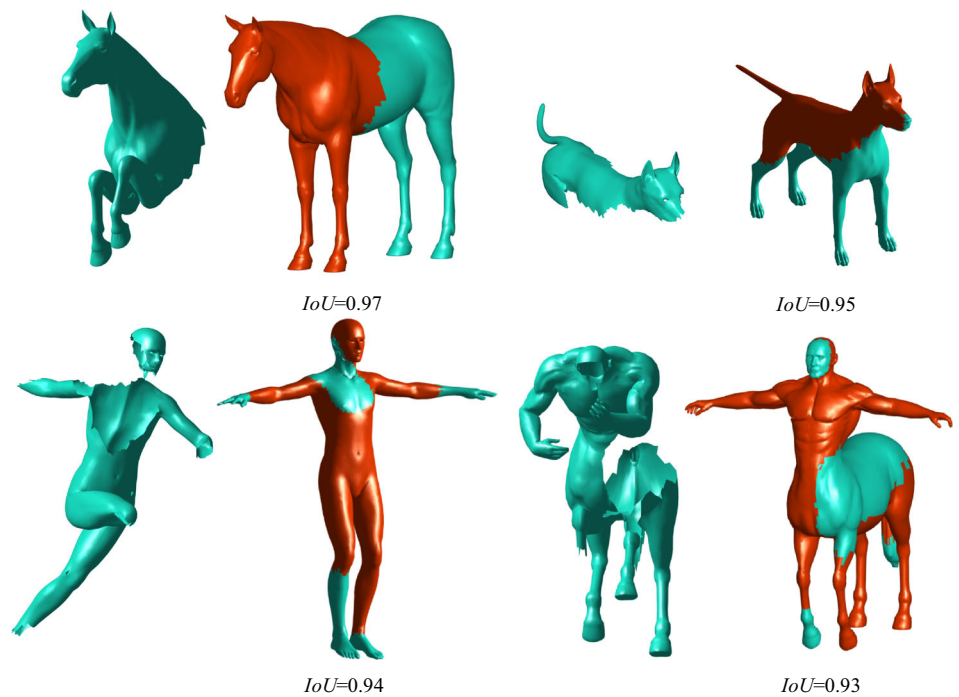


Fig. 4 Multi-part matching with missing parts. The parts $\{X_i\}$ are matched to the reference shape Y

the full shape Y with the eigenvalues of the Dirichlet Laplacian on the partial shape X [57].

3.4.4 Part regularization

We rely upon the eigenvalue alignment to implement the partial shape localization via optimization over the potential function region on the full shape. Along this line of thought, we utilize the eigenvalues of Eq. (22) and boil down the eigenvalue alignment to the following problem:

$$\min_{\mathbf{V} > 0} \left\| \boldsymbol{\mu}(\mathbf{S}_Y^{-1} \mathbf{M}_Y + \text{diag}(\mathbf{V})) - \boldsymbol{\lambda} \right\|_{\omega}^2 \quad (28)$$

where we seek the minimum between the Hamiltonian eigenvalues on the full shape and the Laplacian eigenvalues on the

partial shape to localize the partial shape to the correct region. We use the ω -norm $\|\boldsymbol{\mu} - \boldsymbol{\lambda}\|_{\omega}^2 = \sum_{i=1}^k \frac{1}{\lambda_i^2} (\mu_i - \lambda_i)^2$ (equal to the weighted L2 norm) to make a balanced penalty for all frequencies, promoting the robustness property. Using the ω -norm in Eq. (28) could lead to an accurate shape localization: An optimum can be reached by correctly aligning eigenvalues.

Equation (28) gives us a straightforward way to model part $X' \subseteq Y$ using the potential vector \mathbf{V} , which confines the sought region X' on Y corresponding to the approximately isometric partial shape X . That is, the eigenfunctions with associated eigenvalue $\mu_i < \tau$ are isolated in the area where $\mathbf{V} = 0$. Our goal is to introduce the eigenvalue alignment as a partial shape localization technology to regularize the partial matching in the optimization problem (15). To make Eq. (28) act as a part regularization term R_{part} , we express this equation as

$$R_{\text{part}}(\sigma(z)) = \min_z \left\| \boldsymbol{\mu}(\mathbf{S}_Y^{-1} \mathbf{M}_Y + \text{diag}(\sigma(z))) - \boldsymbol{\lambda} \right\|_{\omega}^2, \quad (29)$$

where $\sigma(z) = \frac{\tau}{2}(\tanh(z) + 1)$ is a saturation function used to restrict the value of $\mathbf{V} = \sigma(z)$ to the range $(0, \tau)$, s.t. $\mu_i < \tau$.

In Fig. 3, we illustrate the localization behavior of the new part regularization term under different inputs. Here, we use the intersection over union (*IoU*) between the located region and the ground truth under each full shape to reflect the localization accuracy. As we show in this experiment, it is sufficient to use fully spectral eigenvalue alignment technology to get a sound localization to the corresponding latent region.

4 Multi-part shape matching

4.1 Optimization problem

As stated before, throughout the paper, we consider the setting where we are given a full model shape Y and a collection of non-overlapping partial shapes $\{X_i\}_{i=1}^m$ that correspond to the non-rigid deformed parts $\{X'_i \subseteq Y\}_{i=1}^m$. Our goal is to match the parts to the full shape utilizing non-rigid mappings $\{T_i : X_i \rightarrow X'_i\}_{i=1}^m$ such that the matching regions $\{X'_i \subseteq Y\}_{i=1}^m$ are non-overlapping and might not cover Y entirely. As we demonstrate in Fig. 4, X'_0 can be seen as a “null segment,” i.e., a missing part.

We propose a practical method of the so-called simultaneous partial functional correspondence (SPFC) based on the previous analysis. Combining the data term (15) and the part regularization term (29), we formulate the above problem of simultaneous multi-part matching as

$$\begin{aligned} & \arg \min_{\mathbf{Q}_i, z_i} \sum_{i=1}^m \alpha \left\| \mathbf{W}_{r \times k} \mathbf{A}_i - \mathbf{Q}_i^T \mathbf{B}_i(\eta(z_i)) \right\|_{2,1} + \sum_{i=1}^m \text{off}(\mathbf{Q}_i^T \Lambda_Y \mathbf{Q}_i) \\ & + \gamma_1 \sum_{i=1}^m R_{\text{part}}(\sigma(z_i)) + \gamma_2 \sum_{i=0}^m R_{\text{Mumford-Shah}}(\eta(z_i)) \quad (30) \\ & \text{s.t.} \begin{cases} \sum_{i=0}^m \eta(z_i) = 1 \\ \text{area}(X_i) = \mathbf{s}_Y^T \eta(z_i) \end{cases}, \end{aligned}$$

where $\eta(z) = \frac{1}{2}(\tanh(z) + 1)$ saturates the part indicator function with values in the range (0, 1) and $\sigma(z) = \tau \eta(z)$,

γ_1 and γ_2 are weights, and \mathbf{s}_Y denotes the vectors of discrete area elements on Y . $R_{\text{Mumford-Shah}}(\eta(z)) = \int_Y \delta(\eta(z) - 0.5) \|\nabla_Y \eta(z)\| dx$ is an intrinsic version of the Mumford–Shah functional [61], measuring the length of the boundary of a part by a saturation function. We find the located region on the model closest to the partial shape and with the shortest boundary. Since the null segment X'_0 does not correspond to any of the X_i 's, it has no data term and does have a Mumford–Shah functional term.

