

ORCA - Online Research @ Cardiff

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository:https://orca.cardiff.ac.uk/id/eprint/134584/

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

Citation for final published version:

Dyke, Roberto M., Lai, Yu-Kun, Rosin, Paul, Zappala, Stefano, Dykes, Seana, Guo, Daoliang, Li, Kun, Marin, Riccardo, Melzi, Simone and Yang, Jingyu 2020. SHREC'20: Shape correspondence with non-isometric deformations. Computers and Graphics 92, pp. 28-43. 10.1016/j.cag.2020.08.008

Publishers page: http://dx.doi.org/10.1016/j.cag.2020.08.008

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.




FISEVIER

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag



Special Section on 3DOR 2020

SHREC'20: Shape correspondence with non-isometric deformations ®



Roberto M. Dyke^{a,1,*}, Yu-Kun Lai^{a,1}, Paul L. Rosin^{a,1}, Stefano Zappalà^a, Seana Dykes^a, Daoliang Guo^b, Kun Li^b, Riccardo Marin^c, Simone Melzi^d, Jingyu Yang^b

- ^a School of Computer Science & Informatics, Cardiff University, UK
- ^b Tianjin University, Tianjin, China
- ^c Department of Computer Science, University of Verona, Italy
- ^d Department of Computer Science, Sapienza University of Rome, Italy

ARTICLE INFO

Article history:
Received 15 May 2020
Revised 26 August 2020
Accepted 28 August 2020
Available online 5 September 2020

Keywords: Non-isometric deformation Shape correspondence Surface registration Benchmark dataset

ABSTRACT

Estimating correspondence between two shapes continues to be a challenging problem in geometry processing. Most current methods assume deformation to be near-isometric, however this is often not the case. For this paper, a collection of shapes of different animals has been curated, where parts of the animals (e.g., mouths, tails & ears) correspond yet are naturally non-isometric. Ground-truth correspondences were established by asking three specialists to independently label corresponding points on each of the models with respect to a previously labelled reference model. We employ an algorithmic strategy to select a single point for each correspondence that is representative of the proposed labels. A novel technique that characterises the sparsity and distribution of correspondences is employed to measure the performance of ten shape correspondence methods.

© 2020 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/)

1. Introduction

With a decade passing since the release of the first Kinect, the last decade has seen a large increase in the number of low-cost 3D capturing devices available. As well as bespoke 3D scanning hardware, a combination of improvements in software solutions for photogrammetry and the pervasiveness of high quality cameras has enabled the creation of vast amounts of 3D data. With the increasing amount of new data captured, demand for methods that provide greater automation to understand relationships between shapes is increasing.

Accurately identifying correspondences between two or more surfaces automatically continues to be a challenging and relevant problem. It provides a basis to enable further analysis and applications in a variety of areas. As discussed by van Kaick et al. [1], the problem of shape retrieval is closely related to shape correspondence, as the correspondence between two shapes may be used to measure their similarity.

Recently, the problem of non-isometric shape correspondence has become increasingly popular. Strictly isometric and near-isometric deformation has been well studied. With state-of-the-art methods [2–4] achieving superior performance in current non-isometric scenarios, there is presently an absence of valuable benchmark datasets for non-isometric shape correspondence.

Additionally, in real-world scenarios where real objects are scanned, existing capturing techniques induce natural geometric errors (e.g., noise, self-occlusions, and fusion between parts—causing topological changes). The limitations and errors exhibited vary according to the particular scanning technique employed. Typically, most benchmark datasets consist of scans from one scanning source. When evaluating a method's performance this makes it unclear how well a method's performance may transfer to other technologies. Recently, Melzi et al. [5] published a dataset that sought to address the issue of incompatibilities between meshes, which arise when working with scans from multiple sources.

For this dataset, we have compiled a small database of quadruped shapes. Establishing correspondence between animals poses a pertinent challenge that—with shapes exhibiting extreme non-isometries—is not currently considered. As discussed in Section 2.1, most existing datasets address the problem of correspondence between humans, which has quite limited applications. From the perspective of comparative anatomy, being able to also establish correspondence with other mammalian

This article has been certified as Replicable by the *Graphics Replicability Stamp Initiative*: http://www.replicabilitystamp.org

^{*} Corresponding author.

E-mail addresses: dykerm@cardiff.ac.uk (R.M. Dyke), laiy4@cardiff.ac.uk (Y.-K. Lai), rosinpl@cardiff.ac.uk (P.L. Rosin).

¹ Track organisers.




Fig. 1. A simple example of a homologous part between quadrupeds. The coloured area represents the matching part for the hind leg. Whilst inter-species, it is possible to intuitively recognise that these areas correspond.

vertebrates—focused on here—may be considered a generalisation of the human correspondence problem.

Previous work [6] has discovered deficiencies when reporting the performance of methods. Current error measurements fail to capture valuable quantitative information about the sparsity and distribution of correspondences. This is further discussed in Section 2.2.

Although topological variations exist, due to the common ancestry of tetrapod mammals, many parts are considered *homologous structures*—to correspond. Fig. 1 illustrates the homologous region of the hind leg between quadrupeds, all comprising primarily of a femur, tibia, fibula, and metatarsal. Therefore establishing a valid correspondence automatically is possible.

Methods capable of accurately finding correspondences between different mammals enable further avenues of research, e.g., statistical, behavioural analysis [7], and generative models [8].

For zoologists that use morphometrics—the study and development of techniques for the quantitative measurement of organisms—sparse, manually placed correspondences between animals are required to conduct statistical shape analysis.

