The Visual Computer (2020) 36:1705–1721
https://doi.org/10.1007/s00371-019-01760-0

SURVEY

Recent advances in shape correspondence

Yusuf Sahillioğlu[1]

Published online: 28 September 2019
© Springer-Verlag GmbH Germany, part of Springer Nature 2019

Abstract
Important new developments have appeared since the most recent direct survey on shape correspondence published almost a decade ago. Our survey covers the period from 2011, their stopping point, to 2019, inclusive. The goal is to present the recent updates on correspondence computation between surfaces or point clouds embedded in 3D. Two tables summarizing and classifying the prominent, to our knowledge, papers of this period, and a large section devoted to their discussion lay down the foundation of our survey. The discussion is carried out in chronological order to reveal the distribution of various types of correspondence methods per year. We also explain our classification criteria along with the most basic solution examples. We finish with conclusions and future research directions.

Keywords Shape correspondence · Shape matching · Survey

1 Introduction

Shape correspondence problem is stated as finding a set of corresponding points between given shapes. In our scope, the shapes are deformable and they are discretizations embedded in 3D in the form of (i) surface mesh structures consisting of vertices, edges, and polygonal faces, or (ii) point clouds. While the common scenario dictates operation on two shapes in isolation, it is also possible to process multiple shapes at once.

This problem is interesting because finding correspondences between shapes is a fundamental operation in many computer graphics and digital geometry processing algorithms such as shape interpolation, shape reconstruction, shape retrieval, texture transfer, segmentation transfer, deformation transfer, symmetry detection, change detection, and statistical shape modeling.

This problem is difficult because it involves understanding the structure of the shapes at both the local and global levels. The understanding requires a search in a huge solution space for the deformable scenario that we address in this survey. Unlike the rigid shape correspondence scenario which deals with a space with low degree-of-freedom consisting only of rotation and translation, the deformable, aka non-rigid, scenario searches over the entire second shape to match a given point in the first shape. The second shape exhibits bending and possible stretching. The space is therefore huge and nonlinear in that we have $O(N!)$ possibilities of matching N points spread on both shapes. As in the case of all difficult problems, one needs to learn from experience by being aware of the existing solutions before developing his own, hence this survey.

Motivation The main motivation of this survey is to wrap up all the interesting and brand new correspondence methods since 2011, which is the endpoint of the closest work [104]. It is exciting to list all these new techniques based on well-established frameworks as well as fundamentally different approaches such as functional maps and deep neural networks. Collectionwise correspondence problem has also been introduced and extensively studied within the 2011–2019 scope of our survey.

2 Comparison with other surveys

Almost a decade later after the most recent shape correspondence survey by van Kaick et al. [104,105], we now present a survey that compactly lists and discusses the new work. We indeed pay special attention to cite shape correspondence works published on or after 2011, and refer the reader to STAR—State of The Art Report [104], or its journal version [105], for the older ones.

✉ Yusuf Sahillioğlu
ys@ceng.metu.edu.tr
https://www.ceng.metu.edu.tr/~ys

[1] Computer Engineering Department, METU, Ankara, Turkey

We also note a survey on mapping from 2014 [53] that deals mostly with parameterization and registration techniques, which constitute a subset of correspondence computation methods. Moreover, 87% of their references are published before 2011. The remaining ones also have no overlap with our references at all.

There is also a 2013 survey [98] that is devoted entirely to 3D surface registration, which is one of the many possible ways to compute correspondences as we will see in Sect. 3.5.1.

Surveys on shape retrieval [54,61] also relate to our survey as the evaluation of shape similarity, a requirement for the retrieval problem, can also be performed through correspondences. In fact, retrieval and correspondence problems are so related that the term *shape matching* is used interchangeably in the literature for these tasks. STAR [12], or its journal version [11], explores shape similarity assessment problem by analyzing and describing shapes using geometrical and topological properties of real functions.

In summary, our survey, similar to the closest work [105], overviews shape correspondence methods directly. In addition to presenting new discoveries since 2011, we also discuss more on parameterization-based and learning-based methods which are lacking in [105].

3 Classification criteria

In the sequel, we categorize shape correspondence methods based on several criteria. Tables 1 and 2 summarize the entire survey by positioning the methods of interest with respect to these criteria. We then elaborate on the criteria in the following subsections. Methods are discussed thoroughly in the guidance of these criteria in Sect. 4.

3.1 Similarity level

Two shapes between which correspondence is sought may be either fully similar or partially similar. For the former, although additional deformations may be (and are likely to be) admitted, it is certainly forbidden for one shape to possess a part that has no semantically meaningful counterpart on the other shape. If such a possession occurs, then we have the latter case to deal with, which is more challenging (Fig. 1).

The level of challenge increases further for the latter under arbitrary scaling of shapes. The former is not affected from this issue since scale can always be normalized prior to the matching process by using a global intrinsic property, such as maximum geodesic distance, that is certainly compatible between fully similar shapes.

3.2 Deformation type

Two shapes to be matched may differ by deformations most common of which is isometry. Under isometric deformations pairwise geodesic distances over the surface are preserved (Fig. 1a), a constraint simplifies the correspondence searching process. Isometry is an important clue for shape correspondence; not only since most real-world deformations are isometric, but also because semantically similar shapes have similar metric structures. Articulation/bending and rigid transformations are isometries.

As soon as we have stretching and/or squeezing involved, we have a non-isometric deformation, a mild example of which is depicted in Fig. 1b. Some severely non-isometric pairs are human versus gorilla and cat versus giraffe, where semantic similarity is preserved but pairwise geodesics are not.

3.3 Shape processing

The default setting in shape correspondence is matching two shapes in isolation, which we call pairwise processing. When shapes come as part of a collection, one may want to process all shapes at the same time, which is likely to bring more consistency and accuracy to the results (Fig. 2). We refer to this setting as collectionwise processing and list the alias terms as groupwise correspondence and multiple shape correspondence.

3.4 Output density

Shape correspondence methods aim to produce a mapping between some or all of the surface points of the two given shapes. Points can be continuous or discretized to mesh vertices. Continuous case, also referred to as sub-vertex resolution, uses barycentric coordinates of the enclosing triangle to define the match. If the output involves only some featured surface points, then we have a sparse correspondence, otherwise we have a dense one (Fig. 3). While sparse correspondence is quick and useful for initializing dense pipelines such as real-world scan registrations, dense correspondence is required for globally smooth attribute transfer and shape morphing applications.

3.5 Solution approach

Shape correspondence solutions can be investigated under three groups as follows.

3.5.1 Registration-based solution

Given two shapes, the registration-based scheme aligns them by either registering one shape to the other or registering both

Table 1 Shape correspondence methods classified according to the criteria of our work

Method	Criteria							
	Similarity level		Deformation type		Shape processing		Output density	
	Full	Partial	Isometric	Non-isometric	Pairwise	Collectionwise	Sparse	Dense
[46,60]	✓			✓	✓			✓
[70,87]	✓		✓		✓			✓
[23]	✓		✓		✓		✓	
[101]	✓		✓		✓			✓
[75]		✓	✓		✓		✓	
[102]		✓		✓	✓		✓	
[67]	✓		✓			✓		✓
[66,96]	✓			✓	✓			✓
[79,88]	✓		✓		✓		✓	
[89,103]		✓	✓		✓		✓	
[41,45]	✓			✓		✓		✓
[43,90]	✓		✓		✓			✓
[99,112]	✓		✓		✓		✓	
[72]	✓		✓		✓			✓
[57,74]		✓	✓		✓			✓
[27,73]	✓			✓	✓			✓
[39]	✓			✓		✓		✓
[48]	✓			✓		✓		✓
[59]	✓		✓		✓		✓	
[18,80]	✓			✓	✓			✓
[4,5]	✓			✓	✓			✓
[91,100]		✓	✓		✓			✓
[15,58]		✓	✓		✓			✓
[7,114]		✓		✓	✓		✓	
[92]	✓		✓			✓		✓
[40]		✓		✓		✓		✓
[93]	✓			✓		✓		✓
[47]	✓			✓		✓	✓	
[2,3]	✓			✓	✓			✓
[6]	✓			✓	✓			✓
[17,26]	✓		✓		✓			✓
[86,113]	✓			✓	✓			✓
[49,82]		✓		✓	✓			✓
[10,19]	✓		✓		✓			✓
[16,56]	✓		✓		✓			✓
[95]	✓		✓		✓			✓
[8,116]		✓		✓	✓		✓	
[97]	✓			✓	✓		✓	
[13]		✓		✓	✓			✓
[31]		✓		✓		✓	✓	
[22,32]		✓		✓	✓		✓	
[1,110]	✓		✓		✓			✓
[36,50]	✓		✓		✓			✓
[63]	✓		✓		✓			✓

Table 1 continued

Method	Criteria							
	Similarity level		Deformation type		Shape processing		Output density	
	Full	Partial	Isometric	Non-isometric	Pairwise	Collectionwise	Sparse	Dense
[55,109]	✓		✓		✓			✓
[28,107]	✓			✓	✓			✓
[81]		✓	✓		✓			✓
[69,76]	✓		✓		✓			✓
[106]		✓		✓	✓			✓
[33,62]	✓			✓	✓			✓
[64]	✓			✓	✓			✓
[30,111]	✓			✓	✓		✓	
[21]		✓	✓			✓	✓	
[84]	✓		✓		✓		✓	
[34,35]	✓		✓		✓			✓
[77,78]	✓			✓	✓			✓
[68,108]	✓		✓		✓			✓
[9,65]	✓			✓	✓			✓
[29]	✓			✓	✓			✓
[20,24]	✓			✓	✓			✓
[52]	✓			✓	✓			✓
[94]	✓			✓	✓			✓
[42]	✓			✓		✓		✓
[37]	✓		✓		✓			✓

✓ in the Sparse column means that the method produces a sparse correspondence from parts to parts instead of points to points. Please see Table 2 for continuation

shapes to a common intermediate domain. One can see the former operation as deformation and the latter as parameterization. Once registered, correspondence is derived from the proximity of the aligned shape elements.