Note that the first aggregate $\sum_{i=1}^m \alpha \left\| \mathbf{W}_{r \times k} \mathbf{A}_i - \mathbf{Q}_i^T \mathbf{B}_i(\eta(z_i)) \right\|_{2,1}$ constitutes the data term measuring the proximity of the parts $\mathbf{W}_{r \times k} \mathbf{A}_i$ to the corresponding transformed segments $\mathbf{Q}_i^T \mathbf{B}_i(\eta(z_i))$ on the reference shape, while the second aggregate $\sum_{i=1}^m \text{off}(\mathbf{Q}_i^T \Lambda_N \mathbf{Q}_i)$ and the third aggregate $\sum_{i=0}^m R_{\text{part}}(\sigma(z_i))$ are the regularization terms measuring the sought localization of each part. We use this efficient optimization scheme (30) to perform the matching of multiple shapes simultaneously.

4.1.1 Alternating scheme

The problem (30) is solved by employing the manifold alternating direction method of multipliers (MADMM) algorithm [62] in Algorithm 1. Since the MADMM algorithm does not support constraints inherently, the constraints in the problem (30) were replaced by large quadratic penalties.

Input: partial shapes $\{X_i\}_{i=1}^m$, full shape Y

Output: indicators $\{z_i\}$, matrices $\{\mathbf{Q}_i\}$

Initialize $\eta(z_i) = \mathbf{1}$ (a vector of n_2 ones), $\mathbf{Q}_i = \mathbf{I}_{r \times r}$

Repeat

Q-step: Fix $\{z_i^*\}$ and compute the $\{\mathbf{Q}_i\}$ solving the problem

$$\arg \min_{\mathbf{Q}_i} \sum_{i=1}^m \alpha \left\| \mathbf{W}_{r \times k} \mathbf{A}_i - \mathbf{Q}_i^T \mathbf{B}_i(\eta(z_i^*)) \right\|_{2,1} + \sum_{i=1}^m \text{off}(\mathbf{Q}_i^T \Lambda_N \mathbf{Q}_i).$$

z-step: Fix $\{\mathbf{Q}_i^*\}$ and compute the $\{z_i\}$ solving the problem

$$\arg \min_{z_i} \sum_{i=1}^m \alpha \left\| \mathbf{W}_{r \times k} \mathbf{A}_i - \mathbf{Q}_i^{*T} \mathbf{B}_i(\eta(z_i)) \right\|_{2,1} + \gamma_1 \sum_{i=1}^m R_{\text{part}}(\sigma(z_i)) + \gamma_2 \sum_{i=0}^m R_{\text{Mumford-Shah}}(\eta(z_i))$$

$$\text{s.t.} \sum_{i=0}^m \eta(z_i) = 1, \text{ area}(X_i) = \mathbf{s}_Y^T \eta(z_i).$$

Until convergence.

Algorithm 1. Our pipeline for solving the problem of simultaneous multi-part matching

As noted above, we use the transformation coefficients matrix \mathbf{Q} to calculate the new basis $\widehat{\psi}_i$, from which it is thus natural to obtain functional map \mathbf{C} by deducing \mathbf{B} .

4.2 Post-processing

The functional map pipeline outlined above demonstrates that reconstructing point-wise correspondences is key to the functional correspondence framework. To convert the functional map \mathbf{C} to a higher-quality point-wise dense correspondence, we formulate a novel map refinement approach, which can be seen as a combination of ZoomOut [14] and RPMR [13] algorithm, called Regularized ZoomOut.

For simplicity, we assume that the bases $\{\phi_i\}_{i=1}^k$ and $\{\widehat{\psi}_i\}_{i=1}^k$ are the first k eigenfunctions of the Laplace–Beltrami operators of the partial shape X and the full shape Y under homogeneous Neumann boundary conditions. Φ and $\widehat{\Psi}$ are matrices containing $\{\phi_i\}_{i=1}^k$ and $\{\widehat{\psi}_i\}_{i=1}^k$ as their columns, of sizes $n_1 \times k$ and $n_2 \times k$, respectively.

In this section, we take account of the point-wise map $T : X \rightarrow Y$ by means of a permutation matrix

$$\mathbf{\Pi}(i, j) = \begin{cases} 1 & T(i) = j \\ 0 & T(i) \neq j \end{cases}, \quad (31)$$

where i and j are vertex indices on shape X and Y , respectively. Note that the $n_1 \times n_1$ matrix $\mathbf{\Pi}$ encodes the map T . By analogy to the simple representation of matrix $\mathbf{C} = \langle \langle T_F(\phi_i), \widehat{\psi}_j \rangle \rangle$ on orthonormal bases, the expression for \mathbf{C} built upon T can be compactly written as

$$\mathbf{C} = (\widehat{\Psi})^T \mathbf{\Pi} \Phi. \quad (32)$$

Along the same line of thought, we get the map T from the following problem:

$$T(i) = \arg \min_i \left\| \mathbf{C}(\Phi(i))^T - (\widehat{\Psi}(j))^T \right\|^2, \quad (33)$$

with $\Phi(i)$ and $\widehat{\Psi}(j)$ being the i th row of the matrix Φ and j th row of the matrix $\widehat{\Psi}$, respectively.

Traditionally, we solve the problem (33) by the approximate nearest neighbor (ANN) algorithm or the k -nearest neighbor (k -NN) algorithm. Rodolà et al. [13, 63] considered the point-wise map recovery problem as a point cloud alignment problem and introduced a regularized probabilistic model to solve the problem (33). Combining the RPMR [13] algorithm with the objective function, we formulate Eq. (33) as

$$\begin{aligned} & \arg \min_{\mathbf{\Pi} \in \{0,1\}^{n_1 \times n_1}} D_{\text{KL}}(\mathbf{C}\Phi^T, \widehat{\Psi}^T \mathbf{\Pi}) \\ & + \rho \left\| \Gamma(\mathbf{C}\Phi^T - \widehat{\Psi}^T \mathbf{\Pi}) \right\|^2 \\ & \text{s.t. } \mathbf{\Pi}^T \mathbf{1} = \mathbf{1}. \quad (\text{a vector of } n_1 \text{ ones}), \end{aligned} \quad (34)$$

where D_{KL} denotes the Kullback–Leibler divergence between a continuous GMM distribution (represented by $\mathbf{C}\Phi^T$) and a mixture of Dirac distributions (represented by $\widehat{\Psi}^T$) and $\rho > 0$ controls the regularity of the assignment. Here, $\|\Gamma\|^2$ is the Tikhonov regularizer, where Γ denotes a low-pass operator promoting smooth velocity vectors. Such regularizer is supported by the motion coherence theory (MCT) [64], which states that points close to one another tend to move coherently, and the displacement (i.e., velocity) function between the point sets should be smooth [65]. We employ the expectation maximization (EM) algorithm to solve the problem (34). This publicly available code comes from the coherent point drift (CPD) [65] method.

The key novelty of the ZoomOut method comes from its capability to iterate the refinement procedure to obtain progressively larger functional maps $\{\mathbf{C}_i\}_{i \geq 0}$ until i reaches the threshold r . This way, as the functional map grows, the point-wise map T becomes both more smooth and accurate.

