

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository:<https://orca.cardiff.ac.uk/id/eprint/148411/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Deng, Bailin , Yao, Yuxin, Dyke, Roberto M. and Zhang, Juyong 2022. A survey of non-rigid 3D registration. *Computer Graphics Forum* 41 (2) , pp. 559-589.
10.1111/cgf.14502

Publishers page: <https://doi.org/10.1111/cgf.14502>



Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



A Survey of Non-Rigid 3D Registration

Bailin Deng¹  Yuxin Yao² Roberto M. Dyke³  Juyong Zhang²  †

¹Cardiff University ²University of Science and Technology of China ³Università della Svizzera italiana

Abstract

Non-rigid registration computes an alignment between a source surface with a target surface in a non-rigid manner. In the past decade, with the advances in 3D sensing technologies that can measure time-varying surfaces, non-rigid registration has been applied for the acquisition of deformable shapes and has a wide range of applications. This survey presents a comprehensive review of non-rigid registration methods for 3D shapes, focusing on techniques related to dynamic shape acquisition and reconstruction. In particular, we review different approaches for representing the deformation field, and the methods for computing the desired deformation. Both optimization-based and learning-based methods are covered. We also review benchmarks and datasets for evaluating non-rigid registration methods, and discuss potential future research directions.

CCS Concepts

• **Computing methodologies** → *Shape modeling; Computer vision; Machine learning;*

1. Introduction

Surface registration computes a deformation that aligns a source surface with a target surface (see Fig. 1 for an example). It has many applications in computer graphics and computer vision. One common application is 3D surface acquisition, where an object is scanned from different directions and the scans are aligned via registration to derive the 3D measurements in a common coordinate system. Depending on the application, the surfaces can be aligned by a rigid or non-rigid deformation. In the former case, the whole source surface undergoes a single rotation and translation, which is suitable for the registration of static shapes. In the latter case, different parts of the source surface can undergo different deformations, to account for the non-rigid behavior—such as articulation of the underlying shape. Computing such deformations is a fundamental problem for the acquisition and analysis of non-rigidly deformable objects. In the past decade, with the availability of consumer depth sensors that can measure time-varying surfaces, non-rigid registration has been applied to dynamic shape reconstruction problems such as human performance capture, enabling a wide range of applications in VR, AR, and entertainment.

Non-rigid registration is a challenging problem. First, unlike rigid registration which only involves a single rotation and translation, non-rigid registration often needs to determine a deformation field for the source surface. A proper representation of the deformation field needs to be chosen to ensure sufficient expressiveness for an accurate alignment without incurring excessive computational costs. In addition, 3D measure data often contain noise and outliers,




Figure 1: An example of non-rigid registration, where the source surface is deformed in a non-rigid manner to align with the target.

and the two surfaces to be aligned can have a low ratio of overlap when they cover different parts of an object. These can make it difficult to derive a reliable correspondence between the two surfaces for their alignment. Moreover, in many practical applications, the registration needs to be performed efficiently, preferably in real-time. This survey provides an overview of significant work that addresses these challenges.

1.1. Scope and Related Works

This survey considers non-rigid registration methods for surfaces in 3D spaces, focusing on techniques related to dynamic shape acqui-

† Corresponding author: juyong@ustc.edu.cn (Juyong Zhang)

sition and reconstruction in computer graphics and vision. As such, some shape correspondence methods intended for other applications are not included in this survey. For example, we do not include shape correspondence methods that require the source and target surfaces to have a bijective correspondence, as the inputs to a 3D acquisition system are often partial scans rather than complete shapes. We do include methods that compute correspondence between a partial shape and a complete shape, since such scenarios can occur in a reconstruction system where a template mesh is updated according to input partial scans [GXW*15, LJY*17, YZG*18, SXZ*20]. Our discussion is focused on methods that perform registration using geometry information. Multimodal methods that utilize geometry together with other types of information have been recently reviewed in [ST20] and are outside the scope of this survey.

Methods for 3D registration and correspondence have been reviewed in some existing surveys. Among them, [vKZHC011] and [Sah20] focus on shape correspondence rather than registration. [TCL*13] consider both rigid and non-rigid registration, but it was published almost a decade ago and did not include more recent development in non-rigid registration, such as the use of robust norms to handle outliers and partial overlaps, efficient numerical solvers for non-smooth formulations, and deep learning-based methods. A more recent review of non-rigid registration methods and software solutions is presented in [KBD17], but it focused on applications in medical imaging only. Our survey provides an up-to-date overview of the literature in graphics and vision.

1.2. Organization

We first review different representations of deformation fields in Section 2. Afterwards, methods for determining the deformation/correspondence are reviewed in detail. We divide such methods into two categories: *extrinsic methods* attempt to directly reduce the distance between the source and target surfaces measured in the ambient 3D space, while *intrinsic methods* utilize intrinsic metrics on the surfaces to compute the alignment. These two types of methods are reviewed in Section 3 and Section 4 respectively, each covering both optimization-based and learning-based approaches. Section 5 provides an overview of datasets and benchmarks for evaluating non-rigid registration methods. Finally, we discuss research challenges and future directions in Section 6.

1.3. Notations

In the following discussion, we use \mathcal{X} and \mathcal{Y} to denote the source surface and the target surface in the 3D space, respectively. The surfaces can be represented either explicitly as point clouds or meshes, or implicitly as level sets of signed distance functions. For explicit representations, we use $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ and $\mathbf{Y} = \{\mathbf{y}_1, \dots, \mathbf{y}_n\}$ to denote the point sets on the source surface and the target surface respectively, and use $\widehat{\mathbf{X}} = \{\widehat{\mathbf{x}}_1, \dots, \widehat{\mathbf{x}}_m\}$ to denote the positions of source surface points after the deformation.

2. Representation of the Deformation Field

Many methods compute a deformation field that transforms the source surface to align it with the target surface. The representation

of the deformation field often involves a trade-off between its expressiveness and the computational cost. On one hand, a representation with higher degrees of freedom tends to be more expressive and may achieve better alignment. On the other hand, higher degrees of freedom will require a larger number of variables to be determined, which can increase the computational cost. In this section, we review different representations that have been used in the literature. Since our aim is to categorize different representations, for each type of representation we only cite some representative papers as examples instead of listing all methods that use the representation.

Pointwise position variables One of the simplest ways to define a deformation field is to directly regard the point positions $\widehat{\mathbf{x}}$ on the new shape as the variables and compute them via optimization [LZW*09, HBH11, YKM13]. In practice, due to the physical behavior modeled by the deformation, a vertex can rarely move independently of other vertices, and such pointwise position variables can lead to redundant degrees of freedom. Therefore, the variables are often subject to additional regularization such as local shape preservation [LZW*09, HBH11] and local similarity [YKM13].

Pointwise affine transformation Instead of treating the point positions as variables, some methods define an affine transformation $(\mathbf{A}_i, \mathbf{t}_i)$ on each point of the source surface [ACP03]:

$$\widehat{\mathbf{x}}_i = \mathbf{A}_i \mathbf{x}_i + \mathbf{t}_i, \quad (1)$$

where $\mathbf{A}_i \in \mathbb{R}^{3 \times 3}$ and $\mathbf{t}_i \in \mathbb{R}^3$. Compared with pointwise position variables, this can better model more complex deformations such as local rotations. Based on the observation that the deformations in many non-rigid registration problems are locally rigid, affine transformations are often computed with a constraint that they are close to rigid transformations, *i.e.*, the matrix $\mathbf{A}_i \in \mathbb{R}^{3 \times 3}$ should be close to a rotation matrix [ARV07, YMYK14, YLLL15, LYLG19, YGL*19].

Deformation graph-based methods The methods mentioned previously store the deformation field densely over each sample point on the source surface. This can lead to a large number of variables that incur high computational costs and memory consumption. To reduce the degrees of freedom and achieve better efficiency, some methods adopt an embedded deformation graph [SSP07] to represent the deformation [LSP08, LAGP09, BSB14, CTL15, LLM*16, LJY*17, LJY*17, YDXZ20, LZXH20, ZFA21]. These methods build a deformation graph with nodes embedded on the surface, and with edges indicating the proximity between the nodes (see Fig. 2 for an example). The nodes are typically a subset of the sample points on the source surface. Each node is associated with an affine transformation and influences the deformation in its surrounding region on the surface. For a point \mathbf{x}_i on the surface, the deformation is a convex combination of the transformations induced from the deformation graph nodes that influence \mathbf{x}_i . Specifically, let $\{\mathbf{p}_1, \dots, \mathbf{p}_r\}$ be the set of nodes in the deformation graph. Then the new position of \mathbf{x}_i can be computed as [SSP07]:

$$\widehat{\mathbf{x}}_i = \sum_{\mathbf{p}_j \in \mathcal{I}(\mathbf{x}_i)} \omega_{ij} \cdot (\mathbf{A}_j(\mathbf{x}_i - \mathbf{p}_j) + \mathbf{p}_j + \mathbf{t}_j), \quad (2)$$

where $\mathbf{A}_j \in \mathbb{R}^{3 \times 3}$ and $\mathbf{t}_j \in \mathbb{R}^3$ represent the affine transformation corresponding to the node \mathbf{p}_j , $\mathcal{I}(\mathbf{x}_i)$ is the set of nodes that influence



Figure 2: A deformation graph defined over a mesh surface, with the graph nodes shown in orange.

\mathbf{x}_i , and $\{\omega_{ij} \mid \mathbf{p}_j \in \mathcal{I}\}$ are convex combination weights determined by the distance from \mathbf{x}_i to the nodes in $\mathcal{I}(\mathbf{x}_i)$. Similar to pointwise transformation, the affine transformation is often required to be close to a rigid transformation, which can be enforced via a regularization term when computing the deformation. Alternatively, some methods directly assign a rigid transformation to each node, with the rigid transformation represented using Lie algebra, quaternions or dual quaternions [KCŽO08, DDF*17, LZXH20, LYP*20, ZFA21]. In addition, for smoothness of the deformation, it is a common practice to use a regularization term that enforces consistency between the deformation induced by two neighboring nodes \mathbf{p}_i and \mathbf{p}_j on the same node \mathbf{p}_j [SSP07, LSP08]:

$$E_{\text{smooth}} = \sum_i \sum_{j \in \mathcal{N}_g(i)} \|\mathbf{A}(\mathbf{p}_j - \mathbf{p}_i) + \mathbf{p}_i + \mathbf{t}_i - (\mathbf{p}_j + \mathbf{t}_j)\|_2^2, \quad (3)$$

where $\mathcal{N}_g(i)$ is the index set of neighboring nodes for \mathbf{p}_i .

A deformation graph is typically constructed by selecting a subset of the sample points on the surface as nodes via uniform sampling, and connecting nearby nodes with edges [SSP07]. Since the number of deformation graph nodes is typically much smaller than the number of sample points on the surface, such a representation can significantly reduce the number of variables compared to pointwise transformations. Besides efficiency, the deformation graph also provides more flexibility in the algorithm: a user can change the number of nodes to control the balance between expressiveness and efficiency. It is also possible to adopt a non-uniform sampling that adapts to the semantics of the deformation [YZG*18].

RKHS-based methods Some methods model the deformation with a displacement field $\mathbf{d}(\cdot)$ over the source surface [MSCPn06, MS10, MZT*13, MZY16, MZJZ17]:

$$\hat{\mathbf{x}}_i = \mathbf{x}_i + \mathbf{d}(\mathbf{x}_i). \quad (4)$$

The displacement field $\mathbf{d}(\cdot)$ is often constructed within a reproducing kernel Hilbert space (RKHS) using a Gaussian kernel $\Gamma: \mathbb{R}^3 \times \mathbb{R}^3 \mapsto \mathbb{R}^{3 \times 3}$ with the form

$$\Gamma(\mathbf{x}, \mathbf{y}) = e^{-\beta \|\mathbf{x} - \mathbf{y}\|^2} \cdot \mathbf{I},$$

where β is a parameter and $\mathbf{I} \in \mathbb{R}^{3 \times 3}$ is the identity matrix. The displacement at a point \mathbf{x}_i can then be written as

$$\mathbf{d}(\mathbf{x}_i) = \sum_{j \in \mathcal{I}} \Gamma(\mathbf{x}_i, \mathbf{x}_j) \mathbf{c}_j, \quad (5)$$

where \mathcal{I} is the index set for either all the source surface points or a subset of them, and $\mathbf{c}_j \in \mathbb{R}^3$ are the variables to be determined. The deformation is usually optimized via a probabilistic model, with regularization on the displacement function such as the coherence of motion of nearby points [MSCPn06, MS10].

Patch-based methods Some methods deform the surface in a patch-based manner [HAWG08, CB110, KSI18]. They first segment the surface into a set of patches that are either disjoint or partially overlapping. Each patch is assigned with a separate rigid transformation, which induces deformation for points inside and adjacent to the patch. For globally consistent deformation, the new position of a surface point is computed by combining the deformations induced by the patches, typically using Gaussian weights related to the distance from the point to the patches. Such approaches lead to approximately rigid deformation with each patch. For registration between partially overlapping surfaces, the patch-based representation also helps to extend deformation to regions without correspondence to the target shape [HAWG08].

[XSWL14] also perform registration with a patch-based deformation. Instead of rigid transformation, each patch is deformed by a displacement field using a linear combination of deformation modes learned from simulated data. When determining the deformation, a regularization term is included that penalizes deviation between the displacement fields of two neighboring patches on their common points, to achieve consistency across patches.

[LKSS14] align a depth image of a human face with a database of face shapes in a patch-based manner, to derive a high-resolution face mesh. The input face is divided into several patches corresponding to different facial parts. Each patch is first matched with the faces in the database to determine its new point positions and normals. The positions and normals from different patches are then fused into the complete 3D mesh.

Spline function-based methods Some methods transform the surface using a deformation field defined by a spline function over the 3D space. One common approach is based on the thin plate spline (TPS) [Boo89], which is used to define a displacement field [CR03, DNK12, SK16, FYZ*17, HZD*19]

$$\mathbf{d}(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{b} + \sum_i \mathbf{w}_i Q(\|\mathbf{c}_i - \mathbf{x}\|),$$

where $\{\mathbf{c}_i\}$ is a set of pre-selected control points, and $\mathbf{A} \in \mathbb{R}^{3 \times 3}$, $\mathbf{b} \in \mathbb{R}^3$ and $\mathbf{w}_i \in \mathbb{R}^3$ are parameters that control the displacement field, and $Q(\cdot)$ is a radial basis function. The control points can be either placed on a uniform grid [DNK12], or chosen based on the feature correspondence between the two surfaces [HZD*19].

Another common approach is free-form deformation (FFD) [SP86], which transforms a point $\mathbf{x} = (x_1, x_2, x_3)$ via

$$\hat{\mathbf{X}} = \sum_i \sum_j \sum_k \mathbf{p}_{ijk} B_i(x_1) B_j(x_2) B_k(x_3),$$

where B_i, B_j, B_k are spline basis functions, and $\{\mathbf{p}_{ijk}\}$ are control points that are to be determined. These methods usually select regular grid points as the control points to deform [RSH*99, RS12].

Grid-based methods Some methods define a surface implicitly as a level set of a scalar function [FNT*11, NFS15, ZCS*15, IZN*16, SBCI17]. Typically, the scalar function is a signed distance function (SDF) or a truncated signed distance function (TSDF). The SDF at a point $\mathbf{x} \in \mathbb{R}^3$ can be written as [FNT*11]

$$\text{SDF}(\mathbf{x}) = (\mathbf{x} - \mathbf{x}^*) \cdot \mathbf{n}^*, \quad (6)$$

where \mathbf{x}^* is the closest point to \mathbf{x} on the surface, and \mathbf{n}^* is the unit normal vector of the surface at \mathbf{x}^* . Note that the definition above is applicable also for surfaces represented as point clouds: in this case, $(\mathbf{x} - \mathbf{x}^*) \cdot \mathbf{n}^*$ is the signed distance from \mathbf{x} to the tangent plane at \mathbf{x}^* which is a first-order approximation of the underlying surface [HDD*92]. Considering that points far away from the surface have little effect on the level set, the TSDF truncates the value of $\text{SDF}(\mathbf{x})$ beyond a certain threshold:

$$\text{TSDF}(\mathbf{x}) = \max\{-1, \min\{1, \text{SDF}(\mathbf{x})/t\}\},$$

where t is a threshold parameter. The SDF or TSDF is often discretized using its values evaluated at a set of points sampled as a uniform grid within a volume that encloses the surface. A deformation field defined over the same volume is then used to transform the implicit surface. The deformation field can be represented discretely either on the same grid as the SDF or TSDF [ZCS*15, IZN*16, SBCI17], or on a coarser grid for better efficiency [NFS15]. The deformation of a point on the implicit surface is interpolated from nearby grid points in the deformation field. Alternatively, some methods use such a deformation field to move a set of FFD control points embedded in the same volume, and transform the surface via FFD [FNT*11]. To achieve local rigidity, the deformation field is either represented using a rigid transformation at each grid point [FNT*11, NFS15, ZCS*15, IZN*16], or using a translation at each grid point together with regularization that maintains the distance between points after the deformation [SBCI17].

Prior-based representations Some methods deform the surface in specific ways based on priors about the underlying shape. One example is articulated shapes, which can be deformed using skeleton models [ACP02, PG08, GSDA*09, ZST*10, YGX*17]. For example, in [GSDA*09] the human body motion is modeled using a kinematic skeleton, which consists of straight edges representing the bones, connected at end points representing body joints. The movement of the skeleton is parameterized by a global rigid motion together with the joint angles. The surface is then deformed via skinning [JDKL14], which transfers the skeleton motion to the surface points. Compared with deformation graphs, such skeletons encode a stronger prior of the structure of the underlying shape and can often represent the deformation behavior with fewer degrees of freedom. Some methods also combine a skeleton with a deformation graph or pointwise displacements to model deformations with more detail [BBLR15, YGX*17, YZZ*19, YZG*18, SXZ*20]. Besides skeletons, other representations have been proposed to model articulated structures. For example, the shape of human hands is modeled using a combination of spheres and cylinders in [TST*15], and using sphere-meshes in [TPT16]. The SMPL [LMR*15], a learned

parameterized human body model, has been used for the non-rigid registration of 3D human shapes [BRPMB17, YZG*18, MMRC20].

The shape of human faces is another example where non-rigid registration can benefit from shape priors. Some parameterized models that capture the variation of human face shapes have been used successfully for facial reconstruction and tracking [BTLP16]. For example, the morphable model from [BV99] provides a linear representation for the geometry and texture of the human face, and can be used to reconstruct a neutral face model via alignment to scanned data [WBLP11, IBP15]. Different facial expressions for the same person can be modeled using a linear blendshape model, which can be used for real-time face tracking via registration [WBLP11]. The bilinear face model from [CWZ*14] is used in [AZB15] to align a template face mesh with a point cloud.

Summary and discussion Each representation of deformation has its benefits and limitations. Pointwise position variables and pointwise affine transformations provide an abundance of degrees of freedom to model the deformation in detail, with pointwise affine transformations having the additional benefit of making it easy to model local rotations. However, they require a large number of variables and may induce a high computational cost. Representations based on deformation graphs, RKHS and spline functions can reduce the degrees of freedom needed to model a deformation. This can be suitable for deformations that have a low-dimensional structure. For example, deformation graphs are commonly used for registration between shapes of the same human subject with different poses, where the deformation is primarily induced by the articulation of joints. Meanwhile, such lower-dimensional representations may be less capable of modeling fine-scale variations in the deformation field, and the control structure may need to be carefully selected to account for the topology of the shape and to ensure expressiveness. Patch-based representations can be suitable when the shape consists of multiple regions and the points within the same region undergo similar deformation. Additional care needs to be taken to ensure a smooth transition of the deformation across the boundaries between patches, which can increase the complexity of the algorithm. Grid-based methods provide implicit representations for the surfaces, which can better handle topological changes between the source and the target surfaces. However, they can incur high memory consumption when a high-resolution grid is used. Prior-based representations can effectively utilize the characteristics of specific classes of shapes such as human faces and hands. On the other hand, their performance may degrade for shapes that are not covered by the parametric model.

3. Extrinsic Methods

Given the representation of a deformation field, non-rigid registration amounts to searching the space of deformations for one solution that best aligns the two surfaces. Therefore, it is important to design criteria to measure the alignment quality. A large class of methods measures the alignment error using the distance between two surfaces in the ambient 3D space. These methods are called extrinsic methods in this paper, since they involve the extrinsic distance between surfaces. This section reviews extrinsic methods in detail.

3.1. Optimization-Based Methods

Many methods compute the deformation by minimizing a target function E about the deformation, typically formulated as

$$E = E_{\text{align}} + \alpha E_{\text{reg}}, \quad (7)$$

where E_{align} measures the alignment error between the deformed source surface and the target surface, and E_{reg} is a regularization term that enforces constraints on the deformation such as smoothness. $\alpha > 0$ is a weight that controls the balance between these two terms. Various forms of the alignment term and the regularization term have been proposed in the past. In the following, we review some representative formulations of the terms.

3.1.1. Alignment term

A simple and popular alignment term is based on the distance from each source surface point to the target surface [BM92]:

$$E_{\text{pp}} = \sum_{i=1}^m w_i^d \|\hat{\mathbf{x}}_i - \mathbf{y}_{\rho(i)}\|^2, \quad (8)$$

where $\rho(i)$ is the index of the corresponding point on the target surface for the point $\hat{\mathbf{x}}_i$, and w_i^d is a weight that can be used to control the influence of different points based on the reliability of their correspondence. Such point-to-point distance terms have been used in many registration methods [ACP03, PMG*05, ARV07, LSP08, LZW*09, CZ11, HMS12, RS12, YMYK14]. The correspondence can be identified by either searching for the closest point on the target surface [ACP03, ARV07], formulating the correspondence as optimization variables [LSP08], or matching features such as scene flows [XLC*18].

Some methods measure the alignment using the following function instead [CM92]:

$$E_{\text{ppi}} = \sum_i w_i^d \left(\mathbf{n}_{\rho(i)} \cdot (\hat{\mathbf{x}}_i - \mathbf{y}_{\rho(i)}) \right)^2, \quad (9)$$

where $\mathbf{n}_{\rho(i)}$ is the unit normal of the target surface at the corresponding point $\mathbf{y}_{\rho(i)}$ for \mathbf{x}_i . The term $\left(\mathbf{n}_{\rho(i)} \cdot (\hat{\mathbf{x}}_i - \mathbf{y}_{\rho(i)}) \right)^2$ measures the squared distance from $\hat{\mathbf{x}}_i$ to the tangent plane at $\mathbf{y}_{\rho(i)}$, and is often called the point-to-plane distance. In rigid registration, it is well-known that the ICP algorithm based on the point-to-plane distance can converge faster than the point-to-point distance, as the tangent plane provides a first-order approximation of the local surface shape [PHYH06]. The point-to-plane distance has been used in many non-rigid registration algorithms [LAGP09, XLC*18, LYP*20]. It can be used either as the alignment term alone [DKD*16, YGX*17, DDF*17, LZG18, YZG*18, ZBYX19, LYP*20, LZXH20], or in combination with the point-to-point distance in (8) [LAGP09, CZ11, YKM13, AZB15, GXW*15, TT16, IZN*16, LJY*17, WZX17, XLC*18, LG20, SXZ*20, ZFA21].