The alternative approach is to keep each shape as is and derive the correspondence directly from the pointwise and/or pairwise *similarity* of elements (Sect. 3.5.2). One can argue that this approach is more popular than the registration-based one because it avoids extra deformation and parameterization processing power and errors. The argument is verified by the higher number of checkmarks in the corresponding column of Table 2.

Deformation We give a generic framework to deform one shape toward a data set, which would be the other shape in our case, while preserving its original geometric features. This demand is achieved by minimizing a combination of data and regularization energy terms:

$$E_{\text{def}}(\mathbf{v}) = E_{\text{data}}(\mathbf{v}) + \gamma E_{\text{regularization}}(\mathbf{v}) \tag{1}$$

where $\mathbf{v} \in \mathbb{R}^{n\times 3}$ stores 3 elements per row for each of the n vertices, namely the x-, y-, and z-coordinates. In the most basic Laplacian deformation model, Laplacian coordinate δ_i for vertex v_i encapsulates local geometric information around v_i by approximating mean curvature times the normal direction (Fig. 4a). This model tries the preserve original geometric features by minimizing the difference between the original and the deformed Laplacian coordinates as the shape is pulled toward the data points (Fig. 4b):

$$E_{\text{data}}(\mathbf{v}) = ||\mathbf{v} - \mathbf{c}||^2 \tag{2}$$

where $\mathbf{c} \in \mathbb{R}^{n\times 3}$ stores the data points to get close to, and

$$E_{\text{regularization}}(\mathbf{v}) = ||\mathbf{L}\mathbf{v} - \mathbf{L}\mathbf{v_0}||^2 \tag{3}$$

where $\mathbf{v_0} \in \mathbb{R}^{n\times 3}$ stores the original cooridantes and $\mathbf{L} \in \mathbb{R}^{n\times n}$ is the cotangent Laplacian matrix filled using

$$\mathbf{L}_{ij} = \begin{cases} \frac{-w_{ij}}{\sum_{j\in N(i)} w_{ij}} & v_i \text{ and } v_j \text{ are neighbors,} \\ 1 & i = j, \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

Discussing the shortcomings of this basic model is beyond the scope of this survey. We refer the reader to [85] for derivation and usage of more sophisticated deformation models,

Table 2 Complementing Table 1 with the remaining criteria for the same methods

Method	Criteria Cont'D										
	Solution approach						Speed			Surf topology	
	Based on		Learning-based		Automatic		Fast	Med	Slow	Arbitrary	Sphere
	Registration	Similarity	Yes	No	Fully	Semi					
[46,60]	✓			✓	✓		✓				✓
[70,87]		✓		✓	✓		✓			✓	
[23]		✓		✓	✓			✓		✓	
[101]		✓		✓	✓		✓			✓	
[75]		✓		✓	✓			✓		✓	
[102]		✓	✓		✓				✓	✓	
[67]		✓		✓	✓				✓	✓	
[66,96]		✓		✓	✓		✓			✓	
[79,88]		✓		✓	✓		✓			✓	
[89,103]		✓		✓	✓				✓	✓	
[41,45]		✓		✓	✓			✓			✓
[43,90]		✓		✓	✓		✓			✓	
[99,112]		✓		✓	✓		✓			✓	
[72]		✓		✓		✓	✓			✓	
[57,74]		✓		✓	✓		✓			✓	
[27,73]		✓		✓		✓	✓			✓	
[39]		✓		✓	✓			✓			✓
[48]	✓		✓			✓		✓		✓	
[59]		✓	✓		✓				✓	✓	
[18,80]		✓	✓		✓			✓		✓	
[4,5]	✓			✓		✓	✓				✓
[91,100]		✓		✓	✓			✓		✓	
[15,58]		✓		✓	✓				✓	✓	
[7,114]		✓		✓		✓			✓	✓	
[92]		✓		✓	✓		✓			✓	
[40]		✓		✓	✓				✓	✓	
[93]	✓			✓		✓		✓		✓	
[47]		✓	✓		✓				✓	✓	
[2,3]	✓			✓		✓	✓				✓
[6]	✓			✓		✓			✓		✓
[17,26]	✓			✓	✓		✓			✓	
[86,113]	✓			✓	✓		✓			✓	
[49,82]		✓		✓	✓		✓			✓	
[10,19]		✓		✓	✓		✓			✓	
[16,56]		✓	✓		✓			✓		✓	
[95]	✓			✓	✓			✓		✓	
[8,116]	✓			✓		✓			✓	✓	
[97]		✓		✓	✓				✓	✓	
[13]		✓	✓		✓			✓		✓	
[31]		✓		✓		✓		✓		✓	
[22,32]		✓		✓	✓				✓	✓	
[1,110]	✓			✓	✓		✓			✓	
[36,50]		✓		✓	✓				✓	✓	
[63]	✓			✓	✓				✓	✓	

Table 2 continued

Method	Criteria Cont'D										
	Solution approach						Speed			Surf topology	
	Based on		Learning-based		Automatic		Fast	Med	Slow	Arbitrary	Sphere
	Registration	Similarity	Yes	No	Fully	Semi					
[55,109]		✓	✓		✓				✓	✓	
[28,107]		✓		✓		✓			✓	✓	
[81]		✓		✓	✓			✓		✓	
[69,76]		✓		✓	✓		✓			✓	
[106]		✓		✓	✓			✓		✓	
[33,62]		✓		✓		✓	✓			✓	
[64]	✓		✓		✓			✓			✓
[30,111]		✓	✓		✓				✓	✓	
[21]		✓		✓	✓		✓			✓	
[84]		✓		✓	✓		✓			✓	
[34,35]	✓		✓		✓			✓		✓	
[77,78]		✓		✓	✓		✓			✓	
[68,108]		✓		✓	✓		✓			✓	
[9,65]		✓		✓	✓				✓	✓	
[29]		✓		✓		✓	✓			✓	
[20,24]	✓			✓	✓		✓			✓	
[52]	✓			✓	✓		✓				✓
[94]		✓		✓	✓		✓			✓	
[42]		✓	✓		✓			✓		✓	
[37]		✓	✓		✓			✓		✓	

In order to save space, two methods are merged into one row if they share the same answers to both question here and in Table 1

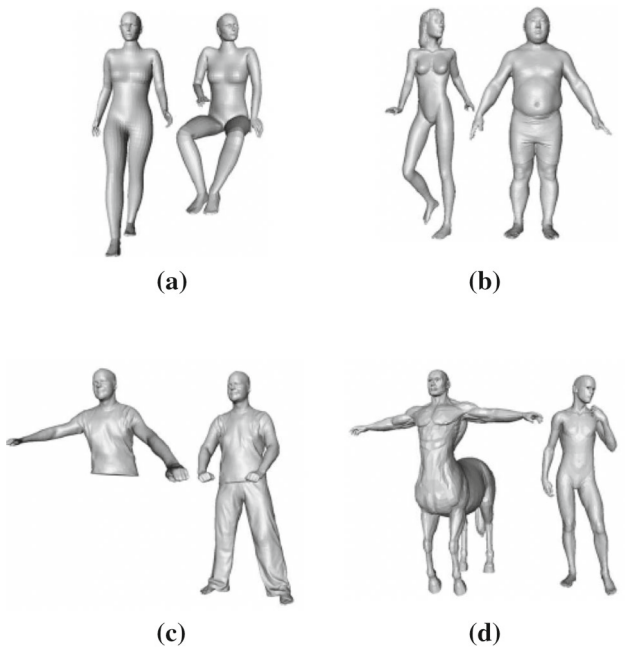

Fig. 1 Fully (**a**, **b**) and partially (**c**, **d**) similar shape pairs. **b**, **d** Involve mild non-isometric deformations between the shapes due to some stretching, whereas the others have purely isometric deformations

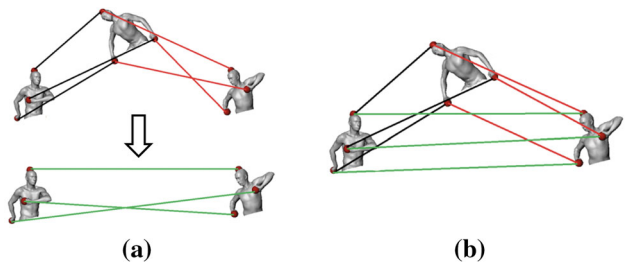

Fig. 2 Applying independent pairwise processings (**a**) lead to a worse total distortion than collectionwise processing does (**b**). If the distortions of the black and red maps are .061 and .063, respectively, then the implied green map at bottom becomes .069, hence a total distortion of .193 (**a**). If all shapes are considered at once, the red map slightly increases to .65, but the saving on the green map is larger, .060, hence a better total of .186 is achieved (**b**)

and to the survey [44] for a thorough overview on the subject.

Parameterization We give a generic framework to parameterize a shape embedded in 3D to a flat domain in 2D. Note that a shape correspondence method performs this process for two shapes and use a common intermediate domain that is not necessarily flat (Sect. 4.5).

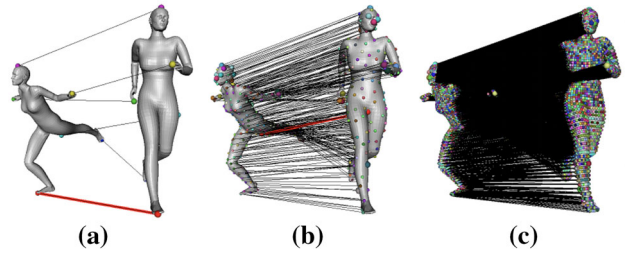

Fig. 3 Sparse (**a**, **b**) and dense (**c**) correspondences computed by Sahillioğlu and Yemez [87]

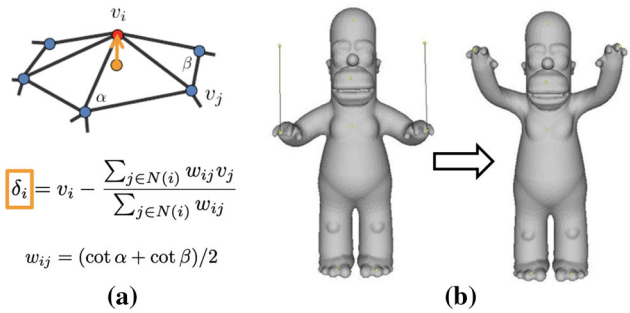

Fig. 4 Deformation in its most basic form using Laplacian coordinates (**a**). Data terms on two points at fingers under the Laplacian regularization perform the deformation (**b**)

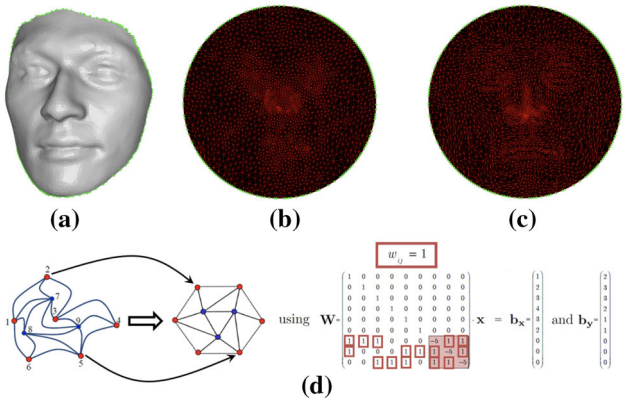

Fig. 5 Parameterization in its most basic form using uniform (**b**) and cotangent (**c**) weights to flatten the 3D surface (**a**) to a disk. In **d**, the toy 2D example at left is parameterized using the solution to the linear system at right

In the most basic disk parameterization scheme, boundary vertices of the surface (green ones in Fig. 5a) are first mapped to the boundary of a disk. Then, the remaining interior vertices are positioned in such a way that each one lands in the center (uniform—Fig. 5b) or weighted average (harmonic—Fig. 5c) of its immediate neighbors, the latter being more respectful to the input geometry. We can solve two linear systems for two coordinates separately to achieve this placement:

$$\mathbf{Wx} = \mathbf{b_x} \quad \text{and} \quad \mathbf{Wy} = \mathbf{b_y} \tag{5}$$

where $\mathbf{b_x} \in \mathbb{R}^{n\times 1}$ stores the x-coordinates of the boundary positions at its top k rows when there are k boundary vertices. Remaining $n-k$ rows, which are set to 0, move the interior vertices accordingly based on the multiplication of the bottom $n-k$ rows of $\mathbf{W} \in \mathbb{R}^{n\times n}$ with $\mathbf{x}$ (Fig. 5d). The y-coordinates are computed similarly.

$$\mathbf{W}_{ij} = \begin{cases} w_{ij} & v_i, v_j \text{ neighbs, } v_i \text{ interior,} \\ -\sum_{k\neq i} w_{ik} & i = j, v_i \text{ interior,} \\ 1 & i = j, v_i \text{ boundary,} \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

Discussing the shortcomings of this basic model and extending it to non-disk topologies are beyond the scope of this survey. We refer to [53] for a comprehensive reading on the subject.

3.5.2 Similarity-based solution

We now investigate the alternative scheme to the registration-based solution, which we call the similarity-based solution. Here, the geometry of the input shapes is not altered in any manner. Instead, we compute geometric invariants, or descriptors, under the appropriate deformation model. Such descriptors can be defined on vertices or between a pair of vertices, the latter being more distinctive and effective. A combination of both pointwise and pairwise terms leads to an energy function whose minimum gives the desired correspondence, or map, $\phi^* : S \to T$:

$$\begin{aligned} E(\phi) = w_1 \sum_{v_i} ||d^S(v_i) - d^T(\phi(v_i))|| \\ + w_2 \sum_{v_i, v_j} |e^S(v_i, v_j) - e^T(\phi(v_i), \phi(v_j))| \end{aligned} \tag{7}$$

where $d^S(.)$ and $e^S(.,.)$ denote the descriptor values on a vertex and between two vertices of the shape S, respectively. Descriptor choice depends on the deformation type, e.g., for isometric deformations a common choice for e is the geodesic distance (Fig. 6a). Once Eq. 7 or a similar energy function is ready, similarity-based solutions strive to minimize it via efficient optimization tools. Similarity of real-valued functions over surfaces (Fig. 6b) has also been a trending topic since 2012 (Sect. 4.2).

3.5.3 Learning-based solution

Recent popularization of deep learning techniques has affected the shape correspondence field significantly. We report a work as learning-based if it learns some sort of prior

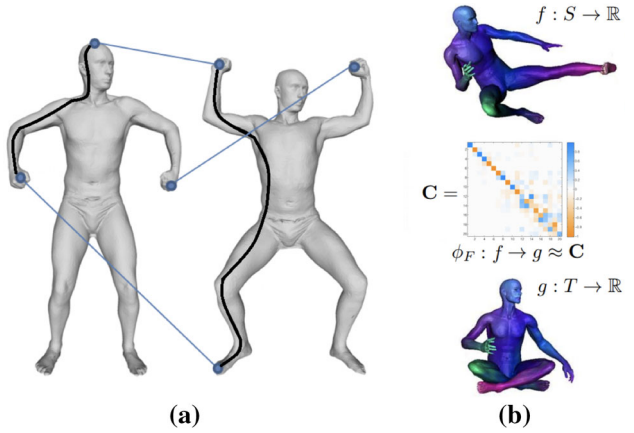

Fig. 6 A bad correspondence composed of three blue matches generates a high energy in Eq. 7 due to incompatible geodesics drawn as black paths (**a**). Matrix representation $\mathbf{C}$ of a functional map ϕ_F (middle) with colors proportional to matrix values. Colors on surfaces are based on function values (**b**)

knowledge from a training set. The speed of such methods is decided based on training and testing times. The corresponding value in Table 2 for a learning-based method is either medium or slow as the training time is naturally longer than the processing time of a non-learning-based method.

3.6 Automatic

A method requiring user interaction, such as manual landmark matches, is labeled as semi-automatic, whereas all others are considered as fully automatic.

3.7 Speed

We tag the methods as fast, medium, or slow according to their reported worst-case time complexities. If this information is missing, then we approximately normalize their reported execution times with a tagged method to make our decision. We also note that dense and sparse methods are treated within their classes, i.e., a dense method is considered fast even if it executes much slower than a sparse method.

3.8 Surface topology

Sometimes correspondence methods require the input surfaces to be at particular topologies, most common of which is genus-zero sphere topology. All other surfaces as well as point cloud inputs fall into the arbitrary topology category in Table 2.

4 The methods

We overview the methods in Tables 1 and 2 in the following subsections that go in chronological order from 2011 to 2019. The prominent papers of each year, in our opinion, are discussed with our criteria (Sect. 3) in mind.

4.1 Year 2011

Kim et al. [46] blend maps obtained by registering shapes to the extended complex plane by means of conformal Möbius transformations. Other than the topology restriction and geodesic centroid approximations that degrade their performance for particular cases, the method produces reliable non-isometric dense maps that are not necessarily onto, i.e., some vertices are left unmatched. Other methods address the isometric correspondence problem by minimizing different variants of Eq. 7 in coarse-to-fine fashion [87] (Fig. 3) or with an integer quadratic programming solver [23]. A coarse correspondence based on entropy-driven planned samples is made dense in one shot using a propagation strategy that respects a geodesic consistency criterion [101]. It is, however, unclear how to connect the decrease in entropy with the stability of shape correspondence. Ovsjanikov et al. [71], on the other hand, identify the samples that make the isometric shape correspondence problem easiest if their matches are known.

In the partial similarity setting, the correspondence-less approach in [75] optimizes the segment-wise similarity over the integration domains by relying on diffusion-based local shape descriptors. It produces part correspondences instead of point correspondences. Knowledge-driven [102] also computes part correspondences where, unlike [75], the input to their system is pre-segmented.

Given a collection of shapes and maps between all pairs, Nguyen et al. [67] compare the composite maps along 3-cycles to identity maps in order to enforce consistency, which in turn improves the input maps. The optimization takes plenty of time but works fine as long as the input maps are sufficiently accurate.

4.2 Year 2012

Ovsjanikov et al. [70] introduce the revolutionary functional map framework which replaces the traditional point-to-point correspondences by matching real-valued functions over surfaces. A point-to-point map $\phi : S \to T$ can also be expressed as a functional map $\phi_F : f \to g$ that associates values of functions $f : S \to \mathbb{R}$ and $g : T \to \mathbb{R}$. This alternative mapping can be concisely represented as a matrix $\mathbf{C} \in \mathbb{R}^{k_S \times k_T}$ (Fig. 6b) using k_S and k_T bases $\{\phi^S\}$ and $\{\phi^T\}$ for functions on shapes S and T, respectively. Namely, $\mathbf{Ca} \approx \mathbf{b}$ where $\mathbf{a}$

and **b** are the vectors of coefficients of f and g in the chosen bases, i.e., $f = \sum_{i=1}^{k_S} a_i \phi_i^S$ and $g = \sum_{i=1}^{k_T} b_i \phi_i^T$.

By providing enough number of **a**, **b** pairs, one can constrain **C** to have a unique solution which can be obtained through a linear solve. This framework is flexible in the sense that it gives user the freedom to choose the basis functions. A common choice for isometric matching is the first hundred Laplace–Beltrami eigenfunctions.

Soft maps [96], as another new map representation, provide a probabilistic relaxation of point-to-point correspondences. In fact, soft maps can be considered as special cases of functional maps, by using the hat function basis defined at each vertex. They basically map probability distributions, instead of functions, between surfaces. Unlike point-to-point maps, functional and soft map matrices are conveniently available for linear-algebraic analysis and manipulation. These frameworks are, however, prone to conversion errors when recovering the point-to-point correspondence from the optimal functional or soft mapping.

Alternative isometric matching frameworks proposed in this year minimize variants of Eq. 7 in probabilistic [88], game-theoretic [79], and combinatorial [89] settings. A non-isometric framework extrapolates the correspondence computed on the automatically extracted symmetry axis curves [60].

Based on initial pairwise maps between some shape pairs, Huang et al. [41] consider consistency of composite maps along 2- and 3-cycles of shapes in order to create soft maps [96] from a set of automatically extracted base shapes to all the shapes in the collection. These maps are also constrained to map neighboring points to neighboring points. Kim et al. [45] produce fuzzy correspondences in the spectral domain that are less accurate mainly because of the missing neighbor-preservation constraint in their system. Both approaches, however, yield much better results than [67] does, especially on datasets where initial maps have moderate or low quality. The main reason of this success is their flexibility to benefit from good parts of various maps, without being forced to use a particular map in its entirety during composition, a feature lacking in [67].

4.3 Year 2013

Huang and Guibas [39] present a theoretically sound collectionwise correspondence method by formulating the cycle-consistency constraint as the solution to a semidefinite program. They cast the problem of estimating cycle-consistent maps to finding the closest positive semidefinite matrix to an input matrix that stores all the initial maps. By jointly solving for segmentation, correspondence, and deformation, Kim et al. [48] achieves higher accuracy for each individual task using user-assisted template initialization.

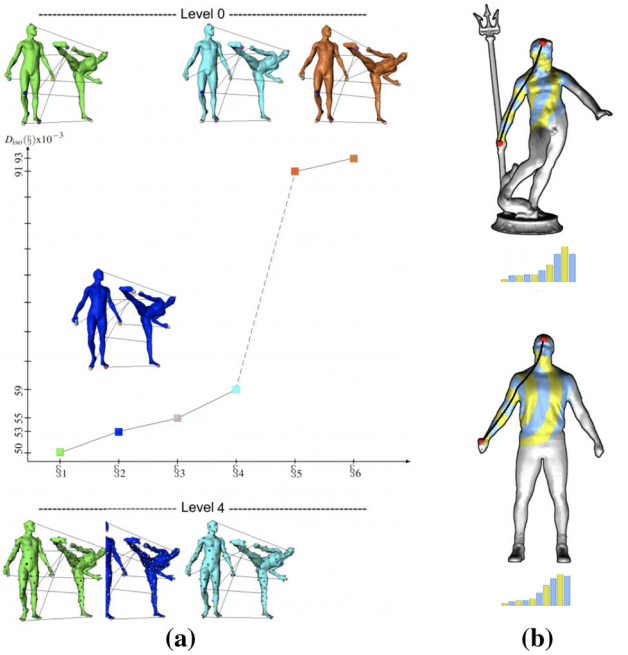

Fig. 7 The orange map after the first significant jump in the map distortion plot (dashed) is not tracked further in denser levels (**a**). The regions captured by two red feature points are compatible (top and bottom), hinting their correspondence. **a** and **b** are taken from [90] and [103], respectively

Symmetric flip problem, i.e., confusion of intrinsically symmetric features, that is inherent to isometric shape matching algorithms is addressed by tracking optimal coarse solutions through denser levels [90] (Fig. 7a) or using a harmonic symmetry-robust descriptor [112]. Ovsjanikov et al. [72] address this symmetric ambiguity problem by performing shape matching in an appropriate quotient space, where the symmetry has been identified and factored out.

The first non-trivial use of sparse models in shape correspondence is seen in [74] while matching partially isometric shapes at dense levels. van Kaick et al. [103] expand isotropic neighborhood around a feature point in an anisotropic way by defining regions of interest with two feature points. This definition facilitates partial matching since potentially extraneous regions of the models are selectively ignored (Fig. 7b). Weighted averages on surfaces [73] enable dense non-isometric correspondence. In particular, they first find the weights of the query point with respect to the anchors and then use these weights at the corresponding anchors on the other shape to complete the matching of the query.

4.4 Year 2014

Litman and Bronstein [59] present a learning scheme for the construction of optimized spectral descriptors for sparse isometric shape correspondence. Another learning-based approach [80] uses the output of their random forest classifier

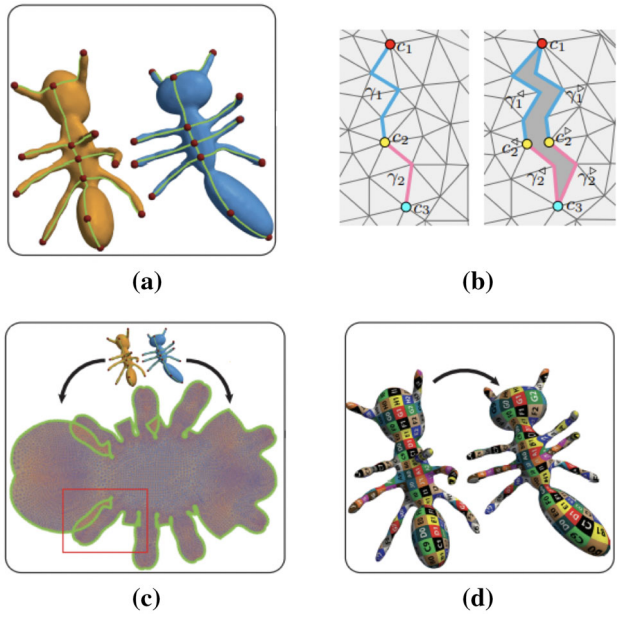

Fig. 8 Spherical surfaces are cut open (**c**) through the cuts designated by landmarks (**a**). Vertices and edges along the cut are duplicated to obtain disk topology (**b**). Flattenings are aligned using landmark correspondences with potential self-overlaps (red rectangle), an issue solved via lifting (**c**). **a**, **c**, **d**, and **b** are taken from [5] and [2], respectively

as a new descriptor that is tuned to the shapes and deformations appearing in their training set. Corman et al. [18], on the other hand, learn the most informative descriptors for isometric matching with a costly training that involves many reference shape pairs with known mappings between them. Non-isometric matching is achieved in [5] by jointly flattening the two surfaces after cutting them to disk topologies consistently (Fig. 8).

Partially isometric shapes are matched with minimum weight perfect matching on similar parts [92]. The same problem is handled in [100] by computing reliable small clusters of locally isometric point correspondences through diffusion pruning [99]. The latter proceeds in a more restricted setting as it requires a query part as input, whereas the former extracts similar parts from global shapes automatically. Both approaches iteratively propagate the computed sparse correspondences to dense ones. Brunton et al. [15] represent partial isometric maps based on equivalence classes of correspondences between pairs of points and their tangent-spaces. The idea is to recover a partial map by an iterative isometric growing procedure starting from sparse pre-matched features over the surfaces. Semi-automatic approaches establish part correspondences between structurally different models based on spatio-structural graphs [7] and relation graphs [114] built per shape.

All collectionwise correspondence efforts before [91] heavily rely on a given initial set of maps between all or some pairs of shapes, and instead of optimizing overall distortion, they rather enforce consistency. Sahillioğlu and Yemez [91] do not require any initial set of maps and explicitly minimize the isometric distortion over all possible pairs of collection shapes via dynamic programming efficiently. Huang et al. [40], on the other hand, construct consistent functional map networks that capture structural similarities within heterogeneous shape collections. Shapira and Ben-Chen [93] differ from other collectionwise methods by their assumption of having two homogeneous shape collections, with good maps within the collection, instead of having a single heterogeneous collection. This allows [93] to assume there exists common structure and use it to align the collections as a whole. To this end, they treat each shape collection as a point sampling from a low-dimensional shape space and use dimensionality reduction techniques based on the intrinsic shape difference distances [83] to embed this point cloud in Euclidean space where they can be aligned via registration. Affordance-based [47] fits a human agent to the models in the extremely diverse collections to establish consistent correspondences.

4.5 Year 2015

Aigerman et al. [2–4] compute harmonic parameterizations of the spherical surfaces into orbifolds, i.e., non-flat/curved surfaces, by minimizing general non-convex energies with L-BFGS optimization. This is in contrast with the flat case (Fig. 5) in which one tries to minimize an energy of a flattening of the surface into the Euclidean plane, where the energy being minimized is convex and quadratic, and hence a linear solve is sufficient. Flat parameterization-based method from the same group [6] improves their previous flat work [5] by bringing invariance to cut placement. This line of works generally requires a few manual landmark correspondences (Fig. 8a) to be able to align the resulting embeddings (Fig. 8c) for full correspondences (Fig. 8d). Instead of common intermediate parameterization domains, Chen and Koltun [17] and Sahillioğlu and Kavan [86] embed one shape to the curved shape of the other directly via convex optimization.

Kovnatsky et al. [49] make the functional maps resilient, to some extent, to missing parts and non-isometries by formulating the computation as a geometric matrix completion. Thanks to completion, this work is advantageous in settings when only scarce corresponding functions is available. Symmetric flip issue arising while converting functional maps to point maps is handled in [19] by computing optimal vector fields entirely within the functional map framework. Carrière et al. [16] map its provable stable multiscale topological signature to a vector, which in turn enables application of learning techniques for matching. A novel spectral embedding that exploits gradient fields provides robust isometric correspondence [95]. The deformation model in [8] allowing both geometric and topological operations, such as part split,

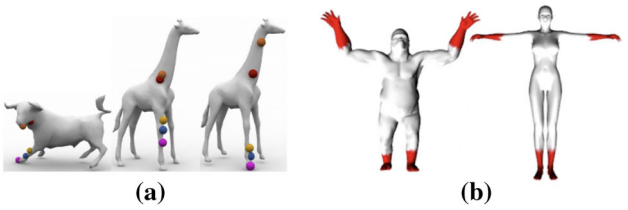

Fig. 9 Point-based (**a**) and part-based (**b**) non-isometric correspondences. Results on the first and second giraffe are obtained via [46] and [97], respectively. **a** and **b** are taken from [97] and [32], respectively

duplication, and merging, leads to part-based correspondence results via a pruned beam search, under the assumptions that input shapes are upright-oriented and pre-segmented.

4.6 Year 2016

Solomon et al. [97] optimize an entropy regularized version of the Gromov–Wasserstein objective function to produce a reasonable correspondence even when the surfaces undergo moderate geometric and topological changes that deviate from isometry. By minimizing a measure of stretch in its objective, Solomon et al. [97] handle non-isometries better than non-isometric methods that are restricted to conformal maps (Fig. 9a). Similar challenging cases are also handled by anisotropic convolutional neural networks [13], orbifold embeddings [4] (Sect. 4.5), and projective analysis [31]. Based on the key observation that points in matching parts often have similar ranks in the sorting of the corresponding feature values, Ganapathi-Subramanian et al. [32] match pairs of parts between two non-isometric shapes (Fig. 9b). Litany et al. [58] solve a non-rigid puzzle which is defined as find a segmentation of the model shape into parts corresponding to (subsets of) the query shapes. This segmentation, realized within the functional map framework, implies dense correspondence over partially similar shapes. Another natural extension to the functional map framework is to additionally compute a smooth interpolation between the given functions, which in turn leads to continuous isometric correspondence [10].

A low-dimensional embedding of shapes symmetrized with respect to the global reflectional symmetry plane produce dense correspondence that alleviates the symmetric flip problem [110]. Another embedding-based effort for dense correspondence treats geodesic distance matching problem in the dual intrinsic spectral domains of the given shapes [1]. Guo et al. [36] establishes isometric correspondence by discovering the articulated rigid parts of the point clouds. Maron et al. [63], on the other hand, use convex semidefinite programming relaxation for isometric matching. Although significantly faster than the standard relaxation, this method is able to deal with the registration of around one hundred points.

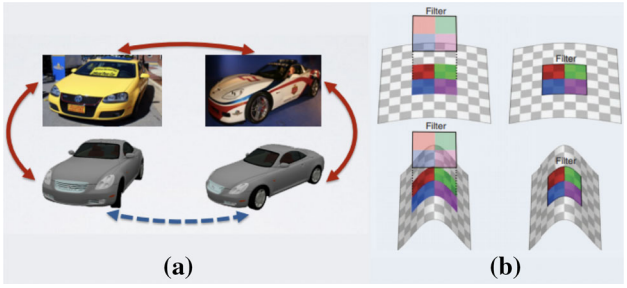

Fig. 10 Cycle of 4 edges (blue and reds) for consistent matches (**a**). Applying the filter to a deforming surface extrinsically (left) and intrinsically (right) gives different results, the latter being invariant to the deformation (**b**). **a** and **b** are taken from [115] and [14], respectively

Learning-based [109] provides impressive dense human body correspondences under large geometric and topological changes. Using a classical extrinsic convolutional neural network architecture to learn invariance to intrinsic pose changes is, however, very costly. In particular, they train on 50M examples in two weeks to create this huge network model that has about 200M parameters. Zhou et al. [115] first retrieve an appropriate 3D model to establish a correspondence 4-cycle and then train the network to minimize the discrepancy between the 3D ground-truth map (Fig. 10a, bottom) and the 3D composed map along the path (Fig. 10a, top) in order to estimate a pixel-wise map.

4.7 Year 2017

Rodolá et al. [82] improve the point-wise map recovery issue of functional maps via isometric alignment between the spectral embeddings of the two shapes to be matched. It is shown and exploited in [81] that the functional map matrix still has a meaningful structure when one of the two shapes has holes or missing parts (Fig. 11a). Consequently, Rodolá et al. [82] additionally improves partial correspondence and also mild non-isometric correspondence (Fig. 11b) as long as the initial functional map is sufficiently good. Alternative recovery methods treat input maps as corrupted versions of the latent correspondence and increase their quality by filtering [107], employ adjoint operators for coupled optimization over the forward and inverse maps [43], or deblur input maps based on a smoothness assumption [27]. The last one is able to produce matchings at sub-vertex resolution, which is more suitable for applications that transport high-frequency data between the shapes, e.g., texture mapping. Improving on [81], partial dense correspondence method [57] presents a purely spectral approach that allows to perform all calculations (except for the initial calculation of the first k Laplacian eigenfunctions) with constant complexity independent of the shape size. Yet another improvement to the functional map framework manages to inject descriptor preservation constraints

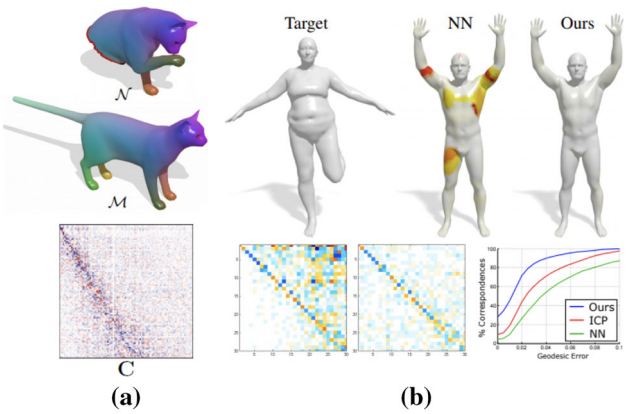

Fig. 11 A partial functional map from $\mathcal{N}$ to $\mathcal{M}$ has a slanted diagonal structure with an angle depending on the area ratio of two surfaces (**a**). A non-isometric functional map yields a point-wise map with the original (middle) and refined (right) recovery methods. Color encodes distance to the ground-truth, increasing from white to red (**b**). **a** and **b** are taken from [81] and [82], respectively

in addition to the original function preservation constraints [69]. To this end, geometric descriptors are represented as linear operators acting on functions through multiplication, rather than as simple scalar-valued signals. Acceleration and partiality improvements to [107] are brought by Vestner et al. [106] through a quadratic assignment problem matching pointwise and pairwise descriptors (Eq. 7) within the functional map framework. Given landmark correspondences or an extrinsic alignment, semi-automatic [62] uses soft maps to minimize the variance of a distribution obtained by matching each source point to a target point with a certain probability. A novel probabilistic biclustering formulation [22] reliably puts similar parts of non-isometric 3D shapes in correspondence.

Sequels to the existing partial correspondence [7,8] (Sects. 4.4 and 4.5) and non-isometric correspondence [3–6] (Sects. 4.4–4.6) methods are proposed by the same research groups. In particular, partial matching method [116] relaxes their previous requirement that asks the input segmentations to be perfect. Non-isometric method [2], on the other hand, complements their previous generalizations of Tutte's algorithm to orbifolds in different geometries. Another non-isometric registration pipeline is based on foliation [113].

Although it is not exceptional by computer vision standards, the size of training set and network in [109] (Sect. 4.6) is complex for computer graphics. Litany et al. [56] produces similar successful results but trains on merely 100 examples and models the network with orders of magnitude less parameters. The main idea is, as depicted in Fig. 10b, to define the convolution filter on the surface intrinsically so that its result does not change when the surface is deformed isometrically. To this effect, functional maps are integrated as differentiable layers into the intrinsic deep learning architecture. Another intrinsic convolutional neural network [64] provides a non-isometric correspondence as long as the shapes come in sphere topology. Using the differentiability feature of the Gromov–Wasserstein objective function utilized in the entropic matching framework of [30,97] models a deep neural network that maps unstructured geometric data to a regular domain. Although it is designed originally for matching local geometric features on real-world depth images, the descriptor learned from millions of correspondence labels found in existing RGB-D reconstructions turns out to be applicable to surface mesh correspondence problem without any modification [111].

Collections with a large presence of outliers are matched consistently by formulating the problem as a series of quadratic programs with sparsity-inducing constraints [21].

4.8 Year 2018

Sahillioğlu [84] observes that genetic algorithm fits well to the sparse isometric correspondence problem where one essentially tries to come up with the best permutation of sample indices in a discrete setting. His method explores the space of permutations wisely without visiting every possibility via specially designed genetic operators. Learning-based [55] establishes dense isometric correspondence by serializing the local neighborhood of vertices. In other words, it avoids commonly used resampling operations that encode neighborhood information in a structured and regular manner while processing unstructured surface meshes. Encoder–decoder deep network architecture in [35], on the other hand, predicts isometric shape correspondence by learning a consistent surface parameterization with a shared template.

Topological [76] and orientation [78] constraints are incorporated into the functional maps framework to promote continuous maps. In addition to continuity, resilience to topological noise and non-isometry is brought as a by-product of the former and latter, respectively. Based on the insight that functional map optimization depends only on inner products between descriptors rather than descriptor values themselves, Wang et al. [108] use kernel functions to efficiently evaluate and preserve such products [68]. Interactivity is also brought to the functional map framework by allowing the user to specify corresponding feature curves on two non-isometric shapes. Feature curve preservation constraints are then incorporated to the framework along with an efficient numerical method to optimize the map with immediate feedback [33], resulting in accurate maps between semantically similar but geometrically different shapes such as horse and elephant. Melzi et al. [65] inject an evolution process scheme into the functional map framework by introducing a multi-scale descriptor per vertex for the functional constraint. By encoding the structure of geodesic neighborhoods of a point across multiple scales, this descriptor enables robust maps

in the presence of mild non-isometric deformations, including missing parts, and topological noise. Another localized spectral analysis, namely the localized manifold harmonics [66], enables faster map computation in the same setting.

4.9 Year 2019

Ezuz et al. [29] formulate a non-isometric correspondence optimization problem which aims to minimize an energy that tries to preserve a given landmark correspondence or functional correspondence while penalizing the Dirichlet energies of both the forward and backward maps. The similar reversibility energy using forward and backward maps is combined with a membrane energy that penalizes area distortion and a bending energy that penalizes misalignment of curvature features in order to obtain more accurate non-isometric maps [28]. The membrane and elastic energies, however, lead to a slower method than [29]. Another non-isometric method [20] deforms one shape with the goal of aligning its Laplacian eigenvalues with the other one's since these values provide a powerful characterization of the shape's geometry. The alignment of the spectra of two shapes makes them more intrinsically isometric and thus facilitates finding an accurate correspondence using an existing isometric technique [25,26,50]. This correspondence then implies the non-isometric map using the identity map between the original and the deformed shapes. Deformation is also used in [52] in order to conformally map each genus-zero shape to the sphere, where the rigid alignment and optical flow produce the correspondence. Dyke et al. [24], on the other hand, deform one shape to the surface of the other under moderate stretching by incorporating an anisotropic deformation estimation into their iterative registration pipeline. Non-isometric correspondence is also obtained by Azencot et al. [9] which iteratively improves the given imperfect initial map by leveraging information from the inverse map. Non-isometric correspondence is finally considered within the functional maps framework by regularizing functional maps based on the resolvent Laplacian commutativity term [77], or by developing a hierarchical pipeline for inference [94].

Learning-based [42] merges two heterogeneous shape collections, where the maps within collection are known, to generate non-isometric functional correspondences for all cross-collection shape pairs. One can argue that the supervised training regime is prohibitive in terms of the amount of the manually annotated data required. Halimi et al. [37] address this issue by proposing the first completely unsupervised learning for dense shape correspondence. The solution is based on a purely geometric criterion that drives the learning process toward distortion-minimizing predictions. Groueix et al. [34] utilize 3-cycle-consistency to define a notion of good correspondences in triplets of objects and uses it as a supervisory signal to train its deformation network in the absence of point correspondence supervision. The deformation network then takes a 3D surface point from one shape and outputs the associated deformed point, which is projected onto the nearest point on the surface of the other shape in order to establish correspondence.

5 Conclusions and future research directions

We can draw the following conclusions based on Tables 1 and 2, and the discussions in Sect. 4. Future research directions are also given along with the items below.

- New compact map representations, namely functional maps [70] and soft maps [96], have proved successful in this decade. In addition to pursuing improvements on each representation separately, it is also a promising direction to combine these functional and probabilistic approaches for alternative representations.
- Learning-based correspondence solutions have proved successful in the time span of our survey. While pure unsupervised approaches, such as [37], can replace massive labeling burden of supervised methods, it is also worth considering combined semi-supervised learning schemes for better accuracy.
- Our chronological ordering in Sect. 4 reveals the uniform distribution of various types of correspondence methods per year except the collectionwise type, which seems to get less attention recently (9 methods in 2011–2014 vs. 3 methods in 2015–2019). Given the increasing availability of digitized 3D models, attention should return to the collectionwise correspondence methods. In particular, novel methods can be developed to avoid repetition of the entire optimization when new models are added to the collection. Automatic augmentation of the shapes in collection to compensate for the missing parts is another direction for future research.
- Building on [47], interaction-based functionality analysis [38] can be investigated further, e.g., with non-human agents, to obtain consistent correspondences through very diverse collections.
- The link between the stability of shape matching and shape symmetries is exploited by only a few methods [19,60,72,90,110,112], and further investigation is necessary.
- Despite the efforts in [15,49,51,57,65,66,75,76,80,97,103,109], finding a meaningful correspondence between shapes under topological noise remains a challenge and should be considered carefully since this type of noise is mostly inevitable while capturing 3D data.
- Lastly, it would be interesting to study the guarantees and the behaviors of correspondence optimization algorithms.

Acknowledgements This work has been supported by TUBITAK under the Project EEEAG-115E471.

References

1. Aflalo, Y., Dubrovina, A., Kimmel, R.: Spectral generalized multi-dimensional scaling. Int. J. Comput. Vis. **118**(3), 380–392 (2016)
2. Aigerman, N., Kovalsky, S., Lipman, Y.: Spherical orbifold Tutte embeddings. ACM Trans. Graph. (Proc. SIGGRAPH) **36**(4), 90 (2017)
3. Aigerman, N., Lipman, Y.: Orbifold Tutte embeddings. ACM Trans. Graph. **34**(6), 190–1 (2015)
4. Aigerman, N., Lipman, Y.: Hyperbolic orbifold Tutte embeddings. ACM Trans. Graph. (Proc. SIGGRAPH Asia) **35**(6), 217 (2016)
5. Aigerman, N., Poranne, R., Lipman, Y.: Lifted bijections for low distortion surface mappings. ACM Trans. Graph. **33**(4), 69 (2014)
6. Aigerman, N., Poranne, R., Lipman, Y.: Seamless surface mappings. ACM Trans. Graph. **34**(4), 72 (2015)
7. Alhashim, I., Li, H., Xu, K., Cao, J., Ma, R., Zhang, H.: Topology-varying 3D shape creation via structural blending. ACM Trans. Graph. (TOG) **33**(4), 158 (2014)
8. Alhashim, I., Xu, K., Zhuang, Y., Cao, J., Simari, P., Zhang, H.: Deformation-driven topology-varying 3D shape correspondence. ACM Trans. Graph. (TOG) **34**(6), 236 (2015)
9. Azencot, O., Dubrovina, A., Guibas, L.: Consistent shape matching via coupled optimization. Comput. Graph. Forum **38**(5), 13–25 (2019)
10. Azencot, O., Vantzos, O., Ben-Chen, M.: Advection-based function matching on surfaces. Comput. Graph. Forum **35**(5), 55–64 (2016)
11. Biasotti, S., Cerri, A., Bronstein, A., Bronstein, M.: Recent trends, applications, and perspectives in 3D shape similarity assessment. Comput. Graph. Forum **35**(6), 87–119 (2016)
12. Biasotti, S., Cerri, A., Bronstein, A.M., Bronstein, M.M.: Quantifying 3D shape similarity using maps: recent trends, applications and perspectives. Eurographics (State of the Art Reports), pp. 135–159 (2014)
13. Boscaini, D., Masci, J., Rodolà, E., Bronstein, M.: Learning shape correspondence with anisotropic convolutional neural networks. In: Advances in Neural Information Processing Systems, pp. 3189–3197 (2016)
14. Bronstein, M.M., Bruna, J., LeCun, Y., Szlam, A., Vandergheynst, P.: Geometric deep learning: going beyond euclidean data. IEEE Signal Process. Mag. **34**(4), 18–42 (2017)
15. Brunton, A., Wand, M., Wuhrer, S., Seidel, H.-P., Weinkauf, T.: A low-dimensional representation for robust partial isometric correspondences computation. Graph. Models **76**(2), 70–85 (2014)
16. Carrière, M., Oudot, S.Y., Ovsjanikov, M.: Stable topological signatures for points on 3D shapes. Comput. Graph. Forum **34**(5), 1–12 (2015)
17. Chen, Q., Koltun, V.: Robust nonrigid registration by convex optimization. In: Proceedings of International Conference on Computer Vision, pp. 2039–2047 (2015)
18. Corman, E., Ovsjanikov, M., Chambolle, A.: Supervised descriptor learning for non-rigid shape matching. In: European Conference on Computer Vision, pp. 283–298 (2014)
19. Corman, E., Ovsjanikov, M., Chambolle, A.: Continuous matching via vector field flow. Comput. Graph. Forum **34**(5), 129–139 (2015)
20. Cosmo, L., Panine, M., Rampini, A., Ovsjanikov, M., Bronstein, M.M., Rodolà, E.: Isospectralization, or how to hear shape, style, and correspondence. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 7529–7538 (2019)
21. Cosmo, L., Rodolà, E., Albarelli, A., Mémoli, F., Cremers, D.: Consistent partial matching of shape collections via sparse modeling. Comput. Graph. Forum **36**(1), 209–221 (2017)
22. Denitto, M., Melzi, S., Bicego, M., Castellani, U., Farinelli, A., Figueiredo, M.A., Kleiman, Y., Ovsjanikov, M.: Region-based correspondence between 3d shapes via spatially smooth biclustering. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 4260–4269 (2017)
23. Dubrovina, A., Kimmel, R.: Approximately isometric shape correspondence by matching pointwise spectral features and global geodesic structures. Adv. Adapt. Data Anal. (AADA) **3**, 203–228 (2011)
24. Dyke, R., Lai, Y.-K., Rosin, P., Tam, G.K.: Non-rigid registration under anisotropic deformations. Comput. Aided Geom. Des. **71**, 142–156 (2019)
25. Dyke, R., Stride, C., Lai, Y., Rosin, P., Aubry, M., Boyarski, A., Bronstein, A., Bronstein, M., Cremers, D., Fisher, M., Groueix, T., Guo, D., Kim, V., Kimmel, R., Lahner, Z., Li, K., Litany, O., Remez, T., Rodolà, E., Russel, B., Sahillioğlu, Y., Slossberg, R., Tam, G., Vestner, M., Wu, Z., Yang, J.: Shape correspondence with isometric and non-isometric deformations. In: Eurographics Workshop on 3D Object Retrieval (2019)
26. Eisenberger, M., Lahner, Z., Cremers, D.: Divergence-free shape correspondence by deformation. Comput. Graph. Forum **38**(5), 1–12 (2019)
27. Ezuz, D., Ben-Chen, M.: Deblurring and denoising of maps between shapes. Comput. Graph. Forum **36**(5), 165–174 (2017)
28. Ezuz, D., Heeren, B., Azencot, O., Rumpf, M., Ben-Chen, M.: Elastic correspondence between triangle meshes. Comput. Graph. Forum **38**(2), 121–134 (2019)
29. Ezuz, D., Solomon, J., Ben-Chen, M.: Reversible harmonic maps between discrete surfaces. ACM Trans. Graph. **38**(2), 15 (2019)
30. Ezuz, D., Solomon, J., Kim, V., Ben-Chen, M.: Gwcnn: a metric alignment layer for deep shape analysis. Comput. Graph. Forum (Proc. SGP) **36**(5), 49–57 (2017)
31. Fish, N., van Kaick, O., Bermano, A., Cohen-Or, D.: Structure-oriented networks of shape collections. ACM Trans. Graph. **35**(6), 171 (2016)
32. Ganapathi-Subramanian, V., Thibert, B., Ovsjanikov, M., Guibas, L.: Stable region correspondences between non-isometric shapes. Comput. Graph. Forum **35**(5), 121–133 (2016)
33. Gehre, A., Bronstein, M., Kobbelt, L., Solomon, J.: Interactive curve constrained functional maps. Comput. Graph. Forum **37**(5), 1–12 (2018)
34. Groueix, T., Fisher, M., Kim, V.G., Russell, B., Aubry, M.: Unsupervised cycle-consistent deformation for shape matching. Comput. Graph. Forum **38**(5), 123–133 (2019)
35. 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 (ECCV), pp. 230–246 (2018)
36. Guo, H., Zhu, D., Mordohai, P.: Correspondence estimation for non-rigid point clouds with automatic part discovery. Vis. Comput. **32**(12), 1511–1524 (2016)
37. Halimi, O., Litany, O., Rodola, E., Bronstein, A.M., Kimmel, R.: Unsupervised learning of dense shape correspondence. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4370–4379 (2019)
38. Hu, R., Savva, M., van Kaick, O.: Functionality representations and applications for shape analysis. Comput. Graph. Forum **37**(2), 603–624 (2018)
39. Huang, Q., Guibas, L.: Consistent shape maps via semidefinite programming. Comput. Graph. Forum (Proc. SGP) **32**(5), 177–186 (2013)

40. Huang, Q., Wang, F., Guibas, L.: Functional map networks for analyzing and exploring large shape collections. ACM Trans. Graph. (TOG) **33**(4), 36 (2014)
41. Huang, Q., Zhang, G., Gao, L., Hu, S., Bustcher, A., Guibas, L.: An optimization approach for extracting and encoding consistent maps in a shape collection. ACM Trans. Graph. (TOG) Proc. SIGGRAPH Asia **31**(6), 167 (2012)
42. Huang, R., Achlioptas, P., Guibas, L., Ovsjanikov, M.: Limit shapes—a tool for understanding shape differences and variability in 3D model collections. Comput. Graph. Forum **38**(5), 187–202 (2019)
43. Huang, R., Ovsjanikov, M.: Adjoint map representation for shape analysis and matching. Comput. Graph. Forum **36**(5), 151–163 (2017)
44. Jacobson, A., Deng, Z., Kavan, L., Lewis, J.P.: Skinning: real-time shape deformation. In: ACM SIGGRAPH 2014 Courses, vol. 24 (2014)
45. Kim, V., Li, W., Mitra, N., DiVerdi, S., Funkhouser, T.: Exploring collections of 3D models using fuzzy correspondences. Proc. SIGGRAPH **31**, 54-1 (2012)
46. Kim, V., Lipman, Y., Funkhouser, T.: Blended intrinsic maps. ACM Trans. Graph. **30**(4), 79 (2011)
47. Kim, V.G., Chaudhuri, S., Guibas, L., Funkhouser, T.: Shape2pose: human-centric shape analysis. ACM Trans. Graph. (TOG) **33**(4), 120 (2014)
48. Kim, V.G., Li, W., Mitra, N.J., Chaudhuri, S., DiVerdi, S., Funkhouser, T.: Learning part-based templates from large collections of 3D shapes. ACM Trans. Graph. (TOG) **32**(4), 70 (2013)
49. Kovnatsky, A., Bronstein, M., Bresson, X., Vandergheynst, P.: Functional correspondence by matrix completion. In: Proceedings of Computer Vision and Pattern Recognition, pp. 905–914 (2015)
50. Küpçü, E., Yemez, Y.: Diffusion-based isometric depth correspondence. Comput. Vis. Image Underst. (2019) **(in press)**
51. Lahner, Z., Rodolà, E., Bronstein, M., Cremers, D., Burghard, O., Cosmo, L., Dieckmann, A., Klein, R., Sahillioğlu, Y.: Shrec'16: matching of deformable shapes with topological noise. In: Proceedings of Eurographics Workshop on 3D Object Retrieval (2016)
52. Lee, S., Kazhdan, M.: Dense point-to-point correspondences between genus-zero shapes. Comput. Graph. Forum **38**(5), 27–37 (2019)
53. Li, X., Iyengar, S.S.: On computing mapping of 3D objects: a survey. ACM Comput. Surv. **47**(2), 34:1–34:45 (2014)
54. Lian, Z., Godil, A., Bustos, B., Daoudi, M., Hermans, J., Kawamura, S., Kurita, Y., Lavoué, G., Van Nguyen, H., Ohbuchi, R., et al.: A comparison of methods for non-rigid 3D shape retrieval. Pattern Recognit. **46**(1), 449–461 (2013)
55. Lim, I., Dielen, A., Campen, M., Kobbelt, L.: A simple approach to intrinsic correspondence learning on unstructured 3D meshes. In: Proceedings of the European Conference on Computer Vision (ECCV) (2018)
56. Litany, O., Remez, T., Rodolà, E., Bronstein, A.M., Bronstein, M.M.: Deep functional maps: structured prediction for dense shape correspondence. In: Proceedings of International Conference on Computer Vision (2017)
57. Litany, O., Rodolà, E., Bronstein, A., Bronstein, M.: Fully spectral partial shape matching. Comput. Graph. Forum **36**(2), 247–258 (2017)
58. Litany, O., Rodolà, E., Bronstein, A., Bronstein, M., Cremers, D.: Non-rigid puzzles. Comput. Graph. Forum (Proc. SGP) **35**(5), 135–143 (2016)
59. Litman, R., Bronstein, A.M.: Learning spectral descriptors for deformable shape correspondence. IEEE Trans. Pattern Anal. Mach. Intell. **36**(1), 171–180 (2014)
60. Liu, T., Kim, V.G., Funkhouser, T.: Finding surface correspondences using symmetry axis curves. Comput. Graph. Forum (Proc. SGP) **31**(5), 1607–1616 (2012)
61. Liu, Z.-B., Bu, S.-H., Zhou, K., Gao, S.-M., Han, J.-W., Wu, J.: A survey on partial retrieval of 3D shapes. J. Comput. Sci. Technol. **28**(5), 836–851 (2013)
62. Mandad, M., Cohen-Steiner, D., Kobbelt, L., Alliez, P., Desbrun, M.: Variance-minimizing transport plans for inter-surface mapping. ACM Trans. Graph. (TOG) **36**(4), 39 (2017)
63. Maron, H., Dym, N., Kezurer, I., Kovalsky, S., Lipman, Y.: Point registration via efficient convex relaxation. ACM Trans. Graph. **35**(4), 73 (2016)
64. Maron, H., Galun, M., Aigerman, N., Trope, M., Dym, N., Yumer, E., Kim, V.G., Lipman, Y.: Convolutional neural networks on surfaces via seamless toric covers. ACM Trans. Graph. **36**(4), 71–1 (2017)
65. Melzi, S., Ovsjanikov, M., Roffo, G., Cristani, M., Castellani, U.: Discrete time evolution process descriptor for shape analysis and matching. ACM Trans. Graph. (TOG) **37**(1), 4 (2018)
66. Melzi, S., Rodolà, E., Castellani, U., Bronstein, M.M.: Localized manifold harmonics for spectral shape analysis. Comput. Graph. Forum **37**(6), 20–34 (2018)
67. Nguyen, A., Ben-Chen, M., Welnicka, K., Ye, Y., Guibas, L.: An optimization approach to improving collections of shape maps. Comput. Graph. Forum (Proc. SGP) **30**(5), 1481–1491 (2011)
68. Nogneng, D., Melzi, S., Rodolà, E., Castellani, U., Bronstein, M., Ovsjanikov, M.: Improved functional mappings via product preservation. Comput. Graph. Forum **37**(2), 179–190 (2018)
69. Nogneng, D., Ovsjanikov, M.: Informative descriptor preservation via commutativity for shape matching. Comput. Graph. Forum (Proc. Eurographics) **36**(2), 259–267 (2017)
70. Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L.: Functional maps: a flexible representation of maps between shapes. ACM Trans. Graph. (TOG) Proc. SIGGRAPH **31**, 30 (2012)
71. Ovsjanikov, M., Huang, Q., Guibas, L.: A condition number for non-rigid shape matching. Comput. Graph. Forum (Proc. SGP) **30**(5), 1503–1512 (2011)
72. Ovsjanikov, M., Mérigot, Q., Pătrucean, V., Guibas, L.: Shape matching via quotient spaces. In: Proceedings of the Eleventh Eurographics/ACMSIGGRAPH Symposium on Geometry Processing, pp. 1–11 (2013)
73. Panozzo, D., Baran, I., Diamanti, O., Sorkine-Hornung, O.: Weighted averages on surfaces. ACM Trans. Graph. (TOG) Proc. SIGGRAPH **32**, 60 (2013)
74. Pokrass, J., Bronstein, A., Bronstein, M., Sprechmann, P., Sapiro, G.: Sparse modeling of intrinsic correspondences. Comput. Graph. Forum **32**(2), 459–468 (2013)
75. Pokrass, J., Bronstein, A.M., Bronstein, M.M.: A correspondence-less approach to matching of deformable shapes. In: Proceedings of Scale Space and Variational Methods (2011)
76. Poulenard, A., Skraba, P., Ovsjanikov, M.: Topological function optimization for continuous shape matching. Comput. Graph. Forum **37**(5), 13–25 (2018)
77. Ren, J., Panine, M., Wonka, P., Ovsjanikov, M.: Structured regularization of functional map computations. Comput. Graph. Forum **38**(5), 39–53 (2019)
78. Ren, J., Poulenard, A., Wonka, P., Ovsjanikov, M.: Continuous and orientation-preserving correspondences via functional maps. In: SIGGRAPH Asia 2018 Technical Papers, p. 248 (2018)
79. Rodolà, E., Bronstein, A., Albarelli, A., Bergamasco, F., Torsello, A.: A game-theoretic approach to deformable shape matching. In: Proceedings of Computer Vision and Pattern Recognition, pp. 182–189 (2012)
80. Rodolà, S. Bulo, E., Windheuser, T., Vestner, M., Cremers, D.: Dense non-rigid shape correspondence using random forests. In:

Proceedings of Computer Vision and Pattern Recognition, pp. 4177–4184 (2014)
81. Rodolà, E., Cosmo, L., Bronstein, M.M., Torsello, A., Cremers, D.: Partial functional correspondence. Comput. Graph. Forum **36**(1), 222–236 (2017)
82. Rodolà, E., Möller, M., Cremers, D.: Regularized pointwise map recovery from functional correspondence. Comput. Graph. Forum **36**(8), 700–711 (2017)
83. Rustamov, R., Ovsjanikov, M., Azencot, O., Ben-Chen, M., Chazal, F., Guibas, L.: Map-based exploration of intrinsic shape differences and variability. ACM Trans. Graph. **32**(4), 72 (2013)
84. Sahillioğlu, Y.: A genetic isometric shape correspondence algorithm with adaptive sampling. ACM Trans. Graph. (TOG) **37**(5), 175 (2018)
85. Sahillioğlu, Y.: A shape deformation algorithm for constrained multidimensional scaling. Comput. Graph. **53**, 156–165 (2015)
86. Sahillioğlu, Y., Kavan, L.: Skuller: a volumetric shape registration algorithm for modeling skull deformities. Med. Image Anal. **23**(1), 15–27 (2015)
87. Sahillioğlu, Y., Yemez, Y.: Coarse-to-fine combinatorial matching for dense isometric shape correspondence. Comput. Graph. Forum (Proc. SGP) **30**(5), 1461–1470 (2011)
88. Sahillioğlu, Y., Yemez, Y.: Minimum-distortion isometric shape correspondence using EM algorithm. IEEE Trans. PAMI **34**(11), 2203–2215 (2012)
89. Sahillioğlu, Y., Yemez, Y.: Scale normalization for isometric shape matching. Comput. Graph. Forum (Proc. Pac. Graph.) **31**(7), 2233–2240 (2012)
90. Sahillioğlu, Y., Yemez, Y.: Coarse-to-fine isometric shape correspondence by tracking symmetric flips. Comput. Graph. Forum **32**(1), 177–189 (2013)
91. Sahillioğlu, Y., Yemez, Y.: Multiple shape correspondence by dynamic programming. Comput. Graph. Forum (Proc. Pac. Graph.) **33**(7), 121–130 (2014)
92. Sahillioğlu, Y., Yemez, Y.: Partial 3D correspondence from shape extremities. Comput. Graph. Forum **33**(6), 63–76 (2014)
93. Shapira, N., Ben-Chen, M.: Cross-collection map inference by intrinsic alignment of shape spaces. Comput. Graph. Forum **33**(5), 281–290 (2014)
94. Shoham, M., Vaxman, A., Ben-Chen, M.: Hierarchical functional maps between subdivision surfaces. Comput. Graph. Forum **38**(5), 55–73 (2019)
95. Shtern, A., Kimmel, R.: Spectral gradient fields embedding for nonrigid shape matching. Comput. Vis. Image Underst. **140**, 21–29 (2015)
96. Solomon, J., Nguyen, A., Butscher, A., Ben-Chen, M., Guibas, L.: Soft maps between surfaces. Comput. Graph. Forum **31**(5), 1617–1626 (2012)
97. Solomon, J., Peyre, G., Kim, V., Sra, S.: Entropic metric alignment for correspondence problems. ACM Trans. Graph. **35**(4), 72 (2016)
98. Tam, G., Cheng, Z., Lai, Y., Langbein, F., Liu, Y., Marshall, D., Martin, R., Sun, X., Rosin, P.: Registration of 3D point clouds and meshes: a survey from rigid to non-rigid. IEEE Trans. Vis. Comput. Graph. **19**(7), 1199–1217 (2013)
99. Tam, G.K., Martin, R.R., Rosin, P.L., Lai, Y.: Diffusion pruning for rapidly and robustly selecting global correspondences using local isometry. ACM Trans. Graph. **33**(1), 1–17 (2014)
100. Tam, G.K., Martin, R.R., Rosin, P.L., Lai, Y.-K.: An efficient approach to correspondences between multiple non-rigid parts. Comput. Graph. Forum **33**, 137–146 (2014)
101. Tevs, A., Berner, A., Wand, M., Ihrke, I., Seidel, H.-P.: Intrinsic shape matching by planned landmark sampling. Comput. Graph. Forum (Proc. Eurographics) **30**(2), 543–552 (2011)
102. van Kaick, O., Tagliasacchi, A., Sidi, O., Zhang, H., Cohen-Or, D., Wolf, L., Hamarneh, G.: Prior knowledge for part correspondence. Comput. Graph. Forum **30**(2), 553–562 (2011)
103. van Kaick, O., Zhang, H., Hamarneh, G.: Bilateral maps for partial matching. Comput. Graph. Forum **32**(6), 189–200 (2013)
104. van Kaick, O., Zhang, H., Hamarneh, G., Cohen-Or, D.: A survey on shape correspondence. In: Proceedings of Eurographics State-of-the-Art Report, pp. 1–24 (2010)
105. van Kaick, O., Zhang, H., Hamarneh, G., Cohen-Or, D.: A survey on shape correspondence. Comput. Graph. Forum **30**(6), 1681–1707 (2011)
106. Vestner, M., Lähner, Z., Boyarski, A., Litany, O., Slossberg, R., Remez, T., Rodola, E., Bronstein, A., Bronstein, M., Kimmel, R., et al.: Efficient deformable shape correspondence via kernel matching. In: 2017 International Conference on 3D Vision (3DV), pp. 517–526 (2017)
107. Vestner, M., Litman, R., Rodolà, E., Bronstein, A., Cremers, D.: Product manifold filter: non-rigid shape correspondence via kernel density estimation in the product space. In: Proceedings of Computer Vision and Pattern Recognition (2017)
108. Wang, L., Gehre, A., Bronstein, M.M., Solomon, J.: Kernel functional maps. Comput. Graph. Forum **37**(5), 27–36 (2018)
109. 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)
110. Yoshiyasu, Y., Yoshida, E., Guibas, L.: Symmetry aware embedding for shape correspondence. Comput. Graph. **60**, 9–22 (2016)
111. Zeng, A., Song, S., Nießner, M., Fisher, M., Xiao, J., Funkhouser, T.: 3DMATCH: learning local geometric descriptors from RGB-D reconstructions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1802–1811 (2017)
112. Zhang, Z., Yin, K., Foong, K.: Symmetry robust descriptor for non-rigid surface matching. Comput. Graph. Forum **32**(7), 355–362 (2013)
113. Zheng, X., Wen, C., Lei, N., Ma, M., Gu, X.: Surface registration via foliation. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 938–947 (2017)
114. Zheng, Y., Cohen-Or, D., Averkiou, M., Mitra, N.J.: Recurring part arrangements in shape collections. Comput. Graph. Forum **33**(2), 115–124 (2014)
115. Zhou, T., Krahenbuhl, P., Aubry, M., Huang, Q., Efros, A.A.: Learning dense correspondence via 3D-guided cycle consistency. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 117–126 (2016)
116. Zhu, C., Yi, R., Lira, W., Alhashim, I., Xu, K., Zhang, H.: Deformation-driven shape correspondence via shape recognition. ACM Trans. Graph. **36**(4), 51 (2017)

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

Yusuf Sahillioğlu is an associate professor in the Department of Computer Engineering at Middle East Technical University. His research interests include digital geometry processing and computer graphics. He has a Ph.D. in computer science from Koç University, Turkey, where he also obtained his M.S. degree. His second computer science M.S. degree is from University of Florida. He received the B.S. degree from the Computer Engineering department of Bilkent University, Turkey. He also worked as a post-doc in the Computer Graphics Lab of University of Pennsylvania.