Our approach is based on the ZoomOut, in which we utilize the RPMR algorithm to recover and refine the point-wise map. To summarize, it includes the following steps:

- (1) Initialize $i = k_0, j = 0$ and $\mathbf{C}_0 = (\text{diag}(1, \dots, 1))_{k_0 \times k_0}$.
- (2) Compute $\mathbf{\Pi}_j$ by alternately using the RPMR algorithm [i.e., utilizing the EM algorithm to solve the problem (34)] or k -NN algorithm. We observed that it typically works better if we used the two algorithms interchangeably in the code loop.
- (3) Compute $\mathbf{C}_j = (\widehat{\Psi}^j)^T \mathbf{\Pi}_j \Phi^i$, where Φ^i and $\widehat{\Psi}^j$ denote the submatrix of Φ and $\widehat{\Psi}$ consisting of the first i columns, respectively. Set $i = i + 1, j = j + 1$.
- (4) Repeat the loop body, until $i > r$.

Figure 5 shows the accuracy compared to the ICP refinement method from the SHREC'16 Partial Correspondence benchmark. We applied our technique and ICP to refine maps between the partial and full shapes and encoded corresponding points in a similar color. Our approach builds upon the feature “upsampling” in the Regularized ZoomOut framework that introduces additional frequencies or equivalently directly adds samples in the spectral domain at every iteration. Initializing the functional map to size 10×10 (second row), we show the point-wise maps visualized via color transfer throughout the upsampling iterations of our approach in the figure. Note that as the correspondence matrix grows, the resulting matrix is more diagonal, and the map becomes more

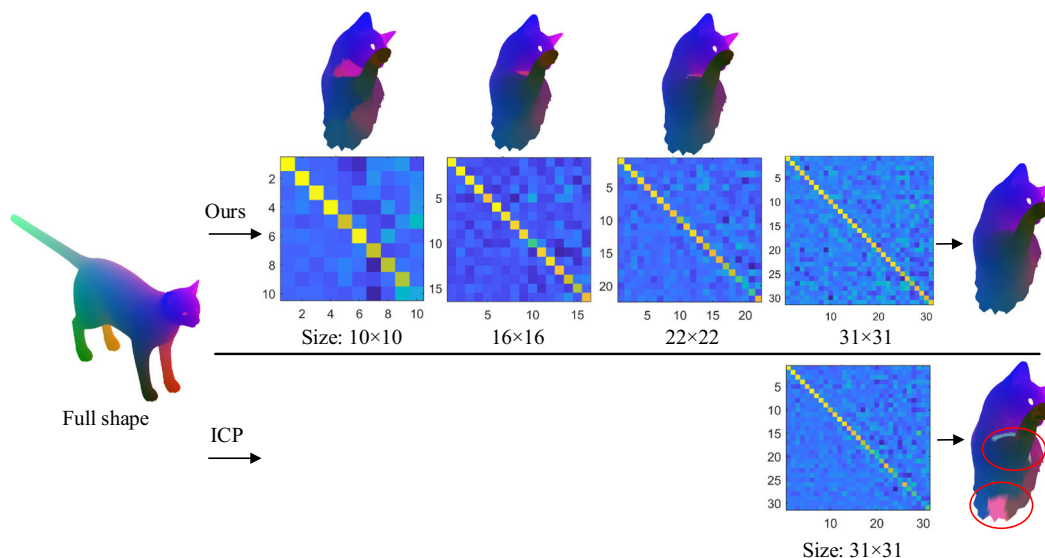


Fig. 5 Comparison of map quality of our method with ICP in 20 outer iterations. Our approach leads to better results, as can be seen, e.g., on the belly and leg

smooth and more semantically accurate. We also show the result of ICP in 20 outer iterations for a map of size 31×31 (bottom row), which nevertheless leads to a low-quality dense point-wise correspondence result. Therefore, we encourage the reader to compare the work of ICP to ours, which provides visual proof that our approach upsampling the map to the same size leads to a significant improvement.

5 Implementation

5.1 Preprocessing

The robustness of the chosen descriptor fields in the shapes directly affects correspondence quality. Compared to other classical local descriptors, such as HKS and WKS, the SHOT [66] descriptor capturing the local normal vector orientations is highly descriptive, computationally efficient, and robust to noise. In all experiments below, we used 352-dimensional SHOT descriptors with ten bins and a SHOT radius roughly chosen to 9% of the maximal pairwise geodesic distance, which means 352 descriptor functions on each shape.

5.2 Eigenvalue number

Roughly speaking, we can estimate the slope of the partial functional map \mathbf{C} as r/k , where r and k denote the eigenvalue numbers of the partial shape X and the full shape Y , respec-

tively. In our experiments, we initialize $k = 90$ and estimate the value of r , directly related to the rank of \mathbf{C} , from

$$r = \max \left\{ i \mid \lambda_i < \max_{j=1}^k \mu_j \right\}, \tag{35}$$

where we write λ_i and μ_j to denote the eigenvalues of the partial shape and the full shape, respectively.

5.3 Step potential

In Sect. 3.4, the potential β confines all eigenfunctions with associated eigenvalues $\mu_i < \tau$ within the region where $\sigma(z)=0$. If $\mu_i > \tau$, it means there is a probability for the particle to penetrate the potential region where $\sigma(z) = \tau$. In our experiments, we set the step of magnitude $\tau = 8\mu_k$, which can get excellent quality of localization.

5.4 Error measure

We use the quantitative criterion introduced in [67] to evaluate point-wise maps' quality. Given the ground-truth match $(x, y^*) \in X \times Y$, the error of the calculated match $(x, y) \in X \times Y$ is computed by the geodesic distance between y and y^* and normalized by the area of Y :

$$e(x) = \frac{d_Y(y, y^*)}{\sqrt{\text{area}(Y)}}, \tag{36}$$

where d_Y is the geodesic distance on Y and $\text{area}(Y)$ is the area of the shape Y .