If the target surface is represented implicitly with an SDF or TSDF and the source surface is represented explicitly, then the distance from a source point \mathbf{x}_i to the target surface can be simply evaluated as $|D(\mathbf{x}_i)|$, where $D(\cdot)$ is the SDF or TSDF. The alignment term can then be defined as

$$E_{\text{align}} = \sum_i w_i^d \phi(|D(\mathbf{x}_i)|), \quad (10)$$

where ϕ is an increasing function on $[0, +\infty)$. This definition has been used in [DDF*17, XSH*20, SXZ*20, LYP*20] for aligning a template mesh with a TSDF.

If both surfaces are represented using SDFs, the alignment term can be defined as the sum of the squared difference between the SDFs of the deformed source surface and the target surface, evaluated over some sample points in an enclosing volume [FNT*11, ZCS*15, SBCH17]. The underlying idea is that if the two surfaces are entirely aligned, their SDFs should be the same everywhere.

Some methods consider the alignment from a probabilistic perspective, typically using mixtures of Gaussians [JV05, MSCPn06, MS10, JV11, MZT*13, GFD14, GF15, MQZ*15, MZY16]. Their alignment terms are formulated as a probabilistic measure, such as the negative log posterior probability [MSCPn06, MS10] and the ℓ_2 difference between probability density functions [JV05, JV11]. Such probabilistic approaches are discussed in detail in Section 3.1.4.

3.1.2. Regularization term

The regularization term can be a weighted sum of multiple terms, each enforcing a different requirement on the deformation field. Below we review the main types of regularization that have been used in the literature.

Smoothness To avoid unnatural deformed shapes, it is a common requirement that the deformation field should be smooth. This can be achieved using different types of regularization terms according to the representation of the deformation. For deformation fields represented as pointwise transformations, the smoothness can be enforced by penalizing the difference between transformations on neighboring points. For example, in [ACP03] the surface is deformed by pointwise affine transformations, and the regularization term for smoothness is defined as

$$E_{\text{smooth}} = \sum_i \sum_{j \in \mathcal{N}(i)} \|\mathbf{T}_i - \mathbf{T}_j\|_F^2, \quad (11)$$

where the matrix \mathbf{T}_i represents the affine transformation on \mathbf{x}_i , $\|\cdot\|_F$ is the Frobenius norm, and $\mathcal{N}(i)$ is the index set of the neighboring vertices for \mathbf{x}_i . Similar smoothness terms have also been used in [PMG*05, ARV07, FHCP18, ZFA21].

For deformations defined using a deformation graph, the smoothness can be enforced via consistency between the deformations on a graph node that is induced by a neighboring node and the node itself, resulting in a smoothness term as shown in (3). This regularization is commonly used in registration methods using deformation graphs [LSP08, LAGP09, GXW*15, CTL15, WZX17, XLC*18, YDXZ20].

For RKHS-based deformation fields, the smoothness is often induced by penalizing a norm of the deformation field within the RKHS. For example, the deformation field in [MSCPn06] is defined using Gaussian kernels as shown in (5), and the regularization is defined as

$$E_{\text{smooth}} = \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{I}} \Gamma(\mathbf{x}_i, \mathbf{x}_j) \mathbf{c}_i \cdot \mathbf{c}_j. \quad (12)$$

Similar regularizations are also used in [JV05, MS10, MZT*13, GFD14, GF15, MZJZ17].

Positional constraints Depending on the application, some points on the deformed surface are required to be close to some target positions. For example, when aligning a template mesh for human bodies or heads to scanned data, the landmark vertices on the deformed template mesh are required to be close to the landmark locations on the data [ACP03, SCO04, ARV07, YLSL10, TT16, LG20, ZFA21]. Such positional constraints can be enforced with the following term:

$$E_{\text{pos}} = \sum_{\mathbf{x}_i \in \mathcal{C}} \|\widehat{\mathbf{x}}_i - \mathbf{m}_i\|^2, \quad (13)$$

where \mathcal{C} is the set of source surface points with positional constraints, and \mathbf{m}_i is the target position for \mathbf{x}_i .

Local shape preservation When deforming the source surface, it is desirable to prevent distortion of the local surface. Depending on the assumption about the deformation, different types of regularizations can be introduced to achieve this goal.

Many methods assume the deformation to be locally rigid, *i.e.*, each local region undergoes an approximately rigid transformation while the whole surface is deformed in a non-rigid way to align with the target. Various forms of regularization have been proposed to achieve such local rigidity. For many methods that assign affine transformations to either the source points or the deformation graph nodes, regularization terms are introduced to require each affine transformation to be close to a rigid transformation. Some methods use the following term [LSP08, LAGP09, GXW*15, CTL15, WZX17, XLC*18]:

$$E_{\text{rigid}} = \sum_i \text{Rot}(\mathbf{A}_i), \quad (14)$$

where $\mathbf{A}_i \in \mathbb{R}^3$ denotes the transformation matrix for the affine transformation at a graph node, and

$$\begin{aligned} \text{Rot}(\mathbf{A}_i) = & (\mathbf{a}_{i,1} \cdot \mathbf{a}_{i,2})^2 + (\mathbf{a}_{i,1} \cdot \mathbf{a}_{i,3})^2 + (\mathbf{a}_{i,2} \cdot \mathbf{a}_{i,3})^2 \\ & + (1 - \mathbf{a}_{i,1} \cdot \mathbf{a}_{i,1})^2 + (1 - \mathbf{a}_{i,2} \cdot \mathbf{a}_{i,2})^2 + (1 - \mathbf{a}_{i,3} \cdot \mathbf{a}_{i,3})^2 \end{aligned}$$

with $\mathbf{a}_{i,1}, \mathbf{a}_{i,2}, \mathbf{a}_{i,3}$ being the three column vectors of \mathbf{A}_i . The function $\text{Rot}(\mathbf{A}_i)$ enforces the condition $\mathbf{A}_i^T \mathbf{A}_i = \mathbf{I}$ by penalizing the deviation between the two matrices using their upper triangular elements. This condition alone does not guarantee \mathbf{A}_i to be a rotation matrix, since a reflection matrix \mathbf{A}_i with orthonormal columns and $\det(\mathbf{A}_i) = -1$ can still satisfy the condition. In [DKD*16], an additional term $(\det(\mathbf{A}_i) - 1)^2$ is introduced to avoid such reflections. Alternatively, some methods enforce local rigidity by penalizing the deviation between the matrix \mathbf{A}_i and its closest rotation matrix [YLLL15, GLLY17, LYL19, YGL*19, YDXZ20]. A similar approach is adopted in [Sah15] for the registration of tetrahedral meshes, by penalizing the deviation between the deformation gradient of each tetrahedron in the source mesh and its closest rotation matrix; an additional Tikhonov regularizer is also introduced to penalize the change of point positions between compared with the previous iteration, which helps to prevent element inversions and handle large deformations.

Since local rigidity implies that the distance metric is also preserved locally, some methods enforce the condition by penalizing the change of distance between each point and its neighboring

points [WJH*07, SWG08, IZN*16, YGL*19, LG20]:

$$E_{\text{ARAP}} = \sum_{i=1}^m \gamma_i \sum_{j \in \mathcal{N}(i)} \gamma_{ij} \|\mathbf{R}_i(\mathbf{x}_j - \mathbf{x}_i) - (\widehat{\mathbf{x}}_j - \widehat{\mathbf{x}}_i)\|_2^2, \quad (15)$$

where \mathbf{R}_i is an auxiliary variable for a local rotation, and γ_i, γ_{ij} are weights. Unlike previous formulations of local rigidity, this regularization term can be applied to general deformation fields, not just those represented with affine transformations.

The local rigidity conditions require the deformation to be *near-isometric*. Other regularizations based on weaker conditions have also been proposed to allow for *non-isometric* deformations. Since the Laplacian at each point encodes the local geometry [Ale03], some methods penalize the change of the Laplacians after the deformation (potentially up to rotation and scaling) [LZW*09, HBH11, AZB15, GF15]. Meanwhile, in [YKM13, APL14, JQL*17, JYZ*19] the deformation is required to be locally close to a similarity transformation. In other works, [YMYK14, WLLY19] require the deformation to be locally as conformal as possible, while [WAO*09] introduce a regularization for local volume preservation.

Parameter regularization In some methods that rely on a parameterized shape model for the deformation, regularization terms for the shape parameters are introduced to obtain more natural results. In [BWP13], the registration for facial tracking uses a linear blendshape model to represent facial expressions, and the target function includes a regularization term that penalizes the ℓ_2 -norm of the blendshape coefficients. In [TST*15], a parameterized hand model is aligned with the scanned data, using a regularization term that penalizes the deviation between the hand parameters and a PCA subspace that indicates feasible poses.

Other regularizations Other types of regularization can be used according to the specific need of an application. For example, the registration between human body motion frames in [LJY*17] introduces a regularization for the temporal coherence of deformation between adjacent frames. For the registration of hand shapes in [TST*15], regularization terms are introduced to prevent self-collision and avoid impossible poses.

3.1.3. Improving robustness

Many alignment terms and regularization terms are formulated in a sum-of-squares form, which can be the squared ℓ_2 norm of an error vector. The minimization of such an ℓ_2 -norm will attempt to reduce the magnitude of all the error components, which may not be suitable for some cases. For example, when the two point sets contain outliers or overlap partially, some source points may not correspond to any point on the target surface. If we minimize the point-to-point distance (8) with the same weights for all points, then reducing the distance for the points without true correspondence may result in an erroneous alignment. One way to address this issue is to choose a proper weight w_i^d to reduce the influence of source points without reliable correspondence on the target surface. This can be done either by setting the weights according to specific quality criteria of the correspondence, or by jointly optimizing the weights such that they are adjusted automatically [ARV07, LSP08, CZ11, ZCS*15, LLM*16, IZN*16, KSI18].

Another popular approach is to perform sparsity optimization on the distance values. The idea is to minimize a robust norm that can reduce the magnitudes of many error components while allowing some components to gain large values, thus promoting the sparsity of the error vector [BJMO12]. This can be applied to not only the alignment term, but also a regularization term such as the smoothness of the transformation. Sparsity optimization for the regularization error is beneficial in some applications, because this would allow a deformation that is less smooth in some local regions such as the joints for an articulated motion.

A popular choice of sparsity-inducing norms is the ℓ_p -norm with $0 \leq p \leq 1$. In [GXW*15], the ℓ_0 -norm is applied to the smoothness term for the affine transformations defined over a deformation graph. [YLLL15] apply the ℓ_1 -norm to the smoothness term for the affine transformations on each point. In [JYZ*19], the ℓ_1 -norm is applied to the alignment term instead. [LYLG19] and [YGL*19] apply the ℓ_1 -norm to both the smoothness term and the alignment term.

Other robust norms have also been used in the literature. The Tukey function is applied to the point-to-plane alignment term in [NFS15]. The Geman-McClure function is used in the alignment terms in [YZG*18], [XSH*20] and [SXZ*20]. The Huber- L_1 loss is applied to the smoothness term in [JYZ*19] and [ZFA21]. The Welsch function is used for both the alignment term and the smoothness term in [YDXZ20].

3.1.4. Probabilistic models

The optimization-based methods above penalize the geometric distance between the two surfaces. Another class of methods models the alignment accuracy from a probabilistic perspective instead.

A popular probabilistic approach is coherent point drift (CPD) and its variants. The original CPD method [MSCPn06] considers the source point set as the centroids of equally-weighted Gaussians with equal isotropic covariance matrices in a Gaussian Mixture Model (GMM), and considers the target point set as the data points. Using an RKHS-based displacement field (5) to deform the source surface, they align the two surfaces by minimizing the posterior probability, which is equivalent to minimizing the following alignment term

$$E_{\text{align}} = - \sum_{j=1}^n \log \sum_{i=1}^m \exp\left(-\frac{1}{2} \left\| \frac{\mathbf{y}_j - \hat{\mathbf{x}}_i}{\sigma} \right\|^2\right), \quad (16)$$

where σ is the standard deviation parameter for the Gaussians. The above measure is combined with a regularization term for the smoothness of the displacement to derive the target function. In [MS10], an additional uniform distribution is introduced to the GMM to account for noise and outliers. In [GFD14] and [GF15], the CPD framework is used with additional regularization terms that preserve local structures.

Local structures are also utilized in [MZY16] for registration. Similar to CPD, they use a GMM to model the relationship between the two point sets. Unlike the CPD that uses equal membership probabilities in the mixture model, they use local shape features to find correspondence between point sets and determine the membership probabilities. The correspondence and the probabilities are continuously updated during the optimization according to the deformation. [MZT*13] and [MQZ*15] adopt a similar strategy and




Figure 3: Sparse initial correspondence between two models, computed by matching SHOT descriptors [STDS14] followed by pruning the correspondence [TMRL14a].

use feature-based correspondence to facilitate registration. Instead of maximum likelihood estimation, they update the deformation according to the correspondence via L_2E [Sco01], a robust estimator related to minimizing the L_2 distance to the ground-truth density function.

[Hir21] formulates the CPD problem in a Bayesian setting and proposes a solver with guaranteed convergence as well as an acceleration scheme. The Bayesian formulation was also used in [Hir20] to accelerate non-rigid registration, by first performing registration on a downsampled source point set and then interpolating the deformation using Gaussian process regression.

Different from CPD, in [JV05] and [JV11] each of the two point sets is modeled as a mixture of Gaussians, and the registration is formulated as a problem of aligning the two mixtures. They measure the alignment error using the L_2 distance between the probability density functions of the two Gaussian mixtures. Using either a TPS or RKHS representation for the deformation field, the above measure is combined with a regularization term for the deformation field, which is then minimized to derive the deformation.

3.1.5. Initialization

Typically, numerical optimization only finds a local minimum near the initial solution. Therefore, it is often necessary to find a proper initialization in order to achieve the desired result. A simple and popular approach is to first compute a rigid transformation and (optionally) a uniform scaling that roughly align the two surfaces [LSP08, LZW*09, YKM13, FHCP18, YDXZ20, SK21], using the ICP algorithm or its variants.

When there is a significant non-rigid deformation between the two surfaces, more sophisticated initialization may be needed. Many methods initialize the optimization using sparse correspondence that is computed by matching features on the two surfaces. These features are often identified using shape descriptors.

[CZ08] construct correspondence using spin-images [Joh97], a

descriptor for local geometry. They randomly subsample the points on the two surfaces to compute their spin-images, and find the point pairs whose spin-images are significantly more similar than others.

[GTRS16] extract 3D keypoints on the point clouds with an Intrinsic Shape Signature descriptor [Zho09], [RBMB08] compute Persistent Feature Histograms (PFH) at the keypoints, and then compare the PFHs to determine initial correspondence. [MZY16, MZJZ17] adopt Fast Point Feature Histograms (FPFH) [RBB09] descriptors, which simplify the calculation of PFH and achieve comparable discriminative power.

[YLLL15, LYL19, YGL*19, YDXZ20] find the initial correspondence using the SHOT descriptor [STDS14], which is more robust to noisy data. They first match the closest SHOT descriptors, and then use a diffusion pruning method [TMRL14a] to prune the correspondences with inconsistent geodesic distances on the source surface and the target surface. Further discussion of pruning techniques may be found in Section 4.1.1.

Besides the 3D geometry itself, auxiliary information can also be used to identify initial correspondence. For example, for input surfaces in RGB-D format, some methods first use compute SIFT features [Low04, LZW*09, IZN*16, LG20, ZFA21] on the RGB images to identify corresponding pixels; the correspondence is then transferred to the 3D surfaces using the association between the RGB and the depth data [ZST*10, MQZ*15, IZN*16].

3.1.6. Numerical solvers

The optimization problem for non-rigid registration is typically non-linear and non-convex, and the solution is often computed using an iterative numerical solver. Problems with a differentiable target function can be solved using gradient-based solvers. Each iteration of such solvers typically amounts to using the gradient information to determine a search direction, followed by a line search along the direction to find new variable values that lower the target function. Examples of such solvers include quasi-Newton methods such as BFGS or L-BFGS [JS06], which are popular for unconstrained optimization problems [ACP02, ACP03, WAO*09, MZT*13, YHL*16, YDXZ20]. If the target function has a sum-of-squares form, then the problem can also be solved using specialized non-linear least squares (NLLS) solvers such as Gauss-Newton or Levenberg-Marquardt [WJH*07, LSP08, LAGP09, CZ11, SK13, YMYK14, BSB14, ZCS*15, CLS20, ZFA21]. Such quasi-Newton and NLLS solvers can achieve a super-linear convergence rate [JS06], and there are established open-source implementations such as Ceres Solver for Levenberg-Marquardt [AMO]. For a complicated target function, the derivation and computation of its gradient can be non-trivial and error-prone. This issue is mitigated by automatic differentiation tools provided by software libraries such as Ceres Solver, which only require the user to implement the target function and can automatically compute its derivatives.

Despite the popularity of gradient-based solvers, they are not applicable to all problems. For example, robust norms can lead to a non-smooth target function that cannot be handled using gradient-based solvers. Such non-smooth problems are often solved using a first-order method [Bec17] instead. The main idea is to introduce auxiliary variables to reformulate an equivalent optimization

problem, whose solution can be computed using a method that involves simple sub-problems that are easy to solve. For example, block coordinate descent [Tse01] is used to solve the reformulated problem in [GXW*15], while ADMM [BPC*11] is used in [YLLL15, LYL19, YGL*19]. One benefit of first-order solvers is that their sub-problems often have closed-form solutions that can be computed in parallel, allowing them to benefit from multi-core systems to improve efficiency.

The idea of decomposing a problem into sub-problems with closed-form solutions is not limited to non-smooth problems. In [ARV07], it is noted that the target function can be minimized in closed-form by solving a linear system if the point correspondence is fixed. Thus they alternate between point correspondence search and linear solve, resulting in a procedure similar to the ICP algorithm for rigid registration. Such ICP-like solvers are also used in [YKM13, YMYK14, IZN*16]. For methods based on probabilistic models, the problem is often solved using an EM algorithm consisting of alternating updates with closed-form steps [MSCP06, MS10, GF15, MZJZ17, KSI18, GT19, Hir21].

3.2. Learning-Based Methods

With a suitable target function, an optimization-based registration method tends to work well when the two surfaces are close. For problems that involve large deformations, optimization methods may converge to an undesirable local minimum if a reliable initial correspondence is not provided. Such a correspondence is not always easy to compute: closest point correspondence is sensitive to the relative position between the surfaces, while shape descriptors may still be ambiguous. For such challenging cases, prior knowledge learned from large-scale data can be beneficial. With the fast development of deep learning in recent years, various learning-based methods have been proposed for non-rigid registration. These methods train a model—typically a neural network—using a set of example data, such that the trained model can reliably predict the deformation or correspondence for unseen data. The training is done by searching for model parameters to minimize a loss function that is evaluated using the training data. In the following, we will review some representative exiting works.

3.2.1. Supervised methods

A supervised method trains the model using data with ground-truth labels, such as the correspondence or the deformation field. The aim is to train a model that can predict the unknown ground truth for unseen data with reasonably high accuracy.

Shape correspondence Some methods focus on finding correspondence between shapes. [WHC*16] propose a method to compute dense correspondence between human body shapes. They use a *convolutional neural network* (CNN) to extract the feature descriptors of the depth image and establish the correspondences by nearest neighbor search in the space of descriptors. [TCM*21] propose a transformer-based architecture to register 3D point clouds and obtain their correspondence, using a surface attention mechanism that adapts to the point cloud density and the underlying geometry.

Learnable correspondence with deformation graph Some methods integrate the learned correspondence and combine it with a deformation graph, to deform the source surface and align with the target surface by optimization. [BZTN20] adopt an encoder-decoder architecture to determine the correspondence between two RGB-D frames, by predicting for each source point a heatmap of correspondence probability on the target frame. The correspondence, identified with the highest-probability point on the target, is then used to optimize a deformation graph-based transformation to align the surfaces. They also use a semi-supervised process to collect a large amount of data for training and testing. Concurrently, [LBZ*20] propose an end-to-end method to perform non-rigid RGB-D tracking. They train a fully convolutional network to extract feature maps for both the source and the target frames, which are used to define a feature term in an optimization problem for a deformation graph-based transformation that aligns the frames. They solve the optimization using Gauss-Newton, and train an encoder-decoder network to produce PCG preconditioners that accelerate the solver's convergence. Further, [BPZ*20] propose a dense correspondence prediction module for the source RGB image and the target RGB image. Through this module, they obtain the dense correspondences and weights to measure the accuracy of these correspondences. Under these correspondences, they use the deformation graph with a differentiable solver to align the input point clouds, which are obtained by depth images. They can obtain better correspondences and reliable weights by end-to-end training.

Learnable displacement field In [SGT*19], two input point clouds are first converted to voxel grid representations, and then aligned using a voxel displacement field predicted by a network pipeline. The pipeline consists of a displacement estimation stage that predicts large global deformation and a refinement stage that recovers small displacements, each implemented with a UNet-style 3D convolutional network. The model is trained on the FLAME [LBB*17], DFAUST [BRPMB17], thin plate [GSVS18], and cloth [BFS18] datasets.

Registration with parametric representation Recently, many methods based on parametric shape models have emerged. Some methods use the SMPL model [LMR*15] as the template for the human shape, then optimize SMPL parameters to align with any human body point clouds, including the human body in clothes [GFK*18, WXZ*20, WGT21]. In [GFK*18], an auto-encoder network is used to predict the correspondence between a template shape and an input human shape, which is then used to perform matching between two human shapes with the template mesh as the proxy. The network is trained using synthetic data generating from SMPL with estimated parameters from the SURREAL dataset [VRM*17]. [WXZ*20] take a sequence of point clouds as the input, and propose a spatial-temporal attention convolution to directly regress the vertex coordinates of a fitted mesh model in a coarse-to-fine manner. The model is trained using the scanned data sequences and corresponding SMPL models from the DFAUST dataset [BRPMB17]. In [WGT21], the SMPL model is fitted to the input point cloud of a dressed human. They use the occupancy functions to represent the canonical space, and propose a piecewise transformation field to deform the source point into the canonical space. The CAPE dataset [MYR*20] is used for training.

Loss function design Different loss function terms have been adopted in supervised methods to learn the correspondence or deformation. When ground-truth correspondence is available in the training data, some methods introduce function terms that require the predicted correspondence to be close to the ground truth. For example, [BZTN20] use a probability heatmap to represent point correspondence, and compute the binary cross-entropy loss and negative log-likelihood loss between the predicted and ground-truth heatmaps to penalize their deviation. Other methods require some points on the predicted deformed source surface to be close to their ground-truth positions, using a loss function term that penalizes the difference [GFK*18, BPZ*20, WXZ*20, LBZ*20, TCM*21]. In [SGT*19], a coarse alignment prediction module is trained with a loss function that penalizes the deviation from the ground-truth displacement, while a refinement module is trained using an alignment loss similar to (8) to refine the registration.

