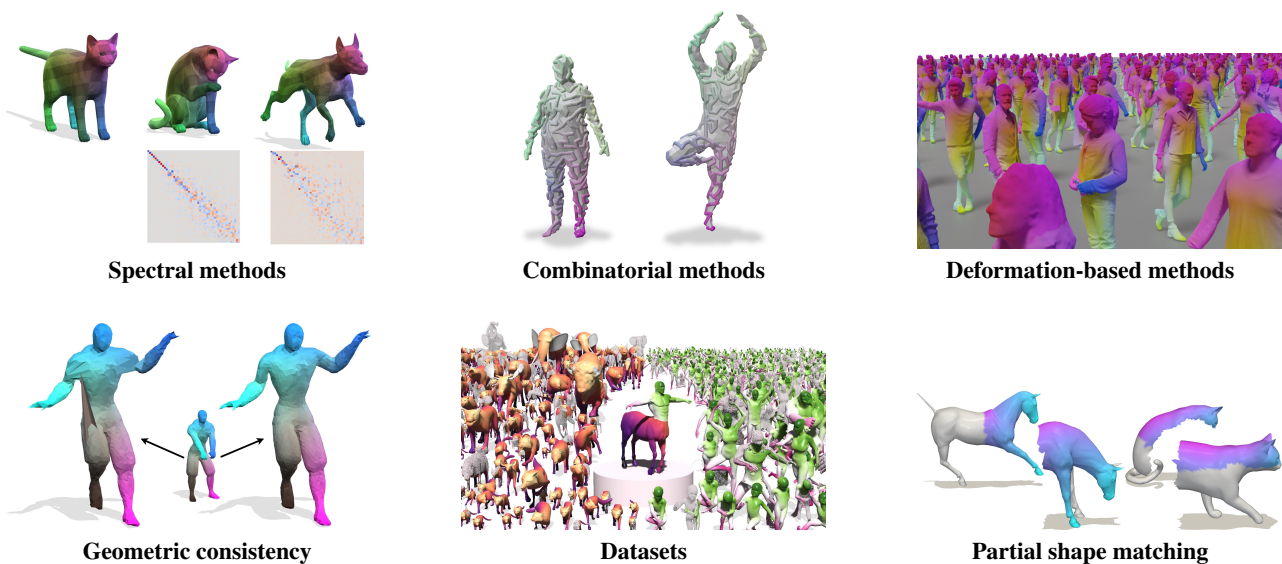


# Non-Rigid 3D Shape Correspondences: From Foundations to Open Challenges and Opportunities

A. Zhuravlev<sup>1\*</sup>  L. Bastian<sup>6,7,8\*</sup>  D. Cao<sup>2,3</sup>  N. El Amrani<sup>2,3</sup>  P. Roetzer<sup>2,3</sup>  V. Ehm<sup>7,8</sup>  R. Marin<sup>7,8</sup>  H. Nishizawa<sup>4</sup>   
S. Morishima<sup>5</sup>  C. Theobalt<sup>1</sup>  N. Navab<sup>7,8</sup>  D. Cremers<sup>7,8</sup>  F. Bernard<sup>2,3</sup>  Z. Lähler<sup>2,3</sup>  V. Golyanik<sup>1</sup> 

<sup>1</sup>MPI for Informatics <sup>2</sup>University of Bonn <sup>3</sup>Lamarr Institute <sup>4</sup>Waseda University <sup>5</sup>Waseda Research Institute for Science and Engineering  
<sup>6</sup>Imperial College London <sup>7</sup>Technical University of Munich <sup>8</sup>Munich Center for Machine Learning \*Equal Contribution



**Figure 1:** This STAR surveys modern approaches to non-rigid 3D shape correspondence, spanning spectral, combinatorial, and deformation-based methods. We further highlight emerging opportunities, such as zero-shot matching and open challenges involving partiality. Image sources: [EEAX\*25, MCPM24, CRB24].

## Abstract

Estimating correspondences between deformed shape instances is a long-standing problem in computer graphics; numerous applications, from texture transfer to statistical modelling, rely on recovering an accurate correspondence map. Many methods have thus been proposed to tackle this challenging problem from varying perspectives, depending on the downstream application. This state-of-the-art report is geared towards researchers, practitioners, and students seeking to understand recent trends and advances in the field. We categorise developments into three paradigms: spectral methods based on functional maps, combinatorial formulations that impose discrete constraints, and deformation-based methods that directly recover a global alignment. Each school of thought offers different advantages and disadvantages, which we discuss throughout the report. Meanwhile, we highlight the latest developments in each area and suggest new potential research directions. Finally, we provide an overview of emerging challenges and opportunities in this growing field, including the recent use of vision foundation models for zero-shot correspondence and the particularly challenging task of matching partial shapes.

## CCS Concepts

• **Computing methodologies** → **Shape analysis; Mesh geometry models; Neural networks;**

## 1. Introduction

From character expression transfer to the study of variability in human anatomy, non-rigid shape correspondence is the cornerstone of numerous applications. It is an inescapable step when transferring or comparing information between 3D shapes, from texture to deformation, and segmentation to motion. At scale, correspondences across shape collections enable parametric shape models of various types, such as bodies, faces, hands, and organs [BRLB14, LMR\*15, RTB17, PCG\*19]. These models are indispensable for numerous downstream tasks: animation pipelines [IMH05], 3D reconstruction systems [LMR\*15], and morphometric analysis of organs [EACB24] are all enabled by parametric models whose power comes from accurate pointwise correspondences. Developing non-rigid shape correspondence methods is, therefore, of fundamental importance and an active area of ongoing research.

But such an important task does not come without challenges. Data scarcity is a hindrance to development in the field, as geometric 3D acquisitions require specialised technologies, modelling is generally done by expert artists, and annotations necessitate extensive manual work, which can be too ambiguous for shape collections with large variations. In the rigid case, a correspondence model can be reduced to a simplified parametrisation with six degrees of freedom [BM92], but this does not hold for *non-rigid deformations*, which are often too complex to obtain an analytical formulation [VLB\*17].

Even the widely studied near-isometric problem setting (e.g., humans simulated in different poses) still requires a deep understanding of geometrical structures to, for instance, distinguish intrinsic symmetries [BRLB14, WWEAB25]. As structural differences grow, such as between animals of different species, so does the semantic ambiguity in the matching (how can one map the neck of a giraffe to that of a dog?) [ZKJB17]. Moving to real-world acquisitions, issues such as noise, holes, and partiality are common and severely degrade the performance even of the best available methods. Such issues continue to motivate scholars to develop robust, widely applicable shape correspondence tools.

Despite the longstanding nature of the shape correspondence problem, recent years have seen continued progress, even in classical settings such as matching two geometrically complete and clean meshes. However, researchers have also explored new relevant directions such as extreme non-isometry, zero-shot matching, partiality, and topological changes. These challenges inspire new paradigms, such as the application of *foundation model* features in 3D geometric contexts. Methods for *partial-to-partial* and *multi-shape matching* address more complex correspondence settings beyond the widely considered *full-to-full* pairwise setting.

Finally, the progress in non-rigid shape correspondence fields owes much to the parallel evolution of benchmarks. Early evaluations focused on synthetic near-isometric shape pairs with uniform meshing, while recent ones have shifted attention toward non-isometry, partiality, real-world noise, and irregular meshing, progressively closing the gap between laboratory and real-world settings. The scope of shape correspondence is as wide as ever, with exciting application domains for interdisciplinary approaches.

### 1.1. Scope of this STAR

This state-of-the-art report is motivated by recent developments in shape correspondence, both in advancing classical foundations and introducing new paradigms. Our primary goals are to: (i) summarise the main directions of current research in shape correspondence, and (ii) highlight open problems and future avenues of research. Our review mainly focuses on works published from 2022 until 2025, including earlier works for background and historical context.

Our report primarily focuses on correspondence problems for 3D meshes, which remain the dominant representation in computer graphics. Meshes incorporate rich geometric structure into their surface representation, which has inspired numerous innovative methods that leverage the wealth of knowledge in geometry processing and numerical optimisation. While point clouds lack such structural information, a review of shape correspondence would be incomplete without mentioning several recent influential approaches as well. Registration of neural implicit representations [PFS\*19, MON\*19] is out of our scope in this paper. Moreover, we do not consider rigid correspondence, as the complexity is vastly reduced when considering transformation models with only six degrees of freedom. Image-based feature matching is also beyond the scope of this report; we refer the readers to [MJF\*21] for an overview.

### 1.2. Related surveys

Previous surveys on shape correspondence have taken various perspectives. [Sah20] covers methods published between 2011 and 2019, many of which provide the foundation for approaches discussed in our report. The more recent survey reviews correspondence methods up to 2021, with an emphasis on techniques applicable for dynamic shape acquisition and reconstruction [DYDZ22]. We build on both, extend the timeline, introduce novel paradigms, and cover directions that were not fully addressed, such as combinatorial methods. [VKZHCO11] summarises shape correspondence methods developed before 2011. An earlier review on functional maps also offers an excellent theoretical primer [OCB\*16], but focuses on pre-2016 methods and thus excludes the deep learning approaches that dominate the recent literature.

### 1.3. Structure

This STAR is organised as follows. In Sec. 2, we introduce the mathematical foundations and problem setting of non-rigid 3D shape correspondence. The recent body of methods is primarily covered in the next three sections: Sec. 3 reviews spectral methods, built on the functional map framework; Sec. 4 covers combinatorial approaches that formulate correspondence as a discrete optimisation problem. Sec. 5 focuses on deformation-based methods that align shapes through spatial transformations.

Following the main body of works, we discuss recent emerging trends that expand the scope of correspondence beyond classical settings in Sec. 6. Suitable datasets and benchmarks, including recently published ones, are summarised in Sec. 7. Finally, we outline practical applications of shape correspondence in Sec. 8, highlight open challenges and discuss future research directions in Sec. 9.

Symbol	Domain	Description
$\mathcal{M}, \mathcal{N}$		3D shapes with $n_{\mathcal{M}}$ and $n_{\mathcal{N}}$ vertices
$X_{\mathcal{M}}$	$\mathbb{R}^{n_{\mathcal{M}} \times 3}$	Vertex positions of $\mathcal{M}$
$E_{\mathcal{M}}$	$\subseteq \{1, \dots, n_{\mathcal{M}}\}^2$	Edges of $\mathcal{M}$
$T_{\mathcal{M}}$	$\subseteq \{1, \dots, n_{\mathcal{M}}\}^3$	Faces (triangles) of $\mathcal{M}$
$f_{\mathcal{M}}$	$\mathbb{R}^{n_{\mathcal{M}} \times d}$	Vertex-wise features on $\mathcal{M}$
$L_{\mathcal{M}}$	$\mathbb{R}^{n_{\mathcal{M}} \times n_{\mathcal{M}}}$	Discrete Laplace–Beltrami operator on $\mathcal{M}$
$\phi_{\mathcal{M}}$	$\mathbb{R}^{n_{\mathcal{M}} \times k}$	Eigenfunctions of $L_{\mathcal{M}}$ (spectral basis)
$C_{\mathcal{M}\mathcal{N}}$	$\mathbb{R}^{k \times k}$	Functional map from $\mathcal{M}$ to $\mathcal{N}$
$\Pi_{\mathcal{M}\mathcal{N}}$	$\{0, 1\}^{n_{\mathcal{M}} \times n_{\mathcal{N}}}$	Point-wise map from $\mathcal{M}$ to $\mathcal{N}$
$\mathcal{E}(\cdot)$		Energy or cost function to be minimized
$\mathcal{L}(\cdot)$		Loss for neural network training

**Table 1:** Summary of notation used throughout the report.


## 2. Background & fundamentals

This section provides an overview of the shape correspondence problem setting. We begin by discussing how surfaces can be characterised locally using *feature descriptors*, which is essential for estimating correspondence between shapes. These descriptors serve as an initial, local similarity measure but are often insufficient on their own due to ambiguity and limited discriminative power. Additional *structural priors* can be introduced to constrain the solution space and improve the faithfulness of the recovered correspondences. These priors provide the conceptual foundation for the three methodological categories discussed in later sections. However, no type of structural prior has emerged as a silver bullet; each has its *limitations* and is best suited to specific settings. The notation we adopt throughout this report is summarised in Table 1.


### 2.1. The shape correspondence problem

In the context of 3D geometry processing, a shape is typically defined as a two-dimensional surface embedded in three-dimensional space. Such a surface is often discretised as a triangular mesh, consisting of vertices  $X_{\mathcal{M}}$ , edges  $E_{\mathcal{M}}$ , and faces  $T_{\mathcal{M}}$ . When a surface triangulation is not given, the vertices  $X_{\mathcal{M}}$  can be considered as an unstructured point cloud. Given two shapes, denoted  $\mathcal{M}$  and  $\mathcal{N}$ , the goal of shape correspondence is to determine which points on  $\mathcal{M}$  match semantically or geometrically to points on  $\mathcal{N}$ . This relationship can be expressed as a point-wise map  $\Pi_{\mathcal{M}\mathcal{N}} : \mathcal{M} \rightarrow \mathcal{N}$ , assigning to each point on the source shape a corresponding point on the target shape.

Throughout the report, we use the terms *correspondence* and *matching* interchangeably. We focus primarily on the dense correspondences setting, which establishes a matching for every point on the source shape to the target shape (formally, we consider  $n_{\mathcal{M}} \leq n_{\mathcal{N}}$  with  $\Pi_{\mathcal{M}\mathcal{N}}$  encoding an injective map, i.e.,  $\Pi_{\mathcal{M}\mathcal{N}}\mathbf{1} = \mathbf{1}$  and  $\Pi_{\mathcal{M}\mathcal{N}}^T\mathbf{1} \leq \mathbf{1}$ , where we denote the all-ones vector as  $\mathbf{1}$ ). Sparse, landmark-based approaches can also be useful in certain applications, specifically as they can be annotated more easily than dense maps (see Sec. 4). The deformation-based family of works also considers the registration settings, where a deformation field is recovered in the ambient domain without first matching points locally (see Sec. 5). When two shapes are aligned in this manner, dense correspondence maps can typically be recovered via a heuristic.



(a) Feature functions on source and target shape.



(b) Visualised correspondences.

**Figure 2: Visualising and matching features between shapes.** (a) Per-vertex feature functions are computed on two non-rigidly deformed shapes visualised by colour. When these features are sufficiently descriptive (as evidenced by similar regions coloured identically), a correspondence mapping can be derived. (b) Dense correspondences (visualised by transferring colours from one to the other shape and by connecting lines) computed between two non-rigidly deformed shapes, demonstrating geometric and semantic alignment. Image source: [GRE\*23].

### 2.2. Characterising a shape’s surface

A central goal in correspondence estimation is to quantify the similarity between shapes in a tractable form. Given shapes  $\mathcal{M}$  and  $\mathcal{N}$ , per-vertex feature functions  $f_{\mathcal{M}} \in \mathbb{R}^{n_{\mathcal{M}} \times d}$  and  $f_{\mathcal{N}} \in \mathbb{R}^{n_{\mathcal{N}} \times d}$ , also referred to as descriptors, can be specified based on a heuristic or estimated in a data-driven manner. These descriptors are intended to encode geometric or semantic properties of the shapes, see Fig. 2a for an illustration. Then, for any point on shape  $\mathcal{M}$  and a ground-truth point-wise map  $\Pi_{\mathcal{M}\mathcal{N}} \in \{0, 1\}^{n_{\mathcal{M}} \times n_{\mathcal{N}}}$ , the descriptor preservation principle assumes that the corresponding point on  $\mathcal{N}$  should have similar descriptors:

$$f_{\mathcal{M}} \approx \Pi_{\mathcal{M}\mathcal{N}} f_{\mathcal{N}}. \quad (1)$$

Several types of descriptors have been proposed for both structured and unstructured shape representations. They can be broadly categorised by their complexity and level of abstraction.

The most basic descriptors are derived directly from *geometric properties* of the shapes themselves. These include both extrinsic properties, such as 3D vertex coordinates [TSDS10] or surface normals [STDS14], and intrinsic properties, such as local curvature [MDSB03, WSSC11]. While the resulting *hand-crafted* descriptors capture local geometric information, they each have in-

herent limitations. Extrinsic descriptors such as SHOT [TSDS10] characterise surfaces through local coordinate systems but are not isometrically invariant. Conversely, intrinsic descriptors such as the Heat and Wave Kernel Signatures [SOG09, ASC11] naturally inherit the Laplace-Beltrami operators' isometric invariance but are not symmetry aware, can change drastically under non-isometric deformations, and remain sensitive to noise. Despite these limitations, these descriptors can still yield effective matches when combined with strong geometric constraints on the point-wise map [WSSC11]. Moreover, they are easily applicable in various domains, in contrast to the more discriminative learned descriptors.

Recent approaches increasingly rely on *learning* geometric descriptors by optimising for correspondence on collections of shapes [LRR\*17]. These models can transform raw geometric inputs (e.g. vertex positions) or axiomatic descriptors (e.g. HKS, WKS, SHOT) to more expressive, domain-specific features [SACO22, SDB\*20]. Learned descriptors can overcome many of the limitations of hand-crafted features and adapt to challenging non-isometric or partial matching scenarios. In practice, however, learning sufficiently discriminative feature descriptors remains challenging. Surface ambiguities, noise, symmetries, shape partiality, and discretisation artefacts can yield descriptors with numerous plausible matching candidates on the target shape.

While obtaining informative descriptors is essential for solving the correspondence problem, recovering a point-wise map (visualised in Fig. 2b) is non-trivial even from descriptors learned by state-of-the-art models. A naive approach to obtain a correspondence map would be to match each point on shape  $\mathcal{M}$  to the most similar descriptor on shape  $\mathcal{N}$  via nearest-neighbour search:

$$\arg \min_{\Pi_{\mathcal{M}\mathcal{N}} \in \{0,1\}^{n_{\mathcal{M}} \times n_{\mathcal{N}}}} \|f_{\mathcal{M}} - \Pi_{\mathcal{M}\mathcal{N}} f_{\mathcal{N}}\|_{\tilde{F}}^2. \quad (2)$$

However, this formulation imposes no constraints on the resulting map. Inaccuracies or ambiguities in the feature descriptors representation can lead to errancies: for example, neighbouring points on  $\mathcal{M}$  may be mapped to distant locations on  $\mathcal{N}$ , yielding an undesirable *geometric inconsistency*. To overcome this, many recent methods impose additional assumptions or structural priors to obtain more accurate correspondences.

### 2.3. Methodological Overview

Solving the matching problem in Eq. (2), while straightforward, typically yields point-wise maps with errant feature pairings. To address this, correspondence optimisation is typically augmented with prior knowledge about the structure and geometry of the shapes, constraining solutions to more plausible point-wise maps. We group recent developments into three main categories:

**Spectral methods.** A common assumption is that the desired map is smooth and can be compactly represented in a low-frequency basis. The resulting insight that such a smooth map can be represented in a low-frequency function space has led to the rise of the seminal *functional maps* framework [OBCS\*12]. Most commonly used to construct this function space are the first  $k$  eigenfunctions of the Laplace-Beltrami operator (LBO), which form an orthonormal basis that captures the intrinsic geometry of the surface and

remains invariant under isometric deformations [BN03]. The key insight of the functional map framework is that instead of directly matching points between shapes, one can align the eigenfunctions  $\phi = [\phi_1, \phi_2, \dots, \phi_k] \in \mathbb{R}^{n \times k}$  for a pair of shapes  $\mathcal{M}$  and  $\mathcal{N}$  in spectral space through a low-rank linear map  $C_{\mathcal{M}\mathcal{N}} \in \mathbb{R}^{k \times k}$ . Given eigenfunctions  $\phi_{\mathcal{M}}$  and  $\phi_{\mathcal{N}}$  on shapes  $\mathcal{M}$  and  $\mathcal{N}$ , and a pointwise correspondence  $\Pi_{\mathcal{N}\mathcal{M}}$ , the relationship between them is captured as:

$$\Pi_{\mathcal{N}\mathcal{M}} \phi_{\mathcal{M}} \approx \phi_{\mathcal{N}} C_{\mathcal{M}\mathcal{N}}. \quad (3)$$

Intuitively, this aligns the spectral bases between shapes by transferring an indicator function from shape  $\mathcal{M}$  through the point-wise map, yielding a function on  $\mathcal{N}$  whose spectral coefficients are determined by the functional map  $C_{\mathcal{M}\mathcal{N}}$ .

The LBO eigenfunctions are famously invariant under non-isometric deformation, and robust even under mild deviations from isometry [RWP06]. Numerous recent works leverage the insight that shape descriptors, whether hand-crafted or learned, can be projected onto the basis functions to establish correspondences in spectral space. The functional map  $C$  naturally regularises the point-wise map due to its low dimensionality and the smoothness of the LBO eigenfunctions. Structural priors have also been derived for  $C_{\mathcal{M}\mathcal{N}}$ , encouraging desirable properties such as promoting approximately isometric mappings or bijectivity [RSO19]. The use of such geometric priors has led to unsupervised paradigms that enable state-of-the-art pipelines to learn robust and discriminative feature descriptors from collections of shapes [CRB23].

However, the low-rank spectral representation inherently compresses the recovered map, limiting the ability to represent high-frequency geometric details. Moreover, the quality of the LBO eigenfunctions degrades in the presence of non-isometries [BXNL24] and holes or partiality [EEAX\*25], and extracting a faithful point-wise map from a functional map remains a non-trivial step that can introduce additional error [VM23]. These trade-offs are discussed in detail in Sec. 3.

**Combinatorial methods.** Solving the naive nearest-neighbour matching in Eq. (2) yields an initial point-wise map, but imposes no structure on the solution. A natural extension is the Linear Assignment Problem (LAP) which considers recovering a bijective mapping (i.e.,  $n_{\mathcal{M}} = n_{\mathcal{N}}$ , obtainable through re-meshing or sub-sampling) that minimises a linear dissimilarity cost  $D \in \mathbb{R}^{n_{\mathcal{M}} \times n_{\mathcal{N}}}$  over the set of permutation matrices  $\mathbb{P}$ :

$$\arg \min_{\Pi \in \mathbb{P}} \langle D, \Pi \rangle \quad \text{where, e.g.} \quad D_{ij} = \|f_{\mathcal{M}}(i) - f_{\mathcal{N}}(j)\|^2. \quad (4)$$

This formulation has been extensively studied and can be solved efficiently through the Auction algorithm [Ber92]. Nevertheless, it still relies entirely on feature similarity and neglects neighbourhood relationships between features. This can lead to *geometrically inconsistent* matchings where neighbouring points on  $\mathcal{M}$  map to distant locations on  $\mathcal{N}$ .

On the other hand, the Quadratic Assignment Problem (QAP) incorporates a metric that is preserved between matching pairs of points (geodesic distances are commonly chosen) but is famously NP-hard [RPW94], making solutions computationally intractable for most meshes of practical size. Many recent combinatorial shape

matching methods, therefore, seek a middle ground, incorporating geodesic constraints by exploiting the surface’s discrete structure while avoiding the full quadratic cost. Recent approaches have improved scalability through convex relaxations [KMDL19, DLB\*23], proposed heuristics [VLB\*17, HFH\*22, RSCB22], or additional constraints like neighbourhood preservation [WSSC11, ERE\*24, RB24].

Despite these scalability challenges, the combinatorial family of methods discussed in Sec. 4 has recently gained traction, as they allow enforcing that solutions have certain geometric properties by constraining the feasible set. Beyond neighbourhood preservation, this also extends to cycle consistency across shape collections [BTL\*23, KSH\*25]. However, combinatorial methods require reasonably discriminative feature representations as input and are typically non-differentiable, thus serving as a complement rather than a replacement to learned methods.

**Deformation-based methods.** This class of methods directly recovers a global deformation in the ambient embedding space without first decomposing the problem into one of matching local elements. In general, one may seek to recover a global deformation field  $\mathcal{F} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$  that aligns the source shape  $\mathcal{M}$  to the target shape  $\mathcal{N}$ :

$$\mathcal{F}(X_{\mathcal{M}}) \approx \Pi_{\mathcal{M}\mathcal{N}} X_{\mathcal{N}}. \quad (5)$$

Pioneered by iterative closest point (ICP) [BM92] and coherent point drift (CPD) [MS10], this class of methods remains widely applicable in settings where obtaining discriminative feature descriptors is challenging. As the general formulation is highly ill-posed, the deformation field  $\mathcal{F}$  is typically regularised with priors that induce spatial smoothness, or require that it is locally rigid. Recent methods have emphasised more robust regularisation techniques [YDXZ23, YDHZ25], parametrising the deformation with neural networks [LH22b, SDM\*24], and the incorporation of parametric templates into the optimisation or learning problem [FKL\*23, LFC\*25].

The deformation-based methods discussed in Sec. 5 operate in the ambient Euclidean space and are thus broadly applicable to various 3D data representations, including unstructured point clouds, without requiring manifold connectivity or consistent topology. This extrinsic paradigm is particularly suited to settings where noise, topological changes, or incomplete data pose challenges for methods requiring clean surface geometry. Recent works have further extended the framework with robust norms [YDXZ23], parametric templates [FKL\*23, LFC\*25], and learned deformation priors [LH22b, MCPM24]. However, the joint recovery of correspondence and deformation is inherently ill-posed, and alternating optimisation schemes inherited from ICP [BM92] remain sensitive to initialisation and prone to local minima.

## 2.4. Challenges across matching scenarios

Choosing a suitable correspondence method for a given application largely depends on the characteristics of the shapes being matched. As shape correspondence methods are applied to an increasingly wide range of geometric inputs, no single approach has remained universally applicable. In practice, the assumptions underlying a

given method may hold only approximately, or be entirely violated, depending on the specific matching scenario. Therefore, it is essential to consider the nature of the shapes, particularly their degree of similarity, partiality, and the availability of training data (with or without annotations), when deciding on the appropriate methodology.

A key factor is the nature of the deformation. When the shapes are represented as water-tight meshes and near-isometric, as in cases of simulated pose variation or mild articulation, most structural and spectral assumptions of matching methods remain valid. In contrast, non-isometric scenarios (e.g., matching across different animal species) are more challenging. As geometric similarity decreases, the availability of sufficiently discriminative features decreases as well, requiring stronger constraints in correspondence optimisation, such as strictly preserving local neighbourhoods. Especially difficult non-isometric cases occur when shapes deviate from idealised, clean, and manifold conditions. Real-world 3D data often contains topological noise, including non-uniform sampling, small holes, or self-intersections. Such violations of smooth manifold structure can compromise correspondence methods that depend on local geometric information, resulting in unstable or inconsistent matches between shapes.

A further challenge arises in the case of *partiality*, where one or both shapes are incomplete. In full-to-full matching, it is generally assumed that every point of one shape has a corresponding partner on the other. In the case of partial shape matching, this assumption is often grossly violated. In the partial-to-full matching scenario, a mapping can be defined from the partial shape to the full one, but not the other way round, since not every region of the full shape is observed in the partial shape. This becomes even more challenging in the partial-to-partial case, where both shapes are incomplete. Here, it is necessary not only to establish correspondences but also to identify the overlapping region between the two shapes. These difficulties demand the development of more sophisticated matching strategies, which we emphasise in Section 6.4 as the problem is largely unresolved.


## 2.5. Correspondence evaluation

Once a correspondence map  $\Pi_{\mathcal{M}\mathcal{N}} \in \{0, 1\}^{n_{\mathcal{M}} \times n_{\mathcal{N}}}$  between a pair of shapes is recovered, it can be evaluated through various metrics.

**Accuracy.** When working with standard synthetic datasets, ground-truth correspondence matrices  $\Pi_{\mathcal{M}\mathcal{N}}^*$  are generally available. This allows one to directly evaluate the accuracy of the predicted map  $\Pi_{\mathcal{M}\mathcal{N}}$ . Let  $D_{\mathcal{N}} \in \mathbb{R}^{n_{\mathcal{N}} \times n_{\mathcal{N}}}$  be the pairwise geodesic distance matrix on the target shape  $\mathcal{N}$ , and  $\text{Area}(\mathcal{N})$  the total surface area of the target shape. Then, the vector of normalised geodesic errors over all vertices in  $\mathcal{M}$  is defined as [KLF11]:

$$e = \frac{1}{\sqrt{\text{Area}(\mathcal{N})}} \cdot \text{diag}\left(\Pi_{\mathcal{M}\mathcal{N}} D_{\mathcal{N}} \Pi_{\mathcal{M}\mathcal{N}}^{*\top}\right). \quad (6)$$

Here,  $\Pi_{\mathcal{M}\mathcal{N}} D_{\mathcal{N}} \Pi_{\mathcal{M}\mathcal{N}}^{*\top} \in \mathbb{R}^{n_{\mathcal{M}} \times n_{\mathcal{M}}}$  contains the geodesic distances between the predicted and ground-truth correspondences for each source vertex, and  $\text{diag}(\cdot)$  extracts diagonal entries.



**Figure 3:** Qualitative visualisation of *correspondence quality* using different transfer strategies. (a) Color transfer provides an overview of the global correspondence structure. (b) High-frequency texture transfer can reveal small local misalignments. (c) Deformation transfer can similarly visualise small topological inconsistencies caused by inaccurate correspondences. Figures adapted from [ZLG25, VOM25].

From the vector of normalised geodesic errors  $e \in \mathbb{R}^{n_{\mathcal{M}}}$ , three standard evaluation metrics are typically computed. *Mean geodesic error (MGE)* over all source vertices is defined as:

$$\text{MGE} = \frac{1}{n_{\mathcal{M}}} \sum_{i=1}^{n_{\mathcal{M}}} e(i). \quad (7)$$

*Percentage of correct keypoints (PCK)* measures the fraction of correspondences whose normalised geodesic error is below a threshold  $\tau > 0$ . Using  $\mathbf{1}[\cdot]$  as the indicator function,

$$\text{PCK}(\tau) = \frac{1}{n_{\mathcal{M}}} \sum_{i=1}^{n_{\mathcal{M}}} \mathbf{1}[e(i) < \tau]. \quad (8)$$

*Area under the PCK curve (AUC)* captures the overall accuracy of the predicted correspondences for varying thresholds, typically up to the maximum threshold  $\tau_{\max}$ . This is defined as:

$$\text{AUC}(\tau_{\max}) = \frac{1}{\tau_{\max}} \int_0^{\tau_{\max}} \text{PCK}(\tau) d\tau. \quad (9)$$

**Smoothness.** Smoothness metrics assess whether neighbouring surface elements on the source and target shapes are consistently mapped under the predicted correspondence. *Geodesic distortion* [Ren21] measures how well intrinsic distances are preserved by comparing pairwise geodesic distances under the correspondence:

$$\sum_{i,j} (D_{\mathcal{M}}(i,j) - D_{\mathcal{N}}(\Pi_{\mathcal{M}\mathcal{N}}(i), \Pi_{\mathcal{M}\mathcal{N}}(j)))^2.$$

The *Dirichlet energy* [MRSO22, ESBC19] can be used to evaluate the smoothness of the induced map, essentially penalising large differences between the mapped positions of neighbouring vertices. *Conformal distortion* [ESBC19, RMOW20] quantifies deviations from local angle preservation, indicating the extent to which small surface elements are stretched or sheared under the mapping. Lower values of these metrics indicate smoother and more geometrically consistent correspondences, which has important ramifications for downstream applications (see Sec. 9).

**Coverage.** Coverage metrics evaluate how completely the target shape is reached by the correspondence and are often used in conjunction with smoothness. A commonly used criterion is the *un-coverage rate*, defined as the fraction of target vertices (or surface area) that are not mapped to by any source vertex under the correspondence. Low un-coverage means that the correspondence utilises the target shape more uniformly, while high un-coverage indicates a collapsing or highly non-injective mapping.

**Bijectivity.** Bijectivity metrics assess whether a correspondence behaves approximately as a one-to-one mapping. This can be evaluated by *inverse consistency*, which computes correspondences in both directions,  $\Pi_{\mathcal{M}\mathcal{N}}$  and  $\Pi_{\mathcal{N}\mathcal{M}}$ , by applying the same model with swapped inputs. The forward and backward maps are then composed (i.e.,  $\Pi_{\mathcal{N}\mathcal{M}} \circ \Pi_{\mathcal{M}\mathcal{N}}$  and  $\Pi_{\mathcal{M}\mathcal{N}} \circ \Pi_{\mathcal{N}\mathcal{M}}$ ), and the average geodesic distance between the composed maps and the identity map  $I$  is measured. Low values indicate approximately bijective and mutually consistent correspondences. In addition to inverse-consistency, bijectivity can be further assessed using *orientation preservation* criteria, such as the triangle inversion rate, which measures the fraction of faces whose orientation is reversed after mapping and indicates locally inconsistent correspondences.

**Visualisation.** In addition to quantitative metrics, correspondences are commonly assessed qualitatively through visualisation. A standard approach is *colour transfer* between shapes, which provides an illustration of the overall correspondence structure but may fail to reveal subtle local errors. More diagnostic visualisations can be obtained by transferring *high-frequency textures*, which are much more sensitive to small-scale errors. Similarly, *deformation transfer* [SP04, VOM25] assesses correspondence quality by propagating the deformation field defined on a source shape to the target, making local mismatches appear as visible topological artefacts.

**Accuracy to downstream application.** While the metrics above provide a comprehensive quantitative assessment of correspondence quality, many methods in the literature optimise primarily for geodesic accuracy or other quantifiable metrics. In practice, however, downstream applications may have additional requirements

that pure vertex-to-vertex or spectral maps may not satisfy. For instance, deformation transfer [SP04] and topology transfer require correspondences that are not only accurate but also smooth and bijective; without such guarantees, transferred deformations can exhibit artefacts (see Fig. 3c). Similarly, statistical shape modelling demands geometrically consistent maps across shape collections to produce meaningful modes of variation [LMR\*15, MBT\*23]. In domain-specific settings such as medical shape analysis, purely geometric descriptors may be insufficient or prohibitively challenging to obtain. Incorporating prior knowledge from domain experts, such as anatomical landmarks or physiological constraints, can thus be essential for incorporating such methods into clinical pipelines [EACB24, BRW23]. These application-specific requirements highlight that correspondence accuracy, while necessary, can be insufficient on its own, motivating the development of methods that jointly account for the demands of the intended downstream application. We revisit this gap between benchmark performance and practical applicability in Sec. 9.

### 3. Spectral methods

Spectral methods have emerged as a powerful framework, relying on intrinsic properties of the shapes that are invariant to isometric deformations. Building on this spectral representation, the functional map framework offers a compact and elegant formulation for expressing correspondences between shapes in the spectral domain, reducing the high-dimensional point-wise matching problem to the estimation of a small matrix of coefficients based on frequency decomposition. In this section, we first provide a theoretical background on the Laplace–Beltrami operator (LBO) and functional maps, later discussing recent classical refinement methods as well as learning-based methods.

#### 3.1. Background

In this subsection, we review the theoretical background for spectral methods.

**Laplace–Beltrami operator.** The Laplace–Beltrami operator is a generalisation of the Laplace operator to Riemannian manifolds [Gug12]. It is defined as the divergence of the gradient of a function  $f$ :


$$\Delta f = \operatorname{div}(\nabla f). \quad (10)$$

There exist many different ways to compute the LBO in the discrete setting [WMKG07]. For triangle meshes, the most common discretisation is based on the cotangent Laplacian matrix  $L_{\mathcal{M}} \in \mathbb{R}^{n_{\mathcal{M}} \times n_{\mathcal{M}}}$ , computed as [PP93]:

$$L_{\mathcal{M}} = A_{\mathcal{M}}^{-1} W_{\mathcal{M}}. \quad (11)$$

Here  $A_{\mathcal{M}}$  is the diagonal matrix of lumped area elements and  $W_{\mathcal{M}}$  is the cotangent weight matrix, defined in the following form:

$$W_{ij} = \begin{cases} \frac{\cot(\alpha_{ij}) + \cot(\beta_{ij})}{2} & \text{if } (i, j) \in \mathcal{E} \\ -\sum_{k \in N(i)} W_{ik} & \text{if } i = j \\ 0 & \text{otherwise,} \end{cases} \quad (12)$$



**Figure 4:** Illustration of the notation used for computing the cotangent weight matrix  $W_{ij}$  and the diagonal area matrix  $A_{ij}$ . The cotangent weight matrix  $W_{ij}$  uses the angles opposite the edge between  $v_i$  and  $v_j$ ;  $T_1$  and  $T_2$  denote the areas of the corresponding triangles. The diagonal area matrix  $A_{ij}$  assigns each vertex one-third of the total area of its adjacent triangles.

$$A_{ij} = \begin{cases} \sum_{k \in N(i)} \frac{T_k}{3} & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

Here  $T_k$  is the area of triangle  $k$  and  $N(i)$  denotes the index of the one-ring neighbour vertices. The notations in Eq. (12) and Eq. (13) are visualized in Fig. 4.

The eigendecomposition of the LBO plays a crucial role in spectral shape analysis. Let  $\Delta_{\mathcal{M}}$  be the LBO on shape  $\mathcal{M}$  and

$$\Delta_{\mathcal{M}} \phi_i = \lambda_i \phi_i, \quad (14)$$

where  $\phi_i \in L^2(\mathcal{M})$  is the  $i$ -th LBO eigenfunction with corresponding eigenvalue  $\lambda_i$ . In the discrete setting, the eigendecomposition can be computed via the generalised eigenvalue problem,

$$W_{\mathcal{M}} \phi_i = \lambda_i A_{\mathcal{M}} \phi_i. \quad (15)$$

Since the LBO is positive semi-definite, its eigenvalues are real and non-negative,

$$0 = \lambda_0 \leq \lambda_1 \leq \dots \lambda_{n_{\mathcal{M}}}. \quad (16)$$

The eigenfunctions of the Laplace–Beltrami operator encode the intrinsic geometry of a surface, their eigenvalues correspond to the frequency, and thus provide a natural basis for representing functions defined on it as well as for frequency filters. This spectral representation serves as the foundation for functional approaches to shape correspondence.

**Functional maps.** The functional map framework was first introduced by Ovsjanikov et al. [OBCS\*12]. Rather than establishing correspondences directly between individual points, it represents them implicitly as linear operators that transfer functions between shapes. Once the functional map is estimated, point-wise correspondences can be recovered, for example, by transferring delta functions. This formulation casts the correspondence problem as solving a compact linear system that maps functions to functions, rather than points to points. Given two sets of pointwise descriptor functions  $f_{\mathcal{M}} \in \mathbb{R}^{n_{\mathcal{M}} \times d}$  and  $f_{\mathcal{N}} \in \mathbb{R}^{n_{\mathcal{N}} \times d}$  and the orthogonal basis sets  $\phi_{\mathcal{M}} \in \mathbb{R}^{n_{\mathcal{M}} \times k}$  and  $\phi_{\mathcal{N}} \in \mathbb{R}^{n_{\mathcal{N}} \times k}$ , which are often chosen as

the Laplace-Beltrami eigenfunctions, on each shape, the functional map  $C_{\mathcal{M}\mathcal{N}} \in \mathbb{R}^{k \times k}$  can be computed as

$$C_{\mathcal{M}\mathcal{N}} = \arg \min_C \|C\phi_{\mathcal{M}}^\dagger f_{\mathcal{M}} - \phi_{\mathcal{N}}^\dagger f_{\mathcal{N}}\|_F^2, \quad (17)$$

where  $\phi_{\bullet}^\dagger = \phi_{\bullet}^\top A_{\bullet}^{-1}$ , which is the pseudo-inverse of  $\phi_{\bullet}$ .

The corresponding point-wise map  $\Pi_{\mathcal{N}\mathcal{M}}$  can be recovered from the functional map based on the following relationship:

$$\Pi_{\mathcal{N}\mathcal{M}}\phi_{\mathcal{N}} \approx \phi_{\mathcal{M}}C_{\mathcal{M}\mathcal{N}}, \quad (18)$$

The left-hand side of the equation aligns the LBO eigenfunctions in the spatial domain, while the right-hand side aligns them in the spectral domain. Fig. 5 visualizes the functional map framework.

### 3.2. Axiomatic methods

Axiomatic methods assume that noisy or incomplete correspondences are obtained either based on handcrafted features [SOG09, BK10, ASC11] or manual annotations and aim to improve the final matching performance by iteratively refining functional maps and point-wise maps as a post-processing step. The key idea relies on the relationship between the two representations, as defined in Eq. (18).

As the simplest axiomatic technique, ICP iteratively converts the functional map to a point-wise map via nearest neighbour search in spectral space [OBCS\*12]. Follow-up works [RMC\*15, RMC17, EBC17] iteratively optimise the functional map and point-wise map by minimising an objective function based on Eq. (18) with additional regularisation terms or constraints encouraging smoothness and bijectivity. BCICP [RPWO18] further promotes correspondence continuity and overall coverage by introducing bijective ICP with additional point-wise correspondence smoothing as well as outlier detection and fix. ZoomOut [MRR\*19] gradually increases the spectral resolution of the functional map during the ICP process. Despite its simplicity, it achieves better performance than previous, more complex refinement methods.

RHM [ESBC19] prompts reversible harmonic maps, thereby resulting in correspondences with lower conformal distortion. FSF [PRM\*21] theoretically analyses the relationship between functional maps and point-wise maps and proposes a point-wise map recovery method based on spectral alignment of adjoint maps. Later, DiscreteOpt [RMWO21] introduces the *proper* functional map space and optimises the functional maps in the proper functional map space, leading to superior performance in the context of accuracy and smoothness. The *proper* functional map space is the set of functional maps that arise from pointwise correspondences:

$$\{C_{\mathcal{M}\mathcal{N}} \mid \exists \Pi_{\mathcal{N}\mathcal{M}}, \text{ s.t. } C_{\mathcal{M}\mathcal{N}} = \phi_{\mathcal{N}}^\dagger \Pi_{\mathcal{N}\mathcal{M}} \phi_{\mathcal{M}}\}. \quad (19)$$

SmoothFM [MRSHO22] extends DiscreteOpt by introducing spatial smoothness based on minimising the Dirichlet energy, similar to RHM [ESBC19].

Instead of using LBO eigenfunctions for functional map computation, several works [PKO22, DCMO22, HSA\*23] propose to use different eigenfunctions to improve the matching performance in terms of non-isometry and orientation awareness. Another line of

Method	Cat.	MS	NI	P	TN
ZoomOut [MRR*19]	A	×	✓	✓	×
ConsistentZoomOut [HRWO20]	A	✓	✓	✓	×
RHM [ESBC19]	A	×	✓	×	×
ComplexFMaps [DCMO22]	A	×	✓	×	×
DiscreteOpt [RMWO21]	A	×	✓	×	×
SmoothFM [MRSHO22]	A	×	✓	×	×
GeomFMaps [DSO20]	S	×	×	×	×
DPFM [APO21]	S	×	×	✓	×
SRFeat [LAO22]	S	×	✓	×	×
EchoMatch [XER*25]	S	×	×	✓	×
DenoisFM [ZLG25]	S	×	✓	×	×
FRIDU [RBCL25]	S	×	×	×	×
DiffuMatch [PLDO25]	S	×	✓	×	×
UnsupFMNet [HLR*19]	U	×	×	×	×
SURFMNet [RSO19]	U	×	×	×	×
DPFM-unsup [APO21]	U	×	×	✓	×
UDMSM [CB22]	U	✓	×	×	×
AttentiveFMaps [LDO22]	U	×	✓	×	×
ULRSSM [CRB23]	U	×	✓	✓	✓
G-msm [ETLTC23]	U	✓	✓	×	✓
HybridFMaps [BXNL24]	U	×	✓	×	✓
DiffZoomOut [MO24]	U	×	✓	×	✓

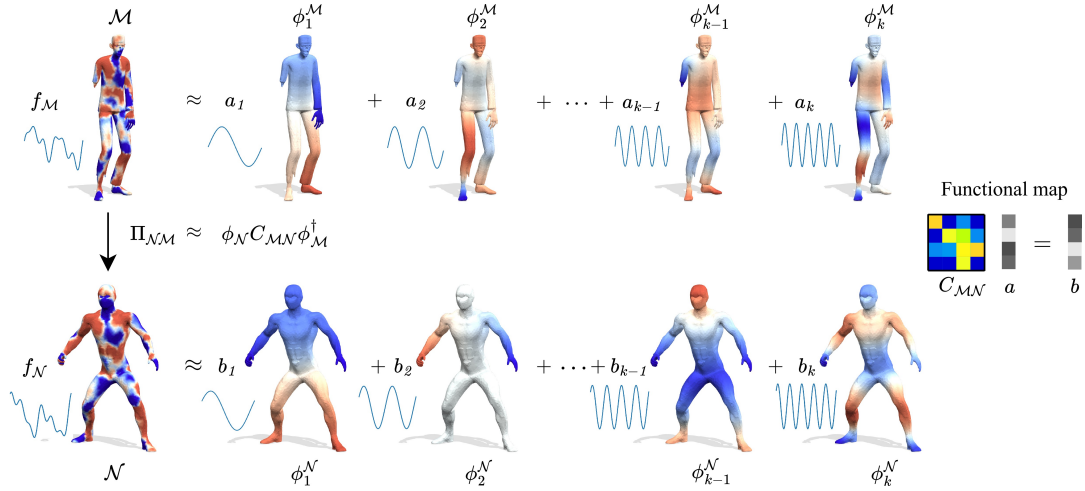
**Table 2:** Overview of *spectral* shape correspondence methods. Each method is categorized by approach (**Cat.**): axiomatic (**A**), supervised (**S**), unsupervised (**U**). The remaining columns indicate supported properties: multi-shape consistency (**MS**), non-isometric deformations (**NI**), partiality (**P**), topological noise (**TN**).

works encourages cycle-consistency [HWG14, BTST19, HRWO20] in the case of processing a collection of shapes.

The study of axiomatic methods establishes a solid theoretical analysis of the relationship between functional maps and point-wise maps. However, in real-world applications, the axiomatic methods often fail due to the violation of the underlying assumptions (e.g., near-isometric deformation). Therefore, it is interesting to combine axiomatic methods with learning-based approaches to improve the matching accuracy. Previous learning-based approaches [HLR\*19, RSO19] typically use them as a post-processing step to improve the matching accuracy. Recent learning-based methods [AO23, CRB23, MO24] are inspired by the insights from those axiomatic methods [RMWO21, MRR\*19]. It is interesting to see that future research in this area integrates and extends existing axiomatic methods within end-to-end learning-based frameworks, combining strong theoretical foundations with data-driven learning to further improve matching accuracy.

### 3.3. Supervised methods

Instead of relying on handcrafted features [SOG09, BK10, ASC11, STDS14] for finding correspondences, learning-based methods attempt to use neural networks to learn the feature descriptors directly from training data. FMNet [LRR\*17] was among the first supervised methods to learn a non-linear transformation of SHOT descriptors [STDS14] based on ground-truth correspondences. Instead of directly obtaining soft correspondences based on feature



**Figure 5:** Visualisation of the **functional map framework** as an analogy of Fourier analysis in signal processing. Given two functions  $f_{\mathcal{M}}$  and  $f_{\mathcal{N}}$  defined on shapes  $\mathcal{M}$  and  $\mathcal{N}$ , respectively, the functional map  $C_{\mathcal{M}\mathcal{N}}$  transfers  $f_{\mathcal{M}}$  from  $\mathcal{M}$  to  $f_{\mathcal{N}}$  on  $\mathcal{N}$ , each function is approximated using the first  $k$  LBO eigenfunctions  $\phi$  of the corresponding shape, which is similar to Fourier transform that uses sinusoidal functions with different frequencies as basis functions. In this representation, the functional map transfers the coefficients  $a$  of the basis functions on  $\mathcal{M}$  to the coefficients  $b$  on  $\mathcal{N}$ , thereby encoding the point-wise correspondences  $\Pi_{\mathcal{N}\mathcal{M}}$  in a compact  $k \times k$  matrix.

similarity, learning-based methods first compute a functional map based on extracted features on each shape  $f_{\mathcal{M}} \in \mathbb{R}^{n_{\mathcal{M}} \times c}$ ,  $f_{\mathcal{N}} \in \mathbb{R}^{n_{\mathcal{N}} \times c}$  by solving the linear least square problem:

$$C_{\mathcal{M}\mathcal{N}} = \operatorname{argmin}_C \mathcal{E}_{\text{data}}(C) + \lambda \mathcal{E}_{\text{reg}}(C). \quad (20)$$

Here, minimizing  $\mathcal{E}_{\text{data}} = \left\| C \phi_{\mathcal{M}}^{\dagger} f_{\mathcal{M}} - \phi_{\mathcal{N}}^{\dagger} f_{\mathcal{N}} \right\|_F^2$  enforces descriptor preservation, while minimizing the regularization term  $\mathcal{E}_{\text{reg}}$  imposes some form of structural properties [OBCS\*12]. After obtaining the functional map, the soft correspondences can be computed as:

$$\Pi_{\mathcal{N}\mathcal{M}}^{\text{soft}} = \operatorname{softmax} \left( \phi_{\mathcal{N}} C_{\mathcal{M}\mathcal{N}} \phi_{\mathcal{M}}^{\dagger} \right). \quad (21)$$

The loss is measured as the geodesic distance error between the predicted soft correspondences and the ground truth. Instead of using geodesic distance, which is computationally expensive, GeomFMaps [DSO20] directly compute the loss in the spectral domain based on the difference between the predicted functional map and the ground-truth functional map. DPFM [APO21] extends this idea to partial shape matching and, for the first time, enables partial-to-partial shape matching by introducing an additional overlap prediction module. Afterwards, SRFeat [LAO22] introduces additional contrastive loss to improve the local discriminative power of learned features and to overcome the limitation that functional maps mainly focus on low-frequency areas. EchoMatch [XER\*25] improves the partial-to-partial matching performance by introducing the concept of correspondence reflection. Correspondence reflection indicates that the valid correspondences for partial shapes should be bijective. Another line of research focuses on learning semantically meaningful basis functions to replace the standard LBO eigenfunctions [MRMO20, HBG\*22] and thereby improve matching performance. Recently, NAM [VOM25] extended the concept

of functional maps to non-linear neural adjoint maps based on multi-layer perceptions.

Despite the appealing matching performance of supervised learning methods, they still face challenges in generalisation. Due to the lack of large annotated training data with ground-truth correspondences across different object categories, supervised methods cannot generalise well to unseen object categories and thus limit their practical applications. Further research in this area should focus on improving generalisation to unseen object categories and on creating large-scale training datasets [EEAX\*25, LZ\*25b].

### 3.4. Unsupervised methods

Due to the lack of large, well-annotated training datasets with ground-truth correspondences, many learning-based methods focus on developing unsupervised training strategies. Starting from UnsupFMNet [HLR\*19], the method proposes an unsupervised loss to train FMNet [LRR\*17] based on geodesic distance preservation:

$$\mathcal{L}_{\text{geo}} = \frac{1}{n_{\mathcal{M}}} \| D_{\mathcal{M}} - \Pi_{\mathcal{M}\mathcal{N}} D_{\mathcal{N}} \Pi_{\mathcal{M}\mathcal{N}}^{\top} \|_F^2, \quad (22)$$

where  $D_{\mathcal{M}} \in \mathbb{R}^{n_{\mathcal{M}} \times n_{\mathcal{M}}}$  and  $D_{\mathcal{N}} \in \mathbb{R}^{n_{\mathcal{N}} \times n_{\mathcal{N}}}$  are the geodesic distance matrices for shape  $\mathcal{M}$  and shape  $\mathcal{N}$ , respectively. Similar to [LRR\*17], the soft correspondences are obtained based on Eq. (21). Later work [ALC20] replaces geodesic distance matrices with heat kernels to be more efficient and robust. Instead of directly regularising the soft correspondences, follow-up works attempt to regularise the more compact functional maps. SURFMNet [RSO19] is one of the first works in this direction. It exploits the structure of the functional map to enable unsupervised learning. Specifically, it uses Laplacian commutativity as the  $\mathcal{E}_{\text{reg}}$  in Eq. (20), which can be expressed as

$$\| C_{\mathcal{M}\mathcal{N}} \Lambda_{\mathcal{M}} - \Lambda_{\mathcal{N}} C_{\mathcal{M}\mathcal{N}} \|_F^2, \quad (23)$$

where  $\Lambda_{\mathcal{M}}$  and  $\Lambda_{\mathcal{N}}$  are the diagonal eigenvalue matrix of shape  $\mathcal{M}$  and  $\mathcal{N}$ , respectively. As for the unsupervised loss terms, SURFM-Net uses bijectivity and orthogonality regularisation to promote the matching to be bijective and area-preserving [OBCS\*12]:

$$\mathcal{L}_{\text{orth}} = \left\| C_{\mathcal{M}\mathcal{N}}^{\top} C_{\mathcal{M}\mathcal{N}} - I \right\|_F^2 + \left\| C_{\mathcal{N}\mathcal{M}}^{\top} C_{\mathcal{N}\mathcal{M}} - I \right\|_F^2, \quad (24)$$

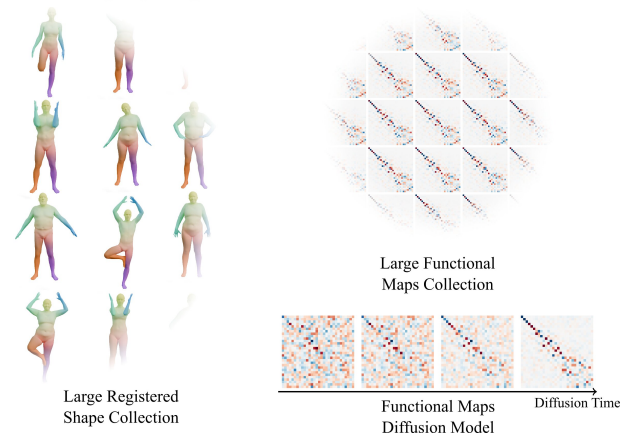
$$\mathcal{L}_{\text{bij}} = \left\| C_{\mathcal{M}\mathcal{N}} C_{\mathcal{N}\mathcal{M}} - I \right\|_F^2 + \left\| C_{\mathcal{N}\mathcal{M}} C_{\mathcal{M}\mathcal{N}} - I \right\|_F^2. \quad (25)$$

Later work [SO20] systematically analyses each component, and DPFM [APO21] extends the idea to partial-to-full shape matching.

Nevertheless, these unsupervised methods still rely on off-the-shelf refinement methods [MRR\*19, OBCS\*12] to obtain better matching results. AttentiveFMaps [LDO22] proposes an attention mechanism to aggregate predicted functional maps with different resolutions and significantly improves the matching performance for non-isometric shapes. ULRSSM [CRB23] utilises the concept of proper functional map space [RMWO21, AO23, SMJ\*23, CRB24] to establish a relationship between the soft correspondences based on feature similarity and the functional maps from Eq. (20) and leads to state-of-the-art performance at the time of publishing for near-isometric, non-isometric as well as topologically noisy shapes. HybridFMaps [BXNL24] extends ULRSSM by incorporating the elastic basis [HSA\*23] to further improve the matching performance, especially for non-isometric deformations. DiffZoomOut [MO24] proposes a differentiable version of ZoomOut [MRR\*19], leading to more efficient inference, especially for high-resolution shapes. Another line of work combines functional maps with spatial alignment, leading to spatially smooth matching [ETLTC20, AO22, CEEA\*24].

Recent unsupervised methods [CRB23, CLB24] have shown significant progress in shape matching for near-isometric, non-isometric shapes, even outperforming some supervised approaches [DSO20]. However, for partial shape matching (particularly the partial-to-partial setting), unsupervised methods struggle to achieve satisfactory results due to the absence of adequate regularization [RCB\*17, EEAX\*25]. Additionally, if significant topological noise arises from self-intersections or poor reconstructions, unsupervised methods are likely to fail. These limitations severely hinder their applicability in real-world scenarios. Future work in this area should focus on improving matching performance for partial and topologically noisy shapes.

**Pointwise map recovery.** An essential step in any spectral pipeline is the recovery of a pointwise map from the functional map. The most straightforward way to achieve this is via nearest-neighbour search in the spectral embedding space [OBCS\*12]. However, this treats each vertex independently, and the resulting maps thus lack guarantees on smoothness or neighbourhood preservation [MRSO22, VM23, RMWO21], which limits their direct applicability in downstream tasks such as deformation transfer [SP04, VOM25] or statistical shape modelling [BBH\*23]. While the refinement methods discussed above mitigate these artefacts [MRR\*19, MRSO22], they operate as post-processing and cannot fully recover information lost due to the low-rank approximation. In Sec. 4, we discuss combinatorial methods that partially



**Figure 6:** Functional maps provide a suitable domain for 2D diffusion models due to their compact matrix representation. Such spectral diffusion models, trained on large collections of registered shapes, introduce a learned prior for functional map solvers. Adapted from [PLDO25].

address this limitation by encoding neighbourhood constraints directly into the matching formulation.

### 3.5. Functional map estimation using image diffusion

Diffusion models have emerged as the dominant paradigm in image synthesis, enabling high-quality generation through iterative denoising [SDWMG15, HJA20, KAAL22, PWG\*24]. They operate by gradually corrupting data with Gaussian noise in a forward process and then learning to reverse this process step by step. A popular approach is Denoising Diffusion Probabilistic Models [HJA20], where a generative model is trained by recovering content from a noise distribution, potentially conditioned on additional data.

Several works have observed that functional maps [OBCS\*12] provide a suitable domain for image diffusion. Given the truncated Laplacian eigenbases  $\phi_{\mathcal{M}}, \phi_{\mathcal{N}} \in \mathbb{R}^{n \times k}$  of shapes  $\mathcal{M}$  and  $\mathcal{N}$ , the resulting functional map  $C_{\mathcal{M}\mathcal{N}} \in \mathbb{R}^{k \times k}$  is a compact matrix, typically with  $k \leq 200$ . This compactness makes it feasible to repurpose architectures originally designed for image generation to predict functional maps.

DenoisFM [ZLG25] trains a diffusion model to directly generate a functional map between a pair of shapes, addressing sign ambiguity of Laplacian eigenfunctions using an unsupervised sign-selection strategy based on learned surface features [ASC11]. FRIDU [RBCL25] presents an approach for functional map refinement, where a diffusion model learns to generate refined maps conditioned on noisy initial maps, guiding learning objectives that encourage plausible output correspondences. DiffuMatch [PLDO25] is trained on the absolute values of functional maps  $|C_{\mathcal{M}\mathcal{N}}|$  to overcome the sign ambiguity of Laplacian eigenbasis, replacing existing axiomatic constraints such as orthogonality Eq. (24) and Laplacian commutativity Eq. (10) with the learned regularizer. These early works demonstrate that diffusion models can serve as a substitute for an explicit functional map solver.

## 4. Combinatorial methods

While the features obtained through neural network architectures based on functional maps are remarkably descriptive, the differentiability requirement on the matching algorithm often yields undesirable matches on disparate but geometrically similar regions of the shape. Another line of work seeks to exploit the manifold structure and geometric coherence of feature matching through powerful combinatorial optimisation techniques.

### 4.1. Background

Combinatorial methods provide powerful tools for shape matching by formulating correspondence problems as discrete optimisation tasks over finite (but potentially very large) solution spaces. As such, combinatorial methods are particularly well-suited for handling discrete shape representations such as meshes and point clouds [Sch03, KV11] and usually seek an optimal assignment between discrete elements of shapes, such as vertex-to-vertex correspondences.

Formally, in the shape matching context, combinatorial methods aim to find an optimal solution  $x^* \in \mathcal{F} \subset \{0, 1\}^n$  (which usually resembles an assignment between surface elements) within a set of feasible solutions  $\mathcal{F}$  while minimizing some (matching cost) function  $\mathcal{E}(x) : \mathcal{F} \rightarrow \mathbb{R}$  such that

$$x^* = \min_{x \in \mathcal{F}} \mathcal{E}(x). \quad (26)$$

Consequently, the advantage of combinatorial formalisms is that they allow (i) to enforce that the solution is guaranteed to have certain properties by constraining the feasible set  $\mathcal{F}$  adequately, and (ii) to add soft constraints and similarity measures by designing appropriate cost functions  $\mathcal{E}(x)$ . In addition to the binary variables  $x \in \mathcal{F}$  used in (26), some methods also introduce additional continuous variables to model deformation alongside matching [BST20, GRE\*23].

Among the most prominent examples of combinatorial matching methods are the linear assignment problem (LAP) and the quadratic assignment problem (QAP). The LAP was introduced to the shape matching community in [BMP01]. Yet, the LAP does not take into account relations between surface elements, such as their geodesic distances, and is therefore not suitable for shape matching when geometric relations must be preserved. On the contrary, the QAP allows for incorporating such relations, but due to the quadratic objective function and resulting NP-hardness [RPW94] it does not scale to practically relevant resolutions. Thus, shape matching approaches based on QAP-formalisms rely on finding solutions via approximations, including convex relaxations [SS05, DML17, BTM18, KMDL19, DLB\*23], and heuristics [DK11, LHP17, VLB\*17, HFH\*22]. To obtain more scalable solutions, other combinatorial shape matching methods exploit the discrete representation of 3D shapes and incorporate neighborhood relations using linear constraints to restrict the feasible set. Among these, one of the most prominent examples is the work by Windheuser et al. [WSSC11] which sparked many follow-up approaches building on similar concepts.

In the following, we summarise recent combinatorial methods. We further categorise approaches into landmark-based formalism,

Method	LB	S2S	MS	Q	LO	RF
Maron et al. [MDK*16]	✓	✓	✗	✗	✗	✓
MINA [BST20]	✓	✗	✗	✗	✗	✗
SIGMA [GRE*23]	✓	✗	✗	✗	✗	✗
Ren et al. [RMWO21]	✗	✓	✗	✗	✗	✗
Pai et al. [PRM*21]	✗	✓	✗	✗	✗	✗
SM-Comb [RSCB22]	✗	✓	✗	✗	✓	✗
DiscoMatch [RAC*24]	✗	✓	✗	✗	✓	✗
SpiderMatch [RB24]	✗	✓	✗	✗	✓	✗
Roetzer et al. [REC*25]	✗	✓	✗	✗	✓	✗
GeCo [RB25]	✗	✓	✗	✗	✓	✗
SuPaMatch [EARB25]	✗	✓	✗	✗	✓	✗
Gao et al. [GLT*21]	✗	✗	✓	✗	✗	✗
Kahl et al. [KSH*25]	✗	✗	✓	✗	✗	✗
Bhatia et al. [BTL*23]	✗	✗	✓	✓	✗	✗
Benkner et al. [SGTM20]	✗	✗	✗	✓	✗	✗
Seelbach et al. [BLG*21]	✗	✗	✗	✓	✗	✗
Meli et al. [MGBM25]	✗	✗	✗	✓	✗	✗


**Table 3:** Overview over *combinatorial* matching methods and classification whether methods are (**LB**) landmark-based, are (**S2S**) Surface-to-Surface methods, are applicable to **MS** multiple shapes, are (**Q**) Quantum, have a (**LO**) linear objective  $f(x)$ , or have a (**RF**) relax the feasible set  $\mathcal{F}$ .

surface-to-surface formalisms, multi-shape matching approaches as well as methods based upon quantum computing. We summarise all methods in Table 3.

### 4.2. Landmark-based approaches

Landmark-based approaches, sometimes also referred to as *sparse* shape matching approaches, focus on a small subset of vertices on each shape, denoted by  $\tilde{X}_{\mathcal{M}} \subset X_{\mathcal{M}}$  and  $\tilde{X}_{\mathcal{N}} \subset X_{\mathcal{N}}$ , see Fig. 7. These vertices, often referred to as *landmarks* or sometimes also as *keypoints*, are intended to capture salient or distinctive locations on the respective shapes  $\mathcal{M}$  and  $\mathcal{N}$ . Generally, such landmarks are not available and are often defined by human annotators. Landmark-based formalisms usually aim to find correspondences only among the landmarks with the underlying assumption that correspondences between shapes can be inferred or approximated effectively from the relationships between their landmarks. Consequently, using landmarks often allows for efficient solvability while still enabling the recovery of dense correspondences via post-processing. Yet, landmark-based approaches assume consistent landmark positions across shapes and thus the solution quality of these methods is heavily based on their position.

Among the landmark-based approaches, one line of work seeks to uncover synergies between shape correspondences and shape reconstruction. In [BST20], Bernard et al. propose a mixed integer programming framework which optimises for a  $m \times m$  permutation matrix  $\Pi$  between landmarks  $\tilde{X}_{\mathcal{M}} \subset X_{\mathcal{M}}$ ,  $\tilde{X}_{\mathcal{N}} \subset X_{\mathcal{N}}$ , and a set of reconstructed landmark vertex-positions  $\tilde{X}_{\mathcal{M}}^{\text{rec}}$ . While  $\tilde{X}_{\mathcal{M}}$  represents the initial positions of the selected landmarks on shape  $\mathcal{M}$ ,  $\tilde{X}_{\mathcal{M}}^{\text{rec}}$  denotes their optimized coordinates in the target space that



**Figure 7:** Visualisation of correspondence between landmarks (in colour) lying on the surface (grey) of two 3D shapes. Image source: [GRE\*23]

best align with the landmarks  $\tilde{X}_{\mathcal{M}}$ . On a high-level the goal is to minimise

$$\min_{\Pi \in \mathbb{P}, \tilde{X}_{\mathcal{M}}^{\text{rec}} \in \mathbb{R}^{m \times r}} \|\tilde{X}_{\mathcal{M}}^{\text{rec}} - \Pi \tilde{X}_{\mathcal{N}}\| + \mathcal{E}_{\text{rec}}(\tilde{X}_{\mathcal{M}}^{\text{rec}}, \mathcal{M}) \quad (27)$$

Here,  $\mathcal{E}_{\text{rec}}(\tilde{X}_{\mathcal{M}}^{\text{rec}}, \mathcal{M})$  is a reconstruction function that aims for a smooth reconstruction of shape  $\mathcal{M}$  using the optimised landmark positions  $\tilde{X}_{\mathcal{M}}^{\text{rec}}$ .

Building on the ideas presented in [BST20], Gao et al. [GRE\*23] propose a similar mixed integer programming formalism and introduce the *projected Laplace-Beltrami* (LBO) operator for the reconstruction, which reads

$$L_{\mathcal{M}}^{\text{proj}} := M^T \Delta_{\mathcal{M}}^{\text{stiff}} M. \quad (28)$$

Here,  $\Delta_{\mathcal{M}}^{\text{stiff}}$  is the stiffness matrix component of the Laplace-Beltrami operator  $L_{\mathcal{M}}$ ,  $M := I - \tilde{X}_{\mathcal{M}}(\tilde{X}_{\mathcal{M}}^T \tilde{X}_{\mathcal{M}})^{-1} \tilde{X}_{\mathcal{M}}^T$  is a projection matrix and  $\tilde{X}_{\mathcal{M}} := (X_{\mathcal{M}}, \mathbf{1})$ . As shown in [GRE\*23], the projected LBO  $L_{\mathcal{M}}^{\text{proj}}$  is scale-invariant and invariant under rigid body transformations. Furthermore, it retains more high-frequency geometric details of the shapes compared to the Laplace-Beltrami operator. This helps to increase reconstruction quality and, in turn, leads to better matching results.

While landmark-based approaches offer improved scalability due to smaller optimisation problems and, with that, greater mathematical modelling flexibility, many argue that full surface-to-surface maps must be considered to compute faithful matchings between shapes.

### 4.3. Surface-to-surface approaches

Surface-to-surface means finding a dense matching for all vertices  $X_{\mathcal{M}}$  of a shape  $\mathcal{M}$ , rather than only for a subset of vertices, as in landmark-based approaches. To find such a matching, some lines of work build on the *Procrustes matching* problem where the goal is to simultaneously find correspondences between two sets of points  $X_{\mathcal{M}}, X_{\mathcal{N}} \in \mathbb{R}^{n \times d}$  (in form of an assignment matrix  $\Pi_{\mathcal{M}\mathcal{N}}$ ), as well as a rigid alignment  $T \in SE(d)$ . This is formulated as the following optimization problem

$$\min_{\Pi_{\mathcal{M}\mathcal{N}} \in \mathbb{P}, T \in SE(d)} \|\tilde{X}_{\mathcal{M}} T - \Pi_{\mathcal{M}\mathcal{N}} \tilde{X}_{\mathcal{N}}\|_F^2, \quad (29)$$

with  $\tilde{X}_{\mathcal{M}} := [X_{\mathcal{M}}, \mathbf{1}_n] \in \mathbb{R}^{n \times (d+1)}$  denoting the homogeneous coordinates of shape  $\mathcal{M}$ , where  $\mathbf{1}_n \in \mathbb{R}^n$  is the all-ones vector. The

homogeneous coordinates  $\tilde{X}_{\mathcal{N}}$  are defined analogously. An iterative approximation to this formulation was popularised by the seminal iterative-closest point (ICP) algorithm [BM92], see also Sec. 5. Other works [MDK\*16] propose convex relaxations for the Procrustes matching problem.

Ren et al. [RMWO21] propose a discrete optimisation approach which computes surface maps between two shapes  $\mathcal{M}$  and  $\mathcal{N}$  by constraining a functional map  $C_{\mathcal{M}\mathcal{N}}$  such that it can be related to a point-wise map  $\Pi_{\mathcal{M}\mathcal{N}}$ . These are then also called *proper functional maps*, see Eq. (19). As such, proper functional maps have been further explored in [PRM\*21] using convex relaxations of point-wise maps, as well as in deep shape matching methods as a loss function [CRB23, AO23].

In contrast, Windheuser et al. [WSSC11] enforce consistent neighbourhood relations of surface elements by utilising linear constraints in an integer linear programming formalism based upon a so-called *product space*. While mathematically elegant, this formalism is hard to solve in practice and thus Windheuser et al. [WSSC11] propose to solve it using linear programming (LP) relaxations and iterative variable fixations. Yet, even with this approximate scheme, it only scales to shapes with around 50 – 100 triangles on modern hardware and with modern LP solvers; see [RB25] for runtime comparisons. This inspired follow-up works [RSCB22, RAC\*24] to develop approximate solvers for this formalism. Both solvers [RSCB22, RAC\*24] build on dual decomposition schemes presented in [LS21, AS22] and run on CPU [RSCB22] and GPU [RAC\*24] respectively. These solvers show significant improvements in scalability, i.e. [RSCB22] scales to shapes with up to a 1000 triangles and [RAC\*24] to 750 triangles. Yet, due to runtime and memory limitations, scalability is nevertheless limited.


This motivated follow-up works to develop new formalisms based on *product graphs*. Product graphs are denoted

$$\mathcal{P} = \mathcal{G}_1 \boxtimes \mathcal{G}_2, \quad (30)$$

where the *strong product* between two graphs  $\mathcal{G}_1$  and  $\mathcal{G}_2$ , see [HIK11] for more details and Fig. 8 for a visualization. In the past, these product graphs were successfully used for dynamic time-warping [SC78], i.e. finding point-to-point correspondence between two contours with start and end, matching a shape to an image [SC10], 2D-to-2D shape matching [STC09], and 2D-to-3D shape matching [LRS\*16]. To obtain an efficient-to-solve formalism for 3D-3D shape matching, SpiderMatch [RB24] changes the representation of shape  $\mathcal{M}$ , i.e. the *source* shape: it uses a single cyclic, self-intersecting curve  $\mathcal{C}$  on the surface of the shape  $\mathcal{M}$  to represent the respective shape. This allows casting the 3D shape matching problem as finding a minimum cost cyclic path in a product graph  $\mathcal{P}$  with additional constraints. As shown in [RB24], this problem is globally optimal and efficiently solvable in practice, scaling to shape resolutions of up to 3500 triangles.

In [REC\*25], authors utilise the curve representation proposed in SpiderMatch [RB24], drop the additional constraints, and show that good shape matching results can be computed when using minimum ratio cycles with local rigidity regularisers [SA07, RLB23], achieving global optimality in practice.

Contrary to these methods, [RB25] argues that representing a



**Figure 8:** Visualisation of the (in-layers-structured) *product graph*  $\mathcal{P}$  between a cyclic chain graph  $\mathcal{C}$  and the graph of a triangle mesh of a 3D shape  $\mathcal{N}$ . Correspondences between  $\mathcal{C}$  and  $\mathcal{N}$  can be computed by finding the shortest cyclic path that visits all layers. Image source: [RB24].

3D shape with a single cyclic curve  $\mathcal{C}$  and matching this curve to the target shape  $\mathcal{N}$  does not lead to a faithful surface-to-surface matching. Therefore, [RB25] proposes to use multiple cyclic curves  $\mathcal{C}_1, \dots, \mathcal{C}_n$  to represent each triangle of the source shape  $\mathcal{M}$ . This allows formulating 3D shape matching via multiple product graphs  $\mathcal{P}_1, \dots, \mathcal{P}_n$  as finding multiple shortest cyclic paths which have to adhere to additional coupling constraints. In turn, this leads to a faithful surface-to-surface matching that is globally optimal in practice and scales to shapes with up to 3000 triangles. Building on this formalism, [EARB25] propose a two-stage approach in which they first transfer patches from source to target shape and subsequently find dense surface-to-surface matchings for each pair of corresponding patches. This allows scaling to shapes with up to 10k triangles without the need for coarse-to-fine strategies. This has the advantage of fewer discretisation artifacts which usually occur in approaches relying on coarse-to-fine strategies.

While [RB25] has shown that multiple coupled problems lead to high-quality matching results, these coupled problems only involve two 3D shapes. As a result, their formalism is not directly applicable to problems that involve collections of shapes, which we discuss next.

#### 4.4. Multi-shape matching approaches

Multi-shape matching approaches simultaneously optimise maps between a collection of shapes. Assuming bijective mappings between shapes, *cycle consistency* emerges as a fundamental property that should be fulfilled in most cases [STCB07, HG13, CRA\*17]. This means that the matching from a vertex on one shape via multiple other shapes back to the first shape must end in the same vertex again. For non-bijective mappings between shapes, the definition of cycle consistency is more intricate [BTGT19]. In particular, considering the bijective case, for a collection of  $N$ -many shapes  $\mathcal{M}_1, \dots, \mathcal{M}_N$ , it must hold that

$$\forall i, j, k \in \{1, \dots, N\} : \Pi_{\mathcal{M}_i \mathcal{M}_j} \Pi_{\mathcal{M}_j \mathcal{M}_k} = \Pi_{\mathcal{M}_i \mathcal{M}_k}. \quad (31)$$

[BTL\*23] enforce cycle consistency constraints by considering triplets of shapes and by proposing a quantum-hybrid approach building on ideas of the famous alpha-expansion algorithm [BVZ01]. [ETLTC23] learns maps between multiple shapes, which are used as edge weights for all pairs  $i, j \in \{1, \dots, N\}$  in the shape graph, to then propagate maps along shortest paths in the shape graph to enforce cycle consistency. [KSH\*25] propose an approximate algorithm for cycle-consistent multi-graph matching, which exploits parallelism, performs local graph matching and maintains cycle consistency at all times.

To avoid enforcing the cycle consistency constraints (31) explicitly, other works revert to shape-to-universe matching where cycle consistency is enforced implicitly [PKS13]. Among these, [GLT\*21] introduces a multi-shape matching formalism for isometric deformed shapes and further proposes an algorithm based upon convex relaxations. Later, the shape-to-universe matching scheme was also employed in unsupervised learning approaches [CB22], see also Table 2.


Multi-shape matching approaches and many other combinatorial methods are most often hard to solve and as a consequence, on classical hardware, their scalability is limited. Motivated by that, other lines of work aim for formalisms which can be employed on quantum hardware which can find global optima naturally in some configurations, such that formalisms might become much more scalable with advances in quantum computing.

#### 4.5. Quantum annealing

Quantum annealers can efficiently solve quadratic unconstrained binary optimisation problems in the form of

$$\min_{x \in \{0,1\}^n} x^\top Q x + q^\top x. \quad (32)$$

Here  $Q \in \mathbb{R}^{n \times n}$  is a symmetric (possibly dense) cost matrix containing pairwise costs and  $q \in \mathbb{R}^n$  a linear bias. Current available systems, for example, the D-Wave line [D-W26], contain several thousand qubits (*quantum bits*). In theory, the number of qubits corresponds directly to  $n$  in (32), however, due to hardware limitations in the connectivity of the annealer, this only holds for certain forms of  $Q$  and in reality, many qubits are blocked when representing a dense  $Q$  matrix [SBLG\*25]. The natural use of binary variables makes quantum annealers primed for solving combinatorial versions like Eq. (26), but the unconstrained setting cannot model many of the desired properties of the solution, such as bijectivity. The first approach encoded the permutation constraint in  $Q$ , but this occupies most of the qubits on fulfilling these constraints and could only correspondence between up to 4 points [SGTM20]. However, by solving the problem iteratively in subproblems based on a version of  $\alpha$ -expansion, quantum annealers can be scaled to several thousand vertices [BLG\*21] and even solve the hard multi-shape case with a linear runtime scaling in the number of shapes [BTL\*23]. The permutation set can also be parametrised by the binary variables directly, which leads to an efficient approximation problem without the need for an iterative approach [MGBM25]. For an introduction to quantum annealing in the context of computer vision and more details on correspondence estimation approaches relying on quantum computa-



**Figure 9: Deformation methods** recover the correspondence between shapes by finding a deformation that aligns them spatially.

tional paradigms, interested readers can refer to the recent survey by Kuete Meli et al. [MWB\*25a].

## 5. Deformation-based methods

So far, the methods that we considered retrieve the correspondence without modifying the input shapes, mainly by relying on intrinsic properties of the geometry. In this section, we instead present methods that recover correspondence by computing a deformation field that aligns a source surface with a target surface in the ambient 3D coordinate space. Unlike *intrinsic* methods, which rely on surface-aware metrics like geodesic distances, this *extrinsic* paradigm fundamentally considers Euclidean metrics and is thus broadly applicable to various 3D data representations, including unstructured point clouds, without requiring manifold connectivity or consistent topology between the shapes being registered. It is worth mentioning that while the majority of such works define the problem in terms of point sets, they are commonly applied to mesh vertices. Surface properties are often used both for regularisation and as a domain for point sampling.

In revisiting the literature of recent years, we categorise deformation works into three main branches: optimisation-, template-, and learning-based. Drawing a clear distinction among them is impossible, as many methods are hybrid, and our categorisation follows what we believe are the main principles behind the paper. We additionally emphasise the type of optimisation that is conducted, feature embedding space, correspondence type, and deformation model (see Table 4). We begin with background on seminal techniques that continue to serve as inspiration.