Contribution The main contributions of this work are as follows:

- Generation of a novel dataset of quadrupeds with sparse ground-truth correspondences labelled by three specialists.
- Development of a new measure to evaluate the coverage of correspondences on a shape's surface—discussed further in Section 5.
- Systematic evaluation of the performance of a selection of recent shape correspondence methods, with additional quantitative insights into performance from our novel measure.

Organisation This report is organised as follows: Section 2 discusses previous works on quadruped benchmarks and their relation to current human body research, as well as discussing correspondence evaluation techniques. Section 3 describes the contents of the dataset, as well as specifying the acquisition techniques used to capture each object. Section 4 describes the correspondence methods and parameters used for this dataset. Section 5 describes the measures used to evaluate the performance of methods, including our novel measure of correspondence coverage. In Section 6 results are presented and discussed. Finally, Section 7 contains concluding thoughts arising from the outputs of this work.

2. Related work

2.1. Correspondence datasets

There has been a great focus on anthropometric—the measurement of humans—surface deformation [9–14], with the many

datasets produced opening up avenues to conduct further research. Because the field of anthropometry may be considered to generalise to the field of morphometry, we shall discuss existing datasets in both.

The notable FAUST dataset [14] contains a total of 300 real scans of 10 humans in 30 poses captured using the 3dMD-body.u System by 3dMD. Subjects were covered in sparse markers to enable shapes to be registered using a novel texture-based technique, which ensured quality alignment in areas with little geometric detail. Ground-truth correspondences between different individuals have been established, with a subset publicly released for training purposes. This was subsequently extended in a dataset that captured human body motion [15], containing 40,000 scans.

Vlasic et al. [12] propose a technique that uses multiple monocular cameras to capture a sequence of images of a human's performance from multiple angles. At each time step the image from each camera is segmented to separate the background from the actor. This is used to produce a silhouette. A template mesh rigged with a skeleton is used in combination with the silhouettes to reconstruct the human's pose. This approach enables all models using the same template to have the same connectivity and thus have dense correspondence. 10 performances have been published, in which sequences consist of between 150 and 250 watertight meshes.

The CAESAR dataset [9] is one of the largest human body datasets. 4431 subjects were scanned in North America, the Netherlands, and Italy using laser scanning. 72 stickers were placed on each subject for use as landmarks; due to initial capturing limitations 110 subjects from the European subset do not have landmarks. Subjects were scanned in three poses: standing, sitting comfortably, and sitting with arms raised.

The CAESAR dataset is not publicly available without purchase of a license. The licensing and copyright of content is an issue present across computer vision, especially in research where it may be unclear as to whether one's work is considered to be for commercial or non-commercial purposes. Whilst still a grey area, subsequent human body datasets have been derived from the dataset [13,16–18]. Loper et al. [17] and Zuffi et al. [8] both use data that is not considered to be in the public domain to develop linear blend skinning algorithms to construct a model of the data. These have been subsequently used to produce synthetic datasets for evaluative and training purposes as ground-truths are easily established [19].

Several synthetic datasets have been derived from the Digital Art Zone (Daz Productions, Inc. or Daz 3D). Daz 3D is a digital model platform, as well as a software development company. They have produced a series of *base* (i.e., template) models which are rigged and include morphs to alter the appearance of a model (e.g., emaciated, muscular, etc.). Many datasets have been derived from these models [18,20,21]. Models that use the same template share the same connectivity, and therefore possess dense ground-truths.

Kim et al. [22] combined three existing datasets [11,20,23], a subset of the shapes are animals, for which a volunteer selected 21 corresponding points on the quadruped shapes. The dataset contains 51 quadruped shapes. Meshes contain between approximately 3,000 and 56,000 faces.

Other datasets that contain a subset of quadruped shapes [24,25] provide ground-truth correspondences between different shapes of the same class but not different mammals, limiting the degree of non-isometric deformation that can be quantitatively evaluated.

Previous SHREC tracks have used datasets that contain non-isometric deformation [5,6], however the degree of non-isometry exhibited is significantly more subtle. This work contains strictly highly non-isometric deformations. Furthermore, Dyke et al. [6]

comprise of scans captured solely with one commercial device, while our new dataset spans a variety of sources. This introduces additional challenging geometric and topological variations.

2.2. Correspondence measures

We shall begin with a few definitions. Let C be a set of correspondences between the source surface X and the target surface Y. For a predicted correspondence between source and target surfaces $c_i = (x_i, y_i) \in X \times Y$, the respective ground-truth correspondence is $(x_i, y_i^*) \in X \times Y$. The distance between the predicted point y_i and the ground-truth point y_i^* is measured using distance function $d_Y(y_i, y_i^*)$, which may represent either the Euclidean or geodesic distance.

Often geodesic distance, or a normalised variant of it, is used to directly measure performance. The average geodesic error may be formulated as

$$E = \frac{1}{|C|} \sum_{i}^{|C|} d_{Y}(y_{i}, y_{i}^{*}).$$
 (1)

Many previous SHREC tracks have principally used geodesic distance as a measure for error [5,26-28]. Note that when measuring the overall error for a collection of target shapes, it may be necessary to normalise the computed error by the surface's properties such as the area of Y or farthest geodesic on Y.