3.2.2. Unsupervised methods

Supervised methods require training data with ground-truth correspondence or deformation, which are not easy to obtain. Thus some methods train the model in an unsupervised manner, using data without ground-truth labels. In this case, the loss function plays an important role in guiding the training of the model. Below we review some loss functions that have been used in the literature.

GMM loss [WCLF19] use a CNN to predict a TPS-based deformation field that aligns two point clouds. To train the model, they use a GMM loss to measure the alignment quality, which has the same form as (16), with the Gaussian standard deviation being a hyper-parameter.

Chamfer distance [WLCF19] use an MLP to predict a displacement field for aligning two point clouds. They train the model using the Chamfer distance to measure the alignment [FSG17]:

$$L_{\text{Chamfer}} = \sum_{\mathbf{x} \in \mathbf{X}} \min_{\mathbf{y} \in \mathbf{Y}} \|\hat{\mathbf{x}} - \mathbf{y}\|_2^2 + \sum_{\mathbf{y} \in \mathbf{Y}} \min_{\mathbf{x} \in \mathbf{X}} \|\hat{\mathbf{x}} - \mathbf{y}\|_2^2.$$

For point clouds with outliers and missing data, they also propose the following clipped Chamfer distance loss:

$$L_{\text{C-Chamfer}} = \sum_{\mathbf{x} \in \mathbf{X}} \max(\min_{\mathbf{y} \in \mathbf{Y}} \|\hat{\mathbf{x}} - \mathbf{y}\|_2^2, c) + \sum_{\mathbf{y} \in \mathbf{Y}} \max(\min_{\mathbf{x} \in \mathbf{X}} \|\hat{\mathbf{x}} - \mathbf{y}\|_2^2, c),$$

where c is a hyper-parameter.

Point-to-point reconstruction loss [ZQZ*21] propose a feature embedding module to extract the deep features of the input point clouds, a correspondence indicator module to learn the correspondences, and a symmetric deformer module to deform the source point cloud to align with the target point cloud and deform the target point cloud to align with the source point cloud. Then they use a point-to-point reconstruction loss to measure the quality of the deformation and the alignment:

$$L_{\text{pp-recon}} = \|\mathbf{X} - \hat{\mathbf{Y}}\|_F^2 + \|\mathbf{Y} - \hat{\mathbf{X}}\|_F^2,$$

where $\hat{\mathbf{X}}$ and $\hat{\mathbf{Y}}$ are the deformed source point cloud and the target point cloud, respectively. This point-to-point reconstruction loss is often easier to train than the Chamfer distance.

Earth mover's distance [WCMN19] use an encoder-decoder network to learn a displacement field to align a source mesh with a target mesh or point cloud. They use the earth mover's distance [FSG17] as part of their alignment measure between the source and target:

$$L_{EMD} = \min_{\phi: \mathcal{X} \mapsto \mathcal{Y}} \sum_{\mathbf{x} \in \mathcal{X}} \|\hat{\mathbf{x}} - \phi(\hat{\mathbf{x}})\|_2,$$

where $\phi: \mathcal{X} \mapsto \mathcal{Y}$ is a bijection. In [WCMN19], this term is used in conjunction with the Chamfer distance to measure the alignment error.

Multi-view projection loss [FZC*21] propose a multi-view projection loss function for aligning point clouds and achieve better results than the Chamfer distance. The idea is to project the 3D shapes onto multiple 2D planes with different orientations, and measure the differences of the depth and mask images between the source shape and the target shapes as an indication of the alignment error. The main difference compared with Chamfer distance and Earth mover's distance is that it provides more directional information. In addition, they propose a recurrent update framework that defines several rigid transformations $\{\mathbf{R}_r\}$ without specific locations to represent the deformation field. To represent the deformation on each source point, they define

$$\hat{\mathbf{x}}_i = \sum_{r=1}^K w_i^r \cdot \mathbf{R}_r,$$

where w_i^r is the skinning weight for \mathbf{x}_i with the constraint:

$$\sum_{r=1}^k w_i^r = 1, \quad \forall i = 1, 2, \dots, m.$$

Their model is trained on the HumanMotion [VBMP08], TOSCA [BBK08] and SURREAL [VRM*17] datasets, and tested on the FaceWareHouse [CWZ*14] and DFAUST [BRPMB17] datasets.

Regularization Similar to optimization-based methods, it is often insufficient to use just an alignment term for the training loss. Some additional regularization terms may be needed to prevent trivial solutions and to produce natural results. For example, in [WCMN19], the training loss includes a symmetry term to promote symmetry for man-made objects, a Laplacian term to preserve local geometric detail, and a local permutation invariant term to prevent self-intersections. In [FZC*21], they include an as-rigid-as-possible term to preserve local shapes, and regularization terms for the translation and the skinning weights, respectively. Some methods include other regularization terms, such as the regularization of the learnable parameters, to make the training easier and to prevent overfitting [FZC*21, BPZ*21].

3.3. Applications and Systems

With the development of the RGB-D data acquisition equipment and technologies, there has been a growing number of works on dynamic reconstruction and tracking based on non-rigid registration.

3.3.1. Template-less methods

DynamicFusion [NFS15] adopts the TSDF to represent the canonical model and support the fusion operation over multiple frames.

Without any prior, they simultaneously fuse geometry information and estimate the deformation field. The system can run in real-time and obtain relatively complete reconstruction results with details. However, it is not robust to fast movements and occluded areas. Further, VolumeDeform [IZN*16] uses a volumetric representation to define the deformation graph and encode the scene's geometry, the accuracy of correspondences is improved by introducing the SIFT feature. [DKD*16] use multi-view capture information and propose a key volume strategy to handle failures in tracking. To improve the accuracy of the correspondences obtained, they use a learning-based method, GPC [WRF*16], and extend it to make use of both the RGB image and the depth image to obtain more robust results within an acceptable time. [GXY*17] use a single RGB-D camera to reconstruct geometry, albedo, and motion in real-time. Observing that the appearance is beneficial to the accuracy of the reconstruction, they combine the shading information to estimate the motion more accurately and also obtain a lighting estimation.

[WZX17] split the input RGB-D sequences into several parts. In each part, they perform the local non-rigid bundle adjustment to register all scans. Then they merge all parts to eliminate the drift problem by global optimization. A final 3D model can be generated by using Poisson surface reconstruction or volumetric fusion. [SBI18] represent each RGB-D frame as a TSDF and deform it to align with the canonical TSDF. They use a Sobolev gradient flow to achieve faster computation and better detail. They additionally compare different regularization strategies, and utilize signature matching to handle rapid motions [SBI21]. [XSWL15] propose a two-stage method. First, they perform CPD to register the previous frame and the current frame. Second, they use a mean shape and a deformation on a low-dimensional space to fuse the information of the new frame. [XXY*18] uses non-rigid registration to reconstruct nearly rigid objects. Due to the noise and distortion caused by camera acquisition, they use a deformed graph with few graph nodes as a deformation representation to reduce the error and improve the reconstruction quality.

3.3.2. Methods based on priors

Considering the prior information of the human body, [BBLR15] use the human parametric model to reconstruct people in motion from monocular RGB-D sequences. Besides estimating the low-dimensional representation, they define a variable-detail shape model and use a low polygon count mesh with a high-resolution texture map and displacement map to represent the details. BodyFusion [YGX*17] embeds the body skeleton model in the human model to enhance the robustness to large motions. They propose a binding term to establish the relationship between the skeleton model and the deformation graph. They improve the performance and robustness of the non-rigid reconstruction and tracking. However, a skeleton model that is too sparse is not robust to handle fast motion or incomplete surfaces. [YZG*18] propose a double-layer model including a body template model and a dense geometry model to address this problem. It makes the system more robust to the fast movement. [ZYL*18] add sparse inertial measurement units (IMUs) to improve the tracking performance. Unstructured-Fusion [XSH*20] uses three deep cameras without any careful pre-calibration to cover the whole human body and achieve real-time high-quality human capture. In addition, they propose a skeleton

warping based non-rigid tracking scheme to calibrate the camera and track at the same time. [TT16] use the first RGB-D frame with facial landmarks and neutral expression to obtain an initial head model, then they use it to track changes in facial expression. [YMO18] combine 3D Morphable Model (3DMM) [BV99] and a head reconstruction framework without prior for accurate head tracking. [TST*15, TPT16, TTR*17] use the prior of a hand to model and track dynamic hand motion in real-time. In [WBLP11] and [BWP13], parametric face models are used for real-time face tracking on commodity RGB-D sensors, allowing the user to control the facial expressions of a digital avatar. Besides these class-specific priors, some methods establish the template model by the first frame or use a pre-scanned model [TZL*12, XLC*18, LZXH20].

3.3.3. Learning-based methods

[DDF*17, BZTN20] use the learning-based technique to obtain a more accurate correspondence. [LBZ*20] improve the registration quality by using the learned feature as an alignment term. [BPZ*20] combine correspondence learning and differentiable solving, and form an end-to-end network. [BPZ*21] propose an implicit neural deformation graph with learnable node positions and influence radii, which can be applied to a sequence of depth camera observations to perform globally-consistent registration and reconstruct a non-rigidly moving object. [SXZ*20] adopts an occupancy network, a pose and shape estimation and semantics networks are used to improve the robustness of the reconstruction. [YZG*21] combine a temporal volumetric fusion technique and deep implicit functions. They first use the multi-view RGB-D input to perform dynamic fusion and obtain a volumetric fusion. They then re-render multi-view RGB-D images from this volume. Finally, they obtain the detailed model through a deep implicit surface reconstruction step. [ZBYX19] perform real-time reconstruction of a hand and deformable objects held by the hand. They first segment the hand, the deformable object and the background, and then use an LSTM-based network to predict the hand pose according to temporal information. Finally, they utilize the predicted pose to perform the frame-by-frame optimized tracking.

3.3.4. Methods combining optimization and learning

Some methods combine optimization-based and learning-based approaches to improve non-rigid reconstruction performance. [BZTN20] first use a learning-based approach to obtain sparse correspondences represented by probability heatmaps. The correspondences are then utilized to formulate an optimization problem for non-rigid registration. [ZBYX19, SXZ*20] first use a learning-based method to predict the parameters of a prior-based shape model to achieve an initial coarse alignment. Afterwards, the registration is further improved by optimizing a detailed deformation field. [BPZ*20] first use a CNN to predict dense correspondences between the two surfaces; the correspondences are then used as constraints to perform an optimization-based registration. [LBZ*20] propose an optimization formulation with an alignment term that is based on learned features and helps the Gauss-Newton solver to avoid local minima; additionally, a learned preconditioner is utilized to improve the convergence of the PCG solver used by the Gauss-Newton method.

3.4. Summary

Tables 1 and 2 summarize the optimization-based and learning-based extrinsic methods discussed in this section, respectively. Papers that involve both optimization-based and learning-based methods are listed in both tables according to their relevant components.

4. Intrinsic Methods

Unlike extrinsic methods that align surfaces in the original \mathbb{R}^3 coordinate space, by the most general definition, intrinsic methods transform a set of surfaces into an alternative representation in another coordinate space—or domain—in which they are aligned. With this in mind, intrinsic methods can be categorized by the characteristics of the domain they are embedded in, *i.e.*, (1) parametric; (2) minimum-distortion; and (3) spectral/functional. Parametrization-based methods compute low-dimensional embeddings, typically in \mathbb{R}^2 . Minimum-distortion refers to methods that minimize some form of intrinsic correspondence measure, the dimensions of which depend on the number of keypoints between surfaces. Spectral/functional methods embed surfaces in a high-dimensional space where the correspondence problem may be solved using established mathematical tools. Furthermore, by avoiding working directly on the graph, the influence of mesh connectivity is lessened greatly.

Much of the related literature only establishes correspondence. Few methods explicitly deform the source shape to the target shape. However, it is possible to incorporate correspondences as hard/soft constraints in a non-rigid registration framework to help align surfaces.

A variety of techniques can be extended to the problem of non-rigid registration. A key deficiency of many approaches is their sensitivity to partial data—a common scenario in registration. This vulnerability is caused by the use of intrinsic measures that are sensitive to geometric and topological noise. Despite this challenge, a variety of methods have been proposed in the literature to address this problem, *e.g.*, minimum-distortion [Sah18]; spectral reconstruction [CPR*19, HZFH20]; learning-based [ENK*21]; hybrid (intrinsic initialization then extrinsic refinement) methods [ELC20, MMRC20].

[Sah20] recently reviewed a broad range of literature that explores the problem of shape correspondence—with a predominant emphasis on intrinsic methods. For this survey, a keen focus is made on literature that seeks to develop robust partial matching techniques. These works contribute towards solving the partial matching problem, which is especially relevant for applications, such as reconstruction, where data is missing or incomplete. One may consider these methods to address similar problems to registration methods through the use of intrinsic data. It is also possible to combine intrinsic and extrinsic techniques to form robust registration pipelines. After a systematic review of material relating to parameterization-based methods, it has been deemed appropriate to limit the scope of this survey to exclude such methods. The reader is directed to [SPR07, HLS07] for a broad overview of parameterization techniques and discussions on early works on cross-parameterization.

Table 2: A summary of learning-based extrinsic methods.

	Objective		Training type			Loss					
	Learnable correspondences	Learnable deformation	Supervised	Semi-supervised	Unsupervised	Ground-truth correspondences	Ground-truth positions	Alignment	Near-isometry	Non-isometry	Other
[WHC* 16]	✓		✓			✓					
[GFK* 18]	✓	✓	✓		✓	✓	✓		✓	✓	
[SGT* 19]		✓	✓				✓	✓			
[WCLF19]	✓				✓						
[WCMN19]		✓			✓			✓			
[WLCF19]		✓			✓			✓			
[ZBYX19, SXZ* 20]		✓	✓								✓
[BPZ* 20]	✓		✓				✓				
[BZTN20]	✓		✓			✓	✓				✓
[LBZ* 20]		✓		✓			✓				
[WXZ* 20]	✓	✓	✓				✓				
[BPZ* 21]		✓			✓			✓	✓		✓
[FZC* 21]		✓			✓			✓	✓		✓
[TCM* 21]	✓	✓	✓				✓				
[ZQZ* 21]	✓	✓			✓			✓			✓

4.1. Optimization-Based methods

The general objective of most intrinsic matching algorithms is to compute a bijective linear map $T : \mathcal{X} \mapsto \mathcal{Y}$ that maps corresponding points between two shapes (\mathcal{X} and \mathcal{Y}). For convenience of notation, in this section, let all shapes comprise of n points (*i.e.*, $m = n$).

Point-to-point mapping Many intrinsic methods seek to approximate T by computing a point-to-point mapping of a permutation matrix $\mathbf{P} \in \mathcal{P}_n$ where $\mathcal{P}_n = \{\mathbf{P} \in \{0, 1\}^{n \times n} \mid \mathbf{P}^\top \mathbf{P} = \mathbf{I}\}$ is the set of permutations. The problem of finding \mathbf{P} is cast as an optimization problem by a given energy function E :

$$\arg \min_{\mathbf{P} \in \mathcal{P}_n} E(\mathbf{P}). \quad (17)$$

In this form, \mathbf{P} is an orthogonal matrix (s.t. $\mathbf{P}^\top \mathbf{P} = \mathbf{I}$), which is naturally bijective. The permutation set is quite large $|\mathcal{P}_n| = n!$. It is therefore impractical to evaluate all permutations, so further heuristics must be employed.

Worse yet, in the case of partial shape matching, the optimal solution is no longer a bijection but either injective or surjective. The solution space is therefore no longer contained within \mathcal{P} but expanded to an even larger set. If the overlap between \mathcal{X} and \mathcal{Y} was known, then it would be possible to simply seek an optimal bijection between the overlapping regions. However, this leads to a chicken-or-the-egg problem, as without some form of mapping it is not possible to determine the overlap.

4.1.1. Minimum-distortion methods

Minimum-distortion methods are a subset of intrinsic techniques that operate principally in the spatial domain and penalize the distortion of intrinsic measures. Intrinsic measures include: *e.g.*, geodesics [BBK06]; conformal maps [LF09, KLCF10, KLF11]; heat kernel maps [OMMG10]; and local feature descriptors [SOG09, BK10, TSDS10, ASC11]. The objective is to then optimize a latent correspondence such that the distortion induced is minimized. There is a substantial amount of literature surrounding this class of approach. The discussions in this survey are therefore limited to methods that consider the problem of partial matching, as these have greater applicability to the problems of registration and reconstruction.

Partial matching The early work [ASP*04] proposes a complete registration pipeline for part-to-full shape matching. To aid the non-rigid ICP method, sparse correspondences are established using a belief propagation algorithm. The algorithm seeks to promote matches between similar keypoint signatures probabilistically. The solution space is constrained by pairwise geodesics, which discard matches that lead to a large metric distortion. This helps to ensure correspondences are locally consistent. The probability-based optimization computes a general mapping, which allows the method to handle partial correspondence. Finally, a non-rigid registration algorithm is applied to align the surfaces. The proposed approach is only suitable for piecewise rigid shape deformation, while most works in the past decade facilitate isometric or near-isometric deformation. [BBK06] extend multi-dimensional scaling to isometrically embed one surface into another. The method is capable of handling partial matching, but is limited to isometric deformation. The method requires points on the input meshes to be sampled coarsely due to the use of geodesics and a quadratic optimization. Given a set of sparse keypoints between a partial source shape and a full target shape, [SY12b] select five of the farthest keypoints on the partial shape, and then perform an exhaustive search to find an optimal permutation of correspondences on the target shape that minimizes a geodesic-based distortion measure. The proposed optimization procedure becomes expensive as the number of keypoints on the full shape is increased. [RBA*12, Sah18] formulate distortion minimization in evolutionary optimization frameworks. [Sah18] employ a genetic algorithm in which permutations of correspondences are represented by chromosomes. Modifying a part of the current correspondence is equivalent to a mutation, and error is measured by geodesic distance. An important advantage of this approach is that it is robust to poor initialization. Similar to the previously discussed works, [Sah18] relies on establishing keypoints consistently between shapes. Uniquely, after each optimization iteration, [Sah18] refines the keypoint locations on the target shape with respect to the pairwise geodesics of candidate correspondences on the source shape. [RBA*12] use an evolutionary game theory-based approach. Rather than determining a bijective correspondence by solving a traditional quadratic assignment problem, they relax this requirement such that not all points must be in correspondence by solving a quadratic semi-assignment problem. [OMMG10] propose a method that relies on constructing heat kernel maps for each shape. To han-

dle partial correspondence, the heat diffusion parameter is decreased, which effectively reduces the size of the neighborhood represented in the map. However, the approach is limited to isometric deformation. [VLB*17] establish a correspondence by solving a linear assignment problem that is refined in an iterative manner. To handle partial correspondence the assignment problem is relaxed by introducing slack variables, which are intended to absorb erroneous candidate correspondences from parts of a full shape that do not overlap. [vKZH13] develop a geodesic-based pairwise descriptor that is designed to be robust for partial matching. When incorporated with distortion-minimizing techniques, the method is demonstrated to be robust to a degree of topological noise. The descriptor requires large geodesics, therefore it is relatively expensive to compute, and is superseded by recent descriptor learning methods.

A common problem amongst most of the described approaches is that many rely on solving complex assignment problems that become very expensive to solve at larger problem sizes. This leads to the use of sparse keypoint sampling techniques, which introduce further complexity to consistently sample shapes. Also, methods that use global measures (e.g., geodesics over the entire shape) are sensitive to topological change in which geodesic paths may change. Furthermore, for partial matching, these methods are only suited to (near-)isometric deformation due to their reliance on consistent geodesics between overlapping regions. [CRB*16] demonstrate that state-of-the-art minimum-distortion methods tend to perform relatively poorly for partial correspondence problems when compared to functional and learning-based approaches.

Sampling Broadly, the partial matching methods, discussed previously, employ computationally demanding techniques to solve general assignment problems. To alleviate some of the computational complexity, sparse sets of keypoints are used. Sampling may be necessary in cases where a method requires the same number of points on each shape. A variety of strategies are employed. [TBW*11] use a probabilistic approach to construct an order to incrementally introduce points that discriminate highly from previously added points, minimizing an entropic measure. [SY11, SY12a] propose a coarse-to-fine strategy, where points that are at shape extremities and high curvature areas are sampled first. Ideally, both shapes will exhibit similar geometric curvature properties. After applying an initial sampling procedure, [Sah18] optimize the location of the sampled points on the target shape—helping to reduce geodesic distortion. [VLB*17] propose a multi-scale/coarse-to-fine approach that uses farthest point sampling [ELPZ97]. A key benefit of this method is that a dense correspondence is retrieved. This is not so critical when used to initialize a registration method; however, if the initial correspondence is too sparse for a particular scenario, such a refinement technique may be necessary. [LJO19] propose a coarsening strategy that seeks to preserve the original shape’s spectral representation. [VLB*17] is adapted to use this technique—producing superior results. [CCS12, NBH18, LLT*20] explore other sampling strategies for surfaces that are able to effectively coarsen and accurately approximate spectral properties of the original surface. When sampling partial shapes, the greatest challenge is selecting consistently localized sample points on the surfaces. Methods that support partial matching generally must incorporate ma-

chinery for this into the matching portion of the correspondence pipeline [SY12b, VLB*17, Sah18].

Correspondence pruning Many non-rigid registration methods require a set of accurate sparse correspondences for initialization [ASP*04, HAWG08, YLLL15, LYL19, YGL*19, YDXZ20]. A simple approach to compute an initial correspondence is to match feature descriptors between shapes by their similarity. In practice, it is challenging to engineer feature descriptors that are suitably discriminative while also being invariant to different types of deformation. It is therefore necessary to remove poor correspondences. A successful intrinsic-based approach has been to prune correspondences that cause a large amount of geodesic distortion with respect to other candidate correspondences (see Fig. 3). An early work, [HAWG08], measures consistency between geodesics c_{ij} as a ratio of their distances. This formulation has the advantage that distortions between distant geodesics do not disproportionately influence the measure.

$$c_{ij} = \min \left\{ \frac{d_{\mathcal{X}}(\mathbf{x}_i, \mathbf{x}_j)}{d_{\mathcal{Y}}(f(\mathbf{x}_i), f(\mathbf{x}_j))}, \frac{d_{\mathcal{Y}}(f(\mathbf{x}_i), f(\mathbf{x}_j))}{d_{\mathcal{X}}(\mathbf{x}_i, \mathbf{x}_j)} \right\}. \quad (18)$$

A threshold τ is set experimentally to 0.7. Correspondence pairs are removed if $c_{ij} \leq \tau$. Using this measure for determining isometry globally is slow due to the cost associated with computing geodesics over large distances. The proposed method is incorporated into a registration pipeline that is suitable for piecewise rigid deformation.