### 5.1. Background

While intrinsic methods come with a number of advantages, such as invariance of the shape’s embedding in the ambient Euclidean domain and strong structural priors, their theoretical premises are also a source of drawbacks. Noise, topological changes, as well as incomplete or cluttered shapes, often introduce enough instability to make purely intrinsic methods inapplicable. On the contrary, we have already seen that considering the extrinsic information of shapes can play an important role, for example, in distinguishing symmetries, and provide semantical nuances difficult to catch on a pure surface level. Along this line, several methods rely on the assumption that the correspondence is equivalent to a certain, generally unknown, deformation. For the sake of simplicity, we formalise the problem in an idealised bijective setting. Considering

two sets of points  $X_M$  and  $X_N$ , and assuming there exist a permutation  $\Pi : X_N \rightarrow \hat{X}_N \in \mathbb{R}^{n_M \times 3}$  that reorganizes their vertices under the desired semanticity, the idea is to recover an equivalent deformation  $\mathcal{F} : X_M \rightarrow \hat{X}_M \in \mathbb{R}^{n_M \times 3}$  such that:

$$\mathcal{F}(X_M) \approx \Pi(X_N). \quad (33)$$

It is important to outline a few observations, which will help us anticipate the challenges addressed in the literature. First, perfect alignment can only occur when we have access to a perfect deformation model, and the two shapes share the same discretisation; this generally holds only for the synthetic settings. Second, given the permutation  $\Pi$  or the deformation  $\mathcal{F}$ , the other can be recovered in a closed form [BM92]; however, the standard case (and the one we will focus on) assumes that both are unknown, an ill-posed problem with many plausible solutions  $\mathcal{F}$  and  $\Pi$ . Finally, it is often the case that only subsets of the points can be correctly aligned; clutter, noise, and outliers often require careful consideration and ad-hoc techniques. For these reasons, the problem is classically phrased as an energy minimisation.

**Energy-based formulation.** Solving Eq. (33) via optimisation requires defining an energy, which aims to align the shapes, but also recovers a plausible deformation. Hence, the energy is generally composed of data ( $\mathcal{E}_{\text{align}}$ ) and regularization ( $\mathcal{E}_{\text{reg}}$ ) terms, with a weighting factor  $\alpha$ :

$$\mathcal{E} = \mathcal{E}_{\text{align}} + \alpha \mathcal{E}_{\text{reg}} \quad (34)$$

**Seminal registration methods.** A survey on non-rigid registration would be incomplete without discussing two seminal works, which still serve as inspiration for many recent methods: Iterative Closest Point (ICP) [BM92] and Coherent Point Drift (CPD) [MS10]. Since both the correspondence and the deformation are unknown, ICP proposes iterating between the two, fixing one and optimising the other. Starting from an initial configuration of the two shapes (i.e.,  $\tilde{\mathcal{F}}(X_M) = X_M$ ), ICP obtains an estimated correspondence  $\tilde{\Pi}$  by Euclidean nearest-neighbour pairing:

$$\tilde{\Pi} = \arg \min_{\Pi} \|\tilde{\mathcal{F}}(X_M) - \Pi(X_N)\|_2^2. \quad (35)$$

Next, the registration is updated  $\tilde{F}$  by minimizing the point-to-point distance with a least-squares objective:

$$\tilde{\mathcal{F}} = \arg \min_{\mathcal{F}} \sum_{i=1}^m \|\mathcal{F}(X_{M_i}) - \tilde{\Pi}(X_{N_i})\|_2^2. \quad (36)$$

The two steps are iterated until convergence. Notice that to avoid trivial solutions,  $\mathcal{F}$  should be restricted to a specific family, which originally was the class of rigid transformations [BM92]. Over the past several decades, numerous non-rigid variations have been proposed. The first non-rigid extension was proposed by Amberg et al. [ARV07], which revises the transformation in terms of local affine transformations, ensuring that changes are minimal between nearby vertices. By fixing the amount of desired stiffness, the authors derive a cost function that still admits a closed-form to recover an optimal iterate.

As opposed to the discrete, nearest-neighbour-based approaches pioneered by ICP, probabilistic methods generate soft matches where each point in the source set  $X_M$  has a certain probability

Method	Category	Correspondence	Embedding Space	Deformation Model	Training
Revisited-CPD [FMT*22]	<b>O</b>	Soft	Euclidean	Coherent	<b>N</b>
Geo-CPD [Hir23]	<b>O</b>	Soft	Manifold	Coherent	<b>N</b>
Tajdari et al. [THYS22]	<b>O</b>	Hard	Euclidean	Affine	<b>N</b>
Yao et al. [YDXZ23]	<b>O</b>	Hard	Euclidean	Locally-affine	<b>N</b>
Zhao et al. [ZJM*24]	<b>O</b>	Soft	Euclidean	Kernel-based	<b>N</b>
Merrouche et al. [MWB25b]	<b>O</b>	Hard	Neural	ARAP	<b>N</b>
SPARE [YDHz25]	<b>O</b>	Hard	Euclidean	ARAP	<b>N</b>
Jiang et al. [JSH23]	<b>H</b>	Neural	Neural	ARAP	<b>S</b>
Chen et al. [CWY*23]	<b>H</b>	Global	Neural	Parametric	<b>S</b>
Qin et al. [QYW*23]	<b>H</b>	Neural	Neural	Parametric	<b>S</b>
ArtEq [FKL*23]	<b>T</b>	Neural	Semantic	Parametric	<b>S</b>
Guo et al. [GPX*24]	<b>H</b>	Hard	Neural	Coherent	<b>S</b>
ETCH [LFC*25]	<b>T</b>	Hard	Manifold	Parametric	<b>S</b>
Jung et al. [JKY*25]	<b>T</b>	Hard	Euclidean	Kernel-based	<b>N</b>
NDP [LH22b]	<b>L</b>	Neural	Neural	Neural Field	<b>S</b>
IFS [SPO22]	<b>L</b>	Neural	Implicit	Neural Field	<b>U</b>
LVD [CPMAMN22]	<b>H</b>	Neural	Neural	Neural Field	<b>S</b>
Sundaraman [SDM*24]	<b>L</b>	Neural	Spectral	Jacobian-based	<b>S</b>
Merrouche [MRWB23]	<b>L</b>	Neural	Neural	Hierarchical	<b>U</b>
Lin et al. [LZZ*23]	<b>L</b>	Neural	Manifold	Neural Field	<b>U</b>
Feng et al. [FCH*23]	<b>L</b>	Neural	Neural	Parametric	<b>S</b>
NICP [MCPM24]	<b>H</b>	Neural	Neural	Neural Field	<b>S</b>

**Table 4:** Overview of *deformation-based* matching methods. Methods are categorized by their approach (*Optimisation*, *Template*, *Learning*, or *Hybrid*), correspondence type (*Hard*, *Soft*, *Neural*, *Global*), feature space, deformation model, and training requirements (*None*, *Supervised*, or *Unsupervised*).

of corresponding to every point in the target set  $X_{\mathcal{N}}$ . The seminal Coherent Point Drift (CPD) algorithm by Myronenko and Song [MS10] reformulates the geometric alignment problem through statistical density estimation. CPD treats the source points as centroids of a Gaussian Mixture Model (GMM), while the target points are considered as observations generated from this GMM. To handle outliers and noise, CPD augments the mixture model with a uniform distribution component, yielding the probability density:

$$p(\mathbf{x}) = w \cdot \frac{1}{n_{\mathcal{N}}} + (1-w) \sum_{i=1}^{n_{\mathcal{M}}} \frac{1}{n_{\mathcal{M}}} \mathcal{N}(\mathbf{x} | \mathcal{F}(x_i), \sigma^2 I) \quad (37)$$

where  $w \in [0, 1]$  represents the assumed proportion of outliers,  $\mathcal{F}$  denotes the transformation function, and  $\sigma^2$  is the isotropic variance of the Gaussian components. The registration problem can then be approximated through maximum likelihood estimation (MLE).


Since the two streams of work rely on different definitions of data fitting, the resulting regularisation techniques also differ. For ICP, the deformation is often chosen within a specific class, such as  $SE(3)$  [BM92] or local Affine transformations [ARV07], which naturally restrict the valid solutions. CPD instead admits free displacements and thus relies on Motion Coherence Theory (MCT), penalising derivatives of all orders of the displacement field and ensuring neighbouring points are deformed consistently (this results in neighbouring points moving *coherently*).

## 5.2. Revisiting classical methods

**Optimisation-based.** The standard point-to-plane metric as used in ICP only considers the normals of the static target surface. The recent SPARE [YDHz25] method thus proposes a symmetrised point-to-plane distance metric. This formulation incorporates normal information from both the deforming source and the target surfaces, providing a more accurate, second-order approximation of the local geometry, which can accelerate convergence and improve final accuracy [YDHz25]. Alternatively, the problem of alignment distribution can be rephrased as a fuzzy clustering problem [ZJM\*24]. Beyond penalising smoothness through motion coherence or stiffness terms, geometrically derived features can provide additional constraints [THYS22, Hir23].

To address noise, outliers, and acquisition artefacts, optimisation can be made more resilient through robust norms such as the  $L_p$  norm with  $p < 2$  (e.g.,  $L_1$ ) for both alignment and regularisation terms [YDXZ20, ZJM\*24]. Unlike the  $L_2$  norm, which quadratically penalises large errors and is thus highly sensitive to outliers, these norms apply a less severe penalty, making the overall cost function more resilient to gross errors. However, many robust norms are non-smooth and non-differentiable, complicating optimisation and inhibiting standard gradient-based solvers. Recent work by Yao et al. [YDXZ23] addresses this through the Majorisation-Minimisation (MM) algorithm with Welsch's function [HW77], enabling closed-form solutions that prove computationally superior to non-smooth proximal methods (such as ADMM [BPC\*11]) while maintaining robustness.

Despite the use of the Fast Gaussian Transform, the speed of



**Figure 10:** An example of deformation from GeoCPD [Hir23]. The source shape (black) is deformed toward the target one (blue), toward multiple steps till final convergence.

the CPD algorithm remains a crucial drawback [GTRS16]. Subsequent works have significantly advanced the regularization framework: Hirose proposed an accelerated version using downsampling [Hir21a] and later introduced a Bayesian formulation that generalizes both CPD and ICP by rephrasing coherence as a distribution prior [Hir21b], while also extending the regularization to use geodesic distances rather than Euclidean ones to better preserve the intrinsic geometry of the surface during deformation [Hir23]. Other extensions have adapted the coherence constraint to non-Euclidean settings [FMT\*22] and incorporated topological preservation through joint optimisation of explicit and implicit representations [MWB25b].

**Use of learned priors.** Learned features can be used to complement non-rigid deformation-based approaches. For instance, Jiang et al. learn deep features via functional maps (see Sec. 3), which induces an initial set of point-to-point correspondences  $\hat{\Pi}$  [JSH23]. This permutation estimate can then be filtered based on geometric consistency and used to recover the deformation function  $\mathcal{F}$ . Guo et al. [GPX\*24] address cloth deformation by learning a diffusion-based prior from real-world clothing deformations. Learned methods can be used to prune correspondences in a manner more suitable for non-rigid deformations [QYW\*23]. As with some functional map-based methods [APO21], transformer-based modules can be useful in guiding overlap prediction to handle deformations between partial shapes [MPS\*23]. Auxiliary tasks can also yield learned feature representations useful for registration [CWY\*23].

### 5.3. Template fitting

For certain classes, we might have access to an archetypal shape that describes the structure, and often comes with information on how shapes of that class are expected to deform. Such a model is referred to as a template, and when it is available, it provides a strong regularisation. Such a template is generally crafted by an artist or domain experts [GLW\*23, DFS99]. The template can also be enriched with a learned parametric deformation space by aligning it to a collection of shapes. A seminal work in this sense is that of Blanz and Vetter [BV99], which aligned a few hundred human faces to learn PCA, both to model the subject's geometry and texture. More recent shape models are the SMPL models for hu-

mans [LMR\*15, PCG\*19], SMAL for the animals [ZKJB17], and MANO for the hands [RTB17], but learning such morphable shapes is also popular in medicine [MLA\*25, MBT\*23, PSE\*26].

While generally such a procedure requires a large amount of curated data [AWLB17, LMR\*15, PMRMB15, ZML\*24], it is worth mentioning that recent approaches also leverage data in the wild, such as images and videos from the web [LLL\*24].

**Fitting.** To take advantage of the available template, its structural properties can be used during the optimisation. For example, humans can be represented as piece-wise rigid deformations, as done in the Stitched Puppet formulation [ZB15]. Recent methods make use of local  $SO(3)$ -equivariance features to represent the local rigidity [LFC\*25, FKL\*23], though in practice, such equivariance assumes little or no noise. Since templates provide a smooth and complete geometry, they can be integrated with intrinsic methods, enhancing their robustness [MMRC20]. It is likely that a template cannot catch all the nuances of a specific class. Hence, it is often the case that to align it to the data, the local details are delegated to free-form offsets. Examples are for skin details [MMRC19], cloths [MYR\*20], and soft-tissues [PMRMB15]. Jung et al. take an approach of jointly optimising a shared template with respect to a collection of input shapes [JKY\*25]. The template fitting can also be enriched by visual semantic information present on the surface, such as textures [ATO\*24, LIPM19]. Recent developments in capturing technologies have also enabled the study of fine-grained template deformation, such as that in fabrics. Modelling wrinkles can be achieved by physics modelling [YCS25] or learned-priors [GPX\*24].

### 5.4. Learning-based methods

With the advent of deep learning, many methods leverage deep neural networks trained to learn a deformation model from a collection of data. 3DCoded [GFK\*18] popularised this line of work, training an autoencoder to deform a template into all the shapes of the training set. In case ground truth registration is available, such a network can be trained in a supervised way using a vertex-to-vertex loss. On the contrary, it can also be trained without supervision, relying on Chamfer Distance. In the follow-up [DGF\*19],

the same authors also investigated the possibility of learning primitive structures together with the deformation module. A revision of the 3DCoded formulation has also been proposed with Transformer architecture [TCM\*21], where the network relies on an attention mechanism that weights points by their area.

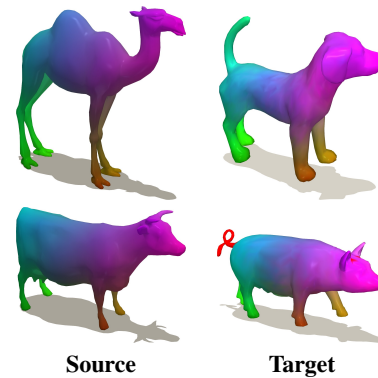
The deformation can be represented globally, for instance, as a set of hierarchical MLPs [LH22b]. Merrouche et al. [MRWB23] suggest training both an association and a deformation network relying on hierarchical features. Temporal network modules such as GRU can be used to represent the deformation incrementally, enhancing robustness [FCH\*23]. A more nuanced geometric view is achieved through Neural Jacobian Fields (NJF) [AGK\*22], which first learns a matrix field in the ambient space and then derives the Jacobian for each triangle via an intrinsic projection. A final global deformation can then be recovered by solving a Poisson equation based on the locally learned Jacobians. Sundararaman et al. build upon NJF by abandoning the need for a global shape encoder, representing shapes' vertices with their local Jacobians, learning a coarse-to-fine deformation entirely based on local properties [SDM\*24].

Particularly convenient is relying on implicit representations, since they provide spatial and differentiable information useful for the fitting. Additionally, they offer a flexible representation that is significantly more robust to noise and partial information. Sundararaman et al. propose representing shapes as neural fields and utilising implicit representations to regularise deformation [SPO22]. On a similar principle, they propose to learn the deformation as local fields [SMRO22]. In IPNET [BSTPM20a], the authors combine a deformable template with a network to predict the occupancy of local body parts, which is then used to guide the SMPL registration. Wang et al. [WGT21] propose PTF, a similar framework that relies on piece-wise implicit representations. LoopReg [BSTPM20b] also proposes a weakly-supervised schema for shape registration.

While these works rely on implicit representations to represent the shapes, a more recent trend is to leverage neural fields to also represent deformations as a velocity field defined all over  $\mathbb{R}^3$ , as proposed in Learned Vertex Descent [CPMAMN22]. The motivation is that they provide a richer representation for the deformation, on which further supervision can be provided. While the original paper utilises the neural field primarily during training, Marin et al. [MCPM24] demonstrate that this formulation can also be beneficial at inference time. In practice, they propose Neural ICP (NICP), a procedure to iteratively refine the neural deformation field by retrieving its induced correspondence. Neural deformation fields have also been defined intrinsically, in the clothing deformation case [LZZ\*23]. Lastly, in the majority of previous works, the correspondence step is delegated to features coming from neural backbones.

## 6. Challenges and emerging trends

In this section, we explore several opportunities that have recently emerged in shape correspondence research. These include leveraging foundation models, preserving local neighbourhoods during matching, and addressing partial shape scenarios. Such directions



**Figure 11: Zero-shot matching** involves establishing correspondences between structurally distinct shape categories (such as different animal species), which often contain unique regions present on only one of the shapes (shown in red). In the absence of training data, feature descriptors from 2D vision models are particularly effective for this task. Shapes are taken from SHREC'20 [DLR\*20] and BeCoS [EEAX\*25] datasets.

extend the capabilities of conventional methods, enabling them to handle more challenging and previously inaccessible settings.

### 6.1. Feature aggregation from foundation models

Axiomatic and learned feature descriptors are fundamental to many correspondence estimation pipelines. However, their key limitations lie in the former's limited expressiveness and the latter's poor generalisation to out-of-distribution shapes. As an alternative, an increasingly popular direction is to utilise features from 2D foundation vision models trained on large-scale image datasets, such as self-supervised transformers [CTM\*21, ODM\*23], image diffusion models [RBL\*22], or segmentation models [KMR\*23]. These distilled descriptors have the potential to generalise zero-shot across shape classes, without category-specific training data.

In a typical pipeline for adapting 2D vision features to 3D shape correspondence, a mesh  $\mathcal{M}$  is projected to image space from  $n$  viewpoints using a projection operator Proj:

$$\text{Proj}_{\mathcal{C}_j}(\mathcal{M}) = I_j \in \mathbb{R}^{H \times W}, \quad (38)$$

where  $H, W$  denote the height and width of the rendered image, with  $\mathcal{C}_j$  representing the  $j^{\text{th}}$  camera producing the image  $I_j$ . Together with a set of additional information  $\mathcal{G}$  such as depth, normal maps, or a text prompt, the rendered image  $I_j$  is processed by a 2D vision model  $\Psi$  to obtain  $d$ -dimensional features  $f_j^{2D}$ ,

$$\Psi(I_j, \mathcal{G}) = f_j^{2D} \in \mathbb{R}^{H \times W \times d}. \quad (39)$$

Known camera parameters are then leveraged to unproject features from image space to mesh vertices  $f_j^{3D}$ . These are aggregated across viewpoints to obtain shape descriptors  $f_{\mathcal{M}}$ ,

$$f_j^{2D} \xrightarrow{\text{Proj}^{-1}} f_j^{3D}, \quad f_{\mathcal{M}} = \frac{1}{n} \sum_{j=1}^n f_j^{3D}. \quad (40)$$

The latter can be passed to a downstream matching pipeline for

computing correspondences, or combined with other axiomatic or learned descriptors.

Among the earliest works to leverage foundation models for 3D shape correspondence, Abdelreheem et al. [AEOW23] use vision-language models [LLSH23, BMR\*20] to infer shape categories and semantic regions, which are localised by a multiview object detector [LZR\*24] and segmented using Segment Anything [KMR\*23]. The resulting masks are aggregated onto the mesh to guide the initialisation of the functional map with WKS descriptors [ASC11], and are then refined via BCICP [RPWO18]. In contrast, Diff3F renders textureless shape views and uses Stable Diffusion [RBL\*22] to synthesise textured images, extracting intermediate features and fusing them with DINO [ODM\*23]. These are unprojected onto the mesh and used in a standard functional map pipeline [OBCS\*12] to compute dense correspondences.


Recent methods extend the early use of 2D vision features by integrating them into more advanced correspondence pipelines. Morreale et al. [MAKM24] employ DINO to co-align shapes by estimating their upright orientation from multiview renderings and to provide semantic descriptors. These features are then integrated into Neural Surface Maps [MAKM21], a framework that computes continuous correspondences between surfaces by encoding them with dedicated neural functions. DenseMatcher [ZJZ\*25] proposes to fuse features from Stable Diffusion and DINO with HKS [TSDS10] descriptors and refine them using a trainable DiffusionNet [SACO22], supervised by a semantic distance loss that encourages feature distances to align with semantic part similarity. In [UEK25], the problem of disambiguating symmetric parts on a shape is addressed by training an autoencoder network to refine Diff3F [DMM24] features through an objective that encourages the preservation of geodesic distances in the learned embedding space.

Several works integrate vision features into deformation-based correspondence pipelines. In SRIF [SGJ\*24], multiview renderings of the source and target meshes are processed by a diffusion-based image morphing model [ZZX\*24], which generates smooth transitional frames between the two shapes. A continuous deformation field is then estimated via a normalising flow [YHH\*19], gradually warping the source mesh to match the target while following the reconstructed trajectory. In DV-Matcher [CJH25], DINO features are combined with positional encodings and processed by a Siamese network with a dual-path attention mechanism, trained to predict the deformation between source and target shapes. Stable-SCore [LZY\*25b] utilises vision features to obtain 2D correspondences between rendered views of meshes, which are then used to guide a mesh deformation model using Neural Jacobian Fields (NJF) [AGK\*22].

We summarise the discussed methods in Table 5. These approaches demonstrate the strong potential of 2D vision models for matching structurally distinct shape categories, such as animal species. However, their limitation lies in the lack of fine-grained geometric awareness, where axiomatic or learned descriptors continue to perform strongly. Moreover, certain shapes may lack sufficiently discriminative textures (e.g., anatomical organs) or may not be sufficiently represented in the training distributions of these models, making them inapplicable. Future research could investigate hybrid strategies that integrate foundation model features

Method	FM	CP
Abdelreheem et al. [AEOW23]	SAM + DINO	SP
Diff3F [DMM24]	SD + DINO	SP
DenseMatcher [ZJZ*25]	SD + DINO	SP
Echomatch [XER*25]	DINO	SP
SRIF [SGJ*24]	DiffMorpher	D
DV-Matcher [CJH25]	DINO	D
Stable-SCore [LZY*25b]	SD + DINO	D
Morreale et al. [MAKM24]	DINO	O
Uzolas et al. [UEK25]	SD + DINO	O

**Table 5:** Overview of methods using *foundation models* for shape correspondence. They are categorized by the foundation model used for descriptor extraction (FM): DINO [CTM\*21, ODM\*23], Stable Diffusion [RBL\*22], SAM [KMR\*23], DiffMorpher [ZZX\*24], as well as the downstream correspondence pipeline (CP): Spectral (SP), Deformation-based (D), or Other (O).



**Figure 12:** Visualisation of the importance of *neighbourhood preservation* for correspondence computation by transferring triangulation from source to target shape via computing matchings. Matchings that violate neighbourhood preservation can lead to arbitrary distortions of triangles. Image source: [RB25].

with other types of descriptors to leverage their complementary strengths, or explore features coming from 3D foundation models [DHL\*25].

## 6.2. Neighbourhood-preservation

Neighbourhood preservation is an important aspect in shape matching as it leads to smoother correspondences (see Fig. 12). Furthermore, many downstream applications rely on the fact that matchings preserve neighbourhoods, including texture transfer, shape reconstruction, and many others. In the literature, neighbourhood preservation is also referred to as *geometric consistency*, a term which was introduced by Windheuser et al. [WSSC11].

Despite its importance, neighbourhood preservation is often overlooked. The resulting optimisation problems are usually harder to solve, and, as a consequence, scalability is limited, as neighbourhood-preserving approaches typically do not scale to shapes with thousands of triangles. For example, the LAP (not neighbourhood preserving) can be solved in polynomial time while the QAP (incorporates neighbourhood relations via costs) is generally NP-hard [RPW94].

Method	IF	S	TF	GNP	P
Takayama [Tak22]	✗	✗	✗	✓	✗
Schmidt et al. [SPK23]	✗	✗	✗	✓	✗
Rorberg et al. [RVBC23]	✗	✓	✗	✗	✗
Born et al. [BSCK21]	✗	✓	✗	✗	✗
Magnet et al. [MRSHO22]	✗	✓	(✓)	✗	✓
Xia et al. [XLGM24]	✗	✓	(✓)	✗	✗
SM-Comb [RSCB22]	✓	✗	✗	✓	(✓)
DiscoMatch [RAC*24]	✓	✗	✗	✓	✗
SpiderMatch [RB24]	✓	(✓)	✗	✗	(✓)
Roetzer et al. [REC*25]	✓	✓	✗	✗	(✓)
GeCo3D [RB25]	✓	✓	✗	✓	✓
SuPaMatch [EARB25]	✓	✓	✗	✓	✓
GC-PSM [ERE*24]	✓	✗	✗	✓	✓
GC-PPSM [EGR*24]	✓	✗	✗	✓	✓
Partial-GeCo [ERBC26]	✓	✓	✗	✓	✓

**Table 6:** Overview of *neighbourhood preserving* approaches and classification if methods are **(IF)** initialisation-free, are **(S)** scalable, are **(TF)** topologically flexible, are **(GNP)** globally neighbourhood preserving, or can solve **(P)** shape matching problems where at least one shape is partial. The scalability of SpiderMatch [RB24] depends on the input representation; see [RB25] for more details. Approaches [MRSHO22, XLGM24] do not show experiments on topologically noisy shapes but would, in theory, support such inputs. The formalism solved by SM-Comb [RSCB22] is not designed for partial-to-full or partial-to-partial shape matching, yet the authors show that it is applicable when closing shapes using triangle fans. Further, the approaches [RB24, REC\*25] do not show partial-to-full experiments, but their formalisms naturally support this case.

Approaches that do incorporate neighbourhood preservation often use combinatorial optimisation techniques. Thus, we pick up many of the aforementioned combinatorial methods in this section. Nevertheless, there are other lines of work which build on other ideas.

In the following, we categorise neighbourhood-preserving methods into local and global approaches; see also Table 6. Local approaches rely on initialisation in the form of sparse landmark correspondences or noisy input maps, whereas global approaches consider neighbourhood relations directly during optimisation of shape correspondences.

### 6.2.1. Local optimisation

Several lines perform local optimisation, or *refinement*, to obtain neighbourhood-preserving maps, or rather to convert existing maps into neighbourhood-preserving ones.

In recent years, Takayama [Tak22] and Schmidt et al. [SPK23] propose using piece-wise linear transformations (which can be seen as piece-wise plane embeddings) to refine sparse landmark correspondences into dense, neighbourhood-preserving matchings by building on so-called intrinsic triangulations, see [SSC19]. Other works use piece-wise Möbius transformations to obtain smoother maps [RVBC23], inspired by [LF09].

For shape pairs with genus greater than zero, [BSCK21] pro-

pose to use homology transfer to obtain high-quality maps. Here, the homology bases and their respective transfer functions act as low-pass filters, making the resulting maps smooth by filtering out local noisy/inconsistent matches. In [MRSHO22], authors propose to use Dirichlet energy minimisation in a functional map framework to obtain smooth maps. Further, [XLGM24] finds inliers of given noisy input correspondences by exploiting neighbourhood information, and subsequently, by utilising these inliers, they refine the outliers.

These refinement methods, by definition, rely on the quality of the input map and, therefore, matchings could in theory be arbitrarily bad. In contrast, global approaches are not dependent on initialisation.

### 6.2.2. Global approaches

Global neighbourhood preserving approaches refer to approaches that are initialisation-free, explicitly consider neighbourhood relations during optimisation, and provide (some form of) optimality certificates. Among the most important approaches in this category is the integer linear programming formalism proposed in [WSSC11] in which neighbourhoods of surface elements are preserved as a hard constraint and, as already mentioned, which shaped the term “geometric consistency”. Yet, this formalism is hard to solve and does not scale to higher shape resolutions. Motivated by this, [RSCB22, RAC\*24] proposed approximate solvers which yield improved scalability.

Other lines of work aim for novel formalisms. Among these, SpiderMatch [RB24] (see Sec. 4) achieves geometric consistency by representing a 3D shape as a cyclic, self-intersecting curve and by preserving intersections in their matching formalism. Building on this idea, [RB25] encodes each triangle of a 3D shape with a single cyclic curve, leading to independent matching problems for each curve. Geometric consistency is then enforced by coupling the problems at opposite edges such that the constraints always lead to a faithful surface-to-surface matching and, with that, to better geometric consistency compared to [RB24]. In addition, this formalism is faster to solve and, in practice, can be solved to global optimality. Motivated by the improved scalability, follow-up work [EARB25] builds on the idea of an alternative shape representation and proposes a two-stage formalism: first, transfer patches from source to target shape and second, find a dense surface-to-surface matching for each pair of corresponding patches. Yet, for both approaches, there are no theoretical guarantees of polynomial runtime, and thus it remains an open research question whether there exists a neighbourhood-preserving shape-matching approach that is globally optimal in polynomial time.

In contrast, [ERE\*24, EGR\*24] propose extensions of the formalism by Windheuser et al. [WSSC11] to partial-to-full [ERE\*24] and partial-to-partial shape matching [EGR\*24]. For partial-to-full shape matching [ERE\*24] relaxes the hard neighbourhood constraints proposed in [WSSC11]. For the more challenging partial-to-partial shape-matching task, [EGR\*24] proposes solving for the mean matching cost rather than the sum of matching costs. This effectively allows for solving for correspondences and unknown overlap at the same time. Yet, it is harder to solve due to the resulting non-linear objective function and requires solving multiple in-

teger linear programs. To address the challenges posed by the non-linear objective, [ERBC26] adapts [RB25] to the partial-to-partial setting and introduces an integer linear programming formalism that explicitly enforces geometric consistency. By integrating a predicted overlap region directly into the objective function, the resulting approach scales significantly better than the previous geometrically consistent method for partial-to-partial shape matching. In Section 6.4, we discuss additional approaches to partial shape matching.

### 6.3. Neural correspondence optimization

The machine learning and computer vision communities alike have become increasingly interested in approximating complex optimisation problems with neural networks [WCK\*25]. Deep learning provides a flexible mechanism for integrating intricate, data-driven priors to approximate what would otherwise be solved iteratively, such as the correspondence between two shapes. As we saw in Sec. 3.5, the optimisation to recover a functional map for a given set of spectral correspondences can be approximated with a diffusion model trained on a large collection of such maps [ZLG25, PLDO25]. [WJL\*24] recently observed that the analogous can be achieved for doubly stochastic “soft” correspondence matrices. This raises the question of whether more complex costs and objectives, such as neighbourhood constraints, can be incorporated into such a formulation, paying the cost of computationally intensive numerical optimisation upfront during training time.

### 6.4. Partial shape matching


As mentioned in previous chapters, the problem of matching two complete shapes has been explored in many works. However, partial shape matching, in which at least one shape is incomplete, has so far received less attention. In the context of partial shape matching, we distinguish between two cases: partial-to-full shape matching, in which only one of the shapes is partially observed, and partial-to-partial shape matching, in which both shapes are incomplete.

#### 6.4.1. Partial-to-full shape matching

In partial-to-full shape matching, we explore the task of a partially observed shape  $\mathcal{N}$  being matched on a full template shape  $\mathcal{M}$  (see Fig. 13 left).

**Partial functional maps.** In one of the first works to address partial-to-full shape matching, the authors observe a partial setting in the functional map context [RCB\*17] for isometric shape matching. For reference, for isometric shapes we observe a close to diagonal structure of the functional map  $C_{\mathcal{M}\mathcal{N}} = \phi_{\mathcal{N}}^\dagger \Pi_{\mathcal{N}\mathcal{M}} \phi_{\mathcal{M}}$  with an optimal point-wise map  $\Pi_{\mathcal{N}\mathcal{M}}$ . In the partial-to-full scenario, this is not the case, as some eigenfunctions of the full shape do not have a corresponding eigenfunction in the partial shape. Therefore, instead of a full diagonal, we observe a slanted diagonal structure with rank  $r$  of the partial functional map  $C_{\mathcal{M}\mathcal{N}}$ . The rank  $r$  can be estimated by comparing the surface areas of the two shapes [Wey11].

Similar to the full-to-full case, the partial functional map has



**Figure 13: Partial shape matching settings:** In the partial-to-full setting (left), a partial shape is matched to a template shape. In comparison, in the partial-to-partial setting, both shapes are only partially observed (right). The grey grid indicates that this part is missing, where red indicates that this part is observed but not in the overlapping region. Shapes are taken from FAUST [BRLB14], KIDS [RRBW\*14] and BeCoS [EEAX\*25] datasets.

Method	Setting	Framework	Type
DPFM [APO21]	P2F,P2P	PFM	S,U
Bensaid et al. [BRGK23]	P2F, P2P	PFM	S
Bensaid et al. [BK24]	P2F	Spectral	A
Bracha et al. [BDK24a]	P2F	Dist	U
Wormhole [BDK24b]	P2F	Dist	U
ULRSSM [CRB23]	P2F	PFM	U
Magnet et al. [MRSHO22]	P2F	Neigh	A
GC-PSM [ERE*24]	P2F	Neigh	A
GeCo3D [RB25]	P2F	Neigh	A
El Amrani et al. [EARB25]	P2F	Neigh	A
EchoMatch [XER*25]	P2P	PFM	S
GC-PPSM [ERE*24]	P2P	Neigh	A
Partial-GeCo [ERBC26]	P2P	Neigh	A

**Table 7: Categorization of Partial Shape Matching Methods:** We categorize partial shape matching methods in their setting (partial-to-full (P2F) or partial-to-partial (P2P)), their underlying framework (Partial Functional Maps, Spectral alignment, Distance-based or Neighbourhood-preserving) and their optimisation type (axiomatic (A), supervised (S), unsupervised (U)).

several key properties. One is semi-orthogonality, which ensures that the map preserves volume when transferring functions from the partial shape to the full one, but not necessarily in the opposite direction. The other is bijectivity, which requires that mapping a function from the partial shape to the full shape and then back recovers the original function. While supervised methods rely on the ground-truth functional map to guide training [APO21], these structural properties provide a foundation for unsupervised learning [APO21, CRB23], including the orthogonality losses

$$\|C_{\mathcal{M}\mathcal{N}}C_{\mathcal{M}\mathcal{N}}^\top - I_r\|_F^2 \text{ and } \|C_{\mathcal{N}\mathcal{M}}^\top C_{\mathcal{N}\mathcal{M}} - I_r\|_F^2, \quad (41)$$

and the bijectivity loss,

$$\|C_{\mathcal{M}\mathcal{N}}C_{\mathcal{N}\mathcal{M}} - I_r\|_F^2. \quad (42)$$

Here,  $I_r$  denotes the identity matrix with only the first  $r$  elements being 1. DPFM [APO21] introduces a cross-attention refinement module that enables feature communication across shapes. Bensaid et al. [BRGK23] propose learning a piecewise-smooth function over the surface: sparse landmark correspondences between a full shape and its partial counterpart are first established via feature similarity, and a neural representation is then optimised to fit these landmarks, enabling smooth interpolation between matched features that serve as anchors.

**Spectral alignment methods.** Rampini et al. [RTO\*19] align spectral representations derived from the Laplace–Beltrami operator (LBO). Building on this idea, Bensaid et al. [BK24] combine the standard LBO with its scale-invariant variant (SI-LBO), leveraging the complementary information from both spectra to better capture semantically meaningful curved regions.

**Distance-based methods.** Another line of work shows that using functional maps in the partial shape matching setting introduces unavoidable errors [BDK24a]. Therefore, the method primarily operates in the spatial domain rather than the spectral domain. The main loss preserves the geodesic distance (surface distance) between points on both shapes. Bracha et al. [BDK24b] tackle the problem that the geodesic distance of a partial shape of two points does not necessarily have the same length as the same two points on a full shape, due to possible holes. Therefore, reliable point pairs are identified and are later used in the loss formulation.

**Neighbourhood-preserving methods.** [MRSHO22, ERE\*24, EARB25, RB25] use integer neighbourhood-preserving techniques to find solutions for the partial-to-full shape matching problem, see also Sec. 6.2.2.

#### 6.4.2. Partial-to-partial shape matching

The more challenging problem in partial shape matching is the partial-to-partial setting, where we have to solve two problems in parallel: finding the overlapping region and determining correspondences in it (see Figure 13 right). Due to its complexity, it remains an underexplored problem and warrants further consideration.

So far, only a few supervised learning-based methods [APO21, BRGK23, XER\*25] exist that tackle the partial-to-partial shape matching problem. Two methods [XER\*25, APO21] mainly operate similarly to the partial-to-full or full-to-full setting, but add an overlap prediction module to determine the overlapping region. For two partial shapes  $\mathcal{M}$  and  $\mathcal{N}$  define a overlap prediction vector  $p_{\mathcal{M}} \in [0, 1]^{n_{\mathcal{M}}}$  and  $p_{\mathcal{N}} \in [0, 1]^{n_{\mathcal{N}}}$ , indicating the probability that a vertex is in the overlapping region or not. These predictions are compared with a binary cross entropy loss (BCE) with the ground truth overlap  $gt_{\mathcal{M}}$  and  $gt_{\mathcal{N}}$ :

$$\mathcal{L}_{\text{overlap}} = \frac{1}{2} [\text{BCE}(p_{\mathcal{M}}, gt_{\mathcal{M}}) + \text{BCE}(p_{\mathcal{N}}, gt_{\mathcal{N}})]. \quad (43)$$

DPFM [APO21] uses the predicted features directly in an

MLP to predict the overlap predictions  $p_{\mathcal{M}}$  and  $p_{\mathcal{N}}$ . In contrast, EchoMatch [XER\*25] introduces a correspondence reflection module that leverages cycle consistency of the point-wise maps,  $P_{\mathcal{M}} = \Pi_{\mathcal{M}\mathcal{N}}\Pi_{\mathcal{N}\mathcal{M}}$ , to verify whether a point returns to its original neighbourhood. It is further demonstrated that image-based DINOv2 features are effective for handling partial shapes. A different strategy is proposed in [BRGK23], where anchor correspondences are first established to enable smooth interpolation. For the partial-to-partial setting, two additional area-based losses are introduced: one to enforce similarity of overlapping regions across shapes, and another to encourage large overlap. [EGR\*24, ERBC26] propose an axiomatic neighbourhood-preserving solution for partial-to-partial shape matching, as mentioned in described 6.2.

**Metric adaptations.** Metrics must be properly adapted for the partial correspondence settings. In the partial-to-full (P2F) case, where the source shape  $\mathcal{M}$  is a partial observation and the target shape  $\mathcal{N}$  is complete, the metrics are computed identically to the full-to-full case. In contrast, the partial-to-partial (P2P) setting assumes that both the predicted and ground-truth correspondences are defined only on the overlapping regions of the source and target shapes. Let  $\Omega^* \subseteq \mathcal{M}$  denote the ground-truth overlap region, and  $\Omega \subseteq \mathcal{M}$  the predicted one. To ensure consistency across partial and full scenarios, all geodesic distances are normalised by the square root of the area of the full target shape,  $\sqrt{\text{Area}(\mathcal{N}_{\text{full}})}$ . Vertices in the ground-truth overlap that were not matched ( $i \in \Omega^* \setminus \Omega$ ), as well as vertices falsely matched outside the overlap ( $i \in \Omega \setminus \Omega^*$ ), are penalised with infinite error:  $e(i) = \infty$ . When evaluating the mean geodesic error, PCK, or AUC, only correspondences with finite error ( $e(i) < \infty$ ) are considered. This ensures that the metrics measure accuracy within the overlap area, while ignoring unmatched or undefined regions. Additional evaluation metrics for the P2P setting are discussed in [EEAX\*25, EGR\*24].

While previously discussed mesh-based approaches rely on explicit surface connectivity, the next section focuses on matching pairs of point clouds, where such connectivity information is unavailable. Structural priors can therefore only approximate connectivity.

#### 6.5. Point cloud matching

Several works have focused specifically on the point cloud setting. Such methods are generally also applicable to structured mesh representations but cannot leverage as strong a structural prior; thus, methods like those in Sec. 3 and 4 are preferred. Nevertheless, unstructured point cloud acquisitions are ubiquitous in real-world settings, and thus these methods remain of fundamental importance.

Point cloud correspondence methods generally rely on strong feature representations for matching because there are few structural priors to exploit. Many require ground-truth supervision; those that don't typically regularise the resulting deformation using the chamfer loss, which requires that the shapes already have a reasonable extrinsic alignment, or using piecewise ARAP deformations. Corrn3d [ZQZ\*21] predicts high-dimensional point-wise features for both point clouds, which are used to regress the permutation matrix between them. This permutation matrix is used to deform the point clouds. DPC [LGAR21] uses point-wise

features for self-construction and cross-construction, where self-construction encourages smooth correspondence maps and cross-construction unique matches. In [NO22], the authors use self- and cross-attention modules to predict a soft correspondence matrix, where, in comparison to CorNet3d and DCP, they also adopt a *dustbin* paradigm to discard correspondences. In [HDZ\*23], the authors use *multi-receptive-field point encoders* with different receptive radii to compute hierarchical features. SE-ORNet [DWL\*23] aligns the orientation of shapes to tackle the symmetry problem in human and animal shapes. In [DLZ24], the authors propose learning a set of templates with explicit structures that can later be used for better shape correspondences. Spectral methods can indeed also be applied in point cloud settings [CB23] due to the development of point cloud Laplacians [SC20], yielding feature encoders that can operate on both mesh and point clouds. The same holds for the Gaussian splatting representation [ZL26]. Alternatively, graphs are also well-suited for such multi-modal encoders [SWC\*22].

While these methods mostly handle the task of matching complete, noisy point clouds, some also target partiality. Leopard [LH22a] uses a Transformer architecture with self and cross-attention layers. While primarily designed for the rigid case, design a deformation module which predicts a warp function for every point in a dense point cloud for non-rigid settings. Training also proceeds in a supervised manner. Later, a rotation invariant method [YQH\*23] was introduced to prevent instabilities for unseen rotations. They introduce a rotation-invariant attention-based encoder-decoder architecture for local geometry and a global transformer with rotation-invariant cross-frame position awareness. Similar to Leopard, they predict a per-point flow to also tackle the deformable case. DiffReg [WJL\*24] builds upon these techniques, but proposes to recover a correspondence map through a diffusion process in the space of double-stochastic matrices.

## 7. Datasets

While large-scale 3D shape datasets such as ShapeNet [CFG\*15] and Objaverse [DSS\*23] have enabled progress in general shape understanding, they typically lack the dense ground-truth correspondences required for evaluating deformable shape matching. In line with prior work in this domain, we instead focus on datasets that contain deformable objects annotated with consistent correspondences across shape instances.

**Near-isometric.** Early 3D shape matching datasets such as the CAESER dataset [RDP99] provided thousands of human body scans annotated with hand-placed landmarks, offering an important benchmark but with limited ground-truth correspondences. To address this limitation, the SCAPE dataset [ASK\*05] introduced 71 registered meshes of a single subject in a single pose, providing dense correspondences but restricting variability since all shapes originate from one individual. The TOSCA dataset [BBK08] expanded diversity by including multiple object categories (e.g., humans, cats and dogs), yet its synthetic nature with fixed topology introduced a domain gap between real-world scans and artist-crafted meshes. Subsequent datasets such as FAUST [BRLB14] improved diversity with 10 individuals captured in multiple poses, using registrations to a common template to establish dense cor-


Setting	Dataset	NS	# US	TN	SC
F2F	FAUST [BRLB14]	✗	100	✗	♀
	SCAPE [ASK*05]	✗	71	✗	♀
	SMAL [ZKJB17]	✓	49	✗	♂/♀
	DT4D [MRSO22, LTT*21]	✗	2198	✗	♀/♂
	TOSCA [BBK08]	✗	82	✗	♀/♂
	KIDS [RRBW*14]	✗	32	✗	♀
	SHREC'20 [DLR*20]	✓	11	✗	♂
	SHREC'19 [MMR*19b]	✗	44	✗	♀
	BeCoS (F2F) [EEAX*25]	✓	2543	✗	♀/♂
	SURREAL [VRM*17]	✗	230K	✗	♀
	TOPKIDS [LRB*16, RRBW*14]	✗	32	✓	♀
	SCAPET [SH22, ASK*05]	✗	71	✓	♀
ExtFAUST [MWB25b, BRLB14]	✗	100	✓	♀	
P2F	SHREC'16 [BBK08, CRB*16]	✗	76	✗	♀/♂
	PFAUST [BDK23, BRLB14]	✗	10	✗	♀
	FARM partial [KMP*21]	✗	5	✗	♀
	PFARM [APO21, KMP*21]	✗	28	✗	♂
	SHREC'20 [DLR*20]	✓	14	✗	♀
BeCoS (P2F) [EEAX*25]	✓	2543	✗	♀/♂	
P2P	CP2P [APO21, CRB*16]	✗	76	✗	♀/♂
	PSMAL [EGR*24, ZKJB17]	✓	43	✗	♂
	BeCoS (P2P) [EEAX*25]	✓	2543	✗	♀/♂

**Table 8:** An overview of 3D deformable shape datasets. The datasets are categorized according to their shape matching setting: full-to-full (F2F), partial-to-full (P2F), and partial-to-partial (P2P). For each dataset, we report several characteristics, including non-isometry (NS), number of unique shapes (#US), topological noise (TN), and type of shape categories (SC), indicating whether the dataset contains humanoid shapes (♀) or animal shapes (♂).

respondences, while DFAUST extended FAUST to dynamic sequences, yielding over 40,000 aligned meshes, though its focus on adults limited generalisation to children. To address this gap, the KIDS dataset [RRBW\*14] introduced eight high-resolution synthetic models (80K vertices) of a child in near-isometric poses.

Other datasets emerged from parametric models such as SMPL [LMR\*15], its variants [RTB17, PCG\*19], and the COMA model [RBSB18], enabling the creation of very large collections such as SURREAL [VRM\*17], which contains 230,000 meshes. However, these datasets suffer from biases introduced by repeated use of identical triangulations across thousands of meshes. More recent efforts, such as DT4D [LTT\*21], extend diversity by providing animatable shapes from multiple categories, such as different animal species, with consistent triangulations across meshes of the same class. Despite these advances, nearly all of these datasets remain constrained to the near-isometric case and, therefore, fail to fully capture the complexities of real-world shape variability. This highlights the need for datasets that provide reliable correspondences under non-isometric deformations.

**Beyond isometry.** A more general and challenging setting for 3D shape matching arises in the non-isometric regime, where datasets need to provide correspondences across fundamentally different



**Figure 14:** The BeCoS dataset [EEAX\*25] connects seven existing 3D shape matching datasets [BRLB14,ZKJB17,ASK\*05,BBK08,LTT\*21,RRBW\*14,DLR\*20], enabling the propagation of dense correspondences across a wide range of categories, including humanoids, cats, elephants, and others. Furthermore, it provides annotations for partial correspondences, supporting both partial-to-full (P2F) and partial-to-partial (P2P) matching settings. Image source: [EEAX\*25].

object categories, such as establishing mappings between a cat and a horse. Constructing such datasets is particularly difficult because non-isometric correspondences are inherently ill-defined. While one may attempt to align overall body structures, it is unclear how to meaningfully associate, for example, the trunk of an elephant with the body of a cat. Early efforts in this direction include the SMAL dataset [ZKJB17], which introduced a parametric model for multiple animal categories trained on 41 scans. While this dataset offers cross-category representations, it relies on fitting a single parametric model across all meshes, which limits its expressiveness and its ability to capture the fine geometric details unique to each class of shapes. Another step forward is SHREC'20 [DLR\*20], which explicitly defines correspondences across different animal categories, such as elephants and cats. However, it only provides sparse correspondences, thereby limiting its utility for dense shape analysis. More recently, Magnet et al. [MRSHO22] addressed this gap by manually establishing dense correspondences across most categories in the DT4D dataset [LTT\*21], thereby creating the first dataset to offer dense non-isometric mappings. Building on this idea, BeCoS [EEAX\*25] further expanded the scope by leveraging seven existing isometric datasets. Instead of introducing new shapes, BeCoS builds upon seven existing isometric datasets and manually generates dense correspondences between a set of selected template meshes, effectively enabling cross-category mappings. This extension expands deformable shape analysis to the non-isometric setting, resulting in a large, interconnected collection of shapes, as illustrated in Fig. 14.

**Challenging connectivity and topological noise.** Another line of research in 3D shape matching has focused on datasets designed to introduce specific challenges, particularly related to mesh connectivity and topological noise. Variations in mesh connectivity can arise from anisotropy, where parts of the same surface are represented with coarse triangulation, while others use much finer discretisation, making correspondence estimation substantially more difficult. To evaluate generalisation to such scenarios, SHREC'19 [MMR\*19b] introduced human scans with altered mesh connectivity. Similarly, Donati et al. [DCO22] created a

challenging setting by remeshing both the SCAPE [ASK\*05] and FAUST [BRLB14] datasets anisotropically.

A related challenge is topological noise, which includes the presence of holes, self-intersections, or other surface artefacts. SHREC'10 [BBC\*10] addressed this by extending TOSCA with artificially injected topological noise, though it inherited TOSCA's broader limitations. Building on this idea, TOPKIDS [LRB\*16] extended the KIDS dataset [RRBW\*14] by synthetically introducing topological perturbations, while TACO [PMM24] extended TOSCA using a non-connectivity-preserving remeshing strategy. In TOCA [PMM24], mesh connectivities were intentionally altered, even among shapes of the same class. Leveraging the 3D alignment between original and remeshed shapes as well as TOSCA's ground-truth correspondences, the authors generated new ground-truth correspondences across the perturbed surfaces. Using a different approach, SCAPET [SH22] and ExtFAUST [BBW\*21,MRWB23] build upon existing datasets, SCAPE [ASK\*05] and FAUST [BRLB14], respectively. These methods modify the topology of the shapes by either adding shortcut tunnels or by merging nearby surface regions, effectively connecting previously separate parts of the mesh. Together, these datasets emphasise the importance of evaluating shape matching methods under realistic noise and connectivity variations, moving beyond the simplified settings of earlier benchmarks.

**Partial 3D shapes datasets.** Beyond the full-to-full (F2F) setting discussed so far, the availability of datasets for the other two key settings, partial-to-full (P2F) and partial-to-partial (P2P), as described in Sec. 6.4.1, is equally important. The first benchmark in this direction was SHREC'16 [CRB\*16], which extended the TOSCA dataset [BBK08] to include partial shapes. This was later followed by PFAUST [BDK23] and PFARM [APO21], based respectively on the FAUST [BRLB14] and FARM [KMP\*21] datasets, while SHREC'20 [DLR\*20] also included a subset of partial-to-full correspondence pairs.

In contrast, P2P shape matching is considerably more challenging, as it requires establishing correspondences between two partial shapes with unknown overlap. The first benchmark for this setting, CP2P [APO21], was introduced by extending the SHREC'16

CUTS dataset [CRB\*16]. More recently, PSMAL [EGR\*24] extended the SMAL dataset [ZKJB17] to incorporate partial-to-partial correspondences across animal categories. However, in both P2F and P2P settings, partiality has traditionally been simulated by artificially removing mesh parts or introducing holes, which does not accurately reflect the challenges of real-world partial observations. To address this limitation, BeCoS [EEAX\*25] introduces a more realistic framework for generating partial shapes through ray-casting, enabling the creation of P2F, P2P, and F2F correspondence pairs within a unified framework.

**Performance of 3D shape matching methods.** Over the past few years, performance on well-established benchmarks has largely saturated, as indicated by the marginal improvements achieved by recent methods (e.g. the mean geodesic error ( $\times 100$ ) on FAUST [BRLB14] has remained at approximately 1.6 [CLB24, BXNL24]). This stagnation highlights the need for more challenging benchmarks, preferably based on real-world shapes obtained from 3D scans, in order to further drive methodological progress. In addition, the reliability of ground-truth annotations in several commonly used datasets, including DT4D [LTT\*21], BeCoS [EEAX\*25], and SHREC'19 [MMR\*19a], may contain errors, suggesting that improvements in numerical accuracy do not necessarily reflect genuine advances in correspondence quality [LDO22]. These challenges are even more pronounced in the partial shape setting, where benchmarks and dedicated methods remain scarce, despite this scenario being the most representative of real-world data. At the time of writing, state-of-the-art results in full-to-full, partial-to-full and partial-to-partial matching scenarios, were obtained by [BXNL24], [CLB24] and [XER\*25], respectively.

## 8. Practical applications

Shape correspondence lies at the core of numerous applications involving 3D meshes. In this section, we highlight some of the most prominent use cases. Examples of these applications are illustrated in Figure 15.

**Information transfer.** A well-known application of shape correspondence is the transfer of annotations between shapes, which enables numerous tasks in graphics, animation, and robotics. A classic example is texture transfer, where a texture defined on a source shape is projected onto a target shape, which is often used to visually inspect the quality of correspondences. Segmentation transfer enables the propagation of part labels across a dataset of registered shapes, allowing a small number of manual annotations to be scaled to a shape collection [ASOW23]. Deformation transfer propagates articulations, such as pose changes in human bodies, and can extend to complete deformation systems, including rigging properties [AGR\*16, MMMC21a] or the mesh structure itself [MMM\*20, SP04, LHP25]. Beyond graphics, shape correspondence finds applications in robotics for transferring contact point annotations between objects to guide robotic arm grasping [ZJZ\*25].

**Shape interpolation.** Shape interpolation (or morphing) aims to generate a continuous sequence of intermediate shapes that

smoothly deform a source geometry into a target. Accurate correspondences help preserve semantic structure and geometric detail throughout the interpolation process. Recent approaches vary in how they incorporate correspondence: some treat it as external supervision or initialization [SCC\*25, SCC\*26], others solve for correspondence and interpolation jointly [ENK\*21], or couple spectral and spatial representations to improve consistency [CEEA\*24].

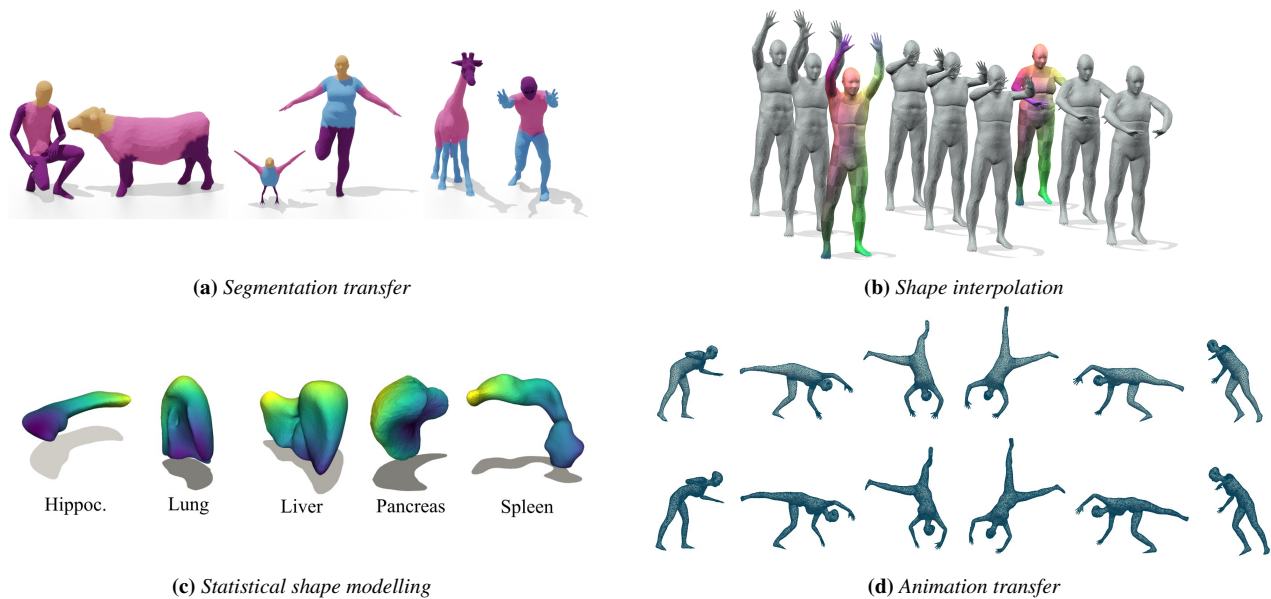
**Statistical shape modelling and medical applications.** Statistical Shape Models (SSMs) capture shape variability within a category by learning low-dimensional representations from populations of 3D shapes. This process relies on dense correspondences to ensure consistent geometry alignment across the population.

Two broad modelling paradigms exist for building SSMs. Template-based methods assume shared mesh topology and model deformation within a fixed coordinate system, with prominent examples including SMPL [LMR\*15], MANO [RTB17], and 3DMM [EST\*20]. Recent extensions adapt this paradigm to animals [LLL\*24] or anatomical skeletons [KZBP22, KWS\*23] by combining template correspondences with learned statistical priors. In contrast, groupwise methods jointly estimate correspondences and a shape model to avoid template bias, using encoder–decoder architectures [IE23] or functional maps [BBH\*23, EACB24]. In both cases, correspondence is the key component enabling statistical analysis.

A major application domain of SSMs is medical data, where correspondence establishes anatomical mappings for disease classification and population-level studies [AE24, IKE25]. Functional map-based frameworks have proven effective for unsupervised correspondence estimation when mesh representations are available [KDMA\*22, BBH\*23, EACB24]. In clinically assisted interventions, however, data may arise from partial point clouds acquired via laparoscopic imaging, ultrasound, X-ray [YSL23, GVS\*24, GVNA25], or time-of-flight (ToF) depth sensors [BWC\*23, WBC\*25]. Template-based deformation approaches have improved robustness in applications including abdominal organs [BRW23], cortical surfaces [BRW24], and orthopaedic fracture reduction planning [YLS\*25]. Atlas construction remains another important research direction [KBFG25]. For a broader overview, we refer to a survey by Zhang et al. [ZLW25] and the extensive MedShapeNet dataset [LZY\*25a].

**Shape retrieval.** Shape retrieval seeks to identify shapes in a database that best match a query despite variations in pose, deformation, or partiality. Recent methods increasingly integrate correspondence estimation into the retrieval process. Dense maps can define similarity metrics directly between query and candidates [DZZ\*23], while keypoint or part-level correspondences support retrieval with subsequent deformation [ZZD\*24, DZW\*24]. Other approaches extend correspondence-based metrics to multimodal retrieval settings [GDN22, GRLD24]. Overall, correspondence helps retrieval focus on geometric rather than purely visual similarity.

**Dynamic reconstruction.** When reconstructing a dynamic scene into a sequence of 3D frames, the underlying problem implicitly



**Figure 15: Practical applications** of shape correspondence span numerous tasks in graphics and geometry processing. **(a)** Segmentation transfer propagates semantic part annotations across a collection of registered shapes. **(b)** Shape interpolation uses correspondences to generate smooth transitions between source and target shapes. **(c)** Statistical shape modelling uses correspondences to build low-dimensional shape spaces, including anatomical models of organs and tissues. **(d)** Animation transfer allows to propagate motion sequences across registered shapes. Figures adapted from [CEEA\*24, AEOW23, EACB24, MMMC21b].

assumes dense correspondence between the frames. Instead of finding correspondences between image keypoints, most recent methods rely on a deformation field between a reconstructed template or an initial frame and subsequent frames [SZT\*23, EL26]. While most methods restrict correspondence to be implicit through the deformation model, it can also be used directly to impose geometric constraints, such as maintaining curvature consistency [COF\*25].

**Data alignment.** Correspondence is of significant importance for annotating 3D data. For example, it is common to decorate acquisition with an aligned template, providing a semantic meaning to otherwise unordered points [BXP\*22, WHG\*24]. Correspondence can also be used to unify different datasets in a common semantic frame [MMR\*19b].

**Measures and shape analysis.** By bringing shapes into correspondence, it is also possible to compare intrinsic properties of objects, especially in terms of measures. For example, correspondence methods are adopted in anthropometry, the science that studies the measurements and proportions of the human body [HLT\*25, ZKA\*22].

## 9. Open challenges

Despite the extensive progress made in non-rigid shape correspondence in recent years, several fundamental challenges remain unresolved and limit its applicability to arbitrary real-world data.

First, existing correspondence methods still struggle with non-trivial partiality and topological changes in the geometry. Both of

these are commonly encountered in 3D scans that contain occlusions or self-intersections. While a few approaches explicitly tackle the partial case (see Sec. 6.4.1 and Sec. 6.4.2), they require extensive training data or rely on hard assumptions to work. Even more challenging is the case of topological changes. The spectral properties can change significantly under different topologies, even if the local geometry is similar [LRB\*16]. This makes the line of spectral methods incredibly sensitive to topological noise, and most papers do not even show results in this direction. In addition, not many datasets, and only very small ones, contain topologically noisy shapes with ground-truth correspondences. For this reason, supervised learning methods cannot be applied in this domain, while unsupervised methods struggle due to violations of structural assumptions arising from changes shape topology.

Another major challenge is the lack of large, diverse datasets. The most famous and commonly used isometric datasets, FAUST [BRLB14] and SCAPE [ASK\*05], are synthetic; they consist only of human shapes and are generated from a fixed template. This leads to near-isometric and very clean meshes, resulting in benchmarks that recent methods have begun to saturate. On the other hand, there is no large collection of alternatives. This is in part due to the challenge of annotating datasets at scale; the largest existing datasets, therefore, still consist of synthetic data. As a result, methods are not sufficiently challenged by the lack of realism in noise and data distributions. The implication that almost all existing data consist of clean shapes, with mostly smooth surfaces and minimal geometric irregularities, is that generalisation to real-world settings is limited. The recent dataset BeCoS [EEAX\*25]

made a step towards realism by simulating extreme partiality using ray casting, but also found that existing methods already struggle under much simpler transformations, such as arbitrary rotations. While no standardised benchmarks exist to test such cases, the findings suggest that the current generation of correspondence methods is unlikely to be widely generalisable in real-world scenarios.

#### Connecting quantitative metrics to downstream applications.

A further challenge that cuts across all three methodological families is the disconnect between what current benchmarks measure and downstream application requirements. Most methods are primarily evaluated on geodesic accuracy, yet applications such as deformation transfer [SP04, VOM25], texture mapping, and statistical shape modelling [LMR\*15, MBT\*23] also require smooth and neighbourhood-preserving maps. While recent developments in unsupervised spectral methods enable the curation of datasets with geometric descriptors, the resulting pointwise maps lack such guarantees (see Sec. 3). Combinatorial methods offer a promising way of estimating smooth maps from these descriptors, but remain limited in scalability [RB25, EARB25]. Closing this gap, either through scalable formulations with geometric guarantees or by employing metrics that go beyond geodesic error, remains an important open problem. In domain-specific settings such as medical image analysis, a further issue is that purely geometric descriptors may be insufficient or prohibitively difficult to obtain. Incorporating prior knowledge from domain experts, such as anatomical landmarks or physiological priors, can improve the ability to produce meaningful correspondences and ultimately drive application impact [BBH\*23, EACB24, BRW23].

**Potential future directions.** Problems with partiality, topological changes, and other noise are closely related to implicit assumptions that most methods make about the solution, such as injectivity. It may be necessary to make a fundamental change in the problem's modelling to overcome these challenges reliably. This holds for both optimisation- and learning-based approaches, which fall in the same categories and rely on the same basic frameworks. To evaluate progress in this direction, new datasets, especially large-scale ones based on real-world data, are needed to capture the diversity of possible data and avoid overfitting to the small data pool. A potential solution might be to incorporate foundation models from other modalities, leveraging their extensive knowledge of the world to compensate for ill-posed settings and a lack of diverse training data. When correspondence methods become robust in real-world scenarios, they will enable a wide range of different applications with noisy data, for example, in medical settings or dynamic reconstruction, that would greatly benefit from correspondence information but cannot provide the clean meshes expected by most methods at the moment.

## 10. Conclusions

The report provides a broad overview of recent advances in non-rigid 3D shape correspondence. We discuss three recent branches of techniques which have been emphasised over the last few years: spectral methods based on functional maps, combinatorial formulations with discrete constraints, and deformation-based techniques. While the field of non-rigid 3D shape correspondence estimation

has seen considerable advances in long-standing paradigms, it is now also fuelled by the recent progress in related computer vision fields and machine learning, including research on foundation models and quantum computing.

We conclude that further progress depends substantially on the availability of new datasets that cover a wide variety of scenarios, especially for learning-based techniques that rely on large training datasets. We also foresee that the outlined open challenges, e.g., partial-to-partial shape matching, multi-shape matching, and handling different kinds of noise in the input data, such as topological changes, will keep the current community and the new generation of researchers curious and excited. The field is continuing to grow, and we are excited for the next generation of robust techniques and tools, broadly applicable across 3D computer vision, graphics, and related fields in diverse scenarios and applications.

## 11. Acknowledgements

This work is supported by the ERC starting grant no. 101160648 (Harmony); the DFG grant LA 5191/2-1; the UK Royal Society grant NIF/R1/254128; the DFG (German Research Foundation) project number 534951134; the ERC advanced grant no. 884679 (SIMULACRON); JSPS KAKENHI No. 21H05054, 24H00742, and 24H00748; the Federal Ministry of Research, Technology and Space of Germany and the state of North Rhine-Westphalia as part of the Lamarr Institute for Machine Learning and Artificial Intelligence.

## References

- [AE24] ADAMS J., ELHABIAN S.: Point2ssm: Learning morphological variations of anatomies from point clouds. In *Proceedings of the International Conference on Learning Representations* (2024). 24
- [AEOW23] ABDELREHEEM A., ELDESOKAY A., OVSJANIKOV M., WONKA P.: Zero-shot 3d shape correspondence. In *Proceedings of the SIGGRAPH Asia Conference* (2023), pp. 1–11. 18, 25
- [AGK\*22] AIGERMAN N., GUPTA K., KIM V. G., CHAUDHURI S., SAITO J., GROUEIX T.: Neural jacobian fields: Learning intrinsic mappings of arbitrary meshes. *ACM Transactions on Graphics* 41, 4 (2022), 109:1–109:17. 17, 18
- [AGR\*16] AVRIL Q., GHAFOURZADEH D., RAMACHANDRAN S., FALLAHDOST S., RIBET S., DIONNE O., DE LASA M., PAQUETTE E.: Animation setup transfer for 3d characters. *Computer Graphics Forum* 35, 2 (2016), 115–126. 24
- [ALC20] AYGÜN M., LÄHNER Z., CREMERS D.: Unsupervised dense shape correspondence using heat kernels. In *Proceedings of the International Conference on 3D Vision* (2020), pp. 573–582. 9
- [AO22] ATTAIKI S., OVSJANIKOV M.: Ncp: Neural correspondence prior for effective unsupervised shape matching. In *Proceedings of the Conference on Neural Information Processing Systems* (2022), vol. 35, pp. 28842–28857. 10
- [AO23] ATTAIKI S., OVSJANIKOV M.: Understanding and improving features learned in deep functional maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), pp. 1316–1326. 8, 10, 12
- [APO21] ATTAIKI S., PAI G., OVSJANIKOV M.: Dpfm: Deep partial functional maps. In *Proceedings of the International Conference on 3D Vision* (2021). 8, 9, 10, 16, 20, 21, 22, 23
- [ARV07] AMBERG B., ROMDHANI S., VETTER T.: Optimal step non-rigid icp algorithms for surface registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2007), pp. 1–8. 14, 15
- [AS22] ABBAS A., SWOBODA P.: Fastdog: Fast discrete optimization on gpu. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022). 12
- [ASC11] AUBRY M., SCHLICKWEI 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* (2011). 4, 8, 10, 18
- [ASK\*05] ANGUELOV D., SRINIVASAN P., KOLLER D., THRUN S., RODGERS J., DAVIS J.: Scape: Shape completion and animation of people. In *Proceedings of the SIGGRAPH Conference* (2005), pp. 408–416. 22, 23, 25
- [ASOW23] ABDELREHEEM A., SKOROKHOV I., OVSJANIKOV M., WONKA P.: Satr: Zero-shot semantic segmentation of 3d shapes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023), pp. 15166–15179. 24
- [ATO\*24] ANTIĆ D., TIWARI G., OZCOMLEKCI B., MARIN R., PONS-MOLL G.: Close: A 3d clothing segmentation dataset and model. In *Proceedings of the International Conference on 3D Vision* (2024), pp. 591–601. 16
- [AWLB17] ACHENBACH J., WALTEMATE T., LATOSCHIK M. E., BOTSCH M.: Fast generation of realistic virtual humans. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology* (2017), pp. 1–10. 16
- [BBC\*10] BRONSTEIN A., BRONSTEIN M. M., CASTELLANI U., DUBROVINA A., GUIBAS L. J., HORAUD R. P., KIMMEL R., KNOS-SOW D., VON LAVANTE E., MATEUS D., ET AL.: Shrec 2010: Robust correspondence benchmark. In *Eurographics Workshop on 3D Object Retrieval* (2010), vol. 10, pp. 087–091. 23
- [BBH\*23] BASTIAN L., BAUMANN A., HOPPE E., BÜRGIN V., KIM H. Y., SALEH M., BUSAM B., NAVAB N.: S3m: Scalable statistical shape modeling through unsupervised correspondences. In *Medical Image Computing and Computer Assisted Intervention* (2023), pp. 459–469. 10, 24, 26
- [BBK08] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: *Numerical geometry of non-rigid shapes*, 1 ed. Springer Publishing Company, Incorporated, 2008. 22, 23
- [BBW\*21] BASSET J., BOUKHAYMA A., WUHRER S., MULTON F., BOYER E.: Neural human deformation transfer. In *Proceedings of the International Conference on 3D Vision* (2021), pp. 545–554. 23
- [BDK23] BRACHA A., DAGÈS T., KIMMEL R.: On partial shape correspondence and functional maps. *CoRR* (2023). 22, 23
- [BDK24a] BRACHA A., DAGÈS T., KIMMEL R.: On unsupervised partial shape correspondence. In *Proceedings of the Asian Conference on Computer Vision* (2024), pp. 4488–4504. 20, 21
- [BDK24b] BRACHA A., DAGÈS T., KIMMEL R.: Wormhole loss for partial shape matching. In *Proceedings of the Conference on Neural Information Processing Systems* (2024). 20, 21
- [Ber92] BERTSEKAS D. P.: Auction algorithms for network flow problems: A tutorial introduction. *Computational Optimization and Applications* 1, 1 (1992), 7–66. 4
- [BK10] BRONSTEIN M. M., KOKKINOS I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2010), pp. 1704–1711. 8
- [BK24] BENSÄID D., KIMMEL R.: A multi-spectral geometric approach for shape analysis. *Journal of Mathematical Imaging and Vision* 66, 4 (2024), 606–615. 20, 21
- [BLG\*21] BENKNER M. S., LÄHNER Z., GOLYANIK V., WUNDERLICH C., THEOBALT C., MOELLER M.: Q-match: Iterative shape matching via quantum annealing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2021). 11, 13
- [BM92] BESL P. J., MCKAY N. D.: Method for registration of 3-d shapes. In *Sensor Fusion IV: Control Paradigms and Data Structures* (1992), vol. 1611, pp. 586–606. 2, 5, 12, 14, 15
- [BMP01] BELONGIE S., MALIK J., PUZICHA J.: Matching shapes. In *Proceedings of the IEEE International Conference on Computer Vision* (2001), vol. 1, pp. 454–461. 11
- [BMR\*20] BROWN T., MANN B., RYDER N., SUBBIAH M., KAPLAN J., DHARIWAL P., NEELAKANTAN A., SHYAM P., SASTRY G., ASKELL A., ET AL.: Language models are few-shot learners. *Proceedings of the Conference on Neural Information Processing Systems* 33 (2020), 1877–1901. 18
- [BN03] BELKIN M., NIYOGI P.: Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation* 15, 6 (2003), 1373–1396. 4
- [BPC\*11] BOYD S., PARIKH N., CHU E., PELEATO B., ECKSTEIN J., ET AL.: Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning* 3, 1 (2011), 1–122. 15
- [BRGK23] BENSÄID D., ROTSTEIN N., GOLDENSTEIN N., KIMMEL R.: Partial matching of nonrigid shapes by learning piecewise smooth functions. *Computer Graphics Forum* 42, 5 (2023), e14913. 20, 21
- [BRLB14] 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* (2014), pp. 3794–3801. 2, 20, 22, 23, 24, 25
- [BRW23] BONGRATZ F., RICKMANN A.-M., WACHINGER C.: Abdominal organ segmentation via deep diffeomorphic mesh deformations. *Scientific Reports* 13, 1 (2023), 18270. 7, 24, 26
- [BRW24] BONGRATZ F., RICKMANN A.-M., WACHINGER C.: Neural deformation fields for template-based reconstruction of cortical surfaces from mri. *Medical Image Analysis* 93 (2024), 103093. 24

- [BSCK21] BORN J., SCHMIDT P., CAMPEN M., KOBBELT L.: Surface map homology inference. *Computer Graphics Forum* 40, 5 (2021), 193–204. [19](#)
- [BST20] BERNARD F., SURI Z. K., THEOBALT C.: Mina: Convex mixed-integer programming for non-rigid shape alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 13826–13835. [11](#), [12](#)
- [BSTPM20a] BHATNAGAR B. L., SMINCHISESCU C., THEOBALT C., PONS-MOLL G.: Combining implicit function learning and parametric models for 3d human reconstruction. In *Proceedings of the European Conference on Computer Vision* (2020), pp. 311–329. [17](#)
- [BSTPM20b] BHATNAGAR B. L., SMINCHISESCU C., THEOBALT C., PONS-MOLL G.: Loopreg: Self-supervised learning of implicit surface correspondences, pose and shape for 3d human mesh registration. *Proceedings of the Conference on Neural Information Processing Systems* 33 (2020), 12909–12922. [17](#)
- [BTGT19] BERNARD F., THUNBERG J., GONCALVES J., THEOBALT C.: Synchronisation of partial multi-matchings via non-negative factorisations. *Pattern Recognition* 92 (2019), 146–155. [13](#)
- [BTL\*23] BHATIA H., TRETSCHK E., LÄHNER Z., BENKNER M. S., MOELLER M., THEOBALT C., GOLYANIK V.: Ccuantum: Cycle-consistent quantum-hybrid matching of multiple shapes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), pp. 1296–1305. [5](#), [11](#), [13](#)
- [BTM18] BERNARD F., THEOBALT C., MOELLER M.: Ds\*: Tighter lifting-free convex relaxations for quadratic matching problems. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2018). [11](#)
- [BTST19] BERNARD F., THUNBERG J., SWOBODA P., THEOBALT C.: Hippo: Higher-order projected power iterations for scalable multi-matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2019), pp. 10284–10293. [8](#)
- [BV99] BLANZ V., VETTER T.: A morphable model for the synthesis of 3d faces. In *Proceedings of the SIGGRAPH Conference* (1999), pp. 187–194. [16](#)
- [BVZ01] BOYKOV Y., VEKSLER O., ZABIH R.: Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23, 11 (2001), 1222–1239. [13](#)
- [BWC\*23] BASTIAN L., WANG T. D., CZEMPIEL T., BUSAM B., NAVAB N.: Disguiser: Holistic face anonymization for the operating room. *International Journal of Computer Assisted Radiology and Surgery* 18, 7 (2023), 1209–1215. [24](#)
- [BXNL24] BASTIAN L., XIE Y., NAVAB N., LÄHNER Z.: Hybrid functional maps for crease-aware non-isometric shape matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 3313–3323. [4](#), [8](#), [10](#), [24](#)
- [BXP\*22] BHATNAGAR B. L., XIE X., PETROV I. A., SMINCHISESCU C., THEOBALT C., PONS-MOLL G.: Behave: Dataset and method for tracking human-object interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 15935–15946. [25](#)
- [CB22] CAO D., BERNARD F.: Unsupervised deep multi-shape matching. In *Proceedings of the European Conference on Computer Vision* (2022), pp. 55–71. [8](#), [13](#)
- [CB23] CAO D., BERNARD F.: Self-supervised learning for multimodal non-rigid 3d shape matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), pp. 17735–17744. [22](#)
- [CEEA\*24] CAO D., EISENBERGER M., EL AMRANI N., CREMERS D., BERNARD F.: Spectral meets spatial: Harmonising 3d shape matching and interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 3658–3668. [10](#), [24](#), [25](#)
- [CFG\*15] CHANG A. X., FUNKHOUSER T., GUIBAS L., HANRAHAN P., HUANG Q., LI Z., SAVARESE S., SAVVA M., SONG S., SU H., XIAO J., YI L., YU F.: Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012* (2015). [22](#)
- [CJH25] CHEN Z., JIANG P., HUANG R.: Dv-matcher: Deformation-based non-rigid point cloud matching guided by pre-trained visual features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2025), pp. 27264–27274. [18](#)
- [CLB24] CAO D., LÄHNER Z., BERNARD F.: Synchronous diffusion for unsupervised smooth non-rigid 3d shape matching. In *Proceedings of the European Conference on Computer Vision* (2024), pp. 262–281. [10](#), [24](#)
- [COF\*25] CHEN D., OBERSON B., FELDMANN I., SCHREER O., HILSMANN A., EISERT P.: Adaptive and temporally consistent gaussian surfels for multi-view dynamic reconstruction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision* (2025), pp. 742–752. [25](#)
- [CPMAMN22] CORONA E., PONS-MOLL G., ALENYA G., MORENO-NOGUER F.: Learned vertex descent: A new direction for 3d human model fitting. In *Proceedings of the European Conference on Computer Vision* (2022), pp. 146–165. [15](#), [17](#)
- [CRA\*17] COSMO L., RODOLÀ E., ALBARELLI A., MÉMOLI F., CREMERS D.: Consistent partial matching of shape collections via sparse modeling. *Computer Graphics Forum* 36, 1 (2017), 209–221. [13](#)
- [CRB\*16] COSMO L., RODOLÀ E., BRONSTEIN M. M., TORSSELLO A., CREMERS D., SAHILLIOĞLU Y., ET AL.: Shrec'16: Partial matching of deformable shapes. *Eurographics Workshop on 3D Object Retrieval* (2016). [22](#), [23](#), [24](#)
- [CRB23] CAO D., ROETZER P., BERNARD F.: Unsupervised learning of robust spectral shape matching. *ACM Transactions on Graphics* 42, 4 (2023), 1–15. [4](#), [8](#), [10](#), [12](#), [20](#)
- [CRB24] CAO D., ROETZER P., BERNARD F.: Revisiting map relations for unsupervised non-rigid shape matching. In *Proceedings of the International Conference on 3D Vision* (2024), pp. 1371–1381. [1](#), [10](#)
- [CTM\*21] CARON M., TOUVRON H., MISRA I., JÉGOU H., MAIRAL J., BOJANOWSKI P., JOULIN A.: Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2021), pp. 9650–9660. [17](#), [18](#)
- [CWY\*23] CHEN G., WANG M., YUAN L., YANG Y., YUE Y.: Re-thinking point cloud registration as masking and reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023), pp. 17717–17727. [15](#), [16](#)
- [D-W26] D-WAVE QUANTUM INC.: D-wave quantum: Quantum computing systems and services. <https://www.dwavequantum.com>, 2026. Accessed: 2026-03-19. [13](#)
- [DCMO22] DONATI N., CORMAN E., MELZI S., OVSJANIKOV M.: Complex functional maps: A conformal link between tangent bundles. *Computer Graphics Forum* 41, 1 (2022), 317–334. [8](#)
- [DCO22] DONATI N., CORMAN E., OVSJANIKOV M.: Deep orientation-aware functional maps: Tackling symmetry issues in shape matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 742–751. [23](#)
- [DFS99] DALE A. M., FISCHL B., SERENO M. I.: Cortical surface-based analysis: I. segmentation and surface reconstruction. *NeuroImage* 9, 2 (1999), 179–194. [16](#)
- [DGF\*19] DEPRELLE T., GROUEIX T., FISHER M., KIM V., RUSSELL B., AUBRY M.: Learning elementary structures for 3d shape generation and matching. *Proceedings of the Conference on Neural Information Processing Systems* 32 (2019). [16](#)
- [DHL\*25] DU K., HU J., LI H., XU H., HUANG H., FU C.-W., LIU S.: Hierarchical neural semantic representation for 3d semantic correspondence. In *Proceedings of the SIGGRAPH Asia Conference* (2025), pp. 1–11. [18](#)

- [DK11] DUBROVINA A., KIMMEL R.: Approximately isometric shape correspondence by matching pointwise spectral features and global geodesic structures. *Advances in Adaptive Data Analysis* 3, 01n02 (2011), 203–228. [11](#)
- [DLB\*23] DRÖGE H., LÄHNER Z., BAHAT Y., MARTORELL NADAL O., HEIDE F., MÖLLER M.: Kissing to find a match: Efficient low-rank permutation representation. In *Proceedings of the Conference on Neural Information Processing Systems* (2023), vol. 36, pp. 48459–48471. [5](#), [11](#)
- [DLR\*20] DYKE R. M., LAI Y.-K., ROSIN P. L., ZAPPALÀ S., DYKES S., GUO D., LI K., MARIN R., MELZI S., YANG J.: Shrec'20: Shape correspondence with non-isometric deformations. *Computers & Graphics* 92 (2020), 28–43. [17](#), [22](#), [23](#)
- [DLZ24] DENG J., LU J., ZHANG T.: Unsupervised template-assisted point cloud shape correspondence network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 5250–5259. [22](#)
- [DML17] DYM N., MARON H., LIPMAN Y.: Ds++: A flexible, scalable and provably tight relaxation for matching problems. *ACM Transactions on Graphics* (2017). [11](#)
- [DMM24] DUTT N. S., MURALIKRISHNAN S., MITRA N. J.: Diffusion 3d features (diff3f): Decorating untextured shapes with distilled semantic features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 4494–4504. [18](#)
- [DSO20] DONATI N., SHARMA A., OVSIANIKOV M.: Deep geometric functional maps: Robust feature learning for shape correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 8592–8601. [8](#), [9](#), [10](#)
- [DSS\*23] DEITKE M., SCHWENK D., SALVADOR J., WEIHS L., MICHEL O., VANDERBILT E., SCHMIDT L., EHSANI K., KEMBHANI A., FARHADI A.: Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), pp. 13142–13153. [22](#)
- [DWL\*23] DENG J., WANG C., LU J., HE J., ZHANG T., YU J., ZHANG Z.: Se-ornet: Self-ensembling orientation-aware network for unsupervised point cloud shape correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), pp. 5364–5373. [22](#)
- [DYDZ22] DENG B., YAO Y., DYKE R. M., ZHANG J.: A survey of non-rigid 3d registration. *Computer Graphics Forum* 41, 2 (2022), 559–589. [2](#)
- [DZW\*24] DI Y., ZHANG C., WANG C., ZHANG R., ZHAI G., LI Y., FU B., JI X., GAO S.: Shapematcher: Self-supervised joint shape canonicalization, segmentation, retrieval and deformation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 21017–21028. [24](#)
- [DZZ\*23] DI Y., ZHANG C., ZHANG R., MANHARDT F., SU Y., RAMBACH J., STRICKER D., JI X., TOMBARI F.: U-red: Unsupervised 3d shape retrieval and deformation for partial point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023), pp. 8884–8895. [24](#)
- [EACB24] EL AMRANI N., CAO D., BERNARD F.: A universal and flexible framework for unsupervised statistical shape model learning. In *Medical Image Computing and Computer Assisted Intervention* (2024), pp. 26–36. [2](#), [7](#), [24](#), [25](#), [26](#)
- [EARB25] EL AMRANI N., ROETZER P., BERNARD F.: High-resolution 3d shape matching with global optimality and geometric consistency. *Computer Graphics Forum* 44, 5 (2025), e70208. [11](#), [13](#), [19](#), [20](#), [21](#), [26](#)
- [EBC17] EZUZ D., BEN-CHEN M.: Deblurring and denoising of maps between shapes. *Computer Graphics Forum* 36, 5 (2017), 165–174. [8](#)
- [EEAX\*25] EHM V., EL AMRANI N., XIE Y., BASTIAN L., GAO M., WANG W., SANG L., CAO D., WEISSBERG T., LÄHNER Z., ET AL.: Beyond complete shapes: A benchmark for quantitative evaluation of 3d shape surface matching algorithms. *Computer Graphics Forum* 44 (2025), e70186. [1](#), [4](#), [9](#), [10](#), [17](#), [20](#), [21](#), [22](#), [23](#), [24](#), [25](#)
- [EGR\*24] EHM V., GAO M., ROETZER P., EISENBERGER M., CREMERS D., BERNARD F.: Partial-to-partial shape matching with geometric consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 27488–27497. [19](#), [21](#), [22](#), [24](#)
- [EL26] EBBED M., LÄHNER Z.: Moangelo: Motion-aware neural surface reconstruction for dynamic scenes. In *Proceedings of the International Conference on 3D Vision* (2026). [25](#)
- [ENK\*21] EISENBERGER M., NOVOTNY D., KERCHENBAUM G., LABATUT P., NEVEROVA N., CREMERS D., VEDALDI A.: Neuromorph: Unsupervised shape interpolation and correspondence in one go. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 7473–7483. [24](#)
- [ERBC26] EHM V., ROETZER P., BERNARD F., CREMERS D.: An integer linear programming approach to geometrically consistent partial-partial shape matching. In *Proceedings of the International Conference on 3D Vision* (2026). [19](#), [20](#), [21](#)
- [ERE\*24] EHM V., ROETZER P., EISENBERGER M., GAO M., BERNARD F., CREMERS D.: Geometrically consistent partial shape matching. In *Proceedings of the International Conference on 3D Vision* (2024), pp. 914–922. [5](#), [19](#), [20](#), [21](#)
- [ESBC19] EZUZ D., SOLOMON J., BEN-CHEN M.: Reversible harmonic maps between discrete surfaces. *ACM Transactions on Graphics* 38, 2 (2019), 1–12. [6](#), [8](#)
- [EST\*20] EGGER B., SMITH W. A. P., TEWARI A., WUHRER S., ZOLLHOEFER M., BEELER T., BERNARD F., BOLKART T., KORTYLEWSKI A., ROMDHANI S., ET AL.: 3d morphable face models—past, present, and future. *ACM Transactions on Graphics* 39, 5 (2020), 1–38. [24](#)
- [ETLTC20] EISENBERGER M., TOKER A., LEAL-TAIXÉ L., CREMERS D.: Deep shells: Unsupervised shape correspondence with optimal transport. In *Proceedings of the Conference on Neural Information Processing Systems* (2020), vol. 33, pp. 10491–10502. [10](#)
- [ETLTC23] EISENBERGER M., TOKER A., LEAL-TAIXÉ L., CREMERS D.: G-msm: Unsupervised multi-shape matching with graph-based affinity priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), pp. 22762–22772. [8](#), [13](#)
- [FCH\*23] FENG W., CAI H., HOU J., DENG B., ZHANG J.: Differentiable deformation graph-based neural non-rigid registration. *Communications in Mathematics and Statistics* 11, 1 (2023), 151–167. [15](#), [17](#)
- [FKL\*23] FENG H., KULITS P., LIU S., BLACK M. J., ABBEVAYA V. F.: Generalizing neural human fitting to unseen poses with articulated se(3) equivariance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023), pp. 7977–7988. [5](#), [15](#), [16](#)
- [FMT\*22] FAN A., MA J., TIAN X., MEI X., LIU W.: Coherent point drift revisited for non-rigid shape matching and registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 1424–1434. [15](#), [16](#)
- [GDN22] GÜMELI C., DAI A., NIESSNER M.: Roca: Robust cad model retrieval and alignment from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 4022–4031. [24](#)
- [GFK\*18] GROUEIX T., FISHER M., KIM V. G., RUSSELL B. C., AUBRY M.: 3d-coded: 3d correspondences by deep deformation. In *Proceedings of the European Conference on Computer Vision* (2018), pp. 235–251. [16](#)
- [GLT\*21] GAO M., LÄHNER Z., THUNBERG J., CREMERS D., BERNARD F.: Isometric multi-shape matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 14183–14193. [11](#), [13](#)
- [GLW\*23] GAO Z., LEI J., WU J., ZHANG J., RUCHAY A., PEZZUOLO A., GUO H.: Automated method for measuring body size parameters of live pigs based on non-rigid registration of point clouds. In *Proceedings of the IEEE International Workshop on Metrology for Agriculture and Forestry* (2023), pp. 472–477. [16](#)

- [GPX\*24] GUO J., PRADA F., XIANG D., ROMERO J., WU C., PARK H. S., SHIRATORI T., SAITO S.: Diffusion shape prior for wrinkle-accurate cloth registration. In *Proceedings of the International Conference on 3D Vision* (2024), pp. 790–799. [15](#), [16](#)
- [GRE\*23] GAO M., ROETZER P., EISENBERGER M., LÄHNER Z., MOELLER M., CREMERS D., BERNARD F.: Sigma: Scale-invariant global sparse shape matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023), pp. 645–654. [3](#), [11](#), [12](#)
- [GRLD24] GAO D., ROZENBERSZKI D., LEUTENEGGER S., DAI A.: Diffcad: Weakly-supervised probabilistic cad model retrieval and alignment from an rgb image. *ACM Transactions on Graphics* **43**, 4 (2024), 1–15. [24](#)
- [GTRS16] GOLYANIK V., TAETZ B., REIS G., STRICKER D.: Extended coherent point drift algorithm with correspondence priors and optimal subsampling. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision* (2016), pp. 1–9. [16](#)
- [Gug12] GUGGENHEIMER H. W.: *Differential geometry*. Courier Corporation, 2012. [7](#)
- [GVNA25] GAFENCU M.-A., VELIKOVA Y., NAVAB N., AZAMPOUR M. F.: Us-x complete: A multi-modal approach to anatomical 3d shape recovery. In *International Workshop on Shape in Medical Imaging* (2025), pp. 218–231. [24](#)
- [GVS\*24] GAFENCU M.-A., VELIKOVA Y., SALEH M., UNGI T., NAVAB N., WENDLER T., AZAMPOUR M. F.: Shape completion in the dark: Completing vertebrae morphology from 3d ultrasound. *International Journal of Computer Assisted Radiology and Surgery* **19**, 7 (2024), 1339–1347. [24](#)
- [HGB\*22] HUANG J., BIRDAL T., GOJCIC Z., GUIBAS L. J., HU S.-M.: Multiway non-rigid point cloud registration via learned functional map synchronization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**, 2 (2022), 2038–2053. [9](#)
- [HDZ\*23] HE J., DENG J., ZHANG T., ZHANG Z., ZHANG Y.: Hierarchical shape-consistent transformer for unsupervised point cloud shape correspondence. *IEEE Transactions on Image Processing* **32** (2023), 2734–2748. [22](#)
- [HFH\*22] HALLER S., FEINEIS L., HUTSCHENREITER L., BERNARD F., ROTHER C., KAINMÜLLER D., SWOBODA P., SAVCHYNSKY B.: A comparative study of graph matching algorithms in computer vision. In *Proceedings of the European Conference on Computer Vision* (2022). [5](#), [11](#)
- [HG13] HUANG Q.-X., GUIBAS L.: Consistent shape maps via semidefinite programming. *Computer Graphics Forum* **32**, 5 (2013), 177–186. [13](#)
- [HIK11] HAMMACK R. H., IMRICH W., KLAVŽAR S.: *Handbook of product graphs*, vol. 2. CRC Press Boca Raton, 2011. [12](#)
- [Hir21a] HIROSE O.: Acceleration of non-rigid point set registration with downsampling and gaussian process regression. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **43**, 8 (2021), 2858–2865. [16](#)
- [Hir21b] HIROSE O.: A bayesian formulation of coherent point drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **43**, 7 (2021), 2269–2286. [16](#)
- [Hir23] HIROSE O.: Geodesic-based bayesian coherent point drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**, 5 (2023), 5816–5832. [15](#), [16](#)
- [HJA20] HO J., JAIN A., ABBEEL P.: Denoising diffusion probabilistic models. *Proceedings of the Conference on Neural Information Processing Systems* **33** (2020), 6840–6851. [10](#)
- [HLR\*19] HALIMI O., LITANY O., RODOLÀ E., BRONSTEIN A. M., KIMMEL R.: Unsupervised learning of dense shape correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 4370–4379. [8](#), [9](#)
- [HLT\*25] HUANG W.-L., LIU J., TASHAYYOD D., KANG J., GANDBAKHCHE A., KAZHDAN M., ARMAND M.: A shape-aware total body photography system for in-focus surface coverage optimization. *IEEE Journal of Biomedical and Health Informatics* (2025). [25](#)
- [HRWO20] HUANG R., REN J., WONKA P., OVSJANIKOV M.: Consistent zoomout: Efficient spectral map synchronization. *Computer Graphics Forum* **39**, 5 (2020), 265–278. [8](#)
- [HSA\*23] HARTWIG F., SASSEN J., AZENCOT O., RUMPF M., BEN-CHEN M.: An elastic basis for spectral shape correspondence. In *Proceedings of the SIGGRAPH Conference* (2023), pp. 1–11. [8](#), [10](#)
- [HW77] HOLLAND P. W., WELSCH R. E.: Robust regression using iteratively reweighted least-squares. *Communications in Statistics - Theory and Methods* **6**, 9 (1977), 813–827. [15](#)
- [HWG14] HUANG Q., WANG F., GUIBAS L.: Functional map networks for analyzing and exploring large shape collections. *ACM Transactions on Graphics* **33**, 4 (2014), 1–11. [8](#)
- [IE23] IYER K., ELHABIAN S. Y.: Mesh2ssm: From surface meshes to statistical shape models of anatomy. In *Medical Image Computing and Computer Assisted Intervention* (2023), pp. 615–625. [24](#)
- [IKE25] IYER K., KARANAM M. S. T., ELHABIAN S.: Mesh2ssm++: A probabilistic framework for unsupervised learning of statistical shape model of anatomies from surface meshes. *arXiv preprint arXiv:2502.07145* (2025). [24](#)
- [IMH05] IGARASHI T., MOSCOVICH T., HUGHES J. F.: As-rigid-as-possible shape manipulation. *ACM Transactions on Graphics* **24**, 3 (2005), 1134–1141. [2](#)
- [JKY\*25] JUNG Y., KIM H., YOON H., HWANG Y., LEE S.: Variable shared template for consistent non-rigid icp. *ACM Transactions on Graphics* **44**, 4 (2025), 1–16. [15](#), [16](#)
- [JSH23] JIANG P., SUN M., HUANG R.: Non-rigid shape registration via deep functional maps prior. *Proceedings of the Conference on Neural Information Processing Systems* **36** (2023), 58409–58427. [15](#), [16](#)
- [KAAL22] KARRAS T., AITTALA M., AILA T., LAINE S.: Elucidating the design space of diffusion-based generative models. *Proceedings of the Conference on Neural Information Processing Systems* **35** (2022), 26565–26577. [10](#)
- [KBFG25] KALAIE S., BULPITT A., FRANGI A. F., GOOYA A.: An end-to-end deep learning generative framework for refinable shape matching and generation. *IEEE Transactions on Medical Imaging* **44**, 8 (2025), 3323–3344. [24](#)
- [KDMA\*22] KLATZOW J., DALMASSO G., MARTÍNEZ-ABADÍAS N., SHARPE J., UHLMANN V.:  $\mu$ match: 3d shape correspondence for biological image data. *Frontiers in Computer Science* **4** (2022). [24](#)
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. *ACM Transactions on Graphics* **30**, 4 (2011), 1–12. [5](#)
- [KMDL19] KUSHINSKY Y., MARON H., DYM N., LIPMAN Y.: Sinkhorn algorithm for lifted assignment problems. *SIAM Journal on Imaging Sciences* (2019). [5](#), [11](#)
- [KMP\*21] KIRGO M., MELZI S., PATANÈ G., RODOLÀ E., OVSJANIKOV M.: Wavelet-based heat kernel derivatives: Towards informative localized shape analysis. *Computer Graphics Forum* **40**, 1 (2021), 165–179. [22](#), [23](#)
- [KMR\*23] KIRILLOV A., MINTUN E., RAVI N., MAO H., ROLLAND C., GUSTAFSON L., XIAO T., WHITEHEAD S., BERG A. C., LO W.-Y., ET AL.: Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023), pp. 4015–4026. [17](#), [18](#)
- [KSH\*25] KAHL M., STRICKER S., HUTSCHENREITER L., BERNARD F., ROTHER C., SAVCHYNSKY B.: Towards optimizing large-scale multi-graph matching in bioimaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2025), pp. 11569–11578. [5](#), [11](#), [13](#)

- [KV11] KORTE B. H., VYGEN J.: *Combinatorial optimization*, vol. 1. Springer, 2011. 11
- [KWS\*23] KELLER M., WERLING K., SHIN S., DELP S., PUJADES S., LIU C. K., BLACK M. J.: From skin to skeleton: Towards biomechanically accurate 3d digital humans. *ACM Transactions on Graphics* 42, 6 (2023), 1–12. 24
- [KZBP22] KELLER M., ZUFFI S., BLACK M. J., PUJADES S.: Osso: Obtaining skeletal shape from outside. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 20492–20501. 24
- [LAO22] LI L., ATTAIKI S., OVSJANIKOV M.: Srfeat: Learning locally accurate and globally consistent non-rigid shape correspondence. In *Proceedings of the International Conference on 3D Vision* (2022), pp. 144–154. 8, 9
- [LDO22] LI L., DONATI N., OVSJANIKOV M.: Learning multi-resolution functional maps with spectral attention for robust shape matching. In *Proceedings of the Conference on Neural Information Processing Systems* (2022), vol. 35, pp. 29336–29349. 8, 10, 24
- [LF09] LIPMAN Y., FUNKHOUSER T.: Möbius voting for surface correspondence. *ACM Transactions on Graphics* 28, 3 (2009), 1–12. 19
- [LFC\*25] LI B., FENG H., CAI Z., BLACK M. J., XIU Y.: Etch: Generalizing body fitting to clothed humans via equivariant tightness. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2025), pp. 8264–8274. 5, 15, 16
- [LGAR21] LANG I., GINZBURG D., AVIDAN S., RAVIV D.: Dpc: Unsupervised deep point correspondence via cross and self construction. In *Proceedings of the International Conference on 3D Vision* (2021), pp. 1442–1451. 21
- [LH22a] LI Y., HARADA T.: Leopard: Learning partial point cloud matching in rigid and deformable scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 5554–5564. 22
- [LH22b] LI Y., HARADA T.: Non-rigid point cloud registration with neural deformation pyramid. *Proceedings of the Conference on Neural Information Processing Systems* 35 (2022), 27757–27768. 5, 15, 17
- [LHP17] LÊ-HUU D. K., PARAGIOS N.: Alternating direction graph matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2017). 11
- [LHP25] LAKSHMIPATHY A. S., HODGINS J. K., POLLARD N. S.: Kinematic motion retargeting for contact-rich anthropomorphic manipulations. *ACM Transactions on Graphics* 44, 2 (2025), 1–20. 24
- [LIPM19] LAZOVA V., INSAFUTDINOV E., PONS-MOLL G.: 360-degree textures of people in clothing from a single image. In *Proceedings of the International Conference on 3D Vision* (2019), pp. 643–653. 16
- [LLL\*24] LI Z., LITVAK D., LI R., ZHANG Y., JAKAB T., RUPPRECHT C., WU S., VEDALDI A., WU J.: Learning the 3d fauna of the web. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 9752–9762. 16, 24
- [LLSH23] LI J., LI D., SAVARESE S., HOI S.: Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *Proceedings of the International Conference on Machine Learning* (2023), pp. 19730–19742. 18
- [LMR\*15] LOPER M., MAHMOOD N., ROMERO J., PONS-MOLL G., BLACK M. J.: Smpl: A skinned multi-person linear model. *ACM Transactions on Graphics* 34, 6 (2015), 248:1–248:16. 2, 7, 16, 22, 24, 26
- [LRB\*16] LÄHNER Z., RODOLÀ E., BRONSTEIN M. M., CREMERS D., BURGHARD O., COSMO L., DIECKMANN A., KLEIN R., SAHILIOGLU Y.: Shrec'16: Matching of deformable shapes with topological noise. In *Eurographics Workshop on 3D Object Retrieval* (2016). 22, 23, 25
- [LRR\*17] LITANY O., REMEZ T., RODOLÀ E., BRONSTEIN A., BRONSTEIN M.: Deep functional maps: Structured prediction for dense shape correspondence. In *Proceedings of the IEEE International Conference on Computer Vision* (2017), pp. 5659–5667. 4, 8, 9
- [LRS\*16] LÄHNER Z., RODOLÀ E., SCHMIDT F. R., BRONSTEIN M. M., CREMERS D.: Efficient globally optimal 2d-to-3d deformable shape matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2016). 12
- [LS21] LANGE J.-H., SWOBODA P.: Efficient message passing for 0–1 ilps with binary decision diagrams. In *Proceedings of the International Conference on Machine Learning* (2021), pp. 6000–6010. 12
- [LTT\*21] LI Y., TAKEHARA H., TAKETOMI T., ZHENG B., NIESSNER M.: 4dcomplete: Non-rigid motion estimation beyond the observable surface. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2021), pp. 12706–12716. 22, 23, 24
- [LZR\*24] LIU S., ZENG Z., REN T., LI F., ZHANG H., YANG J., JIANG Q., LI C., YANG J., SU H., ET AL.: Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *Proceedings of the European Conference on Computer Vision* (2024), pp. 38–55. 18
- [LZY\*25a] LI J., ZHOU Z., YANG J., PEPE A., GSAXNER C., LUIJTEN G., QU C., ZHANG T., CHEN X., LI W., ET AL.: Medshapenet: A large-scale dataset of 3d medical shapes for computer vision. *Biomedical Engineering/Biomedizinische Technik* 70, 1 (2025), 71–90. 24
- [LZY\*25b] LIU H., ZHAN X., YAN Z., LUO Z., WEN Y., HAN X.: Stable-score: A stable registration-based framework for 3d shape correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2025), pp. 917–928. 9, 18
- [LZZ\*23] LIN S., ZHOU B., ZHENG Z., ZHANG H., LIU Y.: Leveraging intrinsic properties for non-rigid garment alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023), pp. 14485–14496. 15, 17
- [MAKM21] MORREALE L., AIGERMAN N., KIM V. G., MITRA N. J.: Neural surface maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 4639–4648. 18
- [MAKM24] MORREALE L., AIGERMAN N., KIM V. G., MITRA N. J.: Neural semantic surface maps. *Computer Graphics Forum* 43, 2 (2024), e15005. 18
- [MBT\*23] MAGNET R., BLOCH K., TAVERNE M., MELZI S., GEOFROY M., KHONSARI R. H., OVSJANIKOV M.: Assessing craniofacial growth and form without landmarks: A new automatic approach based on spectral methods. *Journal of Morphology* 284, 8 (2023), e21609. 7, 16, 26
- [MCPM24] MARIN R., CORONA E., PONS-MOLL G.: Nicp: Neural icp for 3d human registration at scale. In *Proceedings of the European Conference on Computer Vision* (2024), pp. 265–285. 1, 5, 15, 17
- [MDK\*16] MARON H., DYM N., KEZURER I., KOVALSKY S., LIPMAN Y.: Point registration via efficient convex relaxation. *ACM Transactions on Graphics* 35, 4 (2016), 1–12. 11, 12
- [MDSB03] MEYER M., DESBRUN M., SCHRÖDER P., BARR A. H.: Discrete differential-geometry operators for triangulated 2-manifolds. In *Visualization and mathematics III*. Springer, 2003, pp. 35–57. 3
- [MGBM25] MELI N. K., GOLYANIK V., BENKNER M. S., MOELLER M.: Qucoop: A versatile framework for solving composite and binary-parametrised problems on quantum annealers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2025). 11, 13
- [MJF\*21] MA J., JIANG X., FAN A., JIANG J., YAN J.: Image matching from handcrafted to deep features: A survey. *International Journal of Computer Vision* 129, 1 (2021), 23–79. 2
- [MLA\*25] MACCARONE F., LONGARI G., ARRIGONI F., PERUZZO D., MELZI S.: Unraveling brainstem deformations in joubert syndrome: A statistical shape analysis of mri-derived structures. In *Medical Image Computing and Computer Assisted Intervention* (2025), vol. 15973, pp. 659–669. 16
- [MMM\*20] MELZI S., MARIN R., MUSONI P., BARDON F., TARINI M., CASTELLANI U.: Intrinsic/extrinsic embedding for functional remeshing of 3d shapes. *Computers & Graphics* 88 (2020), 1–12. 24

- [MMMC21a] MUSONI P., MARIN R., MELZI S., CASTELLANI U.: A functional skeleton transfer. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 4, 3 (2021), 1–15. [24](#)
- [MMMC21b] MUSONI P., MARIN R., MELZI S., CASTELLANI U.: Reposing and retargeting unrigged characters with intrinsic-extrinsic transfer. In *Smart Tools and Applications in Graphics* (2021). [25](#)
- [MMR\*19a] MELZI S., MARIN R., RODOLÀ E., CASTELLANI U., REN J., POULENARD A., OVSJANIKOV P.: Shrec'19: Matching humans with different connectivity. In *Eurographics Workshop on 3D Object Retrieval* (2019), pp. 1–8. [24](#)
- [MMR\*19b] MELZI S., MARIN R., RODOLÀ E., CASTELLANI U., REN J., POULENARD A., WONKA P., OVSJANIKOV M.: Matching humans with different connectivity. In *Eurographics Workshop on 3D Object Retrieval* (2019). [22](#), [23](#), [25](#)
- [MMRC19] MARIN R., MELZI S., RODOLÀ E., CASTELLANI U.: High-resolution augmentation for automatic template-based matching of human models. In *Proceedings of the International Conference on 3D Vision* (2019), pp. 230–239. [16](#)
- [MMRC20] MARIN R., MELZI S., RODOLÀ E., CASTELLANI U.: Farm: Functional automatic registration method for 3d human bodies. *Computer Graphics Forum* 39, 1 (2020), 160–173. [16](#)
- [MO24] MAGNET R., OVSJANIKOV M.: Memory-scalable and simplified functional map learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 4041–4050. [8](#), [10](#)
- [MON\*19] MESCHEDER L., OECHSLE M., NIEMEYER M., NOWOZIN S., GEIGER A.: Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 4460–4470. [2](#)
- [MPS\*23] MEI G., POIESI F., SALTORI C., ZHANG J., RICCI E., SEBE N.: Overlap-guided gaussian mixture models for point cloud registration. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision* (2023), pp. 4511–4520. [16](#)
- [MRMO20] MARIN R., RAKOTOSAONA M.-J., MELZI S., OVSJANIKOV M.: Correspondence learning via linearly-invariant embedding. *Proceedings of the Conference on Neural Information Processing Systems* 33 (2020), 1608–1620. [9](#)
- [MRR\*19] MELZI S., REN J., RODOLÀ E., SHARMA A., WONKA P., OVSJANIKOV M.: Zoomout: Spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics* 38, 6 (2019), 1–14. [8](#), [10](#)
- [MRSHO22] MAGNET R., REN J., SORKINE-HORNUNG O., OVSJANIKOV M.: Smooth non-rigid shape matching via effective dirichlet energy optimization. In *Proceedings of the International Conference on 3D Vision* (2022), pp. 495–504. [6](#), [8](#), [10](#), [19](#), [20](#), [21](#), [22](#), [23](#)
- [MRWB23] MERROUCHE A., REGATEIRO J. P. C., WUHRER S., BOYER E.: Deformation-guided unsupervised non-rigid shape matching. In *Proceedings of the British Machine Vision Conference* (2023). [15](#), [17](#), [23](#)
- [MS10] MYRONENKO A., SONG X.: Point set registration: Coherent point drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, 12 (2010), 2262–2275. [5](#), [14](#), [15](#)
- [MWB\*25a] MELI N. K., WANG S., BENKNER M. S., SASDELLI M., CHIN T.-J., BIRDAL T., MOELLER M., GOLYANIK V.: Quantum-enhanced computer vision: Going beyond classical algorithms. *arXiv preprint arXiv:2510.07317* (2025). [14](#)
- [MWB25b] MERROUCHE A., WUHRER S., BOYER E.: Matching shapes under different topologies: A topology-adaptive deformation guided approach. *arXiv preprint arXiv:2509.06862* (2025). [15](#), [16](#), [22](#)
- [MYR\*20] MA Q., YANG J., RANJAN A., PUJADES S., PONS-MOLL G., TANG S., BLACK M. J.: Learning to dress 3d people in generative clothing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 6469–6478. [16](#)
- [NO22] NETTO G. M., OLIVEIRA M. M.: Robust point-cloud registration based on dense point matching and probabilistic modeling. *The Visual Computer* 38, 9 (2022), 3217–3230. [22](#)
- [OBCS\*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: A flexible representation of maps between shapes. *ACM Transactions on Graphics* 31, 4 (2012), 1–11. [4](#), [7](#), [8](#), [9](#), [10](#), [18](#)
- [OCB\*16] 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 *Proceedings of SIGGRAPH Asia 2016 Courses* (2016), pp. 1–60. [2](#)
- [ODM\*23] OQUAB M., DARCET T., MOUTAKANNI T., VO H., SZAFRANIEC M., KHALIDOV V., FERNANDEZ P., HAZIZA D., MASSA F., EL-NOUBY A., ET AL.: Dinov2: Learning robust visual features without supervision. *Transactions on Machine Learning Research* (2023). [17](#), [18](#)
- [PCG\*19] PAVLAKOS G., CHOUTAS V., GHORBANI N., BOLKART T., OSMAN A. A. A., TZIONAS D., BLACK M. J.: Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 10975–10985. [2](#), [16](#), [22](#)
- [PFS\*19] PARK J. J., FLORENCE P., STRAUB J., NEWCOMBE R., LOVEGROVE S.: DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 165–174. [2](#)
- [PKO22] PANINE M., KIRGO M., OVSJANIKOV M.: Non-isometric shape matching via functional maps on landmark-adapted bases. *Computer Graphics Forum* 41, 6 (2022), 394–417. [8](#)
- [PKS13] PACHAURI D., KONDOR R., SINGH V.: Solving the multi-way matching problem by permutation synchronization. *Proceedings of the Conference on Neural Information Processing Systems* 26 (2013). [13](#)
- [PLDO25] PIERSON E., LI L., DAI A., OVSJANIKOV M.: Diffumatch: Category-agnostic spectral diffusion priors for robust non-rigid shape matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2025). [8](#), [10](#), [20](#)
- [PMM24] PEDICO S., MELZI S., MAGGIOLI F.: Taco: A benchmark for connectivity-invariance in shape correspondence. *Smart Tools and Applications in Graphics - Eurographics Italian Chapter Conference* (2024). [23](#)
- [PMRMB15] PONS-MOLL G., ROMERO J., MAHMOOD N., BLACK M. J.: Dyna: A model of dynamic human shape in motion. *ACM Transactions on Graphics* 34, 4 (2015), 1–14. [16](#)
- [PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics* 2, 1 (1993), 15–36. [7](#)
- [PRM\*21] PAI G., REN J., MELZI S., WONKA P., OVSJANIKOV M.: Fast sinkhorn filters: Using matrix scaling for non-rigid shape correspondence with functional maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 384–393. [8](#), [11](#), [12](#)
- [PSE\*26] POSNER E., SHTALRID O., ERELL O., NOY D., BOUHNİK M.: A statistical 3d stomach shape model for anatomical analysis. In *Shape in Medical Imaging* (2026), pp. 117–128. [16](#)
- [PWG\*24] PO R., WANG Y., GOLYANIK V., ABERMAN K., BARRON J. T., BERMANO A., CHAN E., DEKEL T., HOLYNSKI A., KANAZAWA A., ET AL.: State of the art on diffusion models for visual computing. *Computer Graphics Forum* 43, 2 (2024). [10](#)
- [QYW\*23] QIN Z., YU H., WANG C., PENG Y., XU K.: Deep graph-based spatial consistency for robust non-rigid point cloud registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), pp. 5394–5403. [15](#), [16](#)
- [RAC\*24] ROETZER P., ABBAS A., CAO D., BERNARD F., SWOBODA P.: Discomatch: Fast discrete optimisation for geometrically consistent

- 3d shape matching. In *Proceedings of the European Conference on Computer Vision* (2024). 11, 12, 19
- [RB24] ROETZER P., BERNARD F.: Spidermatch: 3d shape matching with global optimality and geometric consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 14543–14553. 5, 11, 12, 13, 19
- [RB25] ROETZER P., BERNARD F.: Fast globally optimal and geometrically consistent 3d shape matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2025). 11, 12, 13, 18, 19, 20, 21, 26
- [RBCL25] RIMON A. C., BEN-CHEN M., LITANY O.: Fridu: Functional map refinement with guided image diffusion. *Computer Graphics Forum* 44, 5 (2025), e70203. 8, 10
- [RBL\*22] ROMBACH R., BLATTMANN A., LORENZ D., ESSER P., OMMER B.: High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 10684–10695. 17, 18
- [RBSB18] RANJAN A., BOLKART T., SANYAL S., BLACK M. J.: Generating 3d faces using convolutional mesh autoencoders. In *Proceedings of the European Conference on Computer Vision* (2018), pp. 725–741. 22
- [RCB\*17] RODOLÀ E., COSMO L., BRONSTEIN M. M., TORSOLLO A., CREMERS D.: Partial functional correspondence. *Computer Graphics Forum* 36, 1 (2017), 222–236. 10, 20
- [RDP99] ROBINETTE K. M., DAANEN H., PAQUET E.: The caesar project: A 3-d surface anthropometry survey. In *Proceedings of the International Conference on 3-D Digital Imaging and Modeling* (1999), pp. 380–386. 22
- [REC\*25] ROETZER P., EHM V., CREMERS D., LÄHNER Z., BERNARD F.: Higher-order ratio cycles for fast and globally optimal shape matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2025). 11, 12, 19
- [Ren21] REN J.: *Shape matching and map space exploration via functional maps*. PhD thesis, King Abdullah University of Science and Technology, 2021. 6
- [RLB23] ROETZER P., LÄHNER Z., BERNARD F.: Conjugate product graphs for globally optimal 2d-3d shape matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), pp. 21866–21875. 12
- [RMC\*15] RODOLÀ E., MOELLER M., CREMERS D., ET AL.: Pointwise map recovery and refinement from functional correspondence. In *Vision, Modeling and Visualization* (2015), pp. 25–32. 8
- [RMC17] RODOLÀ E., MOELLER M., CREMERS D.: Regularized pointwise map recovery from functional correspondence. *Computer Graphics Forum* 36, 8 (2017), 700–711. 8
- [RMOW20] REN J., MELZI S., OVSJANIKOV M., WONKA P.: Maptree: Recovering multiple solutions in the space of maps. *ACM Transactions on Graphics* 39, 6 (2020), 264:1–264:17. 6
- [RMWO21] REN J., MELZI S., WONKA P., OVSJANIKOV M.: Discrete optimization for shape matching. *Computer Graphics Forum* 40, 5 (2021), 81–96. 8, 10, 11, 12
- [RPW94] RENDL F., PARDALOS P., WOLKOWICZ H.: The quadratic assignment problem: A survey and recent developments. In *DIMACS Workshop* (1994). 4, 11, 18
- [RPWO18] REN J., POULENARD A., WONKA P., OVSJANIKOV M.: Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics* 37, 6 (2018), 1–16. 8, 18
- [RRBW\*14] RODOLÀ E., ROTA BULO S., 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* (2014), pp. 4177–4184. 20, 22, 23
- [RSCB22] ROETZER P., SWOBODA P., CREMERS D., BERNARD F.: A scalable combinatorial solver for elastic geometrically consistent 3d shape matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 428–438. 5, 11, 12, 19
- [RSO19] ROUFOSSE J.-M., SHARMA A., OVSJANIKOV M.: Unsupervised deep learning for structured shape matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2019), pp. 1617–1627. 4, 8, 9
- [RTB17] ROMERO J., TZIONAS D., BLACK M. J.: Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics* 36, 6 (2017). 2, 16, 22, 24
- [RTO\*19] 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* (2019), pp. 37–46. 21
- [RVBC23] RORBERG S., VAXMAN A., BEN-CHEN M.: Bpm: Blended piecewise möbius maps. *Computer Graphics Forum* 42, 5 (2023), e14914. 19
- [RWP06] REUTER M., WOLTER F.-E., PEINECKE N.: Laplace–beltrami spectra as ‘shape-dna’ of surfaces and solids. *Computer-Aided Design* 38, 4 (2006), 342–366. 4
- [SA07] SORKINE O., ALEXA M.: As-rigid-as-possible surface modeling. In *Symposium on Geometry Processing* (2007). 12
- [SACO22] SHARP N., ATTAIKI S., CRANE K., OVSJANIKOV M.: Diffusionnet: Discretization agnostic learning on surfaces. *ACM Transactions on Graphics* 41, 3 (2022), 1–16. 4, 18
- [Sah20] SAHILIOĞLU Y.: Recent advances in shape correspondence. *The Visual Computer* 36, 8 (2020), 1705–1721. 2
- [SBLG\*25] SEELBACH BENKNER M., LÄHNER Z., GOLYANIK V., KLIESCH M., MOELLER M.: Compensating connectivity restrictions in quantum annealers via splitting and linearization techniques. *arXiv preprint arXiv:2507.12536* (2025). 13
- [SC78] SAKOE H., CHIBA S.: Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 26 (1978). 12
- [SC10] SCHOENEMANN T., CREMERS D.: A combinatorial solution for model-based image segmentation and real-time tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, 7 (2010), 1153–1164. 12
- [SC20] SHARP N., CRANE K.: A laplacian for nonmanifold triangle meshes. *Computer Graphics Forum* 39, 5 (2020), 69–80. 22
- [SCC\*25] SANG L., CANFES Z., CAO D., MARIN R., BERNARD F., CREMERS D.: 4deform: Neural surface deformation for robust shape interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2025), pp. 6542–6551. 24
- [SCC\*26] SANG L., CANFES Z., CAO D., MARIN R., BERNARD F., CREMERS D.: Twosquared: 4d generation from 2d image pairs. In *Proceedings of the International Conference on 3D Vision* (2026). 24
- [Sch03] SCHRIJVER A.: *Combinatorial optimization: Polyhedra and efficiency*. Springer, 2003. 11
- [SDB\*20] SALEH M., DEGHANI S., BUSAM B., NAVAB N., TOMBARI F.: Graphite: Graph-induced feature extraction for point cloud registration. In *Proceedings of the International Conference on 3D Vision* (2020), pp. 241–251. 4
- [SDM\*24] SUNDARARAMAN R., DONATI N., MELZI S., CORMAN E., OVSJANIKOV M.: Deformation recovery: Localized learning for detail-preserving deformations. *ACM Transactions on Graphics* 43, 6 (2024), 1–16. 5, 15, 17
- [SDWGM15] SOHL-DICKSTEIN J., WEISS E., MAHESWARANATHAN N., GANGULI S.: Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the International Conference on Machine Learning* (2015). 10

- [SGJ\*24] SUN M., GUO C., JIANG P., MAO S., CHEN Y., HUANG R.: Srf: Semantic shape registration empowered by diffusion-based image morphing and flow estimation. In *Proceedings of the SIGGRAPH Asia Conference* (2024), pp. 1–11. [18](#)
- [SGTM20] SEELBACH BENKNER M., GOLYANIK V., THEOBALT C., MOELLER M.: Adiabatic quantum graph matching with permutation matrix constraints. In *Proceedings of the International Conference on 3D Vision* (2020), pp. 583–592. [11](#), [13](#)
- [SH22] SAHILLIOĞLU Y., HORSMAN D.: Augmented paths and reodesics for topologically-stable matching. *ACM Transactions on Graphics* 42, 2 (2022), 17:1–17:15. [22](#), [23](#)
- [SMJ\*23] SUN M., MAO S., JIANG P., OVSJANIKOV M., HUANG R.: Spatially and spectrally consistent deep functional maps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023), pp. 14497–14507. [10](#)
- [SMRO22] SUNDARARAMAN R. S., MARIN R., RODOLÀ E., OVSJANIKOV M.: Reduced representation of deformation fields for effective non-rigid shape matching. In *Proceedings of the Conference on Neural Information Processing Systems* (2022), vol. 35, pp. 10405–10420. [17](#)
- [SO20] SHARMA A., OVSJANIKOV M.: Weakly supervised deep functional maps for shape matching. In *Proceedings of the Conference on Neural Information Processing Systems* (2020), vol. 33, pp. 19264–19275. [10](#)
- [SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum* 28, 5 (2009), 1383–1392. [4](#), [8](#)
- [SP04] SUMNER R. W., POPOVIĆ J.: Deformation transfer for triangle meshes. *ACM Transactions on Graphics* 23, 3 (2004), 399–405. [6](#), [7](#), [10](#), [24](#), [26](#)
- [SPK23] SCHMIDT P., PIEPER D., KOBBELT L.: Surface maps via adaptive triangulations. *Computer Graphics Forum* 42, 2 (2023), 103–117. [19](#)
- [SPO22] SUNDARARAMAN R., PAI G., OVSJANIKOV M.: Implicit field supervision for robust non-rigid shape matching. In *Proceedings of the European Conference on Computer Vision* (2022), pp. 344–362. [15](#), [17](#)
- [SS05] SCHELLEWALD C., SCHNÖRR C.: Probabilistic subgraph matching based on convex relaxation. In *Energy Minimization Methods in Computer Vision and Pattern Recognition* (2005). [11](#)
- [SSC19] SHARP N., SOLIMAN Y., CRANE K.: Navigating intrinsic triangulations. *ACM Transactions on Graphics* (2019). [19](#)
- [STC09] SCHMIDT F. R., TÖPPE E., CREMERS D.: Efficient planar graph cuts with applications in computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2009). [12](#)
- [STCB07] SCHMIDT F. R., TÖPPE E., CREMERS D., BOYKOV Y.: Intrinsic mean for semi-metrical shape retrieval via graph cuts. In *Pattern Recognition* (2007). [13](#)
- [STDS14] SALTİ S., TOMBARI F., DI STEFANO L.: Shot: Unique signatures of histograms for surface and texture description. *Computer Vision and Image Understanding* 125 (2014), 251–264. [3](#), [8](#)
- [SWC\*22] SALEH M., WU S.-C., COSMO L., NAVAB N., BUSAM B., TOMBARI F.: Bending graphs: Hierarchical shape matching using gated optimal transport. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 11747–11757. [22](#)
- [SZT\*23] SHAO R., ZHENG Z., TU H., LIU B., ZHANG H., LIU Y.: Tensor4d: Efficient neural 4d decomposition for high-fidelity dynamic reconstruction and rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023). [25](#)
- [Tak22] TAKAYAMA K.: Compatible intrinsic triangulations. *ACM Transactions on Graphics* 41, 4 (2022), 1–12. [19](#)
- [TCM\*21] TRAPPOLINI G., COSMO L., MOSCHELLA L., MARIN R., MELZI S., RODOLÀ E.: Shape registration in the time of transformers. In *Proceedings of the Conference on Neural Information Processing Systems* (2021), vol. 34, pp. 5731–5744. [17](#)
- [THYS22] TAJDARI F., HUYSMANS T., YANG Y., SONG Y.: Feature preserving non-rigid iterative weighted closest point and semi-curvature registration. *IEEE Transactions on Image Processing* 31 (2022), 1841–1856. [15](#)
- [TSDS10] TOMBARI F., SALTİ S., DI STEFANO L.: Unique signatures of histograms for local surface description. In *Proceedings of the European Conference on Computer Vision* (2010), pp. 356–369. [3](#), [4](#), [18](#)
- [UEK25] UZOLAS L., EISEMANN E., KELLNHOFFER P.: Surface-aware distilled 3d semantic features. In *Proceedings of the SIGGRAPH Asia Conference* (2025). [18](#)
- [VKZHCO11] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. *Computer Graphics Forum* 30, 6 (2011), 1681–1707. [2](#)
- [VLB\*17] VESTNER M., LÄHNER Z., BOYARSKI A., LITANY O., SLOSSBERG R., REMEZ T., RODOLÀ E., BRONSTEIN A., BRONSTEIN M., KIMMEL R., CREMERS D.: Efficient deformable shape correspondence via kernel matching. In *Proceedings of the International Conference on 3D Vision* (2017). [2](#), [5](#), [11](#)
- [VM23] VIGANÒ G., MELZI S.: Adjoint bijective zoomout: Efficient upsampling for learned linearly-invariant embedding. In *Proceedings of the Italian Chapter Conference* (2023), pp. 37–46. [4](#), [10](#)
- [VOM25] VIGANÒ G., OVSJANIKOV M., MELZI S.: Nam: Neural adjoint maps for refining shape correspondences. *ACM Transactions on Graphics* 44, 4 (2025), 1–15. [6](#), [9](#), [10](#), [26](#)
- [VRM\*17] VAROL G., ROMERO J., MARTIN X., MAHMOOD N., BLACK M. J., LAPTEV I., SCHMID C.: Learning from synthetic humans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2017), pp. 109–117. [22](#)
- [WBC\*25] WANG T. D., BASTIAN L., CZEMPIEL T., HEILIGER C., NAVAB N.: Beyond role-based surgical domain modeling: Generalizable re-identification in the operating room. *Medical Image Analysis* 105 (2025), 103687. [24](#)
- [WCK\*25] WANG J., CHEN M., KARAEV N., VEDALDI A., RUPPRECHT C., NOVOTNY D.: Vggt: Visual geometry grounded transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2025), pp. 5294–5306. [20](#)
- [Wey11] WEYL H.: Über die asymptotische verteilung der eigenwerte. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse 1911* (1911), 110–117. [20](#)
- [WGT21] WANG S., GEIGER A., TANG S.: Locally aware piecewise transformation fields for 3d human mesh registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 7639–7648. [17](#)
- [WHG\*24] WANG W., HO H.-I., GUO C., RONG B., GRIGOREV A., SONG J., ZARATE J. J., HILLIGES O.: 4d-dress: A 4d dataset of real-world human clothing with semantic annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 550–560. [25](#)
- [WJL\*24] WU Q., JIANG H., LUO L., LI J., DING Y., XIE J., YANG J.: Diff-reg: Diffusion model in doubly stochastic matrix space for registration problem. In *Proceedings of the European Conference on Computer Vision* (2024), pp. 160–178. [20](#), [22](#)
- [WMKG07] WARDETZKY M., MATHUR S., KÄLBERER F., GRINSPUN E.: Discrete laplace operators: No free lunch. In *Symposium on Geometry Processing* (2007), vol. 33, p. 37. [7](#)
- [WSSC11] WINDHEUSER T., SCHLICKWEI U., SCHMIDT F. R., CREMERS D.: Geometrically consistent elastic matching of 3d shapes: A linear programming solution. In *Proceedings of the IEEE International Conference on Computer Vision* (2011). [3](#), [4](#), [5](#), [11](#), [12](#), [18](#), [19](#)
- [WWEAB25] WANG W., WEISSBERG T., EL AMRANI N., BERNARD F.: Symmetry understanding of 3d shapes via chirality disentanglement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2025). [2](#)

- [XER\*25] XIE Y., EHM V., ROETZER P., EL AMRANI N., GAO M., BERNARD F., CREMERS D.: Echomatch: Partial-to-partial shape matching via correspondence reflection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2025), pp. 11665–11675. [8](#), [9](#), [18](#), [20](#), [21](#), [24](#)
- [XLGM24] XIA Y., LU Y., GAO Y., MA J.: Locality preserving refinement for shape matching with functional maps. In *Proceedings of the AAAI Conference on Artificial Intelligence* (2024). [19](#)
- [YCS25] YU B., CORDIER F., SEO H.: Phydeformer: High-quality non-rigid garment registration with physics-awareness. In *Eurographics Workshop on 3D Object Retrieval* (2025). [16](#)
- [YDZH25] YAO Y., DENG B., HOU J., ZHANG J.: Spare: Symmetrized point-to-plane distance for robust non-rigid 3d registration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2025). [5](#), [15](#)
- [YDXZ20] YAO Y., DENG B., XU W., ZHANG J.: Quasi-newton solver for robust non-rigid registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 7600–7609. [15](#)
- [YDXZ23] YAO Y., DENG B., XU W., ZHANG J.: Fast and robust non-rigid registration using accelerated majorization-minimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45, 8 (2023), 9681–9698. [5](#), [15](#)
- [YHH\*19] YANG G., HUANG X., HAO Z., LIU M.-Y., BELONGIE S., HARIHARAN B.: Pointflow: 3d point cloud generation with continuous normalizing flows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2019), pp. 4541–4550. [18](#)
- [YLS\*25] YIBULAYIMU S., LIU Y., SANG Y., ZHU G., LI H., LU H., ZHAO C., WU X., WANG Y.: Sim-to-real transformer-based shape reconstruction for automated orthopedic fracture reduction planning. In *Medical Image Computing and Computer Assisted Intervention* (2025), vol. 15962, pp. 587–597. [24](#)
- [YQH\*23] YU H., QIN Z., HOU J., SALEH M., LI D., BUSAM B., ILIC S.: Rotation-invariant transformer for point cloud matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), pp. 5384–5393. [22](#)
- [YSL23] YANG Z., SIMON R., LINTE C. A.: Learning feature descriptors for pre- and intra-operative point cloud matching for laparoscopic liver registration. *International Journal of Computer Assisted Radiology and Surgery* 18, 6 (2023), 1025–1032. [24](#)
- [ZB15] ZUFFI S., BLACK M. J.: The stitched puppet: A graphical model of 3d human shape and pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2015), pp. 3537–3546. [16](#)
- [ZJM\*24] ZHAO M., JIANG J., MA L., XIN S., MENG G., YAN D.-M.: Correspondence-free non-rigid point set registration using unsupervised clustering analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 21199–21208. [15](#)
- [ZJZ\*25] ZHU J., JU Y., ZHANG J., WANG M., YUAN Z., HU K., XU H.: Densmatcher: Learning 3d semantic correspondence for category-level manipulation from one demo. In *Proceedings of the International Conference on Learning Representations* (2025). [18](#), [24](#)
- [ZKA\*22] ZHAO M., KAWAHARA J., ABHISHEK K., SHAMANIAN S., HAMARNEH G.: Skin3d: Detection and longitudinal tracking of pigmented skin lesions in 3d total-body textured meshes. *Medical Image Analysis* 77 (2022), 102329. [25](#)
- [ZKJB17] ZUFFI S., KANAZAWA A., JACOBS D., BLACK M. J.: 3d menagerie: Modeling the 3d shape and pose of animals. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2017). [2](#), [16](#), [22](#), [23](#), [24](#)
- [ZL26] ZHOU H., LÄHNER Z.: Laplace-beltrami operator for gaussian splatting. In *Proceedings of the International Conference on 3D Vision* (2026). [22](#)
- [ZLG25] ZHURAVLEV A., LÄHNER Z., GOLYANIK V.: Denoising functional maps: Diffusion models for shape correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2025), pp. 26899–26909. [6](#), [8](#), [10](#), [20](#)
- [ZLW25] ZHANG T., LIANG Z., WANG B.: A survey of medical point cloud shape learning: Registration, reconstruction and variation. In *Proceedings of the IEEE International Conference on Bioinformatics and Biomedicine* (2025), pp. 5777–5782. [24](#)
- [ZML\*24] ZUFFI S., MELLBIN Y., LI C., HOESCHLE M., KJELLSTRÖM H., POLIKOVSKY S., HERNLUND E., BLACK M. J.: Varen: Very accurate and realistic equine network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 5374–5383. [16](#)
- [ZQZ\*21] ZENG Y., QIAN Y., ZHU Z., HOU J., YUAN H., HE Y.: Cor-net3d: Unsupervised end-to-end learning of dense correspondence for 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 6052–6061. [21](#)
- [ZZD\*24] ZHANG R., ZHANG C., DI Y., MANHARDT F., LIU X., TOMBARI F., JI X.: Kp-red: Exploiting semantic keypoints for joint 3d shape retrieval and deformation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 20540–20550. [24](#)
- [ZZX\*24] ZHANG K., ZHOU Y., XU X., DAI B., PAN X.: Diffmorpher: Unleashing the capability of diffusion models for image morphing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 7912–7921. [18](#)