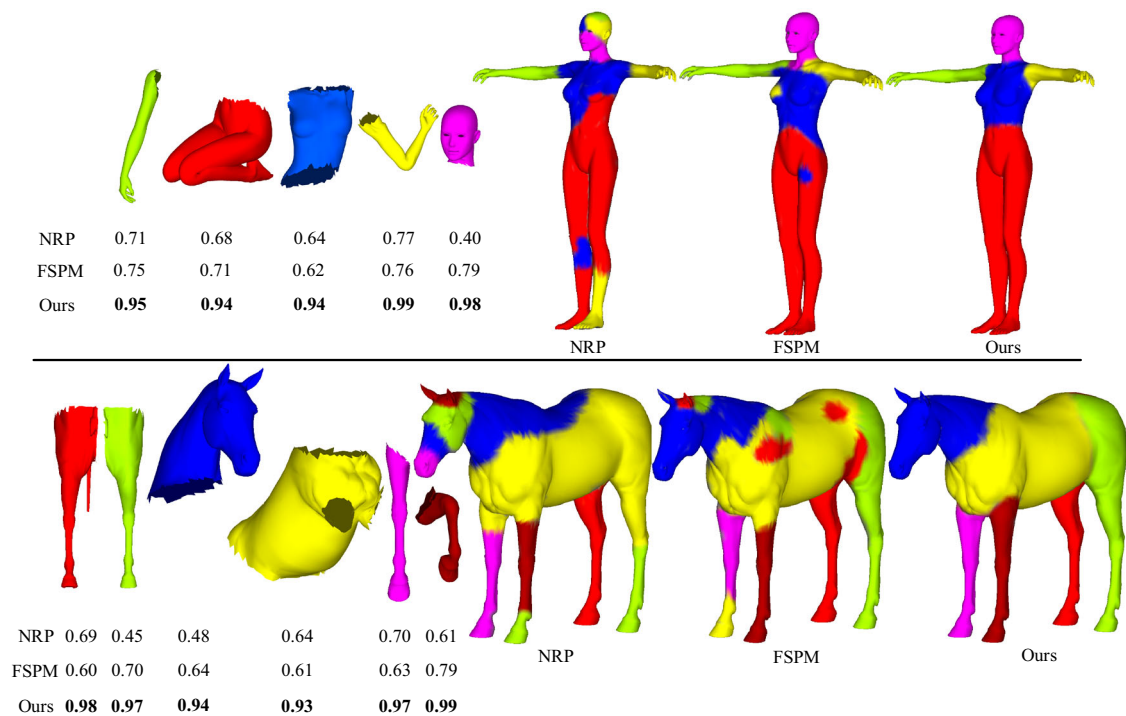


Fig. 6 Comparison between NRP, FSPM, and our method on the SHREC'16 Partial Correspondence benchmark in the multiple part setting. For each example, we show the partial shapes, the NRP solution, the FSPM solution, and our solution. Corresponding regions between

partial and full shapes are indicated with the same color. The numerical score below each partial shape denotes the *IoU* w.r.t. the ground-truth regions from different methods

6 Experimental results

We performed a wide range of real and synthetic data experiments to demonstrate our method's efficacy and carried out qualitative and quantitative comparisons to the state-of-the-art algorithms on several recent benchmarks.

6.1 Datasets

We selected four datasets that have a large variety of objects with ground-truth correspondences.

The SHREC'16 Partial Correspondence benchmark [68] uses shapes from the TOSCA [69] dataset, consisting of nearly isometrically deformed shapes from eight classes, with different removed parts. The benchmark's two types of partiality are cuts (consists of shapes undergoing a single cut) and holes (contains irregular holes and multiple cuts). In each class, the point-wise ground-truth correspondence between the partial and full shape is given.

The SHREC'16 Topology benchmark [70] includes modified shapes from the KIDS set [71], containing 25 shapes of the same class undergoing near-isometric deformations in addition to massive topological shortcuts.

The SCAPE dataset [72] contains 72 clean shapes of scanned humans in different poses.

The FAUST dataset [4] contains 100 human scans belonging to 10 different individuals, partitioned into two classes: 60 requiring intra-subject matching and 40 requiring inter-subject matching.

6.2 Qualitative results

In Fig. 6, we show examples of the correct localization of each part with respect to the full shape. The input data are several non-overlapping pieces taken from the full shape's non-rigid deformations in the SHREC'16 Partial Correspondence benchmark, forming a covering set. We compare the state-of-the-art multi-part shape matching algorithm NRP [42] and partial shape matching algorithm FSPM [12]. Since FSPM is not a partial matching method of shape collections, for fair comparisons, we perform it to each part separately, resulting in corresponding independent partial matching problems. Note that for better comparison results, we report the area ratio *IoU* between the detected region and the ground truth below each partial shape and plot the identified areas corresponding to parts with different colors in the full shapes. Results show that our method has a more superior partial localization performance and is more precise than the one obtained with other approaches. Although there is an explicit model of parts in the architecture, the NRP does not per-

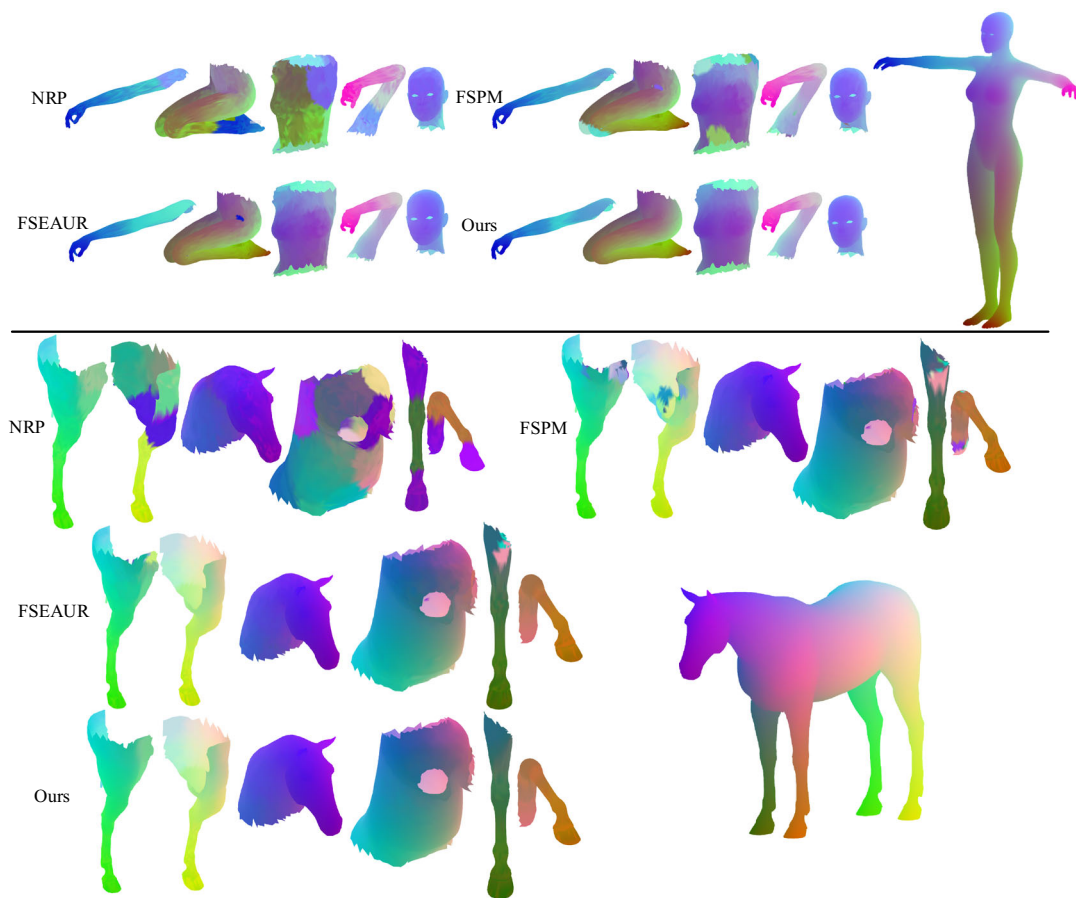


Fig. 7 An example of multi-part dense correspondence was obtained using different methods from the SHREC'16 Partial Correspondence benchmark. Correspondence is visualized via color transfer from the rightmost full shape

form well in the partial localization. Unlike the NRP, FSPM circumvents the need to compute the unknown parts explicitly but lacks partially localized behavior; thus, localization accuracy is low. As noted in our discussion, our localization feature is based on the simultaneous partial functional map framework and the eigenvalue equivalence technique, penalizing the sought localization that lies outside of the correct corresponding region, which intuitively means that no spurious localization should be created.

In Fig. 7, we display examples of the dense correspondence that NRP, FSPM, FSEAU[39], and our algorithm generate, where dense maps between shapes are color-coded such that matching points have the same color. Note that, for FSPM, since each part is matched independently, different parts' correspondence results may cover the wrong areas on the boundary leading to overlapping (e.g., the woman's belly and the horse's four legs). FSEAU has the same problem (e.g., the woman's right arm and the horse's right two legs). Our method leverages the simultaneous functional map algorithm to resolve this problem so that all parts are matched jointly to a whole shape. We exploit it as a regularizing effect on the correspondence to remove this ambiguity. In

the NRP solution, the horse's hind legs are swapped, and the woman and the horse's belly are mapped to the back. Due to symmetry ambiguity, the algorithm fails to produce a good map with consistent orientation. To disambiguate possible symmetric correspondences, we introduce the new part regularization via the eigenvalue equivalence technique for estimating functional maps by promoting the localizing accuracy in a novel way. Moreover, our framework takes Regularized ZoomOut as map refinement, making processing smaller initial functional maps easier to compute and avoiding getting trapped in symmetry ambiguity. We provide comparisons with FSEAU, which is also based on the ZoomOut. We show that our method's post-processing step Regularized ZoomOut leads to a significant boost in performance. Therefore, our method provides the smoothest maps and exhibits the smallest color distortion, close to the ground truth.

In Fig. 8, we use various datasets for the evaluation, including the SHREC'16 Topology benchmark, the FAUST, and the SCAPE datasets. Qualitatively evaluating multi-part dense correspondence on these datasets is challenging since their settings include non-rigid surface deformation, symmetric




Fig. 8 Examples of multi-part dense correspondence obtained with NRP, FSPM, and our method from various datasets. Notice that compared to other competing approaches, our method correctly identifies

non-overlapping subregions on the reference matching all the parts, providing the most accurate and smooth map

ambiguity, and the presence of topological noise. In this experiment, we fragment their model into three parts, give a reference shape, and perform these parts partially matching the reference. We visualize the solution by color coding on the reference and transfer it to the parts using the point-wise maps.

In the SHREC'16 Topology benchmark, the shapes undergo topological changes due to the coalescence of spatially close surface regions, e.g., the hands on the child's face in the figure. As shown in Fig. 8, NRP demonstrates poor performance in this challenging setting. The child's right hand is the same color as his face, and his left hand is mapped to the right hand. Moreover, FSPM is less accurate and more unstable than our method. In contrast with FSEAU, we can obtain high-quality functional maps across each shape collection.

The shapes of FAUST and SCAPE contain artifacts such as scanning noise and symmetric ambiguity. When compared to these state-of-the-art methods, an essential advantage of our

method is the map refinement process. It makes the resulting maps of our approach typically less noisy and more globally consistent, even though it is not designed to optimize this problem, which also helps obtain a more accurate partial functional map.

In Fig. 9, we visualize the match results by transferring a texture from reference to multiple parts via the recovered dense map, which provides good visualization for the resulting maps' local distortion. NRP has a massive distortion on all parts of the woman and the horse and gives a flipped left to right map in the woman's head. The map produced by FSPM has a heavily overall error when compared to the ground truth. Note that the continuity and quality of the maps obtained using our method are significantly better than other techniques. These qualitative results illustrate that our method produces smooth, guaranteed bijective solutions. We also observe that our method is the only algorithm that can reliably recover multiple high-quality non-rigid maps in the collection. The results of FSEAU recover fewer high-




Fig. 9 Comparisons of multi-part matching results are shown by texture transfer

quality maps. For example, the correspondences obtained by FSEAU in the woman's buttocks and the horse's legs show noticeable artifacts.

6.3 Quantitative results

We perform an extensive quantitative evaluation of the proposed method on the SHREC'16 Partial Correspondence benchmark, the SHREC'16 Topology benchmark, the FAUST, and the SCAPE datasets. We compare with the state-of-the-art multi-part matching method NRP and partial correspondence approaches, including PFM, FSPM, and FSEAU.

Cumulative curves are plotted in Fig. 10, reading the fraction (y-axis) of computed correspondences that fall within certain (normalized) geodesic distance to the ground-truth ones (x-axis). Note that when compared to the baseline method NRP, our approach achieves 18.9% improvement on the SHREC'16 Partial Correspondence benchmark, 8.8% on the SHREC'16 Topology benchmark, 4.1% on the FAUST,

and 35.7% on the SCAPE on average. On the SHREC'16 Partial Correspondence and SHREC'16 Topology benchmark, our method (black curve) detects over 90% correct correspondences for a small threshold of 0.05. It converges to finding almost all correct correspondences within geodesic error 0.075. Moreover, even for the FAUST and SCAPE, our solution obtains a higher percentage of correct correspondences within a small geodesic diameter and saturates earlier than other methods.

Although shapes in the SHREC'16 Topology benchmark are still synthetic, self-occlusion poses are merged in the geometry, making it more challenging. We can see that even though the accuracy curve on this dataset does not significantly improve our method over FSPM and FSEAU, our results are more accurate, smoother, and consistent. By nature of the acquisition process, the shapes on the FAUST and SCAPE are affected by several artifacts, resulting in a challenging testbed for shape matching. The low accuracy of FSEAU in SCAPE is motivated by its refinement approach. Our approach significantly outperforms other




Fig. 10 Correspondence quality of different methods on the SHREC' 16 Partial Correspondence benchmark, the SHREC' 16 Topology benchmark, the FAUST dataset, and the SCAPE dataset

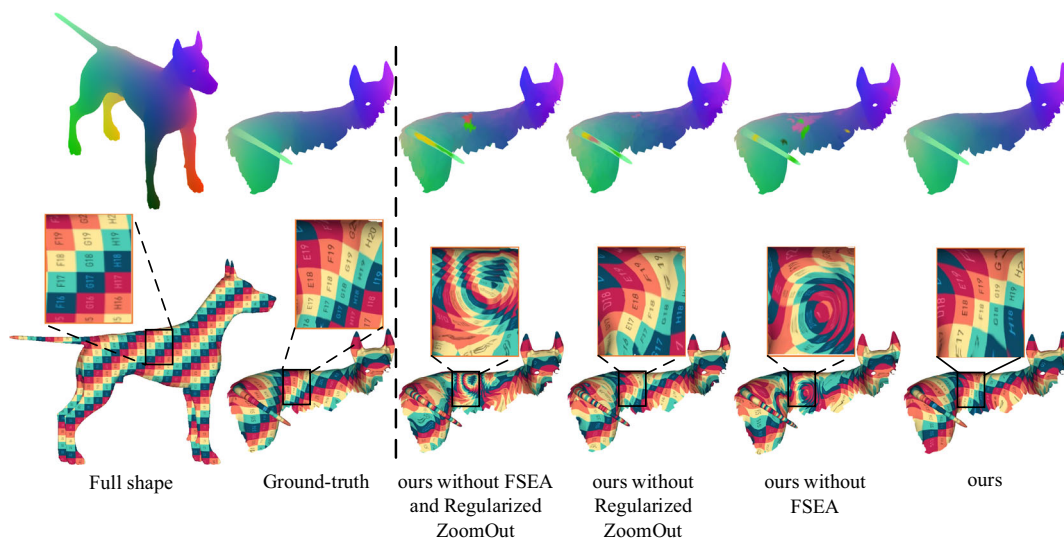


Fig. 11 Ablation study example. Note that we consider four different settings: ours without FSEA and Regularized ZoomOut, ours without Regularized ZoomOut, ours without FSEA, and ours. We then visualize the maps via color (first row) and texture transfer (second row)

methods consistently across these two datasets, whose Regularized ZoomOut refines the FSEAU solutions even further.

6.4 Ablation study

In this section, we present the extensive ablation study of all the vital components of our algorithm, including FSEA and Regularized ZoomOut. We run our algorithm in four different settings on 30 random pairs from the SHREC'16 Partial Correspondence benchmark, where one part of the method changes each time while the rest remains identical. These settings contain our method, ours without FSEA, ours without Regularized ZoomOut (use the ICP algorithm to refine map), and ours without any two components (use the ICP algorithm to refine map). We evaluate the results via visualization of color and texture transfer in Fig. 11. To get a better view, we enlarge the texture details of the corresponding regions in the shapes. The central insight is that our method's different components have an intricate interplay, and the accuracy drops significantly if any part is removed. Our FSEA algorithm can find the correct region corresponding to the partial shape in most cases, even in the presence of large amounts of partiality of the inputs. The figure's fourth and sixth columns show that Regularized ZoomOut makes the map smoother and more accurate semantically.

We assess the effect of the different components of our method shown in Fig. 12. We turn off certain parts of the method and report the correspondence quality in terms of the average geodesic error in % of the diameter obtained by different settings at increasing partiality levels. It demonstrates the importance of all individual blocks and ascertains that all these components are needed to achieve optimal performance with our solution. In particular, the FSEA strategy is vital. Without it, our average geodesic error is over 0.07 at different partiality levels. Remarkably, even when the Regularized ZoomOut strategy is replaced with ICP, the average geodesic error is higher than ours.

6.5 Runtime

We tested our algorithm, NRP, FSPM, and PFM on a workstation with Intel Core i9 3.6 GHz processor and 64 GB memory. We conducted the experiments on 30 random pairs from the SHREC'16 Partial Correspondence benchmark and showed in Fig. 13 a runtime comparison on meshes of sizes 1–10 k vertices. Similar to NRP and PFM, our method used the Mumford–Shah functional term in the optimization problem. The most computationally expensive part of them is the cumbersome optimization over the spectral and spatial domains. Note that our method is significantly faster than PFM and the currently best multi-part matching method NRP. The main reason is that our method's FSEA can quickly find the correct localization of the partial shape and allows the

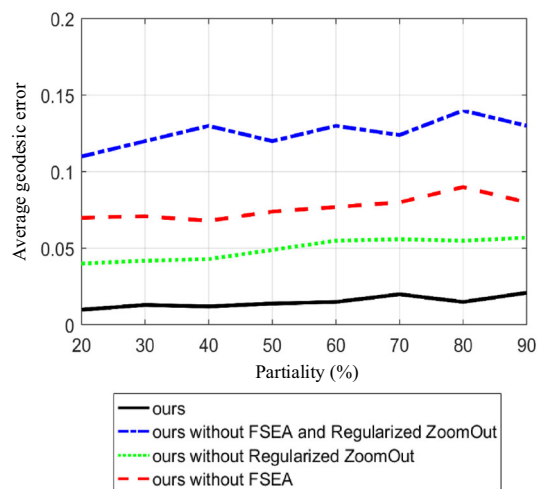


Fig. 12 Comparative results for the different ablations of our method

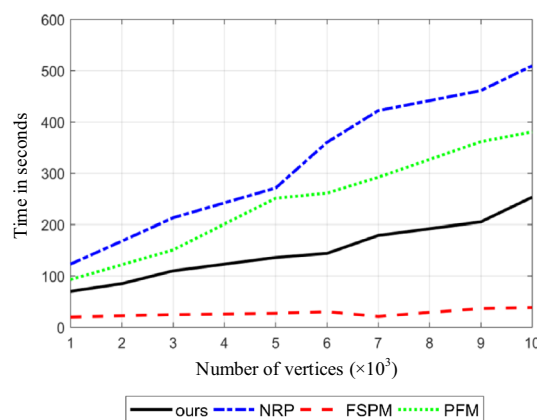


Fig. 13 Runtime of matching shapes with the varying number of vertices using our algorithm, compared to NRP, FSPM, and PFM

optimization to converge fast. At the same time, FSPM performs all the calculations in the spectral domain. Therefore, FSPM is the significant fastest over all methods.

Our technique adopts the Regularized ZoomOut strategy in the post-processing step, and the computational cost is higher than FSPM's refinement method ICP. In the typical case where the number of vertices $n = 10$ k, the CPU implementation of our strategy takes on average 30 s to converge, while ICP only 5 s. Our method's runtime complexity grows linearly with shape size.

7 Discussion and conclusions

This paper formulated a novel approach for computing partial correspondences between non-rigid multiple parts and full shape. In this setting, the parts are jointly matched to the full shape. We firstly introduced a new part regularization term based on Hamiltonian spectrum alignment for the

simultaneous partial functional map framework, promoting the accuracy of region localization directly in the spectral domains. Unlike existing work that utilizes complex regularizers, our technique adopts a simple localization technique calculating entirely in the spectral domain while yielding qualitatively better results at a fraction of the time cost. We then extended the ZoomOut pipeline using the RPMR algorithm, which can be seen as a non-rigid alignment between the two shapes' spectral embeddings, leading to a significant improvement in recovering the point-wise correspondence. We demonstrated our method's advantages by comparing it to several current state-of-the-art methods on well-known test datasets. Contrarily to these approaches, our method can tackle topological noise, strong partiality, missing pieces, and symmetric ambiguity, making it amenable for application in the scenario that frequently occurs when dealing with real data.

Our method still comes with multiple limitations. First, our FSEA strategy needs a larger k of the first few eigenvalues to discriminate parts that only differ at medium and high frequencies. To solve this problem, we have to increase the value of k and bear a higher calculation cost. Second, the post-processing step depends on some parameters that have to be tuned for some shapes. Specifically, we require to identify C_0 in the initialization and the step size of upsampling. Finally, due to the RPMR as a search algorithm, the Regularized ZoomOut obtains an accurate correspondence at a high computational cost.

In the future, we would like to investigate a partial correspondence approach that is suitable for non-isometric pairwise shapes, even multiple shapes. Even though this problem has received surprisingly little interest, we believe that it is a promising direction that makes it possible to accommodate the practical application. Moreover, it would be rewarding to research how deep learning approaches can solve designing a partial functional map.

Acknowledgements This work was supported by the National Natural Science Foundation of China [Grant Number 61862039].

References

1. Van, K.O., Zhang, H., Hamarneh, G., Cohen-Or, D.: A survey on shape correspondence. *Comput. Gr. Forum* **30**(6), 1681–1707 (2011)
2. Aigerman, N., Poranne, R., Lipman, Y.: Seamless surface mappings. *ACM Trans. Gr.* **34**(4), 1–13 (2015)
3. Hasler, N., Stoll, C., Sunkel, M., Rosenhahn, B., Seidel, H.P.: A statistical model of human pose and body shape. *Comput. Gr. Forum* **28**(2), 337–346 (2009)
4. Bogo, F., Romero, J., Loper, M., Black, M.J.: FAUST: dataset and evaluation for 3D mesh registration. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3794–3801 (2014)
5. Gasparetto, A., Torsello, A.: A statistical model of Riemannian metric variation for deformable shape analysis. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1219–1228 (2015)
6. Johnson-Roberson, M., Bryson, M., Friedman, A., Pizarro, O., Troni, G., Ozog, P., Henderson, J.C.: High-resolution underwater robotic vision-based mapping and three-dimensional reconstruction for archaeology. *J. Field Robot.* **34**(4), 625–643 (2017)
7. Papaioannou, G., Karabassi, E.A., Theoharis, T.: Virtual archaeologist: assembling the past. *IEEE Comput. Gr. Appl.* **21**(2), 53–59 (2001)
8. Huang, Q. X., Flöry, S., Gelfand, N., Hofer, M., Pottmann, H.: Reassembling fractured objects by geometric matching. In: *Proceedings of the ACM Siggraph*, pp. 569–578 (2006)
9. Nguyen, A., Ben-Chen, M., Welnicka, K., Ye, Y., Guibas, L.: An optimization approach to improving collections of shape maps. *Comput. Gr. Forum* **30**(5), 1481–1491 (2011)
10. Cosmo, L., Rodolà, E., Albarelli, A., Mémoli, F., Cremers, D.: Consistent partial matching of shape collections via sparse modeling. *Comput. Gr. Forum* **36**(1), 209–221 (2017)
11. Li, H., Sumner, R.W., Pauly, M.: Global correspondence optimization for non-rigid registration of depth scans. *Comput. Gr. Forum* **27**(5), 1421–1430 (2008)
12. Litany, O., Rodolà, E., Bronstein, A.M., Bronstein, M.M.: Fully spectral partial shape matching. *Comput. Gr. Forum* **36**(2), 247–258 (2017)
13. Rodolà, E., Möller, M., Cremers, D.: Regularized point-wise map recovery from functional correspondence. *Comput. Gr. Forum* **36**(8), 700–711 (2017)
14. Melzi, S., Ren, J., Rodolà, E., Sharma, A., Wonka, P., Ovsjanikov, M.: Zoomout: spectral upsampling for efficient shape correspondence. *ACM Trans. Graph.* **38**(6), 155–170 (2019)
15. Biasotti, S., Cerri, A., Bronstein, A.M., Bronstein, M.M.: Recent trends, applications, and perspectives in 3D shape similarity assessment. *Comput. Gr. Forum* **35**(6), 87–119 (2016)
16. Sahillioğlu, Y.: Recent advances in shape correspondence. *Vis. Comput.* **36**(8), 1705–1721 (2020)
17. Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L.: Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.* **31**(4), 30–41 (2012)
18. Pokrass, J., Bronstein, A.M., Bronstein, M.M., Sprechmann, P., Sapiro, G.: Sparse modeling of intrinsic correspondences. *Comput. Gr. Forum* **32**(2pt4), 459–468 (2013)
19. Gasparetto, A., Cosmo, L., Rodolà, E., Bronstein, M. M., Torsello, A.: Spatial maps: From low rank spectral to sparse spatial functional representations. In: *Proceedings of the international conference on 3D vision*, pp. 477–485 (2017)
20. Aflalo, Y., Dubrovina, A., Kimmel, R.: Spectral generalized multi-dimensional scaling. *Int. J. Comput. Vis.* **118**(3), 380–392 (2016)
21. Yang, J., Yan, H.: An algorithm for calculating shape correspondences using functional maps by calibrating base matrix of 3D shapes. *Geomat. Inf. Sci. Wuhan Univ.* **43**(10), 1518–1525 (2018)
22. Eynard, D., Rodolà, E., Glashoff, K., Bronstein, M. M.: Coupled functional maps. In: *Proceedings of the international conference on 3D vision*, pp. 399–407 (2016)
23. Ezuz, D., Solomon, J., Ben-Chen, M.: Reversible harmonic maps between discrete surfaces. *ACM Trans. Gr.* **38**(2), 1–12 (2019)
24. Ezuz, D., Heeren, B., Azencot, O., Rumpf, M., Ben-Chen, M.: Elastic correspondence between triangle meshes. *Comput. Graph. Forum* **38**(2), 121–134 (2019)
25. Melzi, S., Rodolà, E., Castellani, U., Bronstein, M.M.: Localized manifold harmonics for spectral shape analysis. *Comput. Gr. Forum* **37**(6), 20–34 (2018)
26. Wang, Y., Ben-Chen, M., Polterovich, I., Solomon, J.: Steklov spectral geometry for extrinsic shape analysis. *ACM Trans. Gr.* **38**(1), 1–21 (2018)

27. Nogneng, D., Ovsjanikov, M.: Informative descriptor preservation via commutativity for shape matching. *Comput. Gr. Forum* **36**(2), 259–267 (2017)
28. Nogneng, D., Melzi, S., Rodolà, E., Castellani, U., Bronstein, M.M., Ovsjanikov, M.: Improved functional mappings via product preservation. *Comput. Gr. Forum* **37**(2), 179–190 (2018)
29. Aiger, D., Mitra, N. J., Cohen-Or, D.: 4-points congruent sets for robust pairwise surface registration. In: *Proceedings of the ACM Siggraph*, pp. 1–10 (2008)
30. Albarelli, A., Rodolà, E., Torsello, A.: Fast and accurate surface alignment through an isometry-enforcing game. *Pattern Recognit.* **48**(7), 2209–2226 (2015)
31. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Generalized multi-dimensional scaling: a framework for isometry-invariant partial surface matching. *Proc. Natl. Acad. Sci.* **103**(5), 1168–1172 (2006)
32. Bronstein, A.M., Bronstein, M.M., Bruckstein, A.M., Kimmel, R.: Partial similarity of objects, or how to compare a centaur to a horse. *Int. J. Comput. Vis.* **84**(2), 163–183 (2009)
33. Windheuser, T., Schlickewei, U., Schimdt, F.R., Cremers, D.: Large-scale integer linear programming for orientation preserving 3d shape matching. *Comput. Graph. Forum* **30**(5), 1471–1480 (2011)
34. Sahillioğlu, Y., Yemez, Y.: Partial 3-D correspondence from shape extremities. *Comput. Graph. Forum* **33**(6), 63–76 (2014)
35. Masci, J., Boscaini, D., Bronstein, M. M., Vanderghenst, P.: Geodesic convolutional neural networks on Riemannian manifolds. In: *Proceedings of the IEEE international conference on computer vision workshops*, pp. 37–45 (2015)
36. Wei, L., Huang, Q., Ceylan, D., Vouga, E., Li, H.: Dense human body correspondences using convolutional networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1544–1553 (2016)
37. Kovnatsky, A., Bronstein, M. M., Bresson, X., Vanderghenst, P.: Functional correspondence by matrix completion. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 905–914 (2015)
38. Rodolà, E., Cosmo, L., Bronstein, M.M., Torsello, A., Cremers, D.: Partial functional correspondence. *Comput. Gr. Forum* **36**(1), 222–236 (2017)
39. Wu, Y., Yang, J., Zhao, J.: Partial 3D shape functional correspondence via fully spectral eigenvalue alignment and upsampling refinement. *Comput. Gr.* **92**, 99–113 (2020)
40. Torsello, A., Rodolà, E., Albarelli, A.: Multiview registration via graph diffusion of dual quaternions. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2441–2448 (2011)
41. Litany, O., Bronstein, A. M., Bronstein, M. M.: Putting the pieces together: Regularized multi-part shape matching. In: *Proceedings of the European conference on computer vision*, pp. 1–11 (2012)
42. Litany, O., Rodolà, E., Bronstein, A.M., Bronstein, M.M., Cremers, D.: Non-rigid puzzles. *Comput. Gr. Forum* **35**(5), 135–143 (2016)
43. Cosmo, L., Rodolà, E., Masci, J., Torsello, A., Bronstein, M. M.: Matching deformable objects in clutter. In: *Proceedings of the international conference on 3D vision*, pp. 1–10 (2016)
44. Sahillioğlu, Y., Yemez, Y.: Multiple shape correspondence by dynamic programming. *Comput. Gr. Forum* **33**(7), 121–130 (2014)
45. Huang, Q., Wang, F., Guibas, L.: Functional map networks for analyzing and exploring large shape collections. *ACM Trans. Gr.* **33**(4), 1–11 (2014)
46. Yang, J., Lei, M.: Correspondence calculation of model cluster by functional mapping combined with cycle-consistency constraints. *Laser Optoelectron. Progress.* **056**(008), 115–125 (2019)
47. Cohen, A., Ben-Chen, M.: Robust shape collection matching and correspondence from shape differences. *Comput. Gr. Forum* **39**(2), 555–568 (2020)
48. Huang, R., Ren, J., Wonka, P., Ovsjanikov, M.: Consistent zoomout: efficient spectral map synchronization. *Comput. Gr. Forum* **39**(5), 265–278 (2020)
49. Sun, J., Ovsjanikov, M., Guibas, L.: A concise and provably informative multi-scale signature based on heat diffusion. *Comput. Gr. Forum* **28**(5), 1383–1392 (2009)
50. Aubry, M., Schlickewei, U., Cremers, D.: The wave kernel signature: a quantum mechanical approach to shape analysis. In: *Proceedings of the IEEE international conference on computer vision workshops*, pp. 1626–1633 (2011)
51. Cosmo, L., Minello, G., Bronstein, M. M., Rossi, L., Torsello, A.: The average mixing kernel signature. In: *Proceedings of the European conference on computer vision*, pp. 1–17 (2020)
52. Halimi, O., Litany, O., Rodolà, E., Bronstein, A. M., Kimmel, R.: Unsupervised learning of dense shape correspondence. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4370–4379 (2019)
53. Roufousse, J. M., Sharma, A., Ovsjanikov, M.: Unsupervised deep learning for structured shape matching. In: *Proceedings of the IEEE international conference on computer vision*, pp. 1617–1627 (2019)
54. Ren, J., Panine, M., Wonka, P., Ovsjanikov, M.: Structured regularization of functional map computations. *Comput. Gr. Forum* **38**(5), 39–53 (2019)
55. Ezuz, D., Ben-Chen, M.: Deblurring and denoising of maps between shapes. *Comput. Gr. Forum* **36**(5), 165–174 (2017)
56. Kovnatsky, A., Bronstein, M.M., Bronstein, A.M., Glashoff, K., Kimmel, R.: Coupled quasi-harmonic bases. *Comput. Gr. Forum* **32**(2pt4), 439–448 (2013)
57. Rampini, A., Tallini, I., Ovsjanikov, M., Bronstein, A. M., Rodolà, E.: Correspondence-free region localization for partial shape similarity via Hamiltonian spectrum alignment. In: *Proceedings of the international conference on 3D vision*, pp. 37–46 (2019)
58. Choukroun, Y., Shtern, A., Bronstein, A. M., Kimmel, R.: Hamiltonian operator for spectral shape analysis. In: *Proceedings of the IEEE transactions on visualization and computer graphics*, pp. 1–12 (2018)
59. Postolache, E., Fumero, M., Cosmo, L., Rodolà, E.: A parametric analysis of discrete Hamiltonian functional maps. *Comput. Gr. Forum* **39**(5), 103–118 (2020)
60. Rustamov, R. M.: Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In: *Proceedings of the Eurographics Association*, pp. 225–233 (2007)
61. Bronstein, A. M., Bronstein, M. M.: Not only size matters: regularized partial matching of nonrigid shapes. In: *Proceedings of the IEEE computer society conference on computer vision and pattern recognition workshops*, pp. 1–6 (2008)
62. Kovnatsky, A., Glashoff, K., Bronstein, M. M.: MADMM: a generic algorithm for non-smooth optimization on manifolds. In: *Proceedings of the European conference on computer vision*, pp. 680–696 (2016)
63. Rodolà, E., Moeller, M., Cremers, D.: Point-wise map recovery and refinement from functional correspondence. In: *Proceedings of the international symposium on vision, modeling, and visualization*, pp. 25–32 (2015)
64. Yuille, A. L., Grzywacz, N. M.: The motion coherence theory. In: *Proceedings of the IEEE international conference on computer vision*, pp. 26–33 (1988)
65. Myronenko, A., Song, X.: Point set registration: coherent point drift. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(12), 2262–2275 (2010)
66. Tombari, F., Salti, S., Di, S. L.: Unique signatures of histograms for local surface description. In: *Proceedings of the European conference on computer vision*, pp. 356–369 (2010)
67. Kim, V.G., Lipman, Y., Funkhouser, T.: Blended intrinsic maps. *ACM Trans. Gr.* **30**(4), 1–12 (2011)

68. Cosmo, L., Rodolà, E., Bronstein, M. M., Torsello, A., Cremers, D., Sahillioğlu, Y.: Shrec16: partial matching of deformable shapes. In: Proceedings of the Eurographics workshop on 3D object retrieval, pp. 61–67 (2016)
69. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Numerical Geometry of Non-Rigid Shapes. Springer Science & Business Media, Berlin (2008)
70. Länér, Z., Rodolà, E., Bronstein, M. M., Cremers, D., Burghard, O., Cosmo, L., Sahillioğlu, Y.: Shrec16: matching of deformable shapes with topological noise. In: Proceedings of the Eurographics workshop on 3D object retrieval, pp. 55–60 (2016)
71. Rodolà, E., Rota, BuloS., Windheuser, T., Vestner, M., Cremers, D.: Dense non-rigid shape correspondence using random forests. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 4177–4184 (2014)
72. Anguelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., Davis, J.: SCAPE: shape completion and animation of people. ACM Trans. Gr. **24**(3), 408–416 (2005)



Jun Yang is a professor in the Computer Science Department at Lanzhou Jiaotong University (LZJTU) in China, where he directs the Computer Graphics and Computer Vision Lab. He obtained his PhD degree from the School of Information Science and Technology, Southwest Jiaotong University (SWJTU), China. His research area is computer graphics with a focus on geometric shape analysis and processing with emphasis on non-rigid shape matching and comparison.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Yan Wu received his BS and MS degrees in Electronic and Information Engineering from Jiangsu University, Zhenjiang, China, in 2008 and 2011, respectively. He is currently pursuing his PhD in the School of Electronic and Information Engineering, Lanzhou Jiaotong University (LZJTU), Lanzhou, China. His research interests are 3D shape registration, matching and deep learning.