Measures that require each predicted correspondence to be assigned a binary classification $b(y_i) \in \{0, 1\}$ as either *true positive* (TP) or *true negative* (TN), e.g., precision, recall, and specificity, rely on an appropriate classification strategy $b(y_i) = B(y_i, y_i^*)$. A popular strategy is to measure the distance between the predicted and ground-truth correspondence points $d_Y(y_i, y_i^*)$, and points that are below a specified error threshold ϵ are considered correct. This description is the basis of the popular correspondence error measure proposed by Kim et al. [22], where a correspondence is considered to be a TP when $d_Y(y_i, y_i^*) \le \epsilon$. The value of ϵ is increased to measure the number of TP over larger radii, which can be used to produce a curve. This approach fails to characterise the distribution and sparsity of correspondences on shapes where a limited number of ground-truth correspondences are available.

The benchmark protocol described by Kim et al. [22] has been considered the standard error measure for correspondences. The normalised geodesic error may be used to produce further statistics through the use of the area under the curve [6]. For a collection of shapes, Kim et al. [22] also report the average of the maximal geodesic error. Rodolà et al. [29] and Cosmo et al. [30] report the average geodesic error over a dataset of shapes with gradually reducing surface areas to measure the robustness of methods on increasingly partial scans. For functional mapping approaches, Corman et al. [31] measure the quality of ground-truth and predicted functional basis. In the case of registration methods, where one shape is deformed to align with another, it is possible to measure fitting error using the Hausdorff distance [32,33]. van Kaick et al. [1] discuss a variety of other validation methods for shape correspondences.

Qualitative techniques using visual mappings between two shapes in which topological information is transported (e.g., texture transfer) [4,34] are also used. However, these techniques are not an effective way to succinctly summarise the performance of a method on larger datasets or for comparing the performance of multiple methods. Further evaluation may be done by using the proposed algorithm in an application that requires a correspondence mapping (e.g., shape retrieval [35], consistent quadrangulation [4,36,37]).

3. Dataset

For this track we have identified a set of synthetic models and real-world scans of 3D shapes, specifically four-legged animals, and produced a set of ground-truth correspondences. Shapes have been mended to remove major errors such as self-intersecting faces and handles which cause erroneous high genera. Ground-truth correspondences were acquired by asking specialists in geometry processing and animal studies to label the shapes manually using a bespoke labelling tool (see Fig. 2).

Because the dataset includes real-world scans, many of the shapes contain geometric inconsistency and topological change caused by self-contact. The real-scans also contain natural noise, varying triangulation and self-occluded geometry. Some examples of challenging cases are shown in Fig. 3.

The dataset contains 14 models that have been acquired using a variety of techniques (see Table 2). Because the dataset is limited to quadruped mammals, many regions share a similar shape or function, it is therefore possible to establish correspondences between homologous loci with a reasonable degree of accuracy. While the dataset size might initially be considered to be quite small, for the purposes of computing and evaluating corresponding pairs, there are P(14,2)=182 permutations of shape pairs—or 149 pairs when excluding full-to-partial pairs. Our benchmark experiment participants were asked to complete a subset of these pairs comprising of matching pairs of full-to-full and partial-to-full models (Table 1).

The ground-truths for this dataset are acquired using the originally sourced mesh. Three specialists labelled corresponding points on each shape based on a template shape that had initially been labelled with markers. For each point, multiple experts propose a correspondence on the surface and a consensus was found by selecting the medoid point. Approximately 50 marker positions were initially selected on the rhinoceros. The rhinoceros was selected as the template since, although it was reconstructed from multi-view camera array, the shape was subsequently corrected by a professional CGI artist.

For the benchmark, where models have an exceedingly high triangle count, the mesh is simplified to 100,000 triangles. Participants could also submit results using a low-resolution version of the meshes with 20,000 triangles that were also made available.

Ground-truth correspondences were not made available to participants during the track and were solely reserved for evaluative purposes.

3.1. Test sets

Pairs of scans were carefully selected to ensure the non-isometry present in each test-set gradually increased. A description of the contents of each test-set may be found in Table 3.

3.2. Initial correspondences

For many shape correspondence and registration algorithms, a sparse set of correspondences is required for initialisation. A set of high quality sparse correspondences enables subsequent automatic refinement of the estimated non-rigid deformation. However, a poor set of initial correspondences may cause the algorithms to fail. For the purposes of establishing correspondences automatically, it is important to select a robust initial correspondence strategy.

To produce a set of candidate correspondences, SHOT signatures [39] at two radii (2% and 5% of the square root of the total triangle area) and IWKS [40]—a spectral descriptor—were examined, as well as a combination of SHOT and IWKS used together. SHOT was found to produce the most correct correspondences. A




Fig. 2. A screenshot of the software used by specialists to annotate corresponding points between shapes. The rhino on the left was initially labelled, and used as a reference for subsequent animals.




Fig. 3. Illustrations of some of the challenges in our dataset (a) partial scans (green indicates the boundary), (b) significant non-isometric deformations between pairs of models, and (c) topological inconsistencies: inherent to the original object or caused by scanning limitations. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

spectral pruning method proposed by Tam et al. [41] was used to remove noisy candidates and produce a set of globally consistent correspondences. For this method to work optimally input geometry must be locally isometric, however, this was rarely the case in our dataset. Due to memory limitations and computation time, correspondences were computed with the default parameters, except K=5 (which specifies how many initial correspondences are found for each point in the source mesh) and d=0.25 (which corresponds to the local neighbourhood size in diffusion pruning).