[TMRL14a, TMRL14b] adapt (18) to only be computed within a small geodesic disc. The authors note that geodesics to distant points become unreliable due to the presence noise, and under non-isometric deformation or geometric changes such as holes. For queue-based path finding algorithms that are similar to Dijkstra, this is simple to implement: (1) remove nodes from the queue that are beyond the distance threshold; and (2) terminate execution once there are no more unvisited nodes neighboring visited nodes in the queue. [DLRT19] further adapt this algorithm for non-isometric deformation by replacing the conventional geodesic algorithm with an anisotropic geodesic distance algorithm in which the model for anisotropy is updated iteratively. The described methods are sensitive to their initialization, requiring shape descriptors that are invariant to the exhibited deformation.

Optimal transport A variety of works have formulated shape correspondence as an optimal transport problem [SNB*12, RBA*12, SdGP*15, MCSK*17, PRM*21]. These works consider the Gromov-Wasserstein distance, which is a quadratic problem. This can be computed more efficiently using the Sinkhorn algorithm [Cut13], which alternately optimizes the transport of the source and the target. Relevantly, [BPC16] demonstrate that optimal transport can be used for shape completion of voxel-based shapes. However, for intrinsic shape correspondence, these techniques are yet to be successfully applied to partial correspondence.

Non-isometry Predominantly, intrinsic methods in the literature typically consider isometric and near-isometric deformation. The challenge of handling non-isometries has been considered in a select range of works. The types of non-isometric deformation addressed by different techniques varies greatly. [LF09] search for an optimal conformal map. The method is sensitive to deformations




Figure 4: The first eight eigenvectors on \mathcal{X} and \mathcal{Y} in the ascending order of their corresponding eigenvalues. Red indicates where the eigenvector is positive, blue negative, and white is the reference with a value of zero. N.B. the first eigenvalue $\lambda_1^{\mathcal{X}}$ is approximately zero and is therefore omitted.

that distort the map, and, therefore, the method performs poorly on partial matching problems. [KLF11] combine multiple conformal maps to produce a “blended map” that minimizes this distortion. [TMRL14a] seek locally isometric correspondences between shapes, using a geodesic-based global consistency check to prune inconsistent correspondences. This global check can be relaxed to increase the degree of global non-isometry permitted. [DLRT19] extend [TMRL14a] to handle local anisotropy. [SPKS16] solve an optimal transportation problem, that for full-to-full matching, can handle some non-isometric distortion; however, the approach is not guaranteed to obtain the globally optimal solution. [VLB*17] implement a probabilistic framework that is suitable for finding smooth maps. [Sah20] provides further discussions on the development of shape correspondence approaches that handle non-isometries.

Shape recovery Given a template, [DRB*09] consider the problem of reconstructing a deformed triangle mesh by using an optimization-based technique. An objective function comprising of a regularization term and a data term is proposed. The regularization term uses a pairwise geodesic-based distance between the original template and the deformed template, which seeks to reduce distortion. The data term measures the distance between corresponding points on the deformed shape and the target shape.

To summarize, metrics used by minimum-distortion methods, such as geodesics, are highly sensitive to geometric and topological noise. For these methods to cope with partial correspondence, an assumption of (near-)isometric deformation must be made. The sampling strategies discussed exhibit inconsistent results when sampling on partial or full surfaces.

4.1.2. Spectral/Functional mapping methods

Functional mapping fundamentals Rather than directly computing a point-to-point map, the shapes are abstracted by functions defined upon each surface. The functions are computed from the Laplace–Beltrami operator. A common discretization is the cotangent Laplacian [MDSB02] $\mathbf{L} = \mathbf{D}^{-1}\mathbf{M}$, which is composed of $\mathbf{D}^{n \times n}$ a diagonal mass matrix equal to $\frac{1}{3}$ of the total area of the faces within the 1-ring neighborhood of each vertex, and $\mathbf{M}^{n \times n}$ is a matrix of cotangent weights between connected vertices:

$$m_{ij} = \begin{cases} \frac{1}{2}(\cot \alpha_{ij} + \cot \beta_{ij}) & i, j \in E \\ 0 & \text{otherwise,} \end{cases}$$

where α_{ij} and β_{ij} are the internal angles at the vertices opposite edge ij .

The eigenbasis of the intrinsic Laplace–Beltrami operator—denoted by Φ , with corresponding eigenvalues Λ —is often selected as it is invariant under isometric deformation.

$$\mathbf{M}\Phi = \mathbf{D}\Phi\Lambda. \quad (19)$$

Fig. 4 illustrates the characteristics of the eigenvectors. It can be seen that each eigenfunction ϕ_i varies smoothly over the surface. As the eigenvalue increases, the oscillations observed in the eigenfunction are similar between the isometrically deforming shapes, up to their sign (i.e., $\phi_i^{\mathcal{X}} \approx \pm \phi_i^{\mathcal{Y}}$). For the motivations and specific implementation details, please refer to the seminal paper on functional maps [OBCS*12].

Given this functional representation, it is possible to compute




Figure 5: The first eight eigenvectors on shape \mathcal{V} in the ascending order of their corresponding eigenvalues.

a functional correspondence matrix \mathbf{C} that maps corresponding eigenvectors between surfaces, *i.e.*, $\Phi_{\mathcal{Y}}\mathbf{C} = \Phi_{\mathcal{X}}$, where Φ is a matrix of column eigenvectors. N.B. $\Phi_{\mathcal{Y}}$ and $\Phi_{\mathcal{X}}$ are truncated to k eigenvalues. In the example in Fig. 4, this would induce a map that would approximately be an identity matrix \mathbf{I} , where $\phi_i^{\mathcal{Y}}$ maps to $\phi_i^{\mathcal{X}}$.

Since the Laplacian operator is only invariant under isometric deformation, the strong diagonal of \mathbf{C} gradually weakens to greater non-isometries. This effect is most significant in eigenvectors with larger eigenvalues, while smaller eigenvalues are progressively less affected. This leads to a pre-dominantly diagonal structure, with the strongest diagonal found between smaller eigenvalues. $\phi_{3,7,8}^{\mathcal{Y}}$ in Fig. 5 are negative equivalents of $\phi_{3,7,8}^{\mathcal{X}}$ and $\phi_{3,7,8}^{\mathcal{X}}$ in Fig. 4, and $\phi_9^{\mathcal{Y}}$ presents a completely different eigenvector that is not present in any $\phi_i^{\mathcal{Y}}$ and $\phi_i^{\mathcal{X}}$ ($i \leq 9$).

An optimal \mathbf{C} can generally be computed by solving the following linear system:

$$\arg \min_{\mathbf{C}} \|\Phi_{\mathcal{Y}} - \mathbf{C}\Phi_{\mathcal{X}}\|^2 + E_{\text{corr}}. \quad (20)$$

Where E_{corr} imposes some regularization on \mathbf{C} , *e.g.*, promoting Laplacian commutativity by penalizing matches of dissimilar eigenvalues $E_{\text{corr}} = \alpha \|\Lambda_{\mathcal{X}}\mathbf{C} - \mathbf{C}\Lambda_{\mathcal{Y}}\|^2$.

Considering the example of \mathcal{X} and \mathcal{V} , once \mathbf{C} is obtained (such as in Fig. 6), we see that negated corresponding eigenvectors cause negative entries along the diagonal, while eigenvectors that do not correspond have values close to zero.

Much like the permutation matrix \mathbf{P} , for computing bijective correspondence, a key property of \mathbf{C} is that it should be an orthogonal matrix. Since \mathbf{C} is a soft map, it is a member of the larger Stiefel manifold set $\mathcal{S}_{k_1 \times k_2} = \{\mathbf{C} \in \mathbb{R}^{k_1 \times k_2} \mid \mathbf{C}^T \mathbf{C} = \mathbf{I}\}$, where $\mathcal{P} \subset \mathcal{S}$. This can either be incorporated explicitly by limiting the solution space to \mathcal{S} or added as a regularization term [RCB*17] (*e.g.*, $\text{off}(\mathbf{C}^T \mathbf{C}) = \sum_{i \neq j}^k (\mathbf{C}^T \mathbf{C})_{ij}^2$ [CS96]).

For shape correspondence, it is necessary to recover a point-to-point mapping. [OBCS*12] propose the following energy term:

$$\arg \min_{\mathbf{P} \in \mathcal{P}} \|\Phi_{\mathcal{Y}} - \mathbf{P}(\mathbf{C}\Phi_{\mathcal{X}})\|^2, \quad (21)$$

where \mathbf{C} is fixed. This may be considered to be a form of point align-




Figure 6: A functional map \mathbf{C} between \mathcal{X} in Fig. 4 & \mathcal{V} in Fig. 5.

ment in a high-dimensional space where other alignment schemes can be employed [RMC15, RMC17].

Partial matching For partial matching, the requirements of (17) must be relaxed. [RCB*17] propose to optimize a correspondence map by adopting suitable correspondence and part regularization terms, *i.e.*,

$$E = E_{\text{data}} + \underbrace{E_{\text{corr}} + E_{\text{part}}}_{\text{regularization}},$$

where the E_{corr} and E_{part} are solved in alternating steps to optimize the correspondence and the part separately. In the first step, a correspondence is estimated between a part of the full shape $\mathcal{X}' \subseteq \mathcal{X}$ and a partial shape \mathcal{Y} . Given a deformation that is near-isometric, E_{corr} helps constrain the optimizer to guide it towards a sensible solution. The part regularizer seeks to match the area between \mathcal{Y} and \mathcal{X}' , while also minimizing the boundary length ℓ of \mathcal{X}' ,

$$\arg \min_{\mathcal{X}'} E_{\text{data}} + E_{\text{part}},$$

where

$$E_{\text{part}} = \alpha |\text{area}(\mathcal{Y}) - \text{area}(\mathcal{X}')|^2 + \beta \ell(\partial \mathcal{X}').$$

The E_{data} term is modified to optimize the functional map for the subset of eigenvectors $\Phi_{\mathcal{X}'} \subseteq \Phi_{\mathcal{X}}$ for \mathcal{X}' ,

$$\arg \min_{\mathbf{C}} E_{\text{data}} + \alpha E_{\text{corr}},$$

where

$$E_{\text{data}} = \|\Phi_{\mathcal{Y}} - \mathbf{C}\Phi_{\mathcal{X}'}\|^2.$$

In practice, this is implemented using a fuzzy approach. This is done by giving eigenvectors that are not in \mathcal{X}' a low weight, rather than completely removing them.

A fundamental observation by [RCB*17] is that for a given eigenvector $\phi_i^{\mathcal{Y}}$ on \mathcal{Y} , there exists a corresponding $\phi_j^{\mathcal{X}}$ where $i \leq j$. This is illustrated in Fig. 7, where the following pairs of eigenvectors (partially) match: $(\phi_2^{\mathcal{Y}}, \phi_4^{\mathcal{X}})$; $(\phi_3^{\mathcal{Y}}, \phi_5^{\mathcal{X}})$; and $(\phi_5^{\mathcal{Y}}, \phi_9^{\mathcal{X}})$.

The proposed formulation by [RCB*17] is not suitable for part-to-part matching as it relies on an estimate of overlap that assumes the mapping to be injective. [CRM*16] extend [RCB*17] for part-to-part matching by attempting to optimize for both the part \mathcal{Y}' and \mathcal{X}' . To do this, three optimization steps are used: first \mathbf{C} is computed, then \mathcal{X}' is estimated, and finally \mathcal{Y}' is estimated, before repeating all three steps until a convergence criterion is met. Similarly, [LRB*16b] adapt [RCB*17] for matching multiple parts (\mathcal{Y}'_i and \mathcal{X}'_j) between a full shape ($\mathcal{X}'_i \subseteq \mathcal{X}$) and multiple partial shapes ($\mathcal{Y}'_j \subseteq \mathcal{Y}_i$). Unfortunately, the proposed formulations do not scale well, and require optimization in both the spectral domain and the spatial domain.

[LRBB17] propose an approach that is conducted entirely in the spectral domain for computing part-based correspondence. The method builds upon the full-to-full matching method by [KBB*13]. The authors align the eigenbases $\Phi_{\mathcal{X}}$ and $\Phi_{\mathcal{Y}}$ jointly by formulating an optimization in the form of the joint approximate diagonalization problem,

$$\arg \min_{\mathbf{A}, \mathbf{B} \in \mathcal{S}_{n \times n}} \underbrace{\text{off}(\mathbf{A}^\top \mathbf{A}_{\mathcal{X}} \mathbf{A}) + \text{off}(\mathbf{B}^\top \mathbf{A}_{\mathcal{Y}} \mathbf{B})}_{\text{diagonalization}} + \alpha \underbrace{\|\mathbf{P}\Phi_{\mathcal{X}} \mathbf{A} - \mathbf{Q}\Phi_{\mathcal{Y}} \mathbf{B}\|_F^2}_{\text{coupling}}.$$

$\mathbf{P}, \mathbf{Q} \in \{0, 1\}^{l \times n}$ are matrices that map between l initial corresponding vertices and rows from the eigenbases. In practice, a discrete point p is a row in Φ , i.e., $\phi_i(p) \in \mathbb{R}$ for $i = 1, \dots, k$. Effectively, $\mathbf{P}\Phi_{\mathcal{X}}$ and $\mathbf{Q}\Phi_{\mathcal{Y}}$, represent a subset of corresponding rows from each eigenbasis. The diagonalization terms promote \mathbf{A} and \mathbf{B} to be orthogonal matrices. These are used to perform an orthogonal transform, which rotates and flips $\mathbf{P}\Phi_{\mathcal{X}}$ and $\mathbf{Q}\Phi_{\mathcal{Y}}$. Importantly, the inner products are preserved. The coupling term then measures the similarity of corresponding rows between $\mathbf{P}\Phi_{\mathcal{X}}$ and $\mathbf{Q}\Phi_{\mathcal{Y}}$. The aligned eigenbases can now be computed as:

$$\hat{\Phi}_{\mathcal{X}} = \Phi_{\mathcal{X}} \mathbf{A} \quad \text{and} \quad \hat{\Phi}_{\mathcal{Y}} = \Phi_{\mathcal{Y}} \mathbf{B}.$$

This formulation is not suitable for partial matching problems, as all rows in $\mathbf{P}\Phi_{\mathcal{X}}$ and $\mathbf{Q}\Phi_{\mathcal{Y}}$ must be aligned.

For partial correspondence [LRBB17] modify the coupling term to the following:

$$\arg \min_{\mathbf{A}, \mathbf{B} \in \mathcal{S}_{k \times r}} \text{off}(\mathbf{A}^\top \mathbf{A}_{\mathcal{X}} \mathbf{A}) + \text{off}(\mathbf{B}^\top \mathbf{A}_{\mathcal{Y}} \mathbf{B}) + \alpha \|\mathbf{W}_r (\Phi_{\mathcal{X}} \mathbf{A} - \Phi_{\mathcal{Y}} \mathbf{B})\|_{2,1}. \quad (22)$$

The solution space is constrained to a set of rectangular orthogonal matrices $\mathcal{S}_{k \times r}$. The rank $r = \text{rank}(\mathbf{C})$ is based on an estimate of the overlapping area between \mathcal{X} and \mathcal{Y} , where $r \ll k$. A mask

$\mathbf{W}_r = (\mathbf{I}_{r \times r} \quad \mathbf{0}_{r \times k-r})^\top$ is applied. The $\ell_{2,1}$ -norm ($\|\cdot\|_{2,1}$) promotes sparsity in the columns of the masked matrix. The proposed method demonstrates a strictly spectral approach to the partial functional map problem. How to effectively determine a suitable estimate of r for part-to-part matching is unclear. The authors suggest that an optimal r can be estimated by solving (22) for multiple r 's; however, this is an expensive procedure.

[MRR*19] propose a fully-spectral method that uses a coarse-to-fine approach in which the size of \mathbf{C} is increased incrementally. \mathbf{C} initially only maps low-frequency eigenvectors, as the size of \mathbf{C} increases the map is refined by high-frequency eigenvectors. The refinement process comprises of two steps in which the method transforms correspondence data between a functional map \mathbf{C} and a pointwise map \mathbf{P} , i.e., (21). Initially, $\mathbf{C} = \mathbf{I}_{2 \times 2}$ (or larger), and a point-to-point map \mathbf{P} is recovered, using $\Phi_{n \times k_1}^{\mathcal{Y}}$ and $\Phi_{n \times k_2}^{\mathcal{X}}$ pre-computed $n \times k$ eigenbases that are truncated to the first k_1 and k_2 columns respectively (i.e., $\Phi_{n \times k_j}^{\mathcal{Y}} = \{\phi_i^{\mathcal{Y}} \mid i = 1, \dots, k_j\}$, $j = 1, 2$). For full-to-full matching, the value of k_1 and k_2 are then increased by one (or more) before recomputing the functional map,

$$\mathbf{C}_{k_1 \times k_2} = \left(\Phi_{n \times k_1}^{\mathcal{Y}} \right)^+ \mathbf{P} \Phi_{n \times k_2}^{\mathcal{X}}, \quad (23)$$

where $(\cdot)^+$ denotes the Moore-Penrose inverse. Solving for \mathbf{P} and (23) are alternated between until k is reached by either k_1 or k_2 .

To handle partial matching, [MRR*19] propose to compute a rectangular functional map $\mathbf{C}_{k_1 \times k_2}$. An estimate of rank r , based on [RCB*17], is used to determine the shape of \mathbf{C} for each iteration,

$$k_1 = k_1 + 1, \quad k_2 = k_2 + \left\lceil \frac{k_2}{k} (k - r) \right\rceil.$$

This simple approach achieves excellent results for full-to-full matching. However, the method does not fare so well with partial matching. This is because the method still relies on stable eigenvectors, especially in the initial iterations where only low-frequencies are mapped. As illustrated in Fig. 7, even the order of low-frequency eigenvectors that have been independently decomposed can prove to be highly inconsistent for partial matching.

[RTO*19] demonstrate that replacing the typical Laplacian operator with the Hamiltonian operator facilitates the computation of consistent eigenvectors between partial shapes. For the Hamiltonian operator, only the left-hand side of (19) is modified to solve the following generalized eigenvalue problem for the cotangent Laplacian,

$$(\mathbf{M} + \mathbf{D} \text{diag}(\mathbf{v})) \Psi = \mathbf{D} \Psi \Lambda.$$

$\mathbf{v}: \mathcal{X} \mapsto \{0, \tau\}^n$ is a potential function that indicates whether a point $p \in \mathcal{X}$ is inside ($\mathbf{v}(p) = 0$) or outside ($\mathbf{v}(p) = \tau$) of the overlapping region $\mathcal{X}' = \mathcal{X} \cap \mathcal{Y}$ in \mathcal{X} . The effect of adding $\mathbf{D} \text{diag}(\mathbf{v})$ is that rows in Ψ that correspond to a point assigned a value of τ (i.e., $\mathbf{v}(p) = \tau$) become $\Psi(p) \approx \mathbf{0}^\top$. [RTO*19] set $\tau = 10\Lambda_k$ experimentally, Λ_k is the largest eigenvalue of \mathbf{A} .

The advantage of this operator is that given a potential function \mathbf{v} that characterizes $\mathcal{X}' \subseteq \mathcal{X}$, the eigenvectors on \mathcal{X} are localized to \mathcal{X}' . This corresponds with the Laplacian operator—up to the sign ($\phi_i^{\mathcal{X}'} = \pm \psi_i^{\mathcal{X}}$)—when computed on a disconnected component \mathcal{X}' (see Fig. 8).




Figure 7: The first eight eigenvectors on \mathcal{X} and \mathcal{Y} ordered by their corresponding eigenvalues.

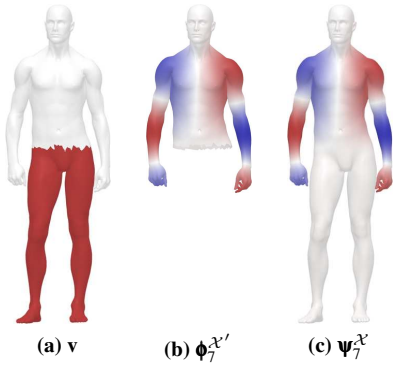


Figure 8: Given two shapes where the overlapping region is known, the seventh eigenvector ($\phi_7^{\mathcal{X}'}$ and $\psi_7^{\mathcal{X}'}$) computed using (b) the Laplacian operator and (c) the Hamiltonian operator $\tau \approx 0.2313$.

To determine \mathcal{X}' , simply \mathbf{v} must be found. An optimization function is formulated to align the Hamiltonian eigenvalues $\lambda_{\mathcal{X}}$ with the Laplacian eigenvalues $\Lambda_{\mathcal{Y}}$,

$$\min_{\mathbf{v} \in \mathbb{R}^n} \left\| \lambda_{\mathcal{X}} (\mathbf{D}_{\mathcal{X}}^{-1} \mathbf{M}_{\mathcal{X}} + \text{diag}(\mathbf{v})) - \Lambda_{\mathcal{Y}} \right\|_w^2. \quad (24)$$

To prevent large eigenvalues disproportionately influencing the result, a weighted ℓ_2 -norm is used,

$$\|\lambda - \Lambda\|_w^2 = \sum_{i=1}^k \frac{1}{\Lambda_i^2} (\lambda_i - \Lambda_i)^2.$$

Once (24) is solved, the eigenvectors should be sufficiently aligned to compute \mathbf{C} , as per (20). The described approach is suitable for isometric partial-to-full matching, additionally the method is capable of handling multiple disconnected non-rigid components. Unfortunately, the non-convex nature of (24) makes this implementation

moderately slow. The method also fails when the number of eigenvectors k is too low.

4.1.3. Spectral/Functional shape recovery methods

Given a mapping, by examining the intrinsic differences between two shapes it is possible to reconstruct a deformed source surface in the pose of the target shape. This section focuses on the problem of reconstructing an extrinsic representation from spectral representations and its application to non-rigid registration.

Shape difference operator reconstruction Given a functional mapping between a source shape and a target shape, [ROA*13] search for a mesh in a shape collection that has the minimal conformal- and equiareal-based shape difference. The method is limited by the collection available, and the complexity of the optimization increases when searching in large shape collections. [BEKB15] consider the problem of finding an embedding of the computed metric that minimizes an energy based on the shape difference operator. The method alternates between minimizing an energy term (based on [ROA*13]) to the Riemannian metric, and generating the extrinsic shape based on the discrete metric. The intrinsic metric does not suitably capture extrinsic curvature, causing ambiguity with possibly multiple isometric solutions. [CSBC*17] address this problem by incorporating extrinsic information into an intrinsic operator. The edge lengths of an offset surface from a dual mesh are used to guide the reconstruction to a better solution. The authors demonstrate how the accuracy of the reconstruction deteriorates as the eigenbasis is truncated. [HRA*19] propose a deep learning framework to recover a shape using shape difference operators that also possess both extrinsic and intrinsic information. A set of training shapes is used to help constrain the solution space. The advantage of learning from shape operators is that the matrices may be treated as a grid structure—a convenient and well-studied domain for deep learning architectures. The method demonstrates superior results compared to [BEKB15, CSBC*17].

Template-less reconstruction Unlike the previously discussed shape recovery works, which rely on shape difference operators, [CPR*19] do not require a mapping to recover a shape as the method deforms a source mesh such that the Laplacian spectra better align. Similarly, [MRC*20, MRC*21] propose a learning-based approach to recover a shape from a given mesh Laplacian. By training a neural network to constrain the solution space to seen examples, the approach efficiently aligns the given spectra. To facilitate partial correspondence within the same class of shapes, [MMC*21] develop a shape union operator, which they incorporate into a reconstruction pipeline [MRC*20]. With this method, reconstructed shapes are in a canonical pose.

Spectral operators The problem of reconstructing a shape based on its spectral representation is not a new problem [KG00]; however, its use for registration has been limited. Given two shapes with an unknown correspondence, [ELC20] align intrinsic and extrinsic shape representations to recover a deformed mesh in a coarse-to-fine manner. A novel operator is proposed that helps constrain the smoothness of the recovered mesh when increasing the number of dimensions of the embedding. In order to align the source and target shapes, the spectral representation is combined with Cartesian coordinates and surface normals. The correspondence and alignment are optimized in alternating steps.

4.2. Learning-Based Methods

Over the past decade, an increasing amount of effort has been placed into developing learning-based techniques that improve or enhance existing intrinsic pipelines. Many techniques seek to develop robust shape descriptors. As input to artificial neural networks, mesh geometry has been adapted to a variety of domains (n -dimensional descriptors [LRR*17], images [WGY*18], voxels [ZSN*17], spectra [LB14], geodesics [MBBV15], and others [CYHH20, BBCV21]). Intrinsic information is also incorporated into the loss function (or equivalent energy) of these techniques (e.g., labels [RBW*14, MBM*17], geodesic distance [LB14], functional map quality [LRR*17, RSO19]).

Learning-based shape descriptors A recent survey of related literature on data-driven descriptors was undertaken recently [RBRY19], so the discussion here is focused on intrinsic methods specifically. A key series of works investigated learning-based methods to develop effective spectral descriptors [LB14, WVR*14, BMM*15, BMR*16]. [LB14] learn low-dimensional class-specific descriptors from the eigenvectors of the Laplacian. The authors observe that the *heat kernel signature* (HKS) [SOG09] and *wave kernel signature* (WKS) [ASC11] may be considered as low-pass and band-pass filters in the spectral domain. A descriptor model that generalizes HKS and WKS is formulated. The training phase attempts to compute coefficients of a cubic spline that optimize the model to produce similar descriptors between corresponding points using a linear discriminant analysis-based approach. For the task of partial matching, the optimized descriptor is found to out-perform HKS & WKS. [WVR*14] propose a similar approach to metric learning to construct an infinite-dimensional descriptor. [BMM*15] exploit a localized spectral CNN. [BMR*16] propose the use of

multiple anisotropic kernels to form a directionally-sensitive descriptor. [BMRB16] incorporate the anisotropic technique as a weighting scheme into a CNN framework. [MBBV15] propose a CNN that uses a geodesic-based local operator based on [KBLB12]. [CRM*16] learn an embedding function to construct descriptors that are robust to cluttered scenes. [MBM*17] present a CNN framework that utilizes an intrinsic local operator that is shown to generalize operators from previous works [MBBV15, BMRB16] as mixtures of Gaussian kernels.

A deficiency of previous chart-based methods is that they require a canonical orientation—well-structured data like digital images naturally have this property, but this is ill-defined on meshes. Previous methods use the maximal response over a set of rotations [MBBV15] and the principal curvature directions [BMRB16], which can become inconsistent. [PO18, WEH20] propose to address this problem by establishing a directional function, based on parallel transport, over the surface. [SRC*20] replace discrete bins used in [PO18] with a 2D rotation through the use of Zernike radial polynomials.

Learning-based functional maps A collection of works have investigated task-driven approaches to learning enhanced descriptors for functional maps [COC15, LRR*17, HLR*19, RSO19, DSO20]. The earliest method [COC15] proposes a shallow learning framework in which spectral descriptors are re-weighted. These re-weighted descriptors are then used to initialize a functional map [OBCS*12], where the difference to a known ground-truth map for training is minimized. [LRR*17] consider a similar approach in a deep learning framework. During the training phase, weights are learned to enhance feature descriptors. The loss function then measures the quality of the subsequent spatial map with geodesic distances to a ground-truth map. The method requires training data with pointwise correspondence, which are often expensive and difficult to acquire. Follow-up works have sought to facilitate unsupervised learning through penalizing geodesic distortion [HLR*19], and functional map heuristics (e.g., orthogonality, bijectivity, and commutativity) [RSO19]. [DSO20] propose a similar framework in which the descriptors are learned from point clouds. While some of these methods do demonstrate promising evaluative results, they rely on simple functional mapping techniques that are known to not be suitable for partial correspondence leading to sub-optimal performance in this scenario. Recently, to address the problem of partial matching, [APO21] propose to incorporate a cross-attention block between sets of learned features for two input shapes. This block enables the communication of features between shapes and causes non-overlapping features to be down-weighted, making it possible to efficiently predict the overlapping region.

Learning-based reconstruction [ETLTC20] incorporate the scheme from [ELC20] to produce a spectral-based learning framework, which—with sufficiently high-dimensional eigenvectors—can approximately reconstruct a target surface. [CNH*20] incorporate geodesic regularization into an auto-encoder framework, leading to low-distortion reconstructions. [ENK*21] adopts a deep neural network architecture to match intrinsic feature descriptors and interpolate extrinsic vertex locations simultaneously unsupervised. The method is not fully intrinsic as it also relies on an as-rigid-as-possible error term. Given a shape's spatial coordinates as input, [BBP*19, GCBZ19] implement auto-encoder architectures that

Table 3: A summary of optimization-based intrinsic methods.

	Sampling strategy	Data term		Regularization constraints					Output			
		Point-to-point	Spectral	Smooth	Point constraints	Near-isometry	Non-isometry	Bijective or part-to-full injective	Partial	Precision		Transform
										Point-to-point	Fuzzy	
[ASP*04]	✓	✓				✓				✓		✓
[BBK06]	✓	✓				✓		✓		✓		✓
[HAWG08]	✓	✓				✓				✓		✓
[DRB*09]		✓				✓				✓		✓
[LF09]		✓		✓	✓					✓		
[OMMG10]	✓	✓				✓				✓		
[KLF11]		✓		✓			✓			✓		
[SY11]	✓	✓				✓				✓		
[TBW*11]	✓	✓				✓				✓		
[OBCS*12]		✓	✓	✓	✓			✓		✓		
[RBA*12]		✓				✓		✓		✓		
[SY12a]	✓	✓				✓				✓		
[SY12b]	✓	✓				✓				✓		
[SNB*12]	✓	✓				✓		✓		✓	✓	
[KBB*13]		✓	✓	✓		✓		✓		✓		
[ROA*13]			✓	✓		✓				✓		✓
[vKZH13]	✓	✓				✓				✓		
[TMRL14a]		✓				✓				✓		
[TMRL14b]		✓				✓				✓		
[RMC15]			✓ ^a	✓			✓			✓		
[SdGP*15]							✓	✓			✓	
[AL16]		✓		✓	✓					✓		
[LRB*16b]		✓	✓	✓		✓		✓		✓		
[SPKS16]				✓				✓			✓	
[LRBB17]		✓	✓	✓		✓		✓		✓		
[MCSK*17]	✓						✓	✓		✓		
[RCB*17]		✓	✓	✓	✓			✓		✓		
[RMC17]			✓ ^a	✓			✓			✓		
[VLB*17]	✓	✓	✓				✓			✓		
[BCK18]	✓	✓		✓ ^b						✓		
[Sah18]	✓	✓				✓				✓		
[CPR*19]	✓		✓ ^a	✓			✓			✓		✓
[DLRT19]		✓					✓			✓		✓
[LJO19]	✓	✓					✓			✓		
[MRR*19]			✓ ^a	✓			✓			✓		
[RTO*19]		✓ ^d		✓			✓			✓ ^c		✓
[ZQW*19]		✓ ^d		✓ ^b			✓			✓		
[ELC20]	✓	✓		✓						✓		✓
[PRM*21]		✓		✓			✓			✓		

^aRefines an initial functional map (e.g., [NO17]); ^bNear-conformal; ^cPart-to-part correspondence; ^dBoundary-to-boundary.

Table 4: A summary of learning-based intrinsic methods.

	Objective		Training type			Loss				
	Learnable correspondences	Learnable deformation	Supervised	Semi-supervised	Unsupervised	Ground-truth correspondences	Ground-truth spectra	Alignment	Functional map heuristics	Non-isometry
[LB14]	✓		✓			✓				
[RBW*14]	✓		✓			✓				
[MBBV15, BMRB16, MBM*17]	✓		✓			✓				✓
[LRR*17]	✓		✓			✓				
[WGY*18]	✓		✓			✓				
[HLR*19]	✓				✓				✓	
[RSO19]	✓				✓				✓	
[DSO20]	✓		✓							
[MRC*20, MRC*21]		✓		✓			✓	✓		
[WRY*20]	✓		✓			✓				
[MMC*21]		✓		✓			✓	✓		

at their core comprise of intrinsic convolutional operators. The operator is composed of neighboring vertices that are traversed in a spiral pattern, which is inherently sensitive to variations in connectivity.

4.3. Summary

The key optimization-based and learning-based works discussed in this section are summarized in Tables 3 and 4.

5. Datasets and Benchmarks

There are many publicly available datasets that can be used for the effective evaluation of non-rigid registration methods. Such datasets may be used to measure the quality of a surface alignment quantitatively (e.g., fitting error [CRS98, RL01], geodesic error [KLF11], coverage error [DLR*20], area ratio [ZWW*10, ZWG*16]) and qualitatively (e.g., model rendering, distortion [OBCCG13], information transport [MMM*19]). Alternatively, applications that may rely on registration (e.g., shape retrieval [FKMS05, ELBS10]) can be used to indirectly measure the performance of registration methods.

Table 5: Datasets for benchmarking correspondence and registration methods.

	Deformation type			Partial correspondence	Scan type			Dense	Ground-truths			Anthropometric	No. of scans	> 25K vertices	
	Near-isometry	Non-isometry	Topological change		Real	Synthetic	Deformed template		Acquisition						Training facility
									Manual	Automatic	Intrinsic				
[RDP99]		✓	✓	✓	✓			✓				4,431 ^a	✓		
[SP04]	✓					✓					partially	551			
[ASK*05]	✓					✓				✓	✓	71	✓		
[Ado]	✓					✓		✓			✓	121 ^b	✓		
[BBK08]	✓					✓					partially	80	✓		
[VBMP08]	✓					✓					partially	1,500			
[HSS*09]		✓				✓		✓			✓	550	✓		
[BBC*10]	✓ ^c			✓		✓					✓	56	✓		
[KLF11]		✓				✓			✓		partially	551			
[BRLB14]		✓	✓	✓		✓				✓	✓	300	✓		
[CWZ*14]		✓				✓					✓	3,000	✓		
[RBW*14]	✓					✓					✓	32	✓		
[YYZ*14]		✓				✓					✓	3,000			
[LRB*16a]	✓		✓	✓		✓					✓	26	✓		
[BRPMB17]		✓	✓	✓		✓					✓	40,000	✓		
[PWH*17]		✓				✓					✓	4,309			
[RCB*17]	✓			✓		✓					partially	1,216	✓		
[RCL*17]	✓			✓		✓					partially	4,258	✓		
[CKPZ18]		✓				✓					✓	1,800,000	✓		
[SAS*18]		✓				✓					✓	400	✓		
[DSL*19]		✓	✓			✓		✓			✓	50	✓		
[MMR*19]		✓	✓			✓					✓	44	✓		
[DLR*20]		✓	✓	✓		✓			✓		✓	14	✓		
[DZL*20]		✓	✓			✓			✓		✓	12	✓		
[CPMK21]		✓		✓		✓					✓	10,200			

^aDataset contains some shapes that are isometric up to scaling; ^bOnly a subset of ground-truths are publicly available; ^c110 of the scans do not have landmarks; ^dBase models.

The reader is directed to [vKZHCO11] for further discussion of other relevant evaluation methods.

State-of-the-art datasets for benchmarking often include ground-truth correspondences, which can be used to measure the root-mean-square error and geodesic error [KLF11] between a predicted ground-truth location and the known ground-truth location on a target surface. This is an effective and commonly used technique to summarize and compare the accuracy of registration methods (*e.g.*, RMSE [SGT*19, YDXZ20, LZXH20], geodesics [CK15, CR16, MDK*16, GFK*18, EC20]). [RBRY19] discuss benchmarks for the purposes of evaluating 3D learning-based shape descriptors, many of which can be—and are—used for evaluating registration techniques.

A common trend has been the use of anthropomorphic shapes, such as human bodies and human faces. An extended overview of public 3D facial datasets is discussed by [ZX18]. Most contain sparse fiducial landmarks (*e.g.*, [PFS*05, SAD*08, GCMB10]) that can be used for evaluating registration methods [GSM15, SK16, FHCP18]. If RGB texture data is acquired, it is possible to establish reliable landmarks automatically [CWZ*14]. Beyond facial datasets, texture data is commonly used for computing ground-truth correspondences reliably [BRLB14, BRPMB17, DSL*19, DZL*20].

Over the past decade, many notable benchmarks have been developed for Shape Retrieval Contest (SHREC) tracks [BBC*10, LRB*16a, RCL*17, DSL*19, MMR*19, DZL*20, DLR*20]. [LRB*16a, RCL*17] introduce challenges such as topological changes and partial correspondence by applying synthetic modifications to models.

Dense ground-truth correspondences for real meshes are often automatically acquired by non-rigidly aligning a template to a given scan using methods such as non-rigid registration [ZKJB17, MMR*19], linear blend skinning [VBMP08, PMRMB15, CPMK21]

or parametric modeling [ASK*05, YYZ*14, XZC18]. The advantage of aligning a template mesh to a set of target scans is that it is possible to compute dense correspondences; however, these correspondences may contain local alignment errors. The alignment can be improved by incorporating texture information [BRLB14, CWZ*14, BRPMB17, CKPZ18, SAS*18] or a sparse set of reliable landmarks [HSS*09].

A key deficiency in existing datasets is the lack of a training facility for learning-based methods. As the popularity of machine learning continues to expand, the necessity for benchmarks that enable fair and consistent comparisons between techniques has increased. Presently, only a few benchmarks address this problem [BRLB14, BRPMB17, RCL*17, SAS*18]. For the benchmark dataset [DSL*19], which does not have a training facility, the learning-based method [GFK*18] is trained on [VRM*17]. Also, [MRC*20, MRC*21] utilize a generative face model [RBSB18] to train a learning-based method for evaluative purposes.

Table 5 provides an overview of datasets used for evaluative purposes in the literature.

6. Directions for Future Works

Input Data Type One important application for non-rigid registration is the acquisition of 3D shapes from scanned data. For existing works, the input data type is mainly 3D geometry such as point clouds or depth maps. Although such data can be captured with depth sensors, currently for the general public it is still easier to capture RGB videos with portable devices. For convenient 3D content acquisition, it is desirable to efficiently reconstruct 3D geometry from monocular RGB video on mobile devices. Recent developments in differentiable rendering [KBM*20] and neural rendering [TFT*20] have enabled 3D reconstruction from monocular

videos [YKM*20, NMOG20, YGKL21, ZYQ21, OPG21], but existing methods are mainly designed for rigid shapes. Their high computational costs also prevent them from being run in real-time on mobile devices. Efficient reconstruction of dynamic 3D shapes from monocular videos will be an interesting research avenue with a significant impact.

Geometric Shape Representation The input 3D data of current registration methods is usually represented as discrete signals. Recently, neural implicit representations [PFS*19, SMB*20, CPM20], which are based on parameterizing a continuous differentiable signal with a neural network, have become a promising alternative to conventional 3D representations such as point clouds and meshes. They are not coupled to the spatial resolution, and have been successfully applied to various tasks like novel view synthesis [MST*20] and multi-view 3D reconstruction [YKM*20]. It must be noted that neural implicit representations are different from the implicit surface representations discussed in Section 2. The neural implicit function is defined over the whole continuous space and is represented in the form of function composition rather than the explicit grid and explicit function values. Adapting existing registration methods to this new representation, as well as designing new methods that can make use of its representational power, is a promising research direction.

Extrinsic methods For registration problems that involve large deformation, optimization-based approaches often require proper initialization to achieve good results. Although they can be initialized using shape descriptors such as FPFH or SHOT, such handcrafted features are inevitably ambiguous when the local shapes are similar. Moreover, such descriptors tend to be sensitive to noise, and may lead to poor registration results on noisy models. Therefore, further research is needed for optimization-based approaches that can handle large deformation between noisy models. In recent years, learning-based approaches have shown promising results for such problems, either as a standalone registration technique or as a way to determine initial correspondence for an optimization method. However, the performance of machine learning models is highly dependent on the quality of the data. Currently, public datasets with accurate correspondence labels are still limited in terms of both their quantities and the types of shapes they cover. As a result, the generalizability of existing machine learning models is still not satisfactory. It will be highly beneficial to construct more labeled datasets with better coverage of shapes, and/or to develop machine learning approaches with better generalization performance.

The use of adaptive weights and robust norms have helped optimization-based methods align shapes that overlap partially. However, their performance is still not satisfactory when the overlap area is small, especially when the alignment involves large deformation and/or the data is noisy. A key issue is that due to the non-rigid nature of the problem, an extrinsic formulation may not effectively distinguish whether a source point has a corresponding point on the target surface. This is a challenging case that can occur in practical applications (*e.g.*, dynamic reconstruction from a small number of scans), and will require further investigation.

Existing shape regularizations in optimization-based methods assume the deformation to be either near-isometric, non-isometric

but with certain characteristics such as conformality, or a combination of them. While such priors can handle deformations with a certain degree of regularity, they are less effective for some complex physically-based deformations such as the ones shown in [DZL*20]. New strategies for handling such complex deformations, such as data-driven approaches to properly model the deformation behavior, will be an interesting research direction.

Intrinsic methods A substantial amount of work has focused on the development of correspondence techniques between humans. Consequently, we observe the strong performance in anthropometric applications.

Many registration methods are sensitive to their initialization, or depend upon invariant point-wise features. While intrinsic methods can also be susceptible to these problems, the use of pairwise measures can alleviate these in some scenarios. Unfortunately, pairwise measures introduce other challenges: for example, geodesic distances are only invariant under isometric deformation. Furthermore, the convenient assumptions of many intrinsic methods can quickly deteriorate when confronted with partial geometry, where certain bijectivity conditions can no longer be satisfied.

For real-time applications, the performance of hand-crafted intrinsic partial matching techniques is inadequate for larger models. Learning-based methods have the potential to achieve superior performance, as they may avoid expensive optimization problems.

Accurately estimating the overlap between the source and target surfaces is instrumental to the accuracy of these methods. The development of overlap estimation techniques (*e.g.*, [APO21]) is a promising avenue of research that may lead to intrinsic methods that handle both partial and non-isometric shapes.

For 3D reconstruction using devices such as hand-held scanners, where the subject is captured in small patches, the inconsistency of standard Laplacian eigenfunctions means that current functional mapping methods are not suitable.

Existing literature on intrinsic shape correspondence tends to focus on a subset of related challenges. The past ten years has seen a focus on relevant problems such as non-isometry of smooth deformations, mesh connectivity and part-to-full correspondence. Topological changes caused by fused geometry during scanning continue to be a challenge for intrinsic optimization methods, while data-driven approaches may help with such scenarios. Furthermore, geometry with extensive and numerous surface discontinuities are understudied (such as that of a tea towel that has been wound up). This type of problem challenges the key underlying assumption about the local smoothness of a mapping. (A variety of challenging deformations are considered in [SPF19, DZL*20].)

Acknowledgements This work was partially supported by the Swiss National Science Foundation (SNSF) under project number 200021-188577, the Guangdong International Science and Technology Cooperation Project (2021A0505030009), the National Natural Science Foundation of China (No. 62122071), the Youth Innovation Promotion Association CAS (No. 2018495), and “the Fundamental Research Funds for the Central Universities” (No. WK3470000021).

References

- [ACP02] ALLEN B., CURLESS B., POPOVIĆ Z.: Articulated body deformation from range scan data. *ACM Transactions on Graphics* 21, 3 (2002), 612–619. 4, 8, 12
- [ACP03] ALLEN B., CURLESS B., POPOVIĆ Z.: The space of human body shapes: reconstruction and parameterization from range scans. *ACM Transactions on Graphics* 22, 3 (2003), 587–594. 2, 5, 6, 8, 12
- [Ado] ADOBE INC.: Mixamo. <https://www.mixamo.com/>, 2008–21
- [AL16] AIGERMAN N., LIPMAN Y.: Hyperbolic orbifold tutte embeddings. *ACM Transactions on Graphics* 35, 6 (Nov. 2016), 217:1–217:14. 20
- [Ale03] ALEXA M.: Differential coordinates for local mesh morphing and deformation. *The Visual Computer* 19, 2 (2003), 105–114. 6
- [AMO] AGARWAL S., MIERLE K., OTHERS: Ceres solver. <http://ceres-solver.org>, 2010–8
- [APL14] AIGERMAN N., PORANNE R., LIPMAN Y.: Lifted bijections for low distortion surface mappings. *ACM Transactions on Graphics* 33, 4 (July 2014), 69:1–69:12. 6
- [APO21] ATTAIKI S., PAI G., OVSJANIKOV M.: DPFM: Deep partial functional maps. In *International Conference on 3D Vision* (Dec. 2021), pp. 175–185. 19, 22
- [ARV07] AMBERG B., ROMDHANI S., VETTER T.: Optimal step non-rigid ICP algorithms for surface registration. In *IEEE Conference on Computer Vision and Pattern Recognition* (2007), IEEE, pp. 1–8. 2, 5, 6, 8, 12
- [ASC11] AUBRY M., SCHLICKEWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *IEEE International Conference on Computer Vision Workshops* (2011), pp. 1626–1633. 13, 19
- [ASK*05] ANGUELOV D., SRINIVASAN P., KOLLER D., THRUN S., RODGERS J., DAVIS J.: SCAPE: Shape completion and animation of people. *ACM Transactions on Graphics* 24, 3 (2005), 408–416. 21
- [ASP*04] ANGUELOV D., SRINIVASAN P., PANG H.-C., KOLLER D., THRUN S., DAVIS J.: The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In *Proceedings of the 17th International Conference on Neural Information Processing Systems* (Cambridge, MA, USA, 2004), NIPS'04, MIT Press, pp. 33–40. 13, 14, 20
- [AZB15] ACHENBACH J., ZELL E., BOTSCH M.: Accurate face reconstruction through anisotropic fitting and eye correction. In *Vision, Modeling & Visualization* (2015), Bommes D., Ritschel T., Schultz T., (Eds.), The Eurographics Association. 4, 5, 6, 12
- [BBC*10] BRONSTEIN A. M., BRONSTEIN M. M., CASTELLANI U., DUBROVINA A., GUIBAS L. J., HORAUD R. P., KIMMEL R., KNOSOW D., LAVANTE E. V., MATEUS D., OVSJANIKOV M., SHARMA A.: SHREC'10 track: Correspondence finding. In *Eurographics Workshop on 3D Object Retrieval* (2010), The Eurographics Association. 21
- [BBCV21] BRONSTEIN M. M., BRUNA J., COHEN T., VELIČKOVIĆ P.: Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478* (May 2021). 19
- [BBK06] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences* 103, 5 (2006), 1168–1172. 13, 20
- [BBK08] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*, 1 ed. Springer, 2008. 10, 21
- [BBLR15] BOGO F., BLACK M. J., LOPER M., ROMERO J.: Detailed full-body reconstructions of moving people from monocular RGB-D sequences. In *IEEE International Conference on Computer Vision* (2015), pp. 2300–2308. 4, 10, 12
- [BBP*19] BOURITSAS G., BOKHNYAK S., PLOUMPIS S., ZAFEIRIOU S., BRONSTEIN M.: Neural 3D morphable models: Spiral convolutional networks for 3D shape representation learning and generation. In *IEEE/CVF International Conference on Computer Vision* (Nov. 2019), IEEE Computer Society, pp. 7212–7221. 19
- [BCK18] BADEN A., CRANE K., KAZHDAN M.: Möbius registration. *Computer Graphics Forum* (2018). 20
- [Bec17] BECK A.: *First-order methods in optimization*. SIAM, 2017. 8
- [BEKB15] BOSCAINI D., EYNARD D., KOUROUNIS D., BRONSTEIN M. M.: Shape-from-operator: Recovering shapes from intrinsic operators. *Computer Graphics Forum* 34, 2 (2015), 265–274. 18
- [BFS18] BEDNARIK J., FUA P., SALZMANN M.: Learning to reconstruct texture-less deformable surfaces from a single view. In *International Conference on 3D Vision* (2018), IEEE, pp. 606–615. 9
- [BJMO12] BACH F., JENATTON R., MAIRAL J., OBOZINSKI G.: Optimization with sparsity-inducing penalties. *Foundations and Trends® in Machine Learning* 4, 1 (2012), 1–106. 7
- [BK10] BRONSTEIN M. M., KOKKINOS I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2010), pp. 1704–1711. 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, International Society for Optics and Photonics, pp. 586–606. 5
- [BMM*15] BOSCAINI D., MASCI J., MELZI S., BRONSTEIN M. M., CASTELLANI U., VANDERGHEYNST P.: Learning class-specific descriptors for deformable shapes using localized spectral convolutional networks. *Computer Graphics Forum* 34, 5 (Aug. 2015), 13–23. 19
- [BMR*16] BOSCAINI D., MASCI J., RODOLÀ E., BRONSTEIN M. M., CREMERS D.: Anisotropic diffusion descriptors. *Computer Graphics Forum* 35, 2 (May 2016), 431–441. 19
- [BMRB16] BOSCAINI D., MASCI J., RODOLÀ E., BRONSTEIN M.: Learning shape correspondence with anisotropic convolutional neural networks. In *Advances in Neural Information Processing Systems* (2016), vol. 29, Curran Associates, Inc. 19, 20
- [Boo89] BOOKSTEIN F.: Principal warps: thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 11, 6 (1989), 567–585. 3
- [BPC*11] BOYD S., PARIKH N., CHU E., PELEATO B., ECKSTEIN J.: Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning* 3, 1 (2011), 1–122. 8
- [BPC16] BONNEEL N., PEYRÉ G., CUTURI M.: Wasserstein barycentric coordinates: Histogram regression using optimal transport. *ACM Transactions on Graphics* 35, 4 (jul 2016), 71:1–71:10. 14
- [BPZ*20] BOŽIČ A., PALAFOX P., ZOLLHÖFER M., DAI A., THIES J., NIESSNER M.: Neural non-rigid tracking. In *Advances in Neural Information Processing Systems* (2020), Larochelle H., Ranzato M., Hadsell R., Balcan M. F., Lin H., (Eds.), vol. 33, Curran Associates, Inc., pp. 18727–18737. 9, 11, 12, 13
- [BPZ*21] BOZIC A., PALAFOX P. R., ZOLLHÖFER M., THIES J., DAI A., NIESSNER M.: Neural deformation graphs for globally-consistent non-rigid reconstruction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 1450–1459. 10, 11, 13
- [BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: FAUST: Dataset and evaluation for 3D mesh registration. In *IEEE Conference on Computer Vision and Pattern Recognition* (June 2014), IEEE. 21
- [BRPMB17] BOGO F., ROMERO J., PONS-MOLL G., BLACK M. J.: Dynamic FAUST: Registering human bodies in motion. In *IEEE Conference on Computer Vision and Pattern Recognition* (July 2017). 4, 9, 10, 12, 21
- [BSB14] BONARRIGO F., SIGNORONI A., BOTSCH M.: Deformable registration using patch-wise shape matching. *Graphical Models* 76, 5 (2014), 554–565. 2, 8, 12

- [BTLP16] BOUAZIZ S., TAGLIASACCHI A., LI H., PAULY M.: Modern techniques and applications for real-time non-rigid registration. In *SIG-GRAPH ASIA 2016 Courses* (2016), SA '16, Association for Computing Machinery, pp. 11:1–11:25. 4
- [BV99] BLANZ V., VETTER T.: A morphable model for the synthesis of 3D faces. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques* (1999), pp. 187–194. 4, 11
- [BWP13] BOUAZIZ S., WANG Y., PAULY M.: Online modeling for realtime facial animation. *ACM Transactions on Graphics* 32, 4 (2013), 40:1–40:10. 6, 11, 12
- [BZTN20] BOZIC A., ZOLLHÖFER M., THEOBALT C., NIESSNER M.: DeepDeform: Learning non-rigid RGB-D reconstruction with semi-supervised data. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 7002–7012. 9, 11, 12, 13
- [CB110] CAGNIART C., BOYER E., ILIC S.: Free-form mesh tracking: a patch-based approach. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2010), IEEE, pp. 1339–1346. 3, 12
- [CCS12] CORSINI M., CIGNONI P., SCOPIGNO R.: Efficient and flexible sampling with blue noise properties of triangular meshes. *IEEE Transactions on Visualization and Computer Graphics* 18 (2012), 914–924. 14
- [CK15] CHEN Q., KOLTUN V.: Robust nonrigid registration by convex optimization. In *IEEE International Conference on Computer Vision* (Dec. 2015), pp. 2039–2047. 21
- [CKPZ18] CHENG S., KOTSIA I., PANTIC M., ZAFEIRIOU S.: 4DFAB: A large scale 4D database for facial expression analysis and biometric applications. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2018), pp. 5117–5126. 21
- [CLS20] CHEBROLU N., LÄBE T., STACHNISS C.: Spatio-temporal non-rigid registration of 3D point clouds of plants. In *IEEE International Conference on Robotics and Automation* (2020), IEEE, pp. 3112–3118. 8
- [CM92] CHEN Y., MEDIONI G.: Object modelling by registration of multiple range images. *Image and vision computing* 10, 3 (1992), 145–155. 5
- [CMPM20] CHIBANE J., MIR A., PONS-MOLL G.: Neural unsigned distance fields for implicit function learning. In *Advances in Neural Information Processing Systems* (2020), Larochelle H., Ranzato M., Hadsell R., Balcan M. F., Lin H., (Eds.), vol. 33, Curran Associates, Inc., pp. 21638–21652. 22
- [CNH*20] COSMO L., NORELLI A., HALIMI O., KIMMEL R., RODOLÀ E.: LIMP: Learning latent shape representations with metric preservation priors. In *Proceedings of the European Conference on Computer Vision* (2020), Springer International Publishing, pp. 19–35. 19
- [COC15] CORMAN É., OVSJANIKOV M., CHAMBOLLE A.: Supervised descriptor learning for non-rigid shape matching. In *Proceedings of the European Conference on Computer Vision* (Mar. 2015), vol. 8928, Springer International Publishing, pp. 283–298. 19
- [CPMK21] CHEN H., PARK H., MACIT K., KAVAN L.: Capturing detailed deformations of moving human bodies. *ACM Transactions on Graphics* 40, 4 (July 2021), 85:1–85:18. 21
- [CPR*19] COSMO L., PANINE M., RAMPINI A., OVSJANIKOV M., BRONSTEIN M. M., RODOLÀ E.: Isospectralization, or how to hear shape, style, and correspondence. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 7521–7530. 11, 19, 20
- [CR03] CHUI H., RANGARAJAN A.: A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding* 89, 2-3 (2003), 114–141. 3, 12
- [CR16] CORRIG J., RANGARAJAN A.: Resonant deformable matching: Simultaneous registration and reconstruction. In *Proceedings of the European Conference on Computer Vision* (Sept. 2016), Springer International Publishing, pp. 51–68. 21
- [CRB*16] COSMO L., RODOLÀ E., BRONSTEIN M. M., TORSELLO A., CREMERS D., SAHILLIOĞLU Y.: Partial matching of deformable shapes. In *Eurographics Workshop on 3D Object Retrieval* (2016), The Eurographics Association. 14
- [CRM*16] COSMO L., RODOLÀ E., MASCI J., TORSELLO A., BRONSTEIN M. M.: Matching deformable objects in clutter. In *International Conference on 3D Vision* (2016), pp. 1–10. 17, 19
- [CRS98] CIGNONI P., ROCCHINI C., SCOPIGNO R.: METRO: Measuring error on simplified surfaces. *Computer Graphics Forum* 17, 2 (1998), 167–174. 20
- [CS96] CARDOSO J.-F., SOULOUMIAC A.: Jacobi angles for simultaneous diagonalization. *SIAM Journal on Matrix Analysis and Applications* 17 (1996), 161–164. 16
- [CSBC*17] CORMAN É., SOLOMON J., BEN-CHEN M., GUIBAS L. J., OVSJANIKOV M.: Functional characterization of intrinsic and extrinsic geometry. *ACM Transactions on Graphics* 36, 2 (Mar. 2017), 14:1–14:17. 18
- [CTL15] CAO V.-T., TRAN T.-T., LAURENDEAU D.: A two-stage approach to align two surfaces of deformable objects. *Graphical Models* 82 (2015), 13–28. 2, 5, 6, 12
- [Cut13] CUTURI M.: Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems* (2013), vol. 26, Curran Associates, Inc. 14
- [CWZ*14] CAO C., WENG Y., ZHOU S., TONG Y., ZHOU K.: Face-Warehouse: A 3D facial expression database for visual computing. *IEEE Transactions on Visualization and Computer Graphics* 20, 3 (2014), 413–425. 4, 10, 12, 21
- [CYHH20] CAO W., YAN Z., HE Z., HE Z.: A comprehensive survey on geometric deep learning. *IEEE Access* 8 (2020), 35929–35949. 19
- [CZ08] CHANG W., ZWICKER M.: Automatic registration for articulated shapes. *Computer Graphics Forum* 27 (2008), 1459–1468. 7
- [CZ11] CHANG W., ZWICKER M.: Global registration of dynamic range scans for articulated model reconstruction. *ACM Transactions on Graphics* 30, 3 (2011), 26:1–26:15. 5, 6, 8, 12
- [DDF*17] DOU M., DAVIDSON P., FANELLO S. R., KHAMIS S., KOWDLE A., RHEMANN C., TANKOVICH V., IZADI S.: Motion2Fusion: Real-time volumetric performance capture. *ACM Transactions on Graphics* 36, 6 (2017), 246:1–246:16. 3, 5, 11, 12
- [DKD*16] DOU M., KHAMIS S., DEGTAREV Y., DAVIDSON P., FANELLO S. R., KOWDLE A., ESCOLANO S. O., RHEMANN C., KIM D., TAYLOR J., KOHLI P., TANKOVICH V., IZADI S.: Fusion4D: Real-time performance capture of challenging scenes. *ACM Transactions on Graphics* 35, 4 (2016), 114:1–114:13. 5, 6, 10, 12
- [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. 20, 21
- [DLRT19] DYKE R. M., LAI Y.-K., ROSIN P. L., TAM G. K.: Non-rigid registration under anisotropic deformations. *Computer Aided Geometric Design* 71 (2019), 142–156. 14, 15, 20
- [DNK12] DOMOKOS C., NEMETH J., KATO Z.: Nonlinear shape registration without correspondences. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 34, 5 (2012), 943–958. 3
- [DRB*09] DEVIR Y. S., ROSMAN G., BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: On reconstruction of non-rigid shapes with intrinsic regularization. In *IEEE International Conference on Computer Vision Workshops* (2009), pp. 272–279. 15, 20
- [DSL*19] DYKE R. M., STRIDE C., LAI Y.-K., ROSIN P. L., AUBRY M., BOYARSKI A., BRONSTEIN A. M., BRONSTEIN M. M., CREMERS D., FISHER M., GROUEIX T., GUO D., KIM V. G., KIMMEL R., LÄHNER Z., LI K., LITANY O., REMEZ T., RODOLÀ E., RUSSELL B. C., SAHILLIOĞLU Y., SLOSSBERG R., TAM G. K. L., VESTNER M., WU Z., YANG J.: Shape correspondence with isometric and non-isometric

- deformations. In *Eurographics Workshop on 3D Object Retrieval* (2019), The Eurographics Association. 21
- [DSO20] DONATI N., SHARMA A., OVSIJANIKOV M.: Deep geometric functional maps: Robust feature learning for shape correspondence. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (June 2020), pp. 8589–8598. 19, 20
- [DZL*20] DYKE R. M., ZHOU F., LAI Y.-K., ROSIN P. L., GUO D., LI K., MARIN R., YANG J.: SHREC 2020 Track: Non-rigid shape correspondence of physically-based deformations. In *Eurographics Workshop on 3D Object Retrieval* (2020), The Eurographics Association. 21, 22
- [EC20] EISENBERGER M., CREMERS D.: Hamiltonian dynamics for real-world shape interpolation. In *Proceedings of the European Conference on Computer Vision* (Oct. 2020), vol. 12349, Springer International Publishing, pp. 179–196. 21
- [ELBS10] ESCOLANO F., LOZANO M. A., BONEV B., SUAU P.: Bypass information-theoretic shape similarity from non-rigid points-based alignment. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2010), pp. 37–44. 20
- [ELC20] EISENBERGER M., LÄHNER Z., CREMERS D.: Smooth shells: Multi-scale shape registration with functional maps. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 12262–12271. 11, 19, 20
- [ELPZ97] ELДАР Y., LINDENBAUM M., PORAT M., ZEEVI Y. Y.: The farthest point strategy for progressive image sampling. *IEEE transactions on image processing* 6, 9 (Sept. 1997), 1305–1315. 14
- [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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (June 2021), pp. 7473–7483. 11, 19
- [ETLTC20] EISENBERGER M., TOKER A., LEAL-TAIXÉ L., CREMERS D.: Deep shells: Unsupervised shape correspondence with optimal transport. In *Advances in Neural Information Processing Systems* (2020), vol. 33, Curran Associates, Inc., pp. 10491–10502. 19
- [FHCP18] FAN Z., HU X., CHEN C., PENG S.: Dense semantic and topological correspondence of 3D faces without landmarks. In *Proceedings of the European Conference on Computer Vision* (Sept. 2018). 5, 7, 12, 21
- [FKMS05] FUNKHOUSER T., KAZHDAN M., MIN P., SHILANE P.: Shape-based retrieval and analysis of 3D models. *Communications of the ACM* 48, 6 (2005), 58–64. 20
- [FNT*11] FUJIWARA K., NISHINO K., TAKAMATSU J., ZHENG B., IKEUCHI K.: Locally rigid globally non-rigid surface registration. In *International Conference on Computer Vision* (2011), IEEE, pp. 1527–1534. 4, 5, 12
- [FSG17] FAN H., SU H., GUIBAS L. J.: A point set generation network for 3D object reconstruction from a single image. In *IEEE Conference on Computer Vision and Pattern Recognition* (2017), pp. 2463–2471. 9, 10
- [FYZ*17] FAN J., YANG J., ZHAO Y., AI D., LIU Y., WANG G., WANG Y.: Convex hull aided registration method (CHARM). *IEEE Transactions on Visualization and Computer Graphics* 23, 9 (2017), 2042–2055. 3, 12
- [FZC*21] FENG W., ZHANG J., CAI H., XU H., HOU J., BAO H.: Recurrent multi-view alignment network for unsupervised surface registration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 10297–10307. 10, 13
- [GCBZ19] GONG S., CHEN L., BRONSTEIN M. M., ZAFEIRIOU S.: SpiralNet++: A fast and highly efficient mesh convolution operator. In *IEEE/CVF International Conference on Computer Vision Workshops* (Oct. 2019), pp. 4141–4148. 19
- [GCMB10] GUPTA S., CASTLEMAN K. R., MARKEY M. K., BOVIK A. C.: Texas 3D face recognition database. In *IEEE Southwest Symposium on Image Analysis Interpretation* (2010), pp. 97–100. 21
- [GF15] GE S., FAN G.: Non-rigid articulated point set registration with local structure preservation. In *IEEE Conference on Computer Vision and Pattern Recognition* (2015), pp. 126–133. 5, 6, 7, 8, 12
- [GFD14] GE S., FAN G., DING M.: Non-rigid point set registration with global-local topology preservation. In *IEEE Conference on Computer Vision and Pattern Recognition* (2014), pp. 245–251. 5, 7, 12
- [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* (Oct. 2018), Springer International Publishing, pp. 235–251. 9, 13, 21
- [GLLY17] GUO D., LI K., LAI Y.-K., YANG J.: Global alignment of deformable objects captured by a single RGB-D camera. In *IEEE International Conference on Multimedia and Expo* (2017), IEEE, pp. 1554–1559. 6, 12
- [GSDA*09] GALL J., STOLL C., DE AGUIAR E., THEOBALT C., ROSENHAHN B., SEIDEL H.-P.: Motion capture using joint skeleton tracking and surface estimation. In *IEEE Conference on Computer Vision and Pattern Recognition* (2009), IEEE, pp. 1746–1753. 4, 12
- [GSM15] GILANI S. Z., SHAFAIT F., MIAN A.: Shape-based automatic detection of a large number of 3D facial landmarks. In *IEEE Conference on Computer Vision and Pattern Recognition* (June 2015), pp. 4639–4648. 21
- [GSVS18] GOLYANIK V., SHIMADA S., VARANASI K., STRICKER D.: HDM-Net: Monocular non-rigid 3D reconstruction with learned deformation model. In *International Conference on Virtual Reality and Augmented Reality* (2018), Springer, pp. 51–72. 9
- [GT19] GAO W., TEDRAKE R.: FilterReg: Robust and efficient probabilistic point-set registration using gaussian filter and twist parameterization. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 11095–11104. 8
- [GTRS16] GOLYANIK V., TAETZ B., REIS G., STRICKER D.: Extended coherent point drift algorithm with correspondence priors and optimal subsampling. In *IEEE Winter Conference on Applications of Computer Vision* (2016), IEEE, pp. 1–9. 8
- [GXW*15] GUO K., XU F., WANG Y., LIU Y., DAI Q.: Robust non-rigid motion tracking and surface reconstruction using L0 regularization. In *IEEE International Conference on Computer Vision* (2015), pp. 3083–3091. 2, 5, 6, 7, 8, 12
- [GXY*17] GUO K., XU F., YU T., LIU X., DAI Q., LIU Y.: Real-time geometry, albedo, and motion reconstruction using a single RGB-D camera. *ACM Transactions on Graphics* 36, 4 (2017), 32:1–32:13. 10, 12
- [HAWG08] HUANG Q.-X., ADAMS B., WICKE M., GUIBAS L. J.: Non-rigid registration under isometric deformations. *Computer Graphics Forum* 27, 5 (2008), 1449–1457. 3, 14, 20
- [HBH11] HUANG P., BUDD C., HILTON A.: Global temporal registration of multiple non-rigid surface sequences. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2011), IEEE, pp. 3473–3480. 2, 6, 12
- [HDD*92] HOPPE H., DEROSE T., DUCHAMP T., McDONALD J., STUETZLE W.: Surface reconstruction from unorganized points. In *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques* (1992), pp. 71–78. 4
- [Hir20] 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 (2020), 2858–2865. 7, 12
- [Hir21] HIROSE O.: A bayesian formulation of coherent point drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43, 7 (2021), 2269–2286. 7, 8
- [HLR*19] HALIMI O., LITANY O., RODOLÀ E., BRONSTEIN A. M., KIMMEL R.: Unsupervised learning of dense shape correspondence. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (June 2019), pp. 4365–4374. 19, 20
- [HLS07] HORMANN K., LÉVY B., SHEFFER A.: Mesh parameterization: Theory and practice. In *ACM SIGGRAPH 2007 Courses* (2007), Association for Computing Machinery, p. 1–115. 11

- [HMS12] HONTANI H., MATSUNO T., SAWADA Y.: Robust nonrigid ICP using outlier-sparsity regularization. In *IEEE Conference on Computer Vision and Pattern Recognition* (2012), IEEE, pp. 174–181. 5, 12
- [HRA*19] HUANG R., RAKOTOSAONA M.-J., ACHLIOPTAS P., GUIBAS L. J., OVSJANIKOV M.: OperatorNet: Recovering 3D shapes from difference operators. In *IEEE/CVF International Conference on Computer Vision* (2019), pp. 8587–8596. 18
- [HSS*09] HASLER N., STOLL C., SUNKEL M., ROSENHAHN B., SEIDEL H.-P.: A statistical model of human pose and body shape. *Computer Graphics Forum* 28, 2 (2009), 337–346. 21
- [HZF*19] HUANG R., ZHAO J., DUAN F., LI X., LIU C., DENG X., PAN Z., WU Z., ZHOU M.: Automatic craniofacial registration based on radial curves. *Computers & Graphics* 82 (2019), 264–274. 3, 12
- [HZFH20] HAMIDIAN H., ZHONG Z., FOTOUHI F., HUA J.: Surface registration with eigenvalues and eigenvectors. *IEEE Transactions on Visualization and Computer Graphics* 26, 11 (2020), 3327–3339. 11
- [IBP15] ICHIM A. E., BOUAZIZ S., PAULY M.: Dynamic 3D avatar creation from hand-held video input. *ACM Transactions on Graphics* 34, 4 (2015), 45:1–45:14. 4
- [IZN*16] INNEMANN M., ZOLLHÖFER M., NIESSNER M., THEOBALT C., STAMMINGER M.: VolumeDeform: Real-time volumetric non-rigid reconstruction. In *Proceedings of the European Conference on Computer Vision* (2016), Springer, pp. 362–379. 4, 5, 6, 8, 10, 12
- [JDKL14] JACOBSON A., DENG Z., KAVAN L., LEWIS J. P.: Skinning: Real-time shape deformation. In *ACM SIGGRAPH 2014 Courses* (2014), SIGGRAPH '14, Association for Computing Machinery, pp. 24:1–24:1. 4
- [Joh97] JOHNSON A. E.: *Spin-images: a representation for 3-D surface matching*. PhD thesis, The Robotics Institute, Carnegie Mellon University, Aug. 1997. 7
- [JQL*17] JIANG T., QIAN K., LIU S., WANG J., YANG X., ZHANG J.: Consistent as-similar-as-possible non-isometric surface registration. *The Visual Computer* 33, 6 (2017), 891–901. 6, 12
- [JS06] JORGE N., STEPHEN J. W.: *Numerical optimization*, second ed. Springer, 2006. 8
- [JV05] JIAN B., VEMURI B. C.: A robust algorithm for point set registration using mixture of gaussians. In *IEEE International Conference on Computer Vision* (2005), IEEE, pp. 1246–1251. 5, 7, 12
- [JV11] JIAN B., VEMURI B. C.: Robust point set registration using gaussian mixture models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33, 8 (2011), 1633–1645. 5, 7, 12
- [JYZ*19] JIANG T., YANG X., ZHANG J., TIAN F., LIU S., XIANG N., QIAN K.: Huber- L_1 -based non-isometric surface registration. *The Visual Computer* 35, 6 (2019), 935–948. 6, 7, 12
- [KBB*13] KOVNATSKY A., BRONSTEIN M. M., BRONSTEIN A. M., GLASHOFF K., KIMMEL R.: Coupled quasi-harmonic bases. *Computer Graphics Forum* 32, 2pt4 (2013), 439–448. 17, 20
- [KBD17] KESZEI A. P., BERKELS B., DESERNO T. M.: Survey of non-rigid registration tools in medicine. *Journal of Digital Imaging* 30, 1 (2017), 102–116. 2
- [KBLB12] KOKKINOS I., BRONSTEIN M. M., LITMAN R., BRONSTEIN A. M.: Intrinsic shape context descriptors for deformable shapes. In *IEEE Conference on Computer Vision and Pattern Recognition* (June 2012), pp. 159–166. 19
- [KBM*20] KATO H., BEKER D., MORARIU M., ANDO T., MATSUOKA T., KEHL W., GAIDON A.: Differentiable rendering: a survey. *CoRR abs/2006.12057* (2020). 21
- [KCŽO08] KAVAN L., COLLINS S., ŽÁRA J., O’SULLIVAN C.: Geometric skinning with approximate dual quaternion blending. *ACM Transactions on Graphics* 27, 4 (2008), 105:1–105:23. 3
- [KG00] KARNI Z., GOTSCHMAN C.: Spectral compression of mesh geometry. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques* (2000), SIGGRAPH '00, ACM Press/Addison-Wesley Publishing Co., p. 279–286. 19
- [KLCF10] KIM V. G., LIPMAN Y., CHEN X., FUNKHOUSER T.: Möbius transformations for global intrinsic symmetry analysis. *Computer Graphics Forum* 29, 5 (2010), 1689–1700. 13
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. *ACM Transactions on Graphics* 30, 4 (2011), 79:1–79:12. 13, 15, 20, 21
- [KSI18] KOZLOV C., SLAVCHEVA M., ILIC S.: Patch-based non-rigid 3D reconstruction from a single depth stream. In *International Conference on 3D Vision* (2018), IEEE, pp. 42–51. 3, 6, 8, 12
- [LAGP09] LI H., ADAMS B., GUIBAS L. J., PAULY M.: Robust single-view geometry and motion reconstruction. *ACM Transactions on Graphics* 28, 5 (2009), 175:1–175:10. 2, 5, 6, 8, 12
- [LB14] LITMAN R., BRONSTEIN A.: Learning spectral descriptors for deformable shape correspondence. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36, 1 (Jan. 2014), 171–180. 19, 20
- [LBB*17] LI T., BOLKART T., BLACK M. J., LI H., ROMERO J.: Learning a model of facial shape and expression from 4D scans. *ACM Transactions on Graphics* 36, 6 (2017), 194:1–194:17. 9
- [LBZ*20] LI Y., BOZIC A., ZHANG T., JI Y., HARADA T., NIESSNER M.: Learning to optimize non-rigid tracking. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 4910–4918. 9, 11, 12, 13
- [LF09] LIPMAN Y., FUNKHOUSER T.: Möbius voting for surface correspondence. *ACM Transactions on Graphics* 28, 3 (July 2009), 72:1–72:12. 13, 14, 20
- [LG20] LI C., GUO X.: Topology-change-aware volumetric fusion for dynamic scene reconstruction. In *Proceedings of the European Conference on Computer Vision* (2020), Springer, pp. 258–274. 5, 6, 8, 12
- [LJO19] LIU H.-T. D., JACOBSON A., OVSJANIKOV M.: Spectral coarsening of geometric operators. *ACM Transactions on Graphics* 38, 4 (July 2019), 105:1–105:13. 14, 20
- [LJY*17] LI Z., JI Y., YANG W., YE J., YU J.: Robust 3D human motion reconstruction via dynamic template construction. In *International Conference on 3D Vision* (2017), IEEE, pp. 496–505. 2, 5, 6, 12
- [LKS14] LIANG S., KEMELMACHER-SHLIZERMAN I., SHAPIRO L. G.: 3D face hallucination from a single depth frame. In *International Conference on 3D Vision* (2014), vol. 1, IEEE, pp. 31–38. 3, 12
- [LLM*16] LIN S., LAI Y.-K., MARTIN R. R., JIN S., CHENG Z.-Q.: Color-aware surface registration. *Computers & Graphics* 58 (2016), 31–42. 2, 6, 12
- [LLT*20] LESCOAT T., LIU H.-T. D., THIERY J.-M., JACOBSON A., BOUBEKEUR T., OVSJANIKOV M.: Spectral mesh simplification. *Computer Graphics Forum* 39, 2 (2020), 315–324. 14
- [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 (Oct. 2015), 248:1–248:16. 4, 9
- [Low04] LOWE D. G.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* 60, 2 (2004), 91–110. 8
- [LRB*16a] LÄHNER Z., RODOLÀ E., BRONSTEIN M. M., CREMERS D., BURGHARD O., COSMO L., DIECKMANN A., KLEIN R., SAHILLIOĞLU Y.: Matching of deformable shapes with topological noise. In *Eurographics Workshop on 3D Object Retrieval* (2016), The Eurographics Association and John Wiley & Sons, Ltd., pp. 55–60. 21
- [LRB*16b] LITANY O., RODOLÀ E., BRONSTEIN A. M., BRONSTEIN M. M., CREMERS D.: Non-rigid puzzles. In *Proceedings of the Symposium on Geometry Processing* (2016), SGP '16, The Eurographics Association, pp. 135–143. 17, 20
- [LRBB17] LITANY O., RODOLÀ E., BRONSTEIN A. M., BRONSTEIN M. M.: Fully spectral partial shape matching. *Computer Graphics Forum* 36, 2 (2017), 247–258. 17, 20

- [LRR*17] LITANY O., REMEZ T., RODOLÀ E., BRONSTEIN A., BRONSTEIN M.: Deep functional maps: Structured prediction for dense shape correspondence. In *IEEE International Conference on Computer Vision* (Oct. 2017), pp. 5660–5668. [19, 20](#)
- [LSP08] LI H., SUMNER R. W., PAULY M.: Global correspondence optimization for non-rigid registration of depth scans. *Computer Graphics Forum* 27 (2008), 1421–1430. [2, 3, 5, 6, 7, 8, 12](#)
- [LYLG19] LI K., YANG J., LAI Y.-K., GUO D.: Robust non-rigid registration with reweighted position and transformation sparsity. *IEEE Transactions on Visualization and Computer Graphics* 25, 6 (2019), 2255–2269. [2, 6, 7, 8, 12, 14](#)
- [LYP*20] LI Z., YU T., PAN C., ZHENG Z., LIU Y.: Robust 3D self-portraits in seconds. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 1344–1353. [3, 5, 12](#)
- [LZG18] LI C., ZHAO Z., GUO X.: ArticulatedFusion: Real-time reconstruction of motion, geometry and segmentation using a single depth camera. In *Proceedings of the European Conference on Computer Vision* (2018), pp. 317–332. [5, 12](#)
- [LZW*09] LIAO M., ZHANG Q., WANG H., YANG R., GONG M.: Modeling deformable objects from a single depth camera. In *IEEE International Conference on Computer Vision* (2009), IEEE, pp. 167–174. [2, 5, 6, 7, 8, 12](#)
- [LZXH20] LI W., ZHAO S., XIAO X., HAHN J.: Robust template-based non-rigid motion tracking using local coordinate regularization. In *IEEE/CVF Winter Conference on Applications of Computer Vision* (2020), pp. 401–410. [2, 3, 5, 11, 12, 21](#)
- [MBBV15] MASCI J., BOSCAINI D., BRONSTEIN M. M., VANDERGHEYNST P.: Geodesic convolutional neural networks on riemannian manifolds. In *IEEE International Conference on Computer Vision Workshops* (Dec. 2015), pp. 832–840. [19, 20](#)
- [MBM*17] MONTI F., BOSCAINI D., MASCI J., RODOLÀ E., SVOBODA J., BRONSTEIN M. M.: Geometric deep learning on graphs and manifolds using mixture model CNNs. In *IEEE Conference on Computer Vision and Pattern Recognition* (July 2017), IEEE Computer Society, pp. 5425–5434. [19, 20](#)
- [MCSK*17] MANDAD M., COHEN-STEINER D., KOBELT L., ALLIEZ P., DESBRUN M.: Variance-minimizing transport plans for inter-surface mapping. *ACM Transactions on Graphics* 36, 4 (jul 2017), 39:1–39:14. [14, 20](#)
- [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), 73:1–73:12. [21](#)
- [MDSB02] MEYER M., DESBRUN M., SCHRÖDER P., BARR A. H.: Discrete differential-geometry operators for triangulated 2-manifolds. In *Visualization and Mathematics III* (2002). [15](#)
- [MMC*21] MOSCHELLA L., MELZI S., COSMO L., MAGGIOLI F., LITANY O., OVSJANIKOV M., GUIBAS L., RODOLÀ E.: Spectral unions of partial deformable 3D shapes. *arXiv preprint arXiv:2104.00514* (2021). [19, 20](#)
- [MMM*19] MELZI S., MARIN R., MUSONI P., CASTELLANI U., TARINI M.: Visual assessments of functional maps. In *Proceedings of the Symposium on Geometry Processing - Posters* (2019), The Eurographics Association. [20](#)
- [MMR*19] 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), The Eurographics Association. [21](#)
- [MMRC20] MARIN R., MELZI S., RODOLÀ E., CASTELLANI U.: FARM: Functional automatic registration method for 3D human bodies. *Computer Graphics Forum* 39, 1 (Feb. 2020), 160–173. [4, 11](#)
- [MQZ*15] MA J., QIU W., ZHAO J., MA Y., YUILLE A. L., TU Z.: Robust L_2E estimation of transformation for non-rigid registration. *IEEE Transactions on Signal Processing* 63, 5 (2015), 1115–1129. [5, 7, 8, 12](#)
- [MRC*20] MARIN R., RAMPINI A., CASTELLANI U., RODOLÀ E., OVSJANIKOV M., MELZI S.: Instant recovery of shape from spectrum via latent space connections. In *International Conference on 3D Vision* (2020), pp. 120–129. [19, 20, 21](#)
- [MRC*21] MARIN R., RAMPINI A., CASTELLANI U., RODOLÀ E., OVSJANIKOV M., MELZI S.: Spectral shape recovery and analysis via data-driven connections. *International Journal of Computer Vision* 129, 10 (2021), 2745–2760. [19, 20, 21](#)
- [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 (Nov. 2019), 155:1–155:14. [17, 20](#)
- [MS10] MYRONENKO A., SONG X.: Point set registration: Coherent point drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, 12 (2010), 2262–2275. [3, 5, 7, 8, 12](#)
- [MSCPn06] MYRONENKO A., SONG X., CARREIRA-PERPIÑAN M.: Non-rigid point set registration: Coherent point drift. In *Advances in Neural Information Processing Systems* (2006), vol. 19, MIT Press, pp. 1009–1016. [3, 5, 7, 8, 12](#)
- [MST*20] MILDENHALL B., SRINIVASAN P. P., TANCIK M., BARRON J. T., RAMAMOORTHY R., NG R.: NeRF: Representing scenes as neural radiance fields for view synthesis. In *Proceedings of the European Conference on Computer Vision* (2020), Vedaldi A., Bischof H., Brox T., Frahm J., (Eds.), vol. 12346 of *Lecture Notes in Computer Science*, Springer, pp. 405–421. [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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 6469–6478. [9](#)
- [MZJZ17] MA J., ZHAO J., JIANG J., ZHOU H.: Non-rigid point set registration with robust transformation estimation under manifold regularization. In *Thirty-First AAAI Conference on Artificial Intelligence* (2017). [3, 5, 8, 12](#)
- [MZT*13] MA J., ZHAO J., TIAN J., TU Z., YUILLE A. L.: Robust estimation of nonrigid transformation for point set registration. In *IEEE Conference on Computer Vision and Pattern Recognition* (2013), pp. 2147–2154. [3, 5, 7, 8, 12](#)
- [MZY16] MA J., ZHAO J., YUILLE A. L.: Non-rigid point set registration by preserving global and local structures. *IEEE Transactions on Image Processing* 25, 1 (2016), 53–64. [3, 5, 7, 8, 12](#)
- [NBH18] NASIKUN A., BRANDT C., HILDEBRANDT K.: Fast approximation of laplace-beltrami eigenproblems. *Computer Graphics Forum* 37, 5 (2018), 121–134. [14](#)
- [NFS15] NEWCOMBE R. A., FOX D., SEITZ S. M.: DynamicFusion: Reconstruction and tracking of non-rigid scenes in real-time. In *IEEE Conference on Computer Vision and Pattern Recognition* (2015), pp. 343–352. [4, 7, 10, 12](#)
- [NMOG20] NIEMEYER M., MESCHEDER L., OECHSLE M., GEIGER A.: Differentiable volumetric rendering: Learning implicit 3D representations without 3D supervision. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 3504–3515. [22](#)
- [NO17] NOGNENG D., OVSJANIKOV M.: Informative descriptor preservation via commutativity for shape matching. *Computer Graphics Forum* 36, 2 (May 2017), 259–267. [20](#)
- [OBCCG13] OVSJANIKOV M., BEN-CHEN M., CHAZAL F., GUIBAS L.: Analysis and visualization of maps between shapes. *Computer Graphics Forum* 32, 6 (Sept. 2013), 135–145. [20](#)
- [OBCS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L. J.: Functional maps: A flexible representation of maps between shapes. *ACM Transactions on Graphics* 31, 4 (July 2012), 30:1–30:11. [15, 16, 19, 20](#)
- [OMMG10] OVSJANIKOV M., MÉRIGOT Q., MÉMOLI F., GUIBAS L. J.: One point isometric matching with the heat kernel. *Computer Graphics Forum* 29, 5 (2010), 1555–1564. [13, 20](#)

- [OPG21] OECHSLE M., PENG S., GEIGER A.: UNISURF: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *IEEE/CVF International Conference on Computer Vision* (2021). 22
- [PFS*05] PHILLIPS P., FLYNN P., SCRUGGS T., BOWYER K., CHANG J., HOFFMAN K., MARQUES J., MIN J., WOREK W.: Overview of the face recognition grand challenge. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2005), vol. 1, pp. 947–954 vol. 1. 21
- [PFS*19] PARK J. J., FLORENCE P., STRAUB J., NEWCOMBE R., LOVEGROVE S.: DeepSDF: Learning continuous signed distance functions for shape representation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 165–174. 22
- [PG08] PEKELNY Y., GOTSMAN C.: Articulated object reconstruction and markerless motion capture from depth video. *Computer Graphics Forum* 27, 2 (2008), 399–408. 4, 12
- [PHYH06] POTTMANN H., HUANG Q.-X., YANG Y.-L., HU S.-M.: Geometry and convergence analysis of algorithms for registration of 3D shapes. *International Journal of Computer Vision* 67, 3 (2006), 277–296. 5
- [PMG*05] PAULY M., MITRA N. J., GIESEN J., GROSS M., GUIBAS L. J.: Example-based 3D scan completion. In *Proceedings of the Symposium on Geometry Processing* (2005), Desbrun M., Pottmann H., (Eds.), The Eurographics Association. 5, 12
- [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 (July 2015), 120:1–120:14. 21
- [PO18] POULENARD A., OVSJANIKOV M.: Multi-directional geodesic neural networks via equivariant convolution. *ACM Transactions on Graphics* 37, 6 (2018), 236:1–236:14. 19
- [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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 384–393. 14, 20
- [PWH*17] PISHCHULIN L., WUHRER S., HELTEN T., THEOBALT C., SCHIELE B.: Building statistical shape spaces for 3D human modeling. *Pattern Recognition* 67 (2017), 276–286. 21
- [RBA*12] RODOLÀ E., BRONSTEIN A. M., ALBARELLI A., BERGAMASCO F., TORSELLO A.: A game-theoretic approach to deformable shape matching. In *IEEE Conference on Computer Vision and Pattern Recognition* (2012), pp. 182–189. 13, 14, 20
- [RBB09] RUSU R. B., BLODOW N., BEETZ M.: Fast point feature histograms (FPFH) for 3D registration. In *IEEE International Conference on Robotics and Automation* (2009), IEEE, pp. 3212–3217. 8
- [RBMBO8] RUSU R. B., BLODOW N., MARTON Z. C., BEETZ M.: Aligning point cloud views using persistent feature histograms. In *IEEE/RSJ international conference on intelligent robots and systems* (2008), IEEE, pp. 3384–3391. 8
- [RBRY19] ROSTAMI R., BASHIRI F. S., ROSTAMI B., YU Z.: A survey on data-driven 3D shape descriptors. *Computer Graphics Forum* 38, 1 (2019), 356–393. 19, 21
- [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* (September 2018). 21
- [RBW*14] RODOLÀ E., BULÒ S., WINDHEUSER T., VESTNER M., CREMERS D.: Dense non-rigid shape correspondence using random forests. In *Computer Vision and Pattern Recognition* (2014), IEEE, pp. 4177–4184. 19, 20, 21
- [RCB*17] RODOLÀ E., COSMO L., BRONSTEIN M. M., TORSELLO A., CREMERS D.: Partial functional correspondence. *Computer Graphics Forum* 36, 1 (2017), 222–236. 16, 17, 20, 21
- [RCL*17] RODOLÀ E., COSMO L., LITANY O., BRONSTEIN M. M., BRONSTEIN A. M., AUDEBERT N., HAMZA A. B., BOULCH A., CASTELLANI U., DO M. N., DUONG A.-D., FURUYA T., GASPARETTO A., HONG Y., KIM J., SAUX B. L., LITMAN R., MASOUMI M., MINELLO G., NGUYEN H.-D., NGUYEN V.-T., OHBUCHI R., PHAM V.-K., PHAN T. V., REZAEI M., TORSELLO A., TRAN M.-T., TRAN Q.-T., TRUONG B., WAN L., ZOU C.: Deformable shape retrieval with missing parts. In *Eurographics Workshop on 3D Object Retrieval* (Apr. 2017), The Eurographics Association, pp. 85–94. 21
- [RDP99] ROBINETTE K. M., DAANEN H., PAQUET E.: The CAESAR project: A 3D surface anthropometry survey. In *International Conference on 3D Digital Imaging and Modeling* (1999), 3DIM '99, IEEE, pp. 380–386. 21
- [RL01] RUSINKIEWICZ S., LEVOY M.: Efficient variants of the ICP algorithm. In *Proceedings Third International Conference on 3-D Digital Imaging and Modeling* (2001), pp. 145–152. 20
- [RMC15] RODOLÀ E., MOELLER M., CREMERS D.: Point-wise map recovery and refinement from functional correspondence. In *Vision, Modeling & Visualization* (2015), The Eurographics Association. 16, 20
- [RMC17] RODOLÀ E., MOELLER M., CREMERS D.: Regularized point-wise map recovery from functional correspondence. *Computer Graphics Forum* 36, 8 (2017), 700–711. 16, 20
- [ROA*13] RUSTAMOV R. M., OVSJANIKOV M., AZENCOT O., BEN-CHEN M., CHAZAL F., GUIBAS L. J.: Map-based exploration of intrinsic shape differences and variability. *ACM Transactions on Graphics* 32, 4 (July 2013), 72:1–72:12. 18, 20
- [RS12] ROUHANI M., SAPP A. D.: Non-rigid shape registration: A single linear least squares framework. In *Proceedings of the European Conference on Computer Vision* (2012), Springer, pp. 264–277. 4, 5, 12
- [RSH*99] RUECKERT D., SONODA L., HAYES C., HILL D., LEACH M., HAWKES D.: Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Transactions on Medical Imaging* 18, 8 (1999), 712–721. 4, 12
- [RSO19] ROUFOSSE J.-M., SHARMA A., OVSJANIKOV M.: Unsupervised deep learning for structured shape matching. In *IEEE/CVF International Conference on Computer Vision* (Oct. 2019), pp. 1617–1627. 19, 20
- [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 *International Conference on 3D Vision* (2019), pp. 37–46. 17, 20
- [SAD*08] SAVRAN A., ALYÜZ N., DIBEKLIOĞLU H., ÇELIKTUTAN O., GÖKBERK B., SANKUR B., AKARUN L.: Bosphorus database for 3D face analysis. In *Biometrics and Identity Management* (2008), Springer Berlin Heidelberg, pp. 47–56. 21
- [Sah15] SAHILLIOĞLU Y.: A shape deformation algorithm for constrained multidimensional scaling. *Computers & Graphics* 53, DEC.PT.B (2015), 156–165. 6, 12
- [Sah18] SAHILLIOĞLU Y.: A genetic isometric shape correspondence algorithm with adaptive sampling. *ACM Transactions on Graphics* 37, 5 (Oct. 2018), 175:1–175:14. 11, 13, 14, 20
- [Sah20] SAHILLIOĞLU Y.: Recent advances in shape correspondence. *The Visual Computer* 36 (Aug. 2020). 2, 11, 15
- [SAS*18] SAINT A., AHMED E., SHABAYEK A. E. R., CHERENKOVA K., GUSEV G., AOUADA D., OTTERSTEN B.: 3DBodyTex: Textured 3D body dataset. In *International Conference on 3D Vision* (2018), pp. 495–504. 21
- [SBCI17] SLAVCHEVA M., BAUST M., CREMERS D., ILIC S.: Killing-Fusion: Non-rigid 3D reconstruction without correspondences. In *IEEE Conference on Computer Vision and Pattern Recognition* (2017), pp. 1386–1395. 4, 5, 12
- [SBI18] SLAVCHEVA M., BAUST M., ILIC S.: SobolevFusion: 3D reconstruction of scenes undergoing free non-rigid motion. In *IEEE Conference on Computer Vision and Pattern Recognition* (2018), pp. 2646–2655. 10, 12

- [SBI21] SLAVCHEVA M., BAUST M., ILIC S.: Variational level set evolution for non-rigid 3D reconstruction from a single depth camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43, 8 (2021), 2838–2850. [10](#), [12](#)
- [Sco01] SCOTT D. W.: Parametric statistical modeling by minimum integrated square error. *Technometrics* 43, 3 (2001), 274–285. [7](#)
- [SCO04] SORKINE O., COHEN-OR D.: Least-squares meshes. In *Proceedings Shape Modeling Applications* (2004), IEEE, pp. 191–199. [6](#)
- [SdGP*15] SOLOMON J., DE GOES F., PEYRÉ G., CUTURI M., BUTSCHER A., NGUYEN A., DU T., GUIBAS L. J.: Convolutional wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Transactions on Graphics* 34, 4 (jul 2015), 66:1–66:11. [14](#), [20](#)
- [SGT*19] SHIMADA S., GOLYANIK V., TRETSCHK E., STRICKER D., THEOBALT C.: DispVoxNets: Non-rigid point set alignment with supervised learning proxies. In *International Conference on 3D Vision* (Sept. 2019), pp. 27–36. [9](#), [13](#), [21](#)
- [SK13] SÁNTA Z., KATO Z.: Correspondence-less non-rigid registration of triangular surface meshes. In *IEEE Conference on Computer Vision and Pattern Recognition* (2013), pp. 2275–2282. [8](#)
- [SK16] SÁNTA Z., KATO Z.: 3D face alignment without correspondences. In *Proceedings of the European Conference on Computer Vision* (Nov. 2016), Springer International Publishing, pp. 521–535. [3](#), [12](#), [21](#)
- [SK21] SAHILLIOĞLU Y., KAVAN L.: Scale-adaptive ICP. *Graphical Models* (2021), 101113. [7](#)
- [SMB*20] SITZMANN V., MARTEL J., BERGMAN A., LINDELL D., WETZSTEIN G.: Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems* (2020), Larochelle H., Ranzato M., Hadsell R., Balcan M. F., Lin H., (Eds.), vol. 33, Curran Associates, Inc., pp. 7462–7473. [22](#)
- [SNB*12] SOLOMON J., NGUYEN A., BUTSCHER A., BEN-CHEN M., GUIBAS L. J.: Soft maps between surfaces. *Computer Graphics Forum* 31, 5 (2012), 1617–1626. [14](#), [20](#)
- [SOG09] SUN J., OVSJANIKOV M., GUIBAS L. J.: A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum* 28, 5 (2009), 1383–1392. [13](#), [19](#)
- [SP86] SEDERBERG T. W., PARRY S. R.: Free-form deformation of solid geometric models. *SIGGRAPH Computer Graphics* 20, 4 (1986), 151–160. [3](#)
- [SP04] SUMNER R. W., POPOVIĆ J.: Deformation transfer for triangle meshes. *ACM Transactions on Graphics* 23, 3 (2004), 399–405. [21](#)
- [SPF19] SCHMIDT F., PHILLIPS F., FLEMING R. W.: Visual perception of shape-transforming processes: ‘shape scission’. *Cognition* 189 (2019), 167–180. [22](#)
- [SPKS16] SOLOMON J., PEYRÉ G., KIM V. G., SRA S.: Entropic metric alignment for correspondence problems. *ACM Transactions on Graphics* 35, 4 (July 2016), 72:1–72:13. [15](#), [20](#)
- [SPR07] SHEFFER A., PRAUN E., ROSE K.: Mesh parameterization methods and their applications. *Foundations and Trends® in Computer Graphics and Vision* 2, 2 (2007), 105–171. [11](#)
- [SRC*20] SUN Z., ROOKE E., CHARTON J., HE Y., LU J., BAEK S.: ZerNet: Convolutional neural networks on arbitrary surfaces via zernike local tangent space estimation. *Computer Graphics Forum* 39, 6 (Sept. 2020), 204–216. [19](#)
- [SSP07] SUMNER R. W., SCHMID J., PAULY M.: Embedded deformation for shape manipulation. *ACM Transactions on Graphics* 26, 3 (2007), 80:1–80:7. [2](#), [3](#)
- [ST20] SAITI E., THEOHARIS T.: An application independent review of multimodal 3D registration methods. *Computers & Graphics* 91 (2020), 153–178. [2](#)
- [STDS14] SALT I., TOMBARI F., DI STEFANO L.: SHOT: Unique signatures of histograms for surface and texture description. *Computer Vision and Image Understanding* 125 (2014), 251–264. [7](#), [8](#)
- [SWG08] SÜSSMUTH J., WINTER M., GREINER G.: Reconstructing animated meshes from time-varying point clouds. *Computer Graphics Forum* 27, 5 (2008), 1469–1476. [6](#)
- [SXZ*20] SU Z., XU L., ZHENG Z., YU T., LIU Y., FANG L.: Robust-Fusion: Human volumetric capture with data-driven visual cues using a RGBD camera. In *Proceedings of the European Conference on Computer Vision* (2020), Springer, pp. 246–264. [2](#), [4](#), [5](#), [7](#), [11](#), [12](#), [13](#)
- [SY11] SAHILLIOĞLU Y., YEMEZ Y.: Coarse-to-fine combinatorial matching for dense isometric shape correspondence. *Computer Graphics Forum* 30, 5 (2011), 1461–1470. [14](#), [20](#)
- [SY12a] SAHILLIOĞLU Y., YEMEZ Y.: Minimum-distortion isometric shape correspondence using EM algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 34, 11 (2012), 2203–2215. [14](#), [20](#)
- [SY12b] SAHILLIOĞLU Y., YEMEZ Y.: Scale normalization for isometric shape matching. *Computer Graphics Forum* 31, 7 (2012), 2233–2240. [13](#), [14](#), [20](#)
- [TBW*11] TEVS A., BERNER A., WAND M., IHRKE I., SEIDEL H.-P.: Intrinsic shape matching by planned landmark sampling. *Computer Graphics Forum* 30, 2 (2011), 543–552. [14](#), [20](#)
- [TCL*13] TAM G. K., CHENG Z.-Q., LAI Y.-K., LANGBEIN F. C., LIU Y., MARSHALL D., MARTIN R. R., SUN X.-F., ROSIN P. L.: Registration of 3D point clouds and meshes: A survey from rigid to nonrigid. *IEEE Transactions on Visualization and Computer Graphics* 19, 7 (2013), 1199–1217. [2](#)
- [TCM*21] TRAPPOLINI G., COSMO L., MOSCHELLA L., MARIN R., MELZI S., RODOLÀ E.: Shape registration in the time of transformers. *NeurIPS 2021* (2021). [8](#), [9](#), [13](#)
- [TFT*20] TEWARI A., FRIED O., THIES J., SITZMANN V., LOMBARDI S., SUNKAVALLI K., MARTIN-BRUALLA R., SIMON T., SARAGIH J. M., NIESSNER M., PANDEY R., FANELLO S. R., WETZSTEIN G., ZHU J., THEOBALT C., AGRAWALA M., SHECHTMAN E., GOLDMAN D. B., ZOLLHÖFER M.: State of the art on neural rendering. *Computer Graphics Forum* 39, 2 (2020), 701–727. [21](#)
- [TMRL14a] TAM G. K., MARTIN R. R., ROSIN P. L., LAI Y.-K.: Diffusion pruning for rapidly and robustly selecting global correspondences using local isometry. *ACM Transactions on Graphics* 33, 1 (2014), 4:1–4:17. [7](#), [8](#), [14](#), [15](#), [20](#)
- [TMRL14b] TAM G. K., MARTIN R. R., ROSIN P. L., LAI Y.-K.: An efficient approach to correspondences between multiple non-rigid parts. *Computer Graphics Forum* 33, 5 (2014), 137–146. [14](#), [20](#)
- [TPT16] TKACH A., PAULY M., TAGLIASACCHI A.: Sphere-meshes for real-time hand modeling and tracking. *ACM Transactions on Graphics* 35, 6 (2016), 222:1–222:11. [4](#), [11](#), [12](#)
- [TSDS10] TOMBARI F., SALT I., DI STEFANO L.: Unique signatures of histograms for local surface description. In *Proceedings of the European Conference on Computer Vision* (2010), Springer Berlin Heidelberg, pp. 356–369. [13](#)
- [Tse01] TSENG P.: Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of Optimization Theory and Applications* 109, 3 (2001), 475–494. [8](#)
- [TST*15] TAGLIASACCHI A., SCHRÖDER M., TKACH A., BOUAZIZ S., BOTSCH M., PAULY M.: Robust articulated-ICP for real-time hand tracking. *Computer Graphics Forum* 34, 5 (2015), 101–114. [4](#), [6](#), [11](#), [12](#)
- [TT16] THOMAS D., TANIGUCHI R.-I.: Augmented blendshapes for real-time simultaneous 3D head modeling and facial motion capture. In *IEEE Conference on Computer Vision and Pattern Recognition* (2016), pp. 3299–3308. [5](#), [6](#), [11](#), [12](#)
- [TTR*17] TKACH A., TAGLIASACCHI A., REMELLI E., PAULY M., FITZGIBBON A.: Online generative model personalization for hand tracking. *ACM Transactions on Graphics* 36, 6 (2017), 243:1–243:11. [11](#), [12](#)
- [TZL*12] TONG J., ZHOU J., LIU L., PAN Z., YAN H.: Scanning 3D full human bodies using kinects. *IEEE Transactions on Visualization and Computer Graphics* 18, 4 (2012), 643–650. [11](#), [12](#)

- [VBMP08] VLASIC D., BARAN I., MATUSIK W., POPOVIĆ J.: Articulated mesh animation from multi-view silhouettes. *ACM Transactions on Graphics* 27, 3 (Aug. 2008), 97:1–97:9. 10, 21
- [vKZH13] VAN KAICK O., ZHANG H., HAMARNEH G.: Bilateral maps for partial matching. *Computer Graphics Forum* 32, 6 (Sept. 2013), 189–200. 14, 20
- [vKZHC011] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. *Computer Graphics Forum* 30, 6 (2011), 1681–1707. 2, 21
- [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 *International Conference on 3D Vision* (2017), pp. 517–526. 14, 15, 20
- [VRM*17] VAROL G., ROMERO J., MARTIN X., MAHMOOD N., BLACK M. J., LAPTEV I., SCHMID C.: Learning from synthetic humans. In *IEEE Conference on Computer Vision and Pattern Recognition* (2017), pp. 4627–4635. 9, 10, 21
- [WAO*09] WAND M., ADAMS B., OVSJANIKOV M., BERNER A., BOKELOH M., JENKE P., GUIBAS L. J., SEIDEL H.-P., SCHILLING A.: Efficient reconstruction of nonrigid shape and motion from real-time 3D scanner data. *ACM Transactions on Graphics* 28, 2 (2009), 15:1–15:15. 6, 8, 12
- [WBLP11] WEISE T., BOUAZIZ S., LI H., PAULY M.: Realtime performance-based facial animation. *ACM Transactions on Graphics* 30, 4 (2011), 77:1–77:10. 4, 11, 12
- [WCLF19] WANG L., CHEN J., LI X., FANG Y.: Non-rigid point set registration networks. *arXiv preprint arXiv:1904.01428* (2019). 9, 13
- [WCMN19] WANG W., CEYLAN D., MECH R., NEUMANN U.: 3DN: 3D deformation network. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 1038–1046. 10, 13
- [WEH20] WIERSMA R., EISEMANN E., HILDEBRANDT K.: CNNs on surfaces using rotation-equivariant features. *ACM Transactions on Graphics* 39, 4 (July 2020), 92:1–92:12. 19
- [WGT21] WANG S., GEIGER A., TANG S.: Locally aware piecewise transformation fields for 3D human mesh registration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 7639–7648. 9
- [WGY*18] WANG H., GUO J., YAN D.-M., QUAN W., ZHANG X.: Learning 3D keypoint descriptors for non-rigid shape matching. In *Proceedings of the European Conference on Computer Vision* (2018), pp. 3–19. 19, 20
- [WHC*16] WEI L., HUANG Q., CEYLAN D., VOUGA E., LI H.: Dense human body correspondences using convolutional networks. In *IEEE Conference on Computer Vision and Pattern Recognition* (2016), pp. 1544–1553. 8, 13
- [WJH*07] WAND M., JENKE P., HUANG Q., BOKELOH M., GUIBAS L. J., SCHILLING A.: Reconstruction of deforming geometry from time-varying point clouds. In *Geometry Processing* (2007), The Eurographics Association. 6, 8, 12
- [WLCF19] WANG L., LI X., CHEN J., FANG Y.: Coherent point drift networks: Unsupervised learning of non-rigid point set registration. *arXiv preprint arXiv:1906.03039* (2019). 9, 13
- [WLLY19] WU Z., LI K., LAI Y.-K., YANG J.: Global as-conformal-as-possible non-rigid registration of multi-view scans. In *IEEE International Conference on Multimedia and Expo* (2019), pp. 308–313. 6, 12
- [WRFR*16] WANG S., RYAN FANELLO S., RHEMANN C., IZADI S., KOHLI P.: The global patch collider. In *IEEE Conference on Computer Vision and Pattern Recognition* (2016), pp. 127–135. 10
- [WRY*20] WANG Y., REN J., YAN D.-M., GUO J., ZHANG X., WONKA P.: MGCN: Descriptor learning using multiscale GCNs. *ACM Transactions on Graphics* 39, 4 (July 2020), 122:1–122:15. 20
- [WVR*14] WINDHEUSER T., VESTNER M., RODOLÀ E., TRIEBEL R., CREMERS D.: Optimal intrinsic descriptors for non-rigid shape analysis. In *Proceedings of the British Machine Vision Conference* (Sept. 2014), BMVA Press. 19
- [WXZ*20] WANG K., XIE J., ZHANG G., LIU L., YANG J.: Sequential 3D human pose and shape estimation from point clouds. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 7275–7284. 9, 13
- [WZX17] WANG K., ZHANG G., XIA S.: Templateless non-rigid reconstruction and motion tracking with a single RGB-D camera. *IEEE Transactions on Image Processing* 26, 12 (2017), 5966–5979. 5, 6, 10, 12
- [XLC*18] XU L., LIU Y., CHENG W., GUO K., ZHOU G., DAI Q., FANG L.: FlyCap: Markerless motion capture using multiple autonomous flying cameras. *IEEE Transactions on Visualization and Computer Graphics* 24, 8 (2018), 2284–2297. 5, 6, 11, 12
- [XSH*20] XU L., SU Z., HAN L., YU T., LIU Y., FANG L.: UnstructuredFusion: realtime 4D geometry and texture reconstruction using commercial RGBD cameras. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42, 10 (2020), 2508–2522. 5, 7, 10, 12
- [XSWL14] XU W., SALZMANN M., WANG Y., LIU Y.: Nonrigid surface registration and completion from RGBD images. In *Proceedings of the European Conference on Computer Vision* (2014), Springer, pp. 64–79. 3, 12
- [XSWL15] XU W., SALZMANN M., WANG Y., LIU Y.: Deformable 3D fusion: From partial dynamic 3D observations to complete 4D models. In *IEEE International Conference on Computer Vision* (2015), pp. 2183–2191. 10, 12
- [XXY*18] XU J., XU W., YANG Y., DENG Z., BAO H.: Online global non-rigid registration for 3D object reconstruction using consumer-level depth cameras. *Computer Graphics Forum* 37, 7 (2018), 1–12. 10, 12
- [XZC18] XU Z., ZHANG Q., CHENG S.: Multilevel active registration for Kinect human body scans: from low quality to high quality. *Multimedia Systems* 24, 3 (2018). 21
- [YDXZ20] YAO Y., DENG B., XU W., ZHANG J.: Quasi-newton solver for robust non-rigid registration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (June 2020), pp. 7597–7606. 2, 5, 6, 7, 8, 12, 14, 21
- [YGKL21] YARIV L., GU J., KASTEN Y., LIPMAN Y.: Volume rendering of neural implicit surfaces. *arXiv preprint arXiv:2106.12052* (2021). 22
- [YGL*19] YANG J., GUO D., LI K., WU Z., LAI Y.-K.: Global 3D non-rigid registration of deformable objects using a single RGB-D camera. *IEEE Transactions on Image Processing* 28, 10 (2019), 4746–4761. 2, 6, 7, 8, 12, 14
- [YGX*17] YU T., GUO K., XU F., DONG Y., SU Z., ZHAO J., LI J., DAI Q., LIU Y.: BodyFusion: Real-time capture of human motion and surface geometry using a single depth camera. In *IEEE International Conference on Computer Vision* (2017), pp. 910–919. 4, 5, 10, 12
- [YHL*16] YIN K., HUANG H., LONG P., GAISSINSKI A., GONG M., SHARF A.: Full 3D plant reconstruction via intrusive acquisition. *Computer Graphics Forum* 35, 1 (2016), 272–284. 8
- [YKM13] YAMAZAKI S., KAGAMI S., MOCHIMARU M.: Non-rigid shape registration using similarity-invariant differential coordinates. In *International Conference on 3D Vision* (2013), IEEE, pp. 191–198. 2, 5, 6, 7, 8, 12
- [YKM*20] YARIV L., KASTEN Y., MORAN D., GALUN M., ATZMON M., BASRI R., LIPMAN Y.: Multiview neural surface reconstruction by disentangling geometry and appearance. In *Annual Conference on Neural Information Processing Systems* (2020). 22
- [YLLL15] YANG J., LI K., LI K., LAI Y.-K.: Sparse non-rigid registration of 3D shapes. *Computer Graphics Forum* 34, 5 (2015), 89–99. 2, 6, 7, 8, 12, 14
- [YLSL10] YEH I.-C., LIN C.-H., SORKINE O., LEE T.-Y.: Template-based 3D model fitting using dual-domain relaxation. *IEEE Transactions on Visualization and Computer Graphics* 17, 8 (2010), 1178–1190. 6, 12

- [YMO18] YU Y., MORA K. A. F., ODOBEZ J.-M.: HeadFusion: 360° head pose tracking combining 3D morphable model and 3D reconstruction. *IEEE transactions on pattern analysis and machine intelligence* 40, 11 (2018), 2653–2667. [11](#), [12](#)
- [YMYK14] YOSHIYASU Y., MA W.-C., YOSHIDA E., KANEHIRO F.: As-conformal-as-possible surface registration. *Computer Graphics Forum* 33, 5 (2014), 257–267. [2](#), [5](#), [6](#), [8](#), [12](#)
- [YYZ*14] YANG Y., YU Y., ZHOU Y., DU S., DAVIS J., YANG R.: Semantic parametric reshaping of human body models. In *International Conference on 3D Vision* (2014), vol. 2, IEEE, pp. 41–48. [21](#)
- [YZG*18] YU T., ZHENG Z., GUO K., ZHAO J., DAI Q., LI H., PONS-MOLL G., LIU Y.: DoubleFusion: Real-time capture of human performances with inner body shapes from a single depth sensor. In *IEEE Conference on Computer Vision and Pattern Recognition* (2018), pp. 7287–7296. [2](#), [3](#), [4](#), [5](#), [7](#), [10](#), [12](#)
- [YZG*21] YU T., ZHENG Z., GUO K., LIU P., DAI Q., LIU Y.: Function4D: Real-time human volumetric capture from very sparse consumer RGBD sensors. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 5746–5756. [11](#), [12](#)
- [YZZ*19] YU T., ZHENG Z., ZHONG Y., ZHAO J., DAI Q., PONS-MOLL G., LIU Y.: SimulCap: Single-view human performance capture with cloth simulation. In *IEEE Conference on Computer Vision and Pattern Recognition* (2019), pp. 5504–5514. [4](#), [12](#)
- [ZBYX19] ZHANG H., BO Z.-H., YONG J.-H., XU F.: InteractionFusion: real-time reconstruction of hand poses and deformable objects in hand-object interactions. *ACM Transactions on Graphics* 38, 4 (2019), 48:1–48:11. [5](#), [11](#), [12](#), [13](#)
- [ZCS*15] ZHANG R., CHEN X., SHIRATORI T., TONG X., LIU L.: An efficient volumetric method for non-rigid registration. *Graphical Models* 79 (2015), 1–11. [4](#), [5](#), [6](#), [8](#), [12](#)
- [ZFA21] ZAMPOGIANNIS K., FERMÜLLER C., ALOIMONOS Y.: Topology-aware non-rigid point cloud registration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43, 3 (2021), 1056–1069. [2](#), [3](#), [5](#), [6](#), [7](#), [8](#), [12](#)
- [Zho09] ZHONG Y.: Intrinsic shape signatures: A shape descriptor for 3D object recognition. In *IEEE International Conference on Computer Vision Workshops* (2009), IEEE, pp. 689–696. [8](#)
- [ZKJB17] ZUFFI S., KANAZAWA A., JACOBS D. W., BLACK M. J.: 3D menagerie: Modeling the 3D shape and pose of animals. In *IEEE Conference on Computer Vision and Pattern Recognition* (2017), pp. 5524–5532. [21](#)
- [ZQW*19] ZHAO J., QI X., WEN C., LEI N., GU X.: Automatic and robust skull registration based on discrete uniformization. In *IEEE/CVF International Conference on Computer Vision* (2019), pp. 431–440. [20](#)
- [ZQZ*21] ZENG Y., QIAN Y., ZHU Z., HOU J., YUAN H., HE Y.: CorrNet3D: Unsupervised end-to-end learning of dense correspondence for 3D point clouds. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 6052–6061. [9](#), [13](#)
- [ZSN*17] ZENG A., SONG S., NIESSNER M., FISHER M., XIAO J., FUNKHOUSER T.: 3DMatch: Learning local geometric descriptors from RGB-D reconstructions. In *IEEE Conference on Computer Vision and Pattern Recognition* (July 2017), IEEE Computer Society, pp. 199–208. [19](#)
- [ZST*10] ZHENG Q., SHARF A., TAGLIASACCHI A., CHEN B., ZHANG H., SHEFFER A., COHEN-OR D.: Consensus skeleton for non-rigid space-time registration. *Computer Graphics Forum* 29, 2 (2010), 635–644. [4](#), [8](#)
- [ZWG*16] ZENG Y., WANG C., GU X., SAMARAS D., PARAGIOS N.: Higher-order graph principles towards non-rigid surface registration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 38, 12 (Dec. 2016), 2416–2429. [20](#)
- [ZWW*10] ZENG Y., WANG C., WANG Y., GU X., SAMARAS D., PARAGIOS N.: Dense non-rigid surface registration using high-order graph matching. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (June 2010), pp. 382–389. [20](#)
- [ZX18] ZHOU S., XIAO S.: 3D face recognition: a survey. *Human-centric Computing and Information Sciences* 8, 1 (2018). [21](#)
- [ZYL*18] ZHENG Z., YU T., LI H., GUO K., DAI Q., FANG L., LIU Y.: HybridFusion: Real-time performance capture using a single depth sensor and sparse IMUs. In *Proceedings of the European Conference on Computer Vision* (2018), pp. 384–400. [10](#), [12](#)
- [ZYQ21] ZHANG J., YAO Y., QUAN L.: Learning signed distance field for multi-view surface reconstruction. In *IEEE/CVF International Conference on Computer Vision* (2021), pp. 6525–6534. [22](#)