The Visual Computer (2026) 42:105
https://doi.org/10.1007/s00371-025-04326-5

RESEARCH

Unsupervised landmark discovery via Karcher means for non-rigid shape correspondence

Haoliang Zhang[1] · Samuel Cheng[1] · Christian El Amm[1]

Received: 24 April 2025 / Accepted: 14 December 2025

Abstract

This paper proposes an unsupervised intrinsic landmark detection method that enhances non-rigid dense shape correspondence. Our approach identifies intrinsic landmarks on near-isometric surfaces, enabling efficient and robust point-wise correspondence. A key observation is that a small set of intrinsic landmarks is sufficient to reconstruct an isometric map, allowing for effective shape matching. Our method eliminates the need for manually defined landmarks by leveraging intrinsic geometric properties. Specifically, we compute the Karcher means of the surface and sequentially select the farthest points from these medians. This provides a stable and repeatable landmark set, which can be used in various shape correspondence frameworks. We evaluate our approach on multiple shape matching benchmarks and show that it consistently improves the accuracy of existing correspondence methods without requiring modifications. The results demonstrate that our method is computationally efficient and can serve as a plug-and-play component in shape matching pipelines. Code is available at https://github.com/ZHLRJ/Non-Rigid-Shape-Correspondence.

Keywords Intrinsic landmarks · Shape correspondence · Non-rigid shape matching

1 Introduction

Given a point on the first shape, how do we find a meaningful correspondence on all other shapes? This "meaningful correspondence" has also been termed registration, alignment, or shape matching. Establishing correspondences provides a valuable initial stage for a diverse array of applications, including texture transfer, shape analysis, co-segmentation, and exploration, just to name a few. To compute this correspondence, one idea is to identify a set of meaningful points on two shapes and then maximize the resemblance between the descriptors associated with these points. With the aid of rich local descriptors, the correspondence problem for 2D images is well solved. In contrast, constructing stable and repeatable descriptors for plain, non-rigid 3D shapes is highly challenging; even just identifying geometrically salient points with reliable descriptors is already quite costly. Instead of using the original representation of a surface, data-driven methods [23] learn descriptors to induce a more accurate correspondence. However, these methods are hindered by the scarcity of annotated 3D shapes and the associated financial costs.

Identifying an isometric mapping between two shapes is challenging due to the vastness of the potential mapping space. Ovsjanikov et al. [31] introduced a functional map-based approach for non-rigid shape matching, which leverages function correspondences to compute a mapping that facilitates the transfer of real-valued functions between shapes. This mapping is then converted into a point-to-point correspondence. A key advantage of this approach is that computations are performed on a predefined functional basis, significantly reducing the dimensionality of the search space. However, while functional map methods are both powerful and efficient, they do not inherently distinguish among multiple plausible solutions. In addition, the process of converting a functional map into a point-to-point correspondence often involves solving complex quadratic optimization problems.

Many shape matching algorithms [9, 14, 30] have shown that a small number of correspondences are sufficient to recover the full isometry. However, in practice, we often

✉ Haoliang Zhang
mars_zhang@ou.edu

Samuel Cheng
samuel.cheng@ou.edu

Christian El Amm
christian-elamm@ouhsc.edu

1 The University of Oklahoma, Norman, OK, USA

Published online: 04 January 2026

lack even a single known correspondence between shapes. Motivated by this, we propose an efficient algorithm that: (1) generates consistently ordered and salient landmark points, (2) matches shapes undergoing near-isometric deformations, and (3) computes dense shape correspondences efficiently.

While landmarks have been used to improve shape mapping in prior work, our goal is to achieve this without data-driven fitting or user-specified inputs. Our algorithm requires no manually registered datasets or predefined correspondences.

In this paper, we present a geometry-based approach for identifying consistent landmarks for shape correspondence. Our method is built upon two key properties:

- We make no assumptions about the availability of input data such as preregistered correspondences. Instead, our approach relies purely on geometric processing: It takes raw shape geometry as input and automatically generates consistent landmarks.
- Our method is general and avoids making specific assumptions about meshes. It can be applied broadly across various shape correspondence algorithms, including functional maps, heat kernel maps, and functional map optimization and refinement methods. This makes it a plug-and-play component for a wide range of shape matching pipelines.

Given a pair of near-isometric meshes, our algorithm automatically generates an ordered set of landmarks to resolve symmetry ambiguities. These consistent intrinsic landmarks can be leveraged in various applications—from improving correspondence accuracy in the functional map framework to enhancing existing optimization methods. To be clear, our method does not resolve ambiguities in perfectly symmetric shapes (e.g., spheres). However, it remains effective on real-world datasets, where exact symmetries are rare. We demonstrate the versatility of our approach across several shape matching algorithms, including functional maps, multiplicative operators, ICP (Iterative Closest Point), and ZoomOut.

2 Related work

The literature on shape matching is extensive [5, 11, 45, 48]. In the following section, we review methods most relevant to ours, focusing on techniques designed to compute dense correspondences between non-rigid shapes or those that can be adapted for this purpose.

Point-wise Methods. Early approaches to non-rigid shape matching focused on directly establishing correspondences between points on two shapes by minimizing an energy function, often based on approximate preservation of geodesic distances [9, 14, 30, 39, 40, 46] or conformality (local angle preservation) [1, 18, 22]. These methods are typically limited to hundreds of points due to high computational cost, a limitation that persists even in recent formulations tackling the quadratic assignment problem [12, 17, 27].

While not scalable to dense matching, these approaches are appealing for their theoretical properties—most notably, that a small number of landmarks can suffice to recover the full isometric map [22, 30]. However, their reliance on user-specified landmarks or training data limits practical applicability. In contrast, our method is fully automatic and produces a consistent ordering of landmarks across similar shapes without requiring any additional input, enhancing downstream correspondence quality.

Functional Maps. Other techniques proposed for non-rigid shape matching fall within the category of methods rooted in the concept of functional maps [31], which has been extended in several follow-up works [19, 34–36]. These methods significantly reduce the number of variables to optimize by selecting truncated Laplacian eigenbases for the function space defined on each shape and representing the mapping as a matrix that signifies a change of basis. However, converting a functional map to a point-wise map can be difficult and error-prone [13, 37]. These methods implicitly impose restrictive assumptions, such as near isometry, on the shapes being analyzed, and often encounter challenges in practice, leading to low-quality outcomes and limited surjectivity. Vestner et al. [49] proposed a filter to improve the quality of the input mapping, and Huang and Ovsjanikov [15] introduced adjoint regularization, both ensuring a bijective correspondence between the shapes. Nevertheless, the application of these techniques is still limited as they rely on handpicked initial matches for initialization and fail when there are no annotated landmarks.

Descriptor Learning. The performance of functional maps is heavily reliant on accurate descriptors and precise hyperparameter tuning [26]. In practice, constructing a stable and repeatable descriptor is highly challenging. Dense, specific dimensional descriptor fields on two shapes are typically employed as corresponding functions. Examples include GPS [38], Heat and Wave Kernel Signatures (HKS and WKS) [4, 44], SHOT [47], and geodesic distance descriptors [42]. Litman and Bronstein [24] introduced a data-driven spectral descriptor model that generalizes the HKS and WKS. Litany et al. [23] further improved descriptor quality by introducing a deep functional maps network (FMNet) for shape correspondence, which takes base descriptors as input and returns matches. FMNet refines the input descriptors by minimizing a geometric loss, yielding soft correspondence. One of the advantages of this supervised method [43] is its versatility, capable of handling different shape discretizations.

However, it has the inherent drawback of requiring training data. It is worth mentioning that, even though they are termed "unsupervised methods" [3, 20], these approaches are still fundamentally data-driven. Within the Ovsjanikov et al. [30] framework, our method consistently enhances the performance of the above models.

Iterative Map Refinement. Intrinsic symmetries, such as left-to-right symmetry, can cause functional maps to yield mixed results from multiple potential solutions. The Iterative Closest Point in spectral embedding was employed in the original functional maps approach [31] for post-processing, effectively mapping vertices from the source shape to the target shape. Typically, the ambiguity arising from symmetry is resolved using user-specified landmarks or training data [7]. However, this diminishes the practical utility of the ensuing pipeline. ZoomOut [28] and subsequent works [16, 21, 36] enhance results quality through the iterative updating of functional and point-to-point maps. Nonetheless, these methods face scalability issues with an increasing number of vertices; for instance, BCICP [34] takes an average of 2 min on meshes with 1, 000 vertices and exhausts memory around 30K vertices. Absent external information, the intensive computation required for basis precomputation, and iterative updating of point-wise maps [26] render these methods impractical for real-time applications.

3 Tool box

The core idea of our method is to improve dense shape matching through ordered intrinsic landmarks, computed automatically without any manual annotations. We first provide a concise overview of the key steps in computing functional maps and their properties, and then describe a set of fundamental building blocks that are integrated to form the final algorithm in Sec. 4.

3.1 Notation

Given two shapes M and N, typically represented as triangle meshes, let $T : N \to M$ be a bijective mapping between the manifolds M and N. Given a function f defined on shape M, $f : M \to \mathbb{R}$, we can transfer f onto N through composition using T, defining $g = f \circ T$. Here, $g : N \to \mathbb{R}$. We use W and A to denote the cotangent weights and area matrices, respectively, which collectively define the positive semi-definite Laplace–Beltrami operator as $L = A^{-1}W$. $d_M(x, y)$ represents the shortest geodesic distance between any two points $x, y \in M$. TM denotes the tangent bundle of M, which is the set of all tangent vectors.

3.2 Basic pipeline

Let M and N contain, respectively, n_M and n_N vertices. The fundamental workflow for computing a map using functional maps typically involves the following main steps, as outlined in Chapter 2 of [32]:

1. We compute a fixed set of basis functions for shapes M and N, storing them as columns in matrices Φ_M and Φ_N, respectively. These matrices are of sizes $n_M \times k_M$ and $n_N \times k_N$, where k_M and k_N are much smaller than n_M and n_N, respectively.
2. We compute q descriptor functions on each shape and project them onto the truncated Laplace–Beltrami basis to obtain coefficient vectors f^i and g^i. These coefficients form the columns of matrices F and G, of size $k_M \times q$ and $k_N \times q$, respectively.
3. Solving the functional map C via
$$\arg\min_C \|CF - G\|^2 + \alpha \|\Delta_N C - C\Delta_M\|^2, \quad (1)$$
where Δ_N and Δ_M are the Laplace–Beltrami operators of the two shapes, represented in their respective bases, and α is a small scalar weight. In some models, the second constraint is omitted.
4. Deriving a point-to-point correspondence from a given functional map matrix C.

Although a small number of Laplace–Beltrami basis functions can be used to represent each shape, the resulting linear transformation matrix between the two function spaces is only an approximation of the true map. However, acquiring a continuous point-to-point map through this workflow has been found to be both difficult and error-prone. In step (3), factors such as descriptor preservation do not disambiguate different possible solutions due to the symmetries of the shapes. It is worth noting that step (4) leads to difficult quadratic optimization problems, making the framework less practical for real-world use. For example, PMF [49] solves an equation at each iteration to determine a bijection, involving high algorithmic complexity and also requires the shapes to possess an identical number of vertices. BCICP runs out of memory for meshes around 30K vertices and takes 1 min per shape pair with 600 vertices (Sect. 5.2 [25]). These two challenges are closely related. One can accelerate the model by obtaining a better C from step 3 and thereby reducing the need for step 4.

Prior work has shown that only a few landmarks are sufficient to recover a full isometric map. Previous methodologies relied on external information, such as user-specified landmarks. In this study, we leverage the intrinsic structure of isometries between a pair of shapes. We show that under

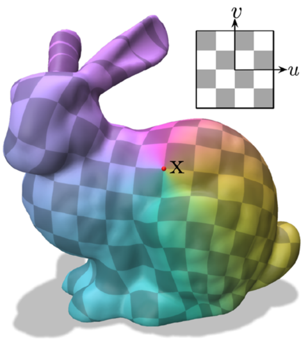

Fig. 1 The logarithmic map refers to a specific 2D parameterization of a surface, relative to a chosen point x. A global logarithmic map facilitates the translation of algorithms from the Euclidean space to curved domains

near-isometric conditions, it is possible to identify significant landmarks on both shapes.

3.3 Geometric medians

To compute ordered intrinsic landmarks, we rely on log and exponential maps to define Karcher means, and on geodesic distances as a core metric. While the previous section outlines the basic functional map framework, these geometric tools are essential for our landmark strategy, enabling the extraction of representative points that respect the intrinsic structure of the shape. We consider several important implementations from geometry processing that are used in our algorithm. To maintain a focused discourse, we restrict our discussion to triangulated surfaces.

Logarithmic Map The logarithmic map is a special 2D parameterization of a surface. At any point x on a smooth surface M, the exponential map [41], denoted as $\exp_x(rv) : T_x M \to M$, provides the point $y \in M$ reached by traversing a geodesic in the direction v (a unit vector) for a distance r. The logarithmic map $\log_x(y)$ is the inverse operation which is summarized in Algorithm 1. Intuitively, we can imagine that exp "wraps" the tangent plane around the surface. Logarithmic map provides a tool of "translation" for curved domains. It enables us to easily compute an average of points on a surface by mapping all points to the tangent plane using the log map (Fig. 1).

Algorithm 1: Logarithmic Map

Input : A point $x \in M$ and a zero direction $H_0 \in T_x M$.
Output: A function $(u, v) : M \to \mathbb{R}^2$
Find the shortest geodesic R from x to target point y.
Compute the angle ϕ between H_0 and R.
Construct $\exp_x(R) = y$; equivalently, $\log_x(y) = R$.
Compute $(u, v) = |R|(\cos\phi, \sin\phi)$ at point y.

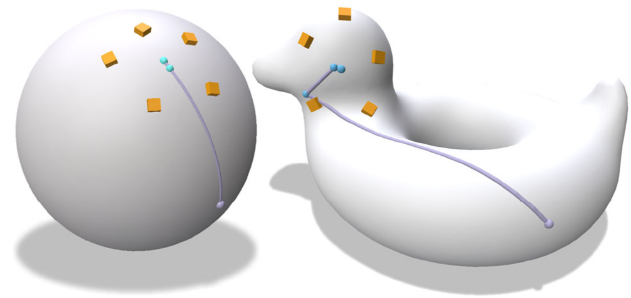

Fig. 2 In the case of the Karcher mean, given a collection of points, we can get the expected solution in just 2–3 steps; here, we show selected points (orange) and the initial guess (purple)

Geometric Medians If we think of the points in the plane, the average center of points can be easily calculated using the Euclidean mean. However, for a curved surface, the average of points may not necessarily lie on the surface. Given a set of points $y_1, \ldots, y_n \in M$, any point that minimizes the energy function can be considered a center.

$$E(x) := \frac{1}{2n} \sum_{i=1}^{n} d(x, y_i)^p. \tag{2}$$

For $p = 2$, centers on the surface are called *Karcher means* or *Frechet means* if unique. For $p = 1$, centers on surfaces are known as *geometric medians*. Both *geometric medians* and *Karcher means* can be efficiently computed via the log map in two steps (as shown in Fig. 2), where $\tau > 0$ is the step size in Algorithm 2.

$$v \leftarrow \frac{1}{n} \sum_{i=1}^{n} \log_{x^k}(y_i), \qquad x^{k+1} \leftarrow \exp_{x^k}(\tau v). \tag{3}$$

In general, an average on a surface is not unique; the resulting points may vary depending on the initialization.

Algorithm 2: Geometric Mean

Input : A set of points $y_1, \ldots, y_n \in M$.
Output: A point $m \in M$
Random initial point m^k from the set M.
while $|v| < \varepsilon$ **do**
  foreach $y_i \in \{y_1, \ldots, y_n\}$ **do**
    Compute the logarithmic map at m^k via Algorithm 1.
    Update vector $v \mathrel{+}= \frac{1}{n} \log_{m^k}(y_i)$.
  Update $m^{k+1} = \exp_{m^k}(\tau v)$, which represents moving forward along the vector v from point m^k for a time duration of τ.

Geodesic Distances In our algorithm, we need to compute geodesic distances between a large number of point pairs,

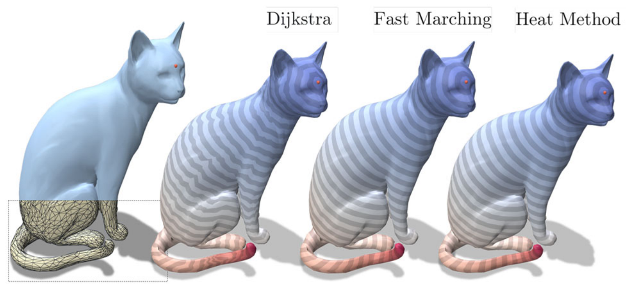

Fig. 3 Comparison of geodesic distance fields computed using Dijkstra, Fast Marching, and the Heat Method. The heat method, based on solving linear elliptic PDEs, produces smoother distance isolines with fewer numerical artifacts, especially on irregular meshes. According to benchmarks by Crane et al. [10], the heat method achieves up to 68× speedup over Fast Marching on large meshes (e.g., 1.45s vs. 98.1s on a 1.6M-triangle model) with comparable or better mean error

making it essential to adopt an accurate and efficient distance computation method. We compare the heat method with Dijkstra and Fast Marching in Fig. 3 to highlight its advantages in smoothness and efficiency, motivating its use in our pipeline. We use heat kernel maps (HKM) to construct robust pointwise descriptors that capture intrinsic geometric information around each landmark. Specifically, for a given shape, we compute the HKM at each selected landmark and treat these as scalar functions defined over the surface. These functions are then projected onto the Laplace–Beltrami eigenbasis to obtain compact functional descriptors. When applied to a pair of shapes, the descriptors derived from corresponding landmarks are used to compute the functional map via least squares fitting. This integration allows HKM to serve as a reliable bridge between intrinsic landmarks and the functional representation, improving the consistency and accuracy of the resulting correspondences.

4 Algorithm

We now have all the ingredients to implement our method using the tools from Sec. 3. Unlike many existing approaches, our method is entirely unsupervised and does not require any manual initialization or user-provided landmarks. Instead, we rely solely on intrinsic geometric information to automatically identify a consistent set of intrinsic landmarks.

In this section, we present our full unsupervised pipeline for intrinsic landmark discovery and dense shape correspondence. The process consists of three main steps: (1) compute pair-wise geodesic distance maps using the heat method; (2) extract intrinsic landmarks via robust geometric center estimation (Sec. 4.1); and (3) use these landmarks to guide downstream correspondence refinement (Sec. 4.2). We now describe each component in detail.

4.1 Ordered landmarks

We refer to ordered landmarks as the selected points ranked by their average geodesic distance from the set of Karcher means. We now introduce our primary algorithm, the *intrinsic landmarks method*, which is summarized in Algorithm 3. The basic idea is to first calculate a set of Karcher means from various starting guesses. To obtain the landmarks, we compute the average geodesic of all vertices relative to this set and sort the list. In Fig. 4, we show the results of Karcher means computation from different initializations, demonstrating the consistency of the selected landmarks across a variety of shapes.

Consider a surface where salient points have been identified and a Karcher mean m has been computed. For any set of salient landmarks $p_1, p_2, \ldots, p_k$ on surface M, the geodesic distance between any two of these landmarks, represented as $d(p_i, p_j)$. This value should be greater than or equal to the maximum geodesic distance between any landmark and the Karcher mean. This can be formulated as (Fig. 5):

$$d(p_i, p_j) \geq \max d(p_l, m), m \in K. \tag{4}$$

Let K denote the set of Karcher means on the surface M. The hypothesis suggests that the Karcher mean serves as a central reference, being closer to each salient point than any pair of salient points are to each other. This hypothesis might not be applicable to non-generic manifolds. The sphere is a classic example of a non-generic manifold, for which counterexamples are easily found. For instance, as shown in Fig. 2, the geodesic distance between two points marked in orange boxes could be less than the distance from any orange box to the Karcher mean (depicted in turquoise). However, most real-world shapes, such as human bodies, possess only reflectional symmetry. Salient points on these shapes are typically stable and reliable in practical applications, as demonstrated in Fig. 6.

Let us first note that incorporating landmark constraints into the functional map formulation does not alter its runtime complexity, as it still reduces to solving a least squares problem [34]. Although we do not provide a formal generalization analysis for Eq. (4), we observe that our algorithm consistently yields stable, ordered landmarks across various shapes. This helps mitigate key factors affecting map quality, all with negligible preprocessing overhead.

For instance, on a mesh with $7K$ vertices, generating 5 ordered landmarks using 10 Karcher mean samples takes approximately 2.37 s for Karcher means computation and 0.015 s for landmark selection. While few existing methods perform analogous preprocessing, our approach remains lightweight and can be readily integrated into existing pipelines—potentially improving correspondence accuracy or accelerating convergence. Benefiting from the efficiency

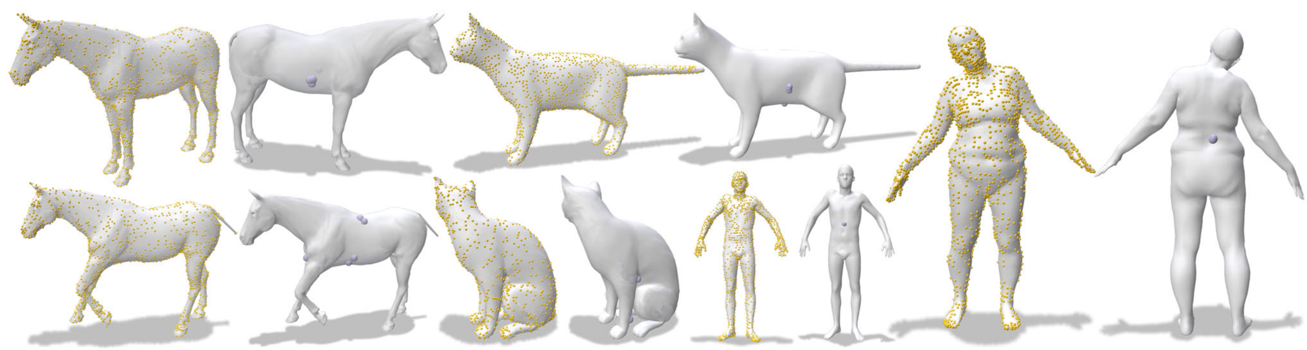

Fig. 4 We compute the *Karcher means* of the yellow points, which were obtained from a substantial number of initial guesses (covering over 90% of all vertices) and are indicated by the purple points. While there can be more than one Karcher mean, we demonstrate here with five center points. In some cases, iterations push the solution toward a local minimum or yield only approximate averages. For human shapes, however, the Karcher means are highly stable

Algorithm 3: Ordered Landmarks

Input : M, n ▷ selected meshes, # landmarks
Output: L ▷ Ordered landmarks
▷ initialization
$K \leftarrow$ KarcherMean(M) ▷ Algorithm 2
$D \leftarrow$ EmptyList() ▷ geodesic to m
foreach *mean* $k \in K$ **do**
  $l_k \leftarrow$ LogarithmicMap (k) ▷ calculate logarithmic map
  foreach *vertex* $v \in M$ **do**
    $D[v] += l_k[v]/n$ ▷ vertex geodesic to mean
  Sort(D) ▷ descending order
▷ build landmarks
while $\#L < n$ ***and*** *notEmpty(D)* **do**
  $v \leftarrow$ Pop (D) ▷ extract maximum geodesic point
  if *isEmpty* (L) ***or*** $d(v, l) \geq D[v] \quad \forall l \in L$ ▷ *Eq. (4)*
  then
    L.append (v)
return L

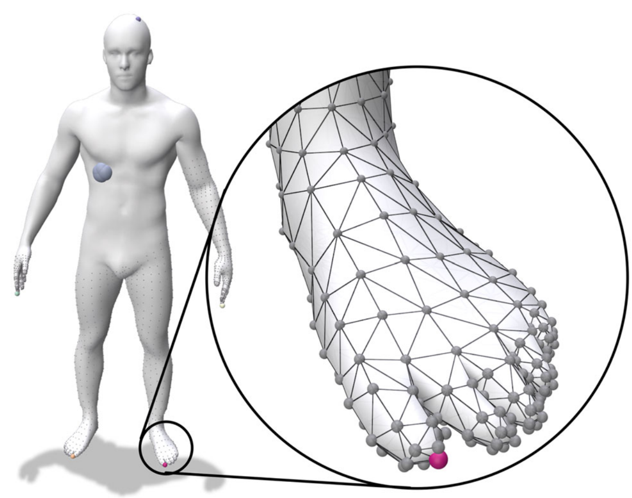

Fig. 5 Selecting the vertices with the greatest geodesic distance from the set of Karcher means, shown in light blue, while also verifying if the point satisfies Eq. (4). In the image, the large purple sphere represents the set of Karcher means. Although the gray points have large geodesic distances to the Karcher mean, they are too close to existing landmarks and thus not selected. The primary landmarks are highlighted in red (visible in the zoomed-out view)

of the heat method, our approach remains lightweight and readily integrates into existing pipelines—potentially improving correspondence accuracy or accelerating convergence. Even on a mesh with 100K vertices, computing a single Karcher mean takes only 1.13 s, confirming the scalability of our method. Compared to prior optimization-based pipelines that require user-specified landmarks or extensive learning-based training, our method offers a significantly faster and fully automated alternative. In practice, we set the number of landmarks n to a small value (typically 3 to 5), which balances coverage and stability across shapes.

On meshes containing approximately $30K$ vertices, the computation of the *Karcher means* can become relatively slow, often taking more than 60 s. In these cases, it may be advisable to initialize the Karcher mean with a smaller set of points. Nonetheless, the efficiency of landmark generation remains unaffected because the *Logarithmic Map* is only applied to a select set of landmarks, which guarantees quick processing.

While this mechanism already performs quite effectively, we can enhance both robustness and accuracy by adjusting the vertex density according to the geometry of the input mesh. The basic idea is to construct an alternative Karcher mean energy for the same piecewise Euclidean surface which leads to better average points for subsequent operations. In this case, we can define a central energy.

$$E(x) := \frac{1}{2n} \sum_{i=1}^{n} \rho(y_i) d(x, y_i)^p. \tag{5}$$

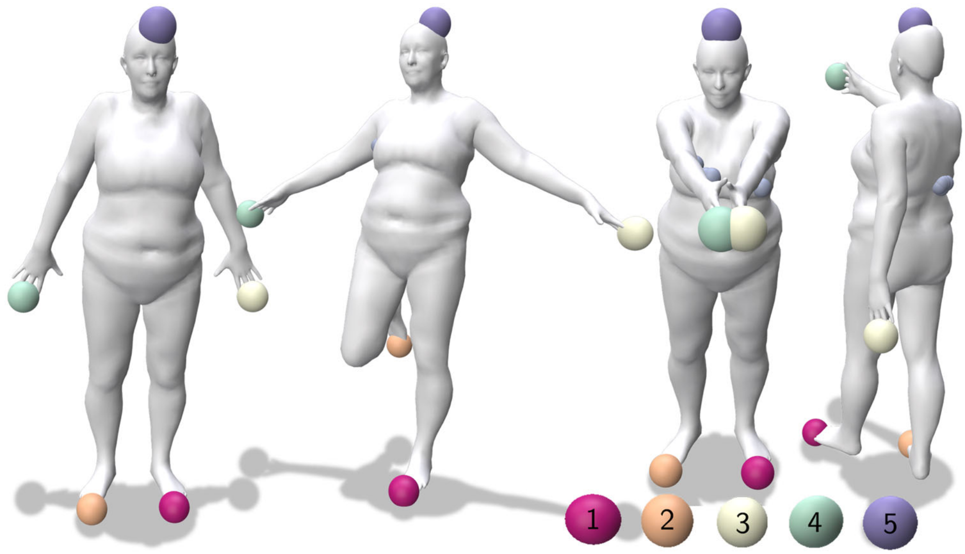

Fig. 6 Our algorithm efficiently provides consistent ordered landmarks, facilitating the most shape correspondence algorithm

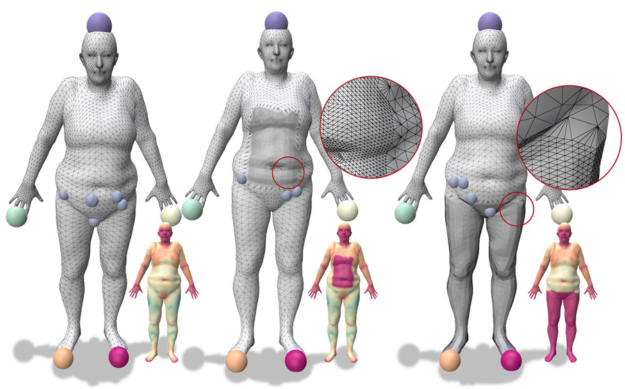

Fig. 7 We subdivided a section of the mesh and compared our method under various scenarios. Our approach demonstrates insensitivity to input triangulation, ensuring that the landmarks remain robust against non-uniform vertex density. The red color represents areas of high vertex density (small vertex dual area). The light blue spheres indicate the Karcher means computed from the input mesh

The only modification to the algorithm involves taking a weighted average across all vertices, where $\rho(y_i) > 0$ denotes the density at each vertex. In our implementation, the weights $\rho(y_i)$ are derived from the vertex dual areas, and they are normalized implicitly during the update step. We use a fixed initial step size $\tau = 1.0$ and a convergence threshold based on local face size (1/3 of the face scale). To ensure stability across different initializations, we run Algorithm 2 multiple times (typically 10) with different random seeds and retain the resulting set of stable Karcher means. This strategy improves robustness and reduces the influence of local minima. In Fig. 7, we construct a stress test by employing an extremely uneven triangulation. Despite these conditions, our algorithm capitalizes on the advantages of area-weighted geodesics to derive more robust landmarks.

4.2 Correspondence augmentation

Previous algorithms that address symmetry ambiguity typically depend on user-specified landmarks or incorporate additional optimization constraints, such as multiplicative operators [29] or orientation-preserving operators [34]. In contrast, our method relies solely on geometric processing to generate ordered landmarks. Our algorithm operates in two distinct stages: Initially, it establishes a select number of landmark correspondences using Algorithm 3, and then, it extends these to a dense map encompassing the entire shape. The heat kernel map theory posits that even a single correspondence can be adequate for isometry recovery, as detailed in Theorem 3.1 of [30]. Furthermore, our algorithm's provision of multiple correspondences significantly bolsters stability.

Perhaps the most closely related work to our method is the study by Ovsjanikov et al [30]. They demonstrated that a single correspondence can be utilized to reconstruct an isometry defined across entire shapes [30]. The original heat kernel map utilizes a related, but fundamentally different technique from our algorithm: The former requires user-specified landmarks, while the latter automatically generates ordered landmarks.

To begin with, it is important to remember that the heat kernel map associated with any point x in M and a predetermined source point p is defined as follows:

$$\Phi_p^M : M \to F, \ \Phi_p^M(x) = k_t^M(p, x), \tag{6}$$

where F is the space of functions. Essentially, the heat kernel $k_t^M(p, x)$ gauges the heat transfer from point p to point x over time t. It is noteworthy that the source point p can be any point on the manifold M. For each point x in M, we discretize $k_t^M(p, x)$ into a collection of time samples $\{t_j\}$. Consequently, for every point, we compute a vector of size $J = |\{t_j\}|$, where the j^{th} component is $k_{t_j}^M(p, x)$. Analogously, for y in N, we calculate $k_{t_j}^M(q, y)$. Subsequently, we evaluate the correspondence by computing the squared distances between the heat kernel maps:

$$E(p, q_i) = \sum_{x \in M} \min_{y \in N} ||\Phi_p^M(x) - \Phi_{q_i}^N(y)||^2. \tag{7}$$

We now understand that certain landmarks $p = \{p_1, p_2, \dots, p_k\} \in M$ correspond to landmarks $q = \{q_1, q_2, \dots, q_k\} \in N$. For the landmark correspondence (p_k, q_k), we can allocate to each point x a combined heat kernel map, where the additional J coordinates represent $k_{t_j}^M(p, x)$. In practice, we incorporate each correspondence (p_k, q_k) by minimizing the combined error:

$$f(x) = \arg\min_{y \in N} \|\Phi_p^M(x) - \Phi_q^N(y)\|. \tag{8}$$

Although the heat kernel map introduces additional coordinates, all the resulting quantities remain commensurable [33]. Consequently, we can employ the same distance metric for their comparison.

5 Results

We have implemented the proposed algorithms in C++. We first detail the competing methods and their implementations. Subsequently, we describe the datasets used and present a series of both quantitative and qualitative evaluations. We demonstrate the efficacy of intrinsic landmarks across a diverse array of shapes and illustrate how our approach significantly enhances the performance of other algorithms.

Description of competing algorithms We conduct a comparative analysis of our method, termed *Intrinsic Landmarks Correspondence (ILC)*, against various competing algorithms.

- ICP: Uses ICP (Iterative Closest Point) to refine soft correspondences, following the approach outlined in the original functional maps framework [31].
- ZoomOut [26, 28]: Improves mappings or correspondences by iteratively upsampling in the spectral domain within the functional map framework.
- Multiplicative Operator Commutativity [29]: Views descriptors as linear operators that act on functions via multiplication, incorporating descriptor preservation constraints formulated via commutativity with the unknown map.
- symmOP + BCICP [34]: Improves the basic ICP by formulating orientation preservation in the spectral domain, enhancing performance in both the spatial and spectral domains.
- (WKS/HKS) [4, 44] Initialization: Employs the basic functional map framework with multiplicative operators to compute initial soft correspondences, employing region-based correspondences computed by an automated algorithm as descriptors.

Datasets We utilize various standard datasets to compare the different techniques: FAUST, TOSCA, and SCAPE.

- FAUST: The dataset [6] comprises 300 triangulated meshes of 10 distinct subjects, with each subject scanned in 30 different poses.
- TOSCA: The TOSCA dataset [8] consists of 80 meshes categorized into nine groups. Each category comprises objects with shared triangulation and an equal number of vertices, numbered compatibly.
- SCAPE [2]: comprises human shapes only. The datasets contains 71 registered meshes of a particular person in different poses.

Metrics We assess various methods using standard metrics for dense shape correspondence [34]. Ground truth correspondences are provided in each dataset. To evaluate the continuity of the dense map $T : N \to M$, we compute the geodesic distance between $T(x)$ and $T^*(x)$ for all $x \in N$, where T^* represents the ground truth map.

We first demonstrate a correspondence augmentation that seamlessly integrates into the basic functional map pipeline, as depicted in Eq. (1). In Fig. 8, we utilized 35 eigenvectors to compute the HKS descriptors for initializing a functional map. The term +Commutative refers to adding the term $w_{d-comm} \sum_i \|C\Gamma_1^i - \Gamma_2^i C\|^2$ [29], while +symmOP denotes the addition of the symmetric operator $w_{orient} \sum_i \|C\Lambda_1^i - \Lambda_2^i C\|^2$ [34] to the basic pipeline. $+ILC$ indicates the inclusion of our five landmarks as mentioned in Sec. 4.2. Here, the multiplicative operators Γ_1^i and Γ_2^i, as well as the orientation-preserving operators Λ_1^i and Λ_2^i, are associated with the i^{th} descriptors. Incorporating the proposed landmarks results in a significant improvement in all three cases.

In addition, we evaluate the utility of our method to enhance various refinement techniques, as discussed. We compare our method to the following:

ICP: We apply ICP to a functional map of dimensions 35×35, estimated using the same pipeline adopted for the initial map.

ZoomOut: We refine the initial functional map and upsample this map with a step size of 1.

From the results in Fig. 9, it is evident that the intrinsic landmark operators are highly beneficial. Our approach

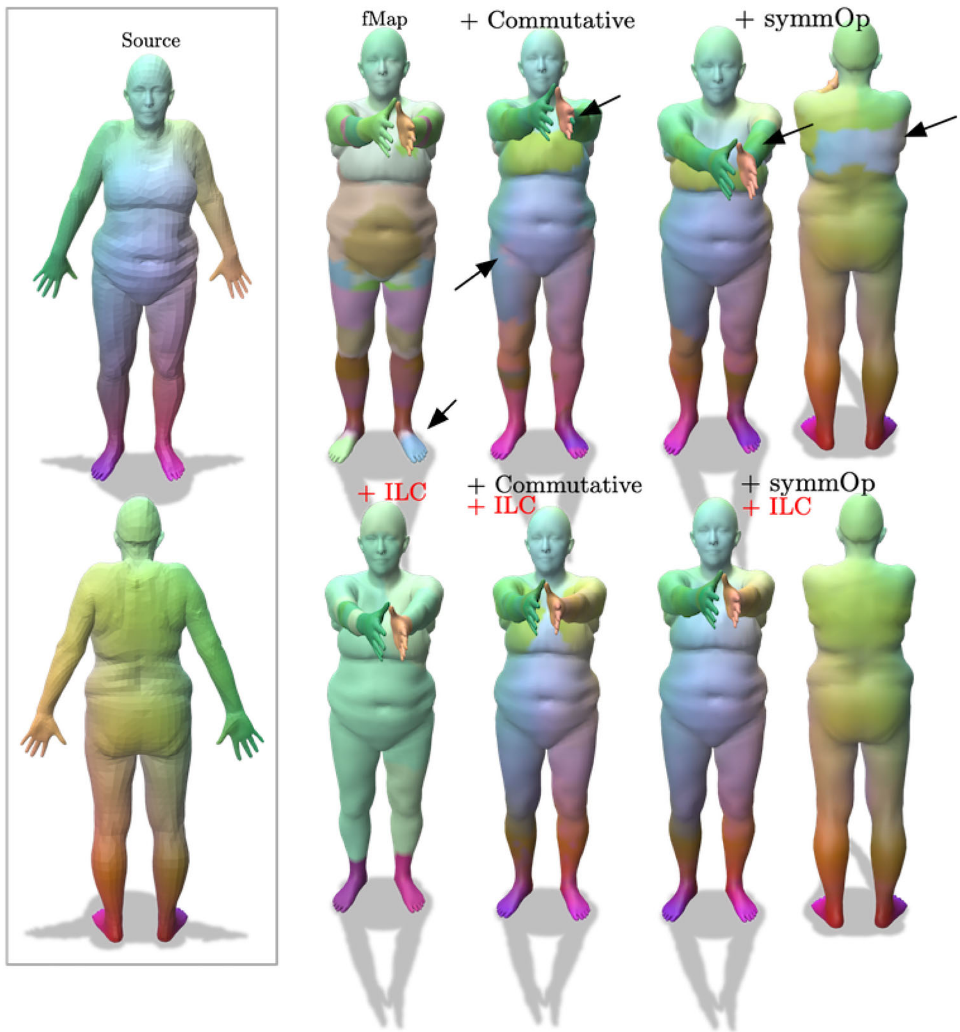

Fig. 8 Comparison of dense correspondence results with and without intrinsic landmarks. The two shapes in the leftmost column serve as the source reference for color transfer. The top row shows the results without using intrinsic landmarks across three functional map frameworks (fMap, + Commutative, + symmOp), where the black arrows highlight significant mapping errors. The bottom row shows the corresponding results enhanced with our intrinsic landmark constraints (indicated by red "+ IL"), leading to consistently improved correspondences across all methods

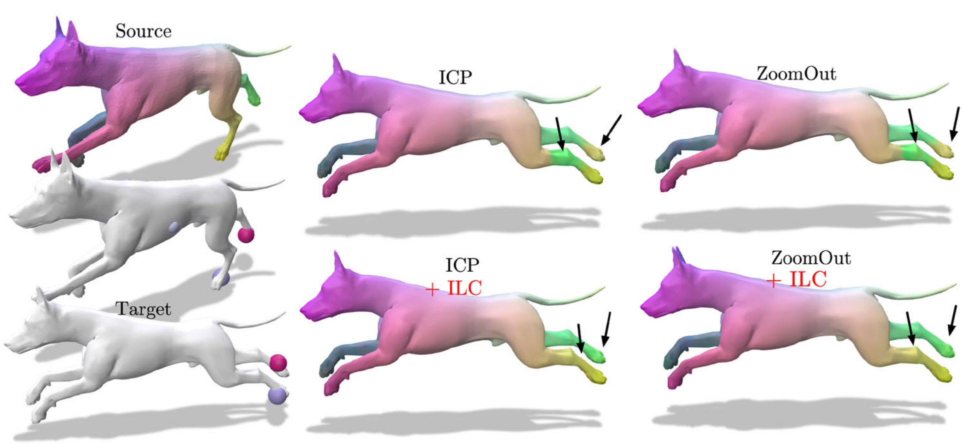

Fig. 9 Qualitative comparison on the TOSCA dataset. The leftmost shapes are the source and target meshes, with two identified intrinsic landmarks indicated by red and purple spheres. The top row shows the correspondence results from ICP and ZoomOut without using intrinsic landmarks, while the bottom row shows the results with our intrinsic landmark constraint (denoted by red "+ IL"). Black arrows highlight areas where the original methods produce visible misalignments, which are corrected by our method. Our approach improves the accuracy of correspondences with minimal preprocessing time (<2 s) and captures finer structural details

Table 1 Quantitative assessment of refinement in shape matching

Method	Average error ($\times 10^{-3}$)	Average runtime (s)
Init	67.5	–
PMF [49]	25.9	311.2
ICP [31]	54.4	12.2
BCICP [34]	22.3	113.5
ZO [28]	20.1	87.4
Scalable ZO [26]	21.7	16.3
Ours + ICP	**47.3**	14.3
Ours + BCICP	**19.2**	115.2
Ours+ ZO	**18.5**	89.3
Ours+ Scalable ZO	**20.9**	18.9

Beginning with a computed 20×20 functional map. The datasets include 190 FAUST pairs and SCAPE
Significance of bold values denote the lowest average error achieved by our method in comparison with the reference methods

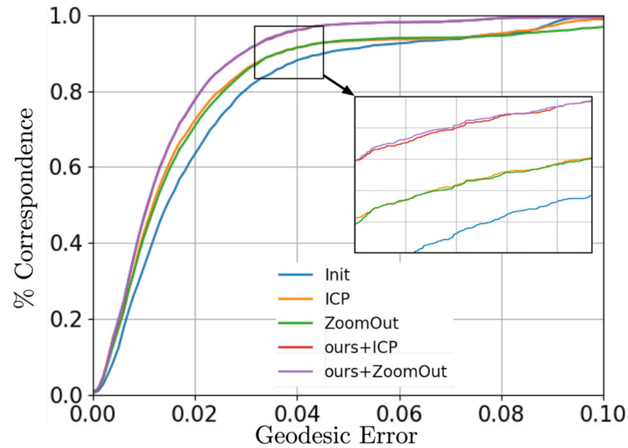

Fig. 10 Evaluation on Dynamic FAUST using geodesic error measures. We present ICP and ZoomOut predictions both before and after applying our augmentation. The initialization parameters are set to $w_{d-comm} = 0.1$ and $w_{orient} = 0.1$

not only achieves quality improvements compared to all raw methods but is also significantly faster. In Fig. 9 and 8, arrows highlight regions where our intrinsic landmark augmentation improves the dense correspondence, especially in symmetric or highly articulated areas.

Many shape refinement methods exhibit poor scalability with the number of vertices. As discussed in Sec. 4.2, the additional landmarks introduced in our method incur nearly negligible computational cost. The main advantage of our approach becomes evident when combined with other correspondence frameworks, as mentioned. Since our algorithm does not entail model training or n-dimensional matrices, its runtime scales linearly with the original number of randomly selected vertices.

Table 1 reports the average error and runtime. Our method achieves a 13.1%, 16.1%, 8.0%, and 3.7% improvement in accuracy over the state of the art while incurring only a couple of seconds in computational cost. Furthermore, our method demonstrates greater versatility, enabling seamless integration into several existing frameworks, in contrast to typical algorithms. In Fig. 10, we present the results of the ICP and ZoomOut methods both with and without our augmentation. It is noteworthy that both methods exhibit improvements when our intrinsic landmarks are applied.

5.1 Ablation study

In this section, we evaluate the robustness and generalization ability of our method under challenging conditions. We first demonstrate that our approach is resilient to changes in mesh connectivity by testing on remeshed datasets. Then, we present representative failure cases to highlight the limitations of our pipeline in practical scenarios with incomplete or partially matching landmarks.

To further evaluate the robustness of our method under varying mesh connectivity, we conduct experiments on the FAUST remeshed dataset [34], where the source and target meshes have same geometry but significantly different triangulations. As shown in Fig. 11, our method consistently identifies semantically meaningful landmarks across differently meshed surfaces. This result highlights the sampling and resolution agnosticism of our pipeline, and demonstrates its reliability in real-world scenarios where mesh quality and connectivity may vary substantially.

Table 2 shows the failure rate of our method. A pair is considered a failure if its result is worse than that of ICP or ZoomOut with WKS initialization. Most failures occur in partial matching scenarios, where missing regions alter the geodesic structure.

Figure 12 illustrates a representative failure case where the left and right feet are incorrectly matched. Although the shapes are complete, the lack of strong extrinsic cues causes the intrinsic landmarks to become symmetrically ambiguous, matching the left foot to the right. This semantic confusion introduces localized distortion. While such ambiguity leads to a 10–20% drop in accuracy, our method still performs comparably or better than baselines without intrinsic landmarks, though it does not match the best-case results achievable under unambiguous geometry.

Limitations and future work

Although our method achieves high-quality results, it has several limitations that we aim to address in future work. First, while intrinsic landmarks often improve correspondence, we cannot guarantee that their inclusion benefits all models. For example, in our experiments with fMap (+ILC), enforcing bijectivity near specific points occasionally degrades performance elsewhere. Extending the current pipeline to provide localized guidance without introduc-

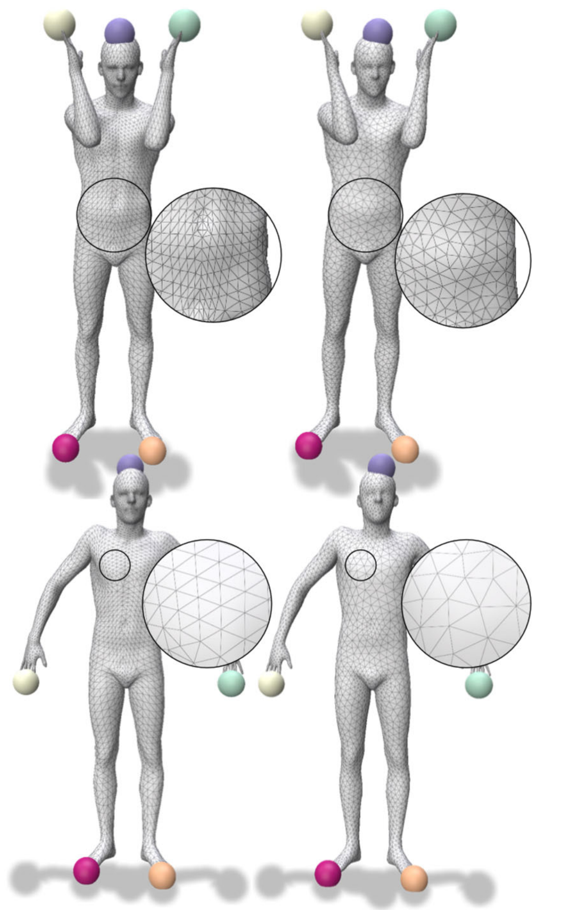

Fig. 11 Our approach produces stable and consistent landmark placement across remeshed surfaces (right), demonstrating sampling and resolution agnosticism. This robustness enables reliable shape matching across a wide range of mesh representations, including those with significantly different vertex counts

Table 2 Failure Rate

Dataset	Compare to ICP + ZoomOut (%)
Dynamic FAUST	8.75
TOSCA	13.19
SCAPE	7.97

The table shows the percentage of shape pairs where our method performs worse than ICP or ZoomOut
Bold values indicate cases where our method outperforms the corresponding baseline methods in terms of average error

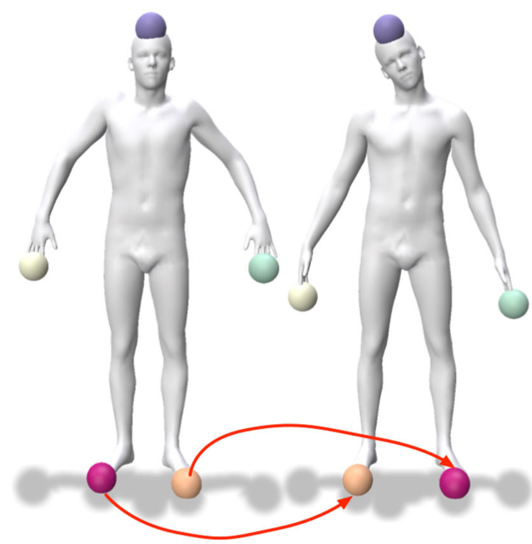

Fig. 12 Failure case under partial matching. The intrinsic landmarks on the feet are mismatched, resulting in incorrect correspondences

ing such side effects remains an open challenge. Second, our method is not directly applicable to partial shape correspondence, as topological differences significantly affect geodesic distances. Thirdly, our method is designed for shapes with approximately preserved intrinsic geometry and is not intended for cross-category landmark discovery, where semantic correspondences are often ill-defined. In future work, we plan to extend our approach to handle partial inputs, enabling more flexible and robust shape matching across a broader range of scenarios.

6 Conclusion

We presented an unsupervised method for computing dense correspondences between non-rigid shapes via automatically generated intrinsic landmarks. Our approach computes a sequence of Karcher means and selects geodesically distributed points to construct an ordered set of landmarks, without requiring any user input or supervision. Built upon the functional map framework of Ovsjanikov et al [30], our method integrates seamlessly into existing pipelines and improves correspondence accuracy with negligible additional cost (2–3 s). The core insight is that a small set of well-ordered intrinsic landmarks can effectively guide dense matching by leveraging only geometric information. Our method demonstrates strong compatibility with a range of existing shape correspondence algorithms and achieves consistent improvements across benchmarks, highlighting its generalization ability across datasets and pipelines.

Author Contributions H.Z. designed the method, implemented the experiments, and wrote the manuscript. S.C. and C.E. supervised the project and provided critical feedback. All authors reviewed and approved the final manuscript.

Data Availability No datasets were generated or analyzed during the current study.

Declarations

Conflict of interest The authors declare no conflict of interest.

References

1. Aigerman, N., Poranne, R., Lipman, Y.: Seamless surface mappings. ACM Trans. Graph., 34 (4), (2015). https://doi.org/10.1145/2766921
2. Anguelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., Davis, J.: Scape: shape completion and animation of people. ACM Trans. Graph. **24**(3), 408–416 (2005). https://doi.org/10.1145/1073204.1073207. (**ISSN 0730-0301**)
3. Attaiki, S., Ovsjanikov, M.: Ncp: neural correspondence prior for effective unsupervised shape matching. In Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22, Red Hook, NY, USA, (2022). Curran Associates Inc. ISBN 9781713871088
4. Aubry, M., Schlickewei, U., Cremers, D.: The wave kernel signature: a quantum mechanical approach to shape analysis . In 2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops), pages 1626–1633, Los Alamitos, CA, USA, Nov. (2011). IEEE Computer Society. https://doi.org/10.1109/ICCVW.2011.6130444
5. Biasotti, S., Cerri, A., Bronstein, A., Bronstein, M.: Recent trends, applications, and perspectives in 3d shape similarity assessment. Comput. Graph. Forum **35**(6), 87–119 (2016). https://doi.org/10.1111/cgf.12734. (**ISSN 0167-7055**)
6. Bogo, F., Romero, J., Loper, M., Black, M.J.: Faust: dataset and evaluation for 3d mesh registration. 2014 IEEE Conference on Computer Vision and Pattern Recognition, pages 3794–3801, (2014). https://api.semanticscholar.org/CorpusID:206592437
7. Boscaini, D., Masci, J., Rodolà, E., Bronstein, M.: Learning shape correspondence with anisotropic convolutional neural networks. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, Advances in Neural Information Processing Systems, volume 29. Curran Associates, Inc., (2016). https://proceedings.neurips.cc/paper_files/paper/2016/file/228499b55310264a8ea0e27b6e7c6ab6-Paper.pdf
8. Bronstein, A., Bronstein, M., Kimmel, R.: Numerical geometry of non-rigid shapes, 1st edn. Springer Publishing Company, Incorporated (2008) ISSN 0387733000
9. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching. Proc. Natl. Acad. Sci. **103**(5), 1168–1172 (2006). https://doi.org/10.1073/pnas.0508601103. https://www.pnas.org/doi/abs/10.1073/pnas.0508601103
10. Crane, K., Weischedel, C., Wardetzky, M.: The heat method for distance computation. Commun. ACM **60**(11), 90–99 (2017). https://doi.org/10.1145/3131280. (**ISSN 0001-0782**)
11. Deng, B., Yao, Y., Dyke, R.M., Zhang, J.: A survey of non-rigid 3d registration. Comput. Graph. Forum **41**(2), 559–589 (2022). https://doi.org/10.1111/cgf.14502
12. Dym, N., Maron, H., Lipman, Y.: DS++: a flexible, scalable and provably tight relaxation for matching problems. ACM Trans. Graph., 36 (6), (2017). https://doi.org/10.1145/3130800.3130826
13. Ezuz, D., Ben-Chen, M.: Deblurring and denoising of maps between shapes. Comput. Graph. Forum **36**(5), 165–174 (2017). https://doi.org/10.1111/cgf.13254. (**ISSN 0167-7055**)
14. Huang, Q.-X., Adams, B., Wicke, M., Guibas, L.J.: Non-rigid registration under isometric deformations. In Proceedings of the Symposium on Geometry Processing, SGP '08, page 1449–1457, Goslar, DEU, (2008). Eurographics Association
15. Huang, R., Ovsjanikov, M.: Adjoint map representation for shape analysis and matching. Computer Graphics Forum **36**(5), 151–163 (2017). https://doi.org/10.1111/cgf.13253. https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13253
16. Huang, R., Ren, J., Wonka, P., Ovsjanikov, M.: Consistent zoomout: efficient spectral map synchronization. Comput. Graph. Forum **39**(5), 265–278 (2020). https://doi.org/10.1111/cgf.14084
17. Kezurer, I., Kovalsky, S.Z., Basri, R., Lipman, Y.: Tight relaxation of quadratic matching. Comput. Graph. Forum **34**(5), 115–128 (2015). https://doi.org/10.1111/cgf.12701
18. Kim, V.G., Lipman, Y., Funkhouser, T.: Blended intrinsic maps. ACM Trans. Graph., 30 (4), (2011). https://doi.org/10.1145/2010324.1964974
19. Kovnatsky, A., Bronstein, M. M., Bronstein, A. M., Glashoff, K., Kimmel, R.: Coupled quasi-harmonic bases. Comput. Graph. Forum, 32 (2pt4): 439–448, 2013. https://doi.org/10.1111/cgf.12064. https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12064
20. Lang, I., Ginzburg, D., Avidan, S., Raviv, D.: Dpc: Unsupervised deep point correspondence via cross and self construction. 2021 International Conference on 3D Vision (3DV), pages 1442–1451, (2020). https://api.semanticscholar.org/CorpusID:239016800
21. Li, L., Donati, N., Ovsjanikov, M.: Learning multi-resolution functional maps with spectral attention for robust shape matching. Adv. Neural. Inf. Process. Syst. **35**, 29336–29349 (2022)
22. Lipman, Y., Funkhouser, T.: Möbius voting for surface correspondence. ACM Trans. Graph., 28 (3), (2009). ISSN 0730-0301. https://doi.org/10.1145/1531326.1531378
23. Litany, O., Remez, T., Rodola, E., Bronstein, A., Bronstein, M.: Deep functional maps: structured prediction for dense shape correspondence . In 2017 IEEE International Conference on Computer Vision (ICCV), pages 5660–5668, Los Alamitos, CA, USA, (Oct. 2017). IEEE Computer Society. https://doi.org/10.1109/ICCV.2017.603. https://doi.ieeecomputersociety.org/10.1109/ICCV.2017.603
24. Litman, R., Bronstein, A.M.: Learning spectral descriptors for deformable shape correspondence . IEEE Trans. Pattern Anal. Mach. Intell., 36 (01): 171–180, (2014). ISSN 1939-3539. https://doi.org/10.1109/TPAMI.2013.148. https://doi.ieeecomputersociety.org/10.1109/TPAMI.2013.148
25. Liu, H.-T.D., Jacobson, A., Ovsjanikov, M.: Spectral coarsening of geometric operators. ACM Trans. Graph., 38 (4), (2019). https://doi.org/10.1145/3306346.3322953
26. Magnet, R., Ovsjanikov, M.: Scalable and efficient functional map computations on dense meshes. Comput. Graph. Forum **42**(2), 89–101 (2023). https://doi.org/10.1111/cgf.14746. https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14746
27. Maron, H., Dym, N., Kezurer, I., Kovalsky, S., Lipman, Y.: Point registration via efficient convex relaxation. ACM Trans. Graph., 35 (4), (2016). ISSN 0730-0301. https://doi.org/10.1145/2897824.2925913
28. Melzi, S., Ren, J., Rodolà, E., Sharma, A., Wonka, P., Ovsjanikov, M.: Zoomout: spectral upsampling for efficient shape correspondence. ACM Trans. Graph., 38 (6), (2019). ISSN 0730-0301. https://doi.org/10.1145/3355089.3356524
29. Nogneng, D., Ovsjanikov, M.: Informative descriptor preservation via commutativity for shape matching. Comput. Graph. Forum **36**(2), 259–267 (2017). https://doi.org/10.1111/cgf.13124. (**ISSN 0167-7055**)
30. Ovsjanikov, M., Mérigot, Q., Mémoli, F., Guibas, L.: One point isometric matching with the heat kernel. Comput. Graph. Forum **29**(5), 1555–1564 (2010). https://doi.org/10.1111/j.1467-8659.2010.01764.x
31. 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), (2012). https://doi.org/10.1145/2185520.2185526
32. Ovsjanikov, M., Corman, E., Bronstein, M., Rodolà, E., Ben-Chen, M., Guibas, L., Chazal, F., Bronstein, A.: Computing and processing correspondences with functional maps. In ACM SIGGRAPH

2017 Courses, SIGGRAPH '17, New York, NY, USA, (2017). Association for Computing Machinery. ISBN 9781450350143. https://doi.org/10.1145/3084873.3084877
33. Ovsjanikovs, M.: Spectral methods for isometric shape matching and symmetry detection. Stanford University, (2011)
34. Ren, J., Poulenard, A., Wonka, P., Ovsjanikov, M.: Continuous and orientation-preserving correspondences via functional maps. ACM Trans. Graph., 37 (6), (2018). https://doi.org/10.1145/3272127.3275040
35. Ren, J., Melzi, S., Ovsjanikov, M., Wonka, P.: Maptree: recovering multiple solutions in the space of maps. ACM Trans. Graph., 39 (6), (2020). https://doi.org/10.1145/3414685.3417800
36. Ren, J., Melzi, S., Wonka, P., Ovsjanikov, M.: Discrete optimization for shape matching. Comput. Graph. Forum **40**(5), 81–96 (2021). https://doi.org/10.1111/cgf.14359. https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14359
37. Rodolà, E., Cosmo, L., Bronstein, M.M., Torsello, A., Cremers, D.: Partial functional correspondence. Comput. Graph. Forum **36**(1), 222–236 (2017). https://doi.org/10.1111/cgf.12797. https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12797
38. Rustamov, R.M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In A. Belyaev and M. Garland, editors, Geometry Processing. The Eurographics Association, (2007). ISBN 978-3-905673-46-3. https://doi.org/10.2312/SGP/SGP07/225-233
39. Sahillioğlu, Y.: A genetic isometric shape correspondence algorithm with adaptive sampling. ACM Trans. Graph., 37 (5), (2018). ISSN 0730-0301. https://doi.org/10.1145/3243593
40. Sahillioğlu, Y., Yemez, Y.: 3d shape correspondence by isometry-driven greedy optimization. 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pages 453–458, (2010). https://api.semanticscholar.org/CorpusID:6988633
41. Schmidt, R., Grimm, C., Wyvill, B.: Interactive decal compositing with discrete exponential maps. ACM Trans. Graph. **25**(3), 605–613 (2006). https://doi.org/10.1145/1141911.1141930. (**ISSN 0730-0301**)
42. Shamai, G., Kimmel, R.: Geodesic distance descriptors . In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 3624–3632, Los Alamitos, CA, USA, July (2017). IEEE Computer Society. https://doi.org/10.1109/CVPR.2017.386. https://doi.ieeecomputersociety.org/10.1109/CVPR.2017.386
43. Sharp, N., Attaiki, S., Crane, K., Ovsjanikov, M.: Diffusionnet: discretization agnostic learning on surfaces. ACM Trans. Graph. (TOG) **41**(3), 1–16 (2022)
44. Sun, J., Ovsjanikov, M., Guibas, L.: A Concise and provably informative multi-scale signature based on heat diffusion. Computer Graphics Forum, (2009). ISSN 1467-8659. https://doi.org/10.1111/j.1467-8659.2009.01515.x
45. Tam, G.K., Cheng, Z.-Q., Lai, Y.-K., Langbein, F., Liu, Y., Marshall, A.D., Martin, R., Sun, X., Rosin, P.: Registration of 3d point clouds and meshes: a survey from rigid to nonrigid. IEEE Trans. Visualization Comput. Graph. **19**(7), 1199–1217 (2013). https://doi.org/10.1109/TVCG.2012.310. (**ISSN 1077-2626**)
46. Tevs, A., Bokeloh, M., Wand, M., Schilling, A., Seidel, H.-P.: Isometric registration of ambiguous and partial data. 2009 IEEE Conference on Computer Vision and Pattern Recognition, pages 1185–1192, (2009). https://api.semanticscholar.org/CorpusID:15366827
47. Tombari, F., Salti, S., Di Stefano, L.: Unique signatures of histograms for local surface description. In Proceedings of the 11th European Conference on Computer Vision Conference on Computer Vision: Part III, ECCV'10, page 356–369, Berlin, Heidelberg, (2010). Springer-Verlag. ISBN 364215557X
48. van Kaick, O., Zhang, H., Hamarneh, G., Cohen-Or, D.: A survey on shape correspondence. Comput. Graph. Forum **30**(6), 1681–1707 (2011). https://doi.org/10.1111/j.1467-8659.2011.01884.x
49. Vestner, M., Litman, R., Rodolà, E., Bronstein, A. M., Cremers, D.: Product manifold filter: non-rigid shape correspondence via kernel density estimation in the product space. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 6681–6690, (2017). https://api.semanticscholar.org/CorpusID:10436114

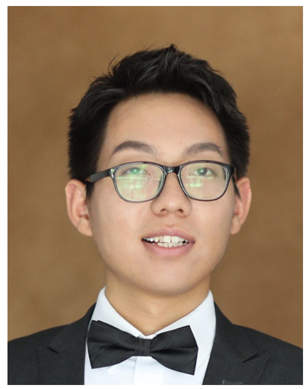

Haoliang Zhang is currently a Ph.D. student in Electrical and Computer Engineering, University of Oklahoma. His research interests include machine learning and discrete differential geometry.

Samuel Cheng received the B.S. degree in Electrical and Electronic Engineering from the University of Hong Kong, and the M.Phil. degree in Physics and the M.S. degree in Electrical Engineering from Hong Kong University of Science and Technology and the University of Hawaii, Honolulu, respectively. He received the Ph.D. degree in Electrical Engineering from Texas A&M University in 2004. He worked in Microsoft Asia, China, and Panasonic Technologies Company, New Jersey, in the areas of texture compression and digital watermarking during the summers of 2000 and 2001. In 2004, he joined Advanced Digital Imaging Research, a research company based in Houston, Texas, as a Research Engineer to perform biomedical imaging research and was promoted to Senior Research Engineer the next year. Since 2006, he joined the School of Electrical and Computer Engineering at the University of Oklahoma and is currently an associate professor and the director of OU-Tulsa Lab of Information and Image Processing. He has been awarded six US patents in miscellaneous areas of signal processing and is the co-recipient of the 2007 IEEE Signal Processing Magazine Best Paper Award. He is a senior member of IEEE and a member of ACM.

Christian El Amm is a Professor in the Department of Surgery at the University of Oklahoma. He specializes in plastic and reconstructive surgery, and pediatric plastic and craniofacial surgery. His areas of expertise include: cleft lip and palate, craniosynostosis, Pierre Robin sequence, midface hypoplasia, fat grafting and liposuction, rhinoplasty, secondary rhinoplasty, facial rejuvenation, body contouring, and breast augmentation or reconstruction. He has practiced at OU Health since 2005 and has received recognition by Castle Connolly Medical Ltd. as a top doctor in 2016 and 2017.