4. Correspondence methods

4.1. Baseline N-ICP

Bouaziz and Pauly [42] describe a naïve non-rigid registration method. The method computes an initial set of correspondences using nearest neighbours, and formulates a data term with point-to-point and point-to-plane metrics. Deformations are regularised by global and local rigidity measures. Local regularisation uses the as-rigid-as-possible formulation, proposed by Sorkine and Alexa

Table 1Summary of models contained in the dataset, as well as a description of the acquisition method used.

Source	Acquisition method	Model name	Original no. of faces	
Sketchfab.com	Metashape by Agisoft ¹	bison*	49,978	
	NextEngine 3D Scanner ²	leopard*	179,680	
	EinScan-S 3D Scanner by SHINING 3D3	giraffe_a*	2,045,246	
	COLMAP ¹ with 20 Canon G16 cameras	rhino*	64,174	
	Sprout Pro by HP ³	pig*	488,892	
	Metashape by Agisoft ¹	elephant_a*	19,998	
	"	dog*	141,698	
	Recap360 by Autodesk1 with an iPhone	hippo#	249,253	
AIM@SHAPE-VISIONAIR	Synthetic	cow*	14,506	
	Unknown	giraffe_b*	18,474	
	Kemelmacher and Basri [38] ⁴	bear#	31,309	
	"	camel_b#	11,698	
	Synthetic	camel_a*	19,536	
	Konica-Minolta Vivid 910 ²	elephant_b*	3,024,588	

¹photogrammetry software ²laser triangulation ³structured light ⁴photometric stereo

Table 2 Description of models remeshed to 20,000 (low-resolution) and 100,000 (high-resolution) faces. L_2 - & L_{∞} -norm errors with respect to diagonal length of each shapes bounding box are reported.

Model name Cleaned no.		High-resolution	model		Low-resolution model			
of faces	No. of faces	L ₂ error (%)	L_{∞} error (%)	No. of faces	L ₂ error (%)	L_{∞} error (%)		
bear	31,309	31,309	0	0	19,999	0.0036	0.0020	
bison	49,978	49,978	0	0	20,000	0.0199	0.2298	
camel_a	19,528	19,528	0	0	19,528	0	0	
camel_b	11,580	11,580	0	0	11,580	0	0	
cow	13,866	13,866	0	0	13,866	0	0	
dog	14,1698	100,000	0.0027	0.0027	20,000	0.0197	0.2305	
elephant_a	19,998	19,998	0	0	19,998	0	0	
elephant_b	3,024,548	100,000	0.0059	0.0590	20,000	0.0219	0.1583	
giraffe_a	1,325,332	100,000	0.0037	0.0329	20,000	0.0168	0.1184	
giraffe_b	18,474	18,474	0	0	18,474	0	0	
hippo	153,714	100,000	0.0026	0.0692	20,000	0.0174	0.2702	
leopard	179,616	100,000	0.0023	0.0498	20,000	0.0175	0.1443	
pig	488,328	100,000	0.0498	0.0214	20,000	0.0147	0.1229	
rhino	63,848	63.848	0	0	20,000	0.0147	0.1872	

Table 3Summary of shape pairs in each respective test-set. All source models in test-set 0 are partial scans and all target models are a full scan. All pairs of models in test-set 1 to 4 are full-to-full scans.

Test-set no.	Source model name	Target model name	Test-set no.	Source model name	Target model name
0	camel_b	camel_a	2	dog	leopard
	hippo	rhino		pig	leopard
	bear	cow		leopard	cow
	camel_b	cow		dog	pig
	bear	leopard		cow	dog
	bear	dog	3	camel_a	cow
	hippo	cow		dog	camel_a
	hippo	dog		rhino	cow
	bear	giraffe_a		pig	elephant_a
	camel_b	elephant_a		bison	elephant_b
1	giraffe_a	giraffe_b	4	rhino	elephant_a
	elephant_a	elaphant_b		cow	elephant_a
	camel_a	giraffe_a		elephant_a	giraffe_a
	giraffe_b	camel_a		cow	giraffe_b
	cow	bison		dog	giraffe_a

[43]. An updated set of correspondences is estimated using the new nearest neighbours based on the present shape deformation. Shapes were registered after being resized by their total triangle area. The terms are combined as the following energy minimisation problem:

$$E_{\text{reg}}(\mathbf{R}, t) = \text{w1}E_{\text{point-to-plane}} + \text{w2}E_{\text{point-to-point}}$$

$$+ \text{w3}E_{\text{global rigidity}} + \text{w4}E_{\text{as-rigid-as-possible}}. \tag{2}$$

The following parameters were used: w1=1, w2=1, w3=1, w4=1000, iter=100.

4.2. Non-rigid registration under anisotropic deformations

Dyke et al. [2] propose a two-stage iterative registration framework. In the first stage a correspondence mapping between surfaces is estimated by applying a variant of non-rigid ICP with an r-ring as-rigid-as-possible constraint for regularisation of larger

^{*} full/watertight mesh # partial mesh

Table 4A taxonomy of the methods evaluated using the described benchmark, complementing the work of Sahillioğlu [54].

Method	Criteria										
	Similarity level		Deformation type		Shape processing		Output density				
	Full	Partial	Isometric	Non-isometric	Pairwise	Collectionwise	Sparse	Dense			
R3DS Wrap 3	✓		✓		✓			✓			
Ezuz and Ben-Chen [46]	✓			✓	✓			✓			
Rodolà et al. [29]	✓		\checkmark		✓			✓			
Ren et al. [47]	✓			✓	✓			✓			
Bouaziz and Pauly [42]	✓		\checkmark		✓			✓			
Li et al. [44]	✓		\checkmark		✓			✓			
Vestner et al. [45]	✓			✓	✓			✓			
Melzi et al. [52]	✓			✓	✓			✓			
Marin et al. [50]	✓			\checkmark	✓			✓			
Method	Criteria conti	nued									
	Solution approach						Speed			Surface top	oology
	Based on		Learning		Automatio	:	Fast	Med	Slow	Arbitrary	Sphere
	Registration	Similarity	Yes	No	Fully	Semi					
R3DS Wrap 3	✓			✓	✓		√			√	
Ezuz and Ben-Chen [46]		✓		✓	✓		✓			\checkmark	
Rodolà et al. [29]		✓		✓	✓			✓		\checkmark	
Ren et al. [47]		✓		✓	✓		✓			\checkmark	
Bouaziz and Pauly [42]	✓			✓	✓		✓			\checkmark	
Li et al. [44]	✓			✓	✓		✓			\checkmark	
Vestner et al. [45]		✓		✓	✓			✓		\checkmark	
Dyke et al. [2]	✓			✓	✓			✓		✓	
Melzi et al. [52]		✓		✓		✓	✓			\checkmark	
Marin et al. [50]		/					/			/	

Table 5A summary of the test-sets completed by participants, which are marked with a tick. The absence of a tick indicates the test-set was not completed.

Method	Test-set						
	0	1	2	3	4		
R3DS Wrap 3	√	✓	✓	✓	√		
Ezuz and Ben-Chen [46]	✓	✓	✓	✓	✓		
Rodolà et al. [29]	✓	✓	✓	✓	✓		
Ren et al. [47]	✓	✓	✓	✓	✓		
Bouaziz and Pauly [42]	✓	✓	✓	✓	✓		
Li et al. [44]	✓	✓	✓	✓	✓		
Vestner et al. [45]	✓	✓	✓	✓	✓		
Dyke et al. [2]	✓	✓	✓	✓	✓		
Melzi et al. [52]			✓				
Marin et al. [50]			✓				

neighbourhoods. The second stage uses the computed mapping to estimate local anisotropy on the surface, represented by a discrete 2-tensor field. The anisotropy map is used to compute anisotropic geodesics for use in an extended spectral diffusion pruning method [41].

The authors observe that non-isometric areas may have few correspondences. This is rectified by interpolating between nearby correspondences that are considered to be good in order to provide correspondences for such problem areas.

The algorithm is initialised using the pre-computed sparse correspondences.

4.3. Robust Non-Rigid Registration with Reweighted Position and Transformation Sparsity

To address large-scale motion in non-rigid deformation, a non-rigid registration method with sparsity-regularised position and transformation constraints is proposed by Li et al. [44]. The distribution of positional errors and transformation differences for typical non-rigid deformation can be well modelled using

the Laplacian distribution, or equivalently, the L_1 -norm should be used to measure both the positional errors and transformation differences. To promote the sparsity, a re-weighted sparse model is adopted, which is solved by the alternating direction method of multipliers (ADMM). The model is robust against outliers as the sparsity terms allow a small fraction of regions with larger deviations. The method is evaluated on both public datasets and real datasets, captured by an RGB-D depth sensor. The results demonstrate that the method obtains better results than other state-of-the-art non-rigid registration and correspondence methods [6,44].

This method requires an initial set of sparse correspondences, the pre-computed correspondences were provided to the participating authors.

4.4. Efficient Deformable Shape Correspondence via Kernel Matching

Vestner et al. [45] consider the importance of certain properties for establishing a quality correspondence mapping, namely: ensuring the predicted mapping is a homeomorphism (bijective, and both itself and its inverse are continuous) and promoting matches of similar points.

The method is controlled primarily by two parameters, α and t. α balances the data and regularisation terms. t is the time parameter for heat diffusion. In practice, changing parameter t changes the influence of distant points during propagation. Vestner et al. [45] state the importance of selecting a large value for t in scenarios where a large amount of noise is present in the initial correspondences.

All shapes were re-scaled to have a similar area. All settings remain as described in the original paper, with the following exceptions $\alpha = 10^{-7}$ and $t = \{500, 323, 209, 135, 88, 57, 37, 24, 15, 10\}.$

4.5. Deblurring and Denoising of Maps between Shapes

Ezuz and Ben-Chen [46] identify that, while versatile, functional maps tend to recover low fidelity correspondences. To address this,



Fig. 4. Exemplar illustration of geodesic distance. For a point at the centre of the leopard's face, the geodesic distance to any given point on the surface is indicated by the local colour value.

they propose an approach that aims to improve the specificity of a given functional map by introducing a novel smoothness prior. The method is also designed to work in cases of highly non-isometric deformation. A regularisation term that promotes smooth mappings, which helps to remove noise from mappings, is incorporated to compute a functional mapping. Precise vertex-to-point correspondences are then recovered using an improved ICP-based recovery method.

Due to memory limitations a subset of 200 correspondences were selected using geodesic-based *farthest point sampling*. For experiments we set the number of basis functions $(k_1 \& k_2)$ to 120. The method requires an initial set of landmarks, the pre-computed SHOT correspondences with diffusion pruning were used for initialisation.




Fig. 5. An example illustrating a barycentric cell (the shaded area) for the vertex in the centre of a one-ring neighbourhood. The area of each vertex A_i —used in Algorithm 1—also corresponds to the area of the vertex's barycentric cell.

4.6. Partial Functional Correspondence

Rodolà et al. [29] propose a functional mapping method that is capable of robustly finding correspondence between non-rigidly deforming partial and full shapes. Observing that the functional mapping between two full near-isometric shapes should be approximately orthogonal and full rank, they investigate how partial shapes deviate from this. The authors take advantage of the low




Fig. 6. Examples of a shape from the dataset segmented using increasingly dense sampling of Voronoi cells. Each sub-figure contains following number of Voronoi cells: (a) two, (b) ten, (c) 25, (d) 50, (e) 150, & (f) 500.




Fig. 7. Examples of points distributed on the surface of a model from this dataset. (a) Bijective mapping, (b) part mapping (75.0%), (c) part mapping (50.0%), (d) sparse correspondence (50.0%), (e) sparse correspondence (25.0%), and (f) sparse correspondence (10.0%).

rank and sloped nature of functional maps in partial cases. Rodolà et al. [29] incorporate novel regularisation terms into a two-step optimisation process. The first step optimises the correspondence of the functional map based on an estimation of how partial the source shape is with respect to the target shape. The second step penalises any change in area and the length of the boundary of a part. This approach is not robust to non-isometric deformations.

Shapes were re-scaled to have surface areas between 1.5×10^4 and 2.0×10^4 . All other parameters remain as per their default, except n_eigen = 100.

4.7. Continuous and Orientation-preserving Correspondences via Functional Maps

Ren et al. [47] seek to address the problem of intrinsic symmetries when estimating correspondence using functional maps. The authors incorporate an orientation preserving constraint term into the optimisation function used to compute a functional mapping, incorporating surface normal information through the use of triple products, although this part of their technique was not used in experiments. The mapping is enhanced with a novel iterative refinement method to further ensure bijectivity and continuity. Outliers are efficiently detected and removed by measuring the Euclidean distance between two pairs of points. Unmapped points are reassigned a correspondence based on neighbouring correspondences. A further step improves point-wise mappings to promote the continuity of correspondences.

The region-level correspondence method of Kleiman and Ovsjanikov [48] was used to establish an initial correspondence between regions using the default parameters, except numComponentsRange = $\{10,9,8,7\}$. An initial functional mapping was computed using Nogneng and Ovsjanikov [49], with $k_1=120$ & $k_2=120$. The refinement method proposed by Ren et al. [47] was run for 10 iterations to recover point-to-point correspondences.

4.8. CMH Connectivity Transfer

We use the CMH framework proposed in Marin et al. [50] and extended on animals in Melzi et al. [51] to establish correspondences by transferring the connectivity. The method relies on extending the standard Laplace-Beltrami Operator (LBO) basis by adding three additional bases that encode extrinsic information of the meshes. This combination of intrinsic and extrinsic information permits fully encoding the geometry of the models without information loss due to a low-pass representation. A functional map is then computed, as proposed by Nogneng and Ovsjanikov [49] using six hand-placed landmarks as probe functions. Finally, the connectivity is transferred using the point-to-point correspondence and refined using an as-rigid-as-possible energy. The match is recovered by finding the nearest neighbour between the target model, and the source connectivity transferred over the model. This method assumes that the target and source shapes share the same pose, and does not use the coherent-point-drift local refinement as proposed in the original paper.




Fig. 8. An example of the coverage measure computed using synthetic correspondences on a real mesh.

4.9. ZoomOut

Similarly to CMH [50], ZoomOut [52] computes standard LBO bases, which it then refines. Correspondence and functional map computations are iterated between, increasing the dimension of the mapping at each step. The following parameters were used: 20 as input and 360 as output for the dimension of the functional map, using an incremental step of 10 and 1,000 samples with farthest point sampling for the correspondence step. As with *CMH*, the connectivity is transferred and the result is refined using as-rigid-as-possible optimisation.

4.10. R3DS Wrap 3.4

Russian3DScanner [53] developed *Wrap 3*, a commercial tool to transfer shape topology through non-rigid registration. The tool uses a variant of coarse-to-fine N-ICP with the facility to provide initial correspondences in the form of hard constraints to further help. This commercial software was found to be highly performant in a previous benchmark with lesser degrees of non-isometry by Dyke et al. [6].

The default parameters were left unchanged, and the method was provided with initial pruned correspondences.

Classification. Based on the comprehensive survey paper of Sahillioğlu [54], all methods have been categorised (see Table 4) based on the criteria described. Please refer to the original survey for the precise definition of each criterion.

Submissions. As shown in Table 5, most methods submitted results for each test-set, except [52] and [50]. This was mainly because these methods were designed to primarily handle cases

where objects have the same genus. In this dataset, test-set 2 does not contain any topological changes.

5. Evaluation

The quality of correspondences is evaluated using two measures. First, a standard error metric is employed, calculating the geodesic distance between predicted correspondences and the ground-truth locations, as per Kim et al. [22]. Then the quantity and uniformity of correspondences are measured using a novel measurement technique.

5.1. Error measure

For convenience, we describe the protocol of Kim et al. [22] here. For an estimated correspondence $(x_i, y_i) \in X \times Y$ and the respective ground-truth correspondence $(x_i, y_i^*) \in X \times Y$. The geodesic distance between the corresponding points on Y is $d_Y(y_i, y_i^*)$ (an example of geodesics is shown in Fig. 4). The area of shape Y is used to normalise the distance. The error of the estimated correspondence is be measured as

$$\varepsilon(x_i) = \frac{d_Y(y_i, y_i^*)}{\operatorname{area}(Y)^{1/2}}.$$
(3)

Cumulative error curves are subsequently produced by counting the number of correspondences with an error $\varepsilon(x_i)$ less than a given threshold of normalised geodesic distance ϵ , i.e., $\varepsilon(x_i) \le \epsilon$.

Algorithm 1: Computation of the coverage measure

```
Input: n – number of seed points/Voronoi cells
          V, F – vertices and triangles of target surface
          A – area of each vertex in surface V, F
          G – geodesic distance map for the surface V, F
          C - binary list with a value of true for vertices with
a correspondence
Output: r - ratio of Voronoi cells containing a
          correspondence
seed \leftarrow random\_integer(1,|V|) // select an initial random
seed point
S \leftarrow \text{fps geodesic}(G, \text{seed}, n) // sample the farthest points on
the surface
Q \leftarrow \text{find\_closest\_point}(G, S) || return classification of each
vertex to the closest point in S
for s \leftarrow S do
    // iterate through each sampled point on S
   if any((Q == s) \cap C) then
       // check if there are any correspondences in a given
       Voronoi cell
       r \leftarrow r + \sum_{i} ((Q_i == s) \cdot A_i) || sum area of the Voronoi
   end
end
                // normalise r by the total surface area
r \leftarrow r / \sum_i A_i
```

5.2. Surface coverage measure

To further assess the performance of methods we develop a coverage measure. The measure is derived by first segmenting a shape's surface into discrete regions and then summing the area of regions that contain a correspondence, this value is then normalised by the shape's total area. When few regions contain a correspondence the resulting value will be low, this indicates potentially poor overall correspondence between surfaces. By being able to numerically summarise the quality of a set correspondences, it is possible to gain valuable quantitative insights of a method's performance over a large dataset. Furthermore, as we demonstrate in this section, by varying the number of regions on a shape, it is possible to gain an even greater understanding of how corresponding points are distributed over the target shape. See Algorithm 1 for a detailed description of the implementation used.

Region segmentation. A set of seed points *S* on the target surface are selected using a geodesic-based farthest point sampling strategy. This helps ensure a reasonably evenly distributed sampling. To obtain discrete regions, a Voronoi segmentation *Q* is subsequently computed using the initial seeds *S*. Fig. 6 illustrates the segmentation of a shape using successively greater numbers of seed points *n*.

Segmentation density. By varying n such that $\{n \in \mathbb{N}^+ \mid n \leq |V|\}$ a reasonably smooth and intuitive performance measure is extracted. When n=1, the shape is unsegmented and a single region covers the whole mesh. If the surface has just one correspondence to any point, the output of Algorithm 1 would be r=1 (i.e., 100%)




Fig. 9. Overall coverage performance of methods that submitted results for all test-sets.




Fig. 10. Cumulative error curves for each test-set, (a) test-set 0, (b) test-set 1, (c) test-set 2 and (d) test-set 3.




Fig. 11. Cumulative error curves for test-set 4.




Fig. 12. Overall performance of methods that submitted completed all test-sets.

coverage for n = 1). While when n = |V|, the barycentric cell of each vertex is a discrete region. To achieve 100% coverage for n = |V|, each vertex must have an associated correspondence.

Precise correspondence. When handling correspondences with sub-vertex accuracy, each correspondence is associated with the closest discrete point on the shapes' surface. The barycentric cell of each vertex is treated as the vertex neighbourhood, as shown in Fig. 5. Points within a given neighbourhood are assigned to that respective vertex.

Segment weighting. Since the initial sampling of seed points does not guarantee that the area of each segmented region is uniform, the regions are weighted by their respective area, i.e., where s identifies a unique seed point/region $\sum_i ((Q_i == s) \cdot A_i)$. Here, the segmentation classification Q is a list of indices where i

refers to a specific vertex and Q_i is the seed point s that is closest to that vertex.

Computation time. Whilst the initial segmentation may be costly to produce, this approach allows the resulting segmentation to be cached and used for further comparison of correspondences with little additional computation. The most costly operation is computing geodesics for the distance map, the complexity of the popular method proposed by Kimmel and Sethian [55] is $O(|V|^2 \log |V|)$.

In Fig. 7, an example shape is used to illustrate different ways in which correspondences may be distributed using a set of synthetic points. Fig. 8 complements this, demonstrating the response of the coverage measure as the number of segments is varied. It can be seen that the characteristics of the shape of each curve vary by the type of correspondence computed.




Fig. 13. Cumulative error curves for each test-set, (a) test-set 0, (b) test-set 1, (c) test-set 2, (d) test-set 3 and (e) test-set 4.

In the case of a bijective mapping, the coverage value will remain at 100% regardless of how fine the segmentation is. Note that this does not mean that the reported correspondences are correct, but that every point on the target surface has a point-to-point correspondence on the source shape. In the case that a part of a shape is not matched (e.g., a leg), the metric will

drop off quickly, with no correspondence, then assuming the rest of the shape is successfully matched, the curve should have a gradient equal to zero. For methods that report evenly spaced sparse correspondences, the coverage should remain high until the frequency of the Voronoi cell samples is sufficient to cover areas in-between the sparse correspondences.

Table 6AUC results from the overall geodesic error of each method.

Method	AUC
R3DS Wrap 3	0.708996
Ezuz and Ben-Chen [46]	0.697713
Rodolà et al. [29]	0.606127
Ren et al. [47]	0.747115
Bouaziz and Pauly [42]	0.743213
Li et al. [44]	0.717810
Vestner et al. [45]	0.610986
Dyke et al. [2]	0.726690

6. Results

In this section we discuss the performance of correspondence methods with respect to the two measures described in Section 5.

6.1. Surface coverage

In Fig. 9 we measure the coverage achieved by each method that completed all test-sets. It is important to note that since a subset of the shapes in the dataset are partial, as the number of Voronoi cells on the surface increases a coverage score of 100% cannot be maintained. The curves of all methods monotonically decrease, this is because all methods report dense correspondences; therefore the sparsity of correspondences is not a factor in these results. Based on the curve characteristics discovered in the example of the coverage measure in Section 5, Vestner et al. [45] exhibit the closest performance to full and dense coverage, with Ren et al. [47] performing second best. This is understandable as both methods promote bijectivity and therefore produce both dense and well distributed correspondences. There is a significant performance gap between the other methods. This may indicate that the other methods do not promote, or do not strongly promote, bijectivity.

Figs. 10 and 11 report the performance of methods measured by coverage in each respective test-set. In test-set 2 (Fig. 10c) results for the methods of Melzi et al. [52] and Marin et al. [50] are shown. Both methods perform exceedingly well with a relatively high level of surface coverage. The results for test-sets 1-4 suggest that most methods failed to establish correspondences for one or more parts of each target shape.

6.2. Geodesic error

Fig. 12 reports cumulative geodesic error curves for methods that have completed all test-sets. Table 6 complements Fig. 12 reporting the *area under the curve* (AUC) of each respective method. Overall, most methods appear to perform similarly, with Ren et al. [47] and Bouaziz and Pauly [42] performing the best. However, also taking into account the results of the coverage measure, Ren et al. [47] produce more desirable results with greater coverage of the target shape.

In Fig. 13 results for each test-set are presented. Results for Melzi et al. [52] and Marin et al. [50] on test-set 2 are shown in Fig. 13c. These methods achieve superior correspondence accuracy in comparison to the fully-automatic methods.

Test-set 0 contains only partial-to-full scans. Bouaziz and Pauly [42] perform particularly well, this may be due to shapes having a similar initial orientation, which is important for N-ICP-based methods. Several methods achieve higher levels of accuracy on

test-set 2, this may be due to this test-set containing little topological change. With the exception of Ren et al. [47], most methods perform poorly on test-set 4. This is likely to be due to the higher degrees of non-isometry exhibited. The performance of Ren et al. [47] may be due in part to the use of a region-level correspondence method [48] for initialisation, since the correspondence method works particularly well on homogeneous shapes.

7. Conclusion

In this paper, a new benchmark dataset of non-isometric non-rigidly deforming shapes has been proposed to evaluate the performance of shape correspondence methods. To ensure greater accuracy, ground-truth correspondences were established by asking multiple specialists to annotate the animals. The performance of a variety of methods was evaluated using this dataset. Whilst traditional measures of correspondence accuracy are useful, they do not show the full picture of a correspondence method's performance. To address this, a new measure of correspondence coverage has been developed. The coverage measure helps quantitatively indicate the sparsity and distribution of correspondences. We find that Ren et al. [47] achieve the greatest accuracy, as well as a high degree of surface coverage, making it the overall best method in this scenario.

Both Melzi et al. [52] and Marin et al. [50] present semiautomatic methods that achieve superior accuracy compared to the fully-automatic methods. Though not evaluated here, in the paper of Groueix et al. [19]—a data-driven shape correspondence method—the authors make use of an out-of-the-box blend skinning model [8] for training. It would be interesting to see how well this method performs on this dataset where there are animals that have not been observed in the trained model.

As the accessibility of 3D scene capturing tools increases, there is a greater need for high-quality datasets that may be used for benchmarking and training purposes. Restrictive copyrights on existing works make the curation of such datasets challenging. Websites such as Sketchfab and AIM@SHAPE-VISIONAIR Shape Repository provide simple licenses that make it clear what the original creator permits and may enable further datasets to be produced.

The coverage measure formulated here offers important insights into the performance of shape correspondence methods. It can be directly incorporated into many existing benchmarks—such as those mentioned in Section 2.1. In the case of [30], where a dense ground-truth correspondence is known between near-isometrically deforming shapes, it is possible to modify the measure to normalise its output coverage based on an approximation of the total corresponding area. This may help ensure a fair comparison of collections of shape pairs.

This report provides a point of reference on the performance of state-of-the-art methods. The results demonstrate that non-isometric deformation remains a challenging shape matching problem, and further research is still required. This dataset could be extended to include a training facility. The performance of learning-based methods on unseen animals would be particularly insightful as a measure of generalisation.

8. Dataset contents

(Fig. 14).




Fig. 14. Illustrations of the contents of the benchmark dataset simplified to consist of a maximum of 100,000 faces.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRediT authorship contribution statement

Roberto M. Dyke: Conceptualization, Methodology, Software, Writing - original draft, Writing - review & editing, Visualization. Yu-Kun Lai: Conceptualization, Methodology, Resources, Writing - original draft, Writing - review & editing, Supervision, Project administration, Funding acquisition. Paul L. Rosin: Conceptualization, Methodology, Resources, Writing - original draft, Writing - review & editing, Supervision, Project administration, Funding acquisition. Stefano Zappalà: Methodology, Software. Seana Dykes: Conceptualization, Methodology, Daoliang Guo: Writing - original draft, Software. Kun Li: Writing - original draft, Software. Riccardo Marin: Writing - original draft, Software. Simone Melzi: Writing - original draft, Software. Jingyu Yang: Writing - original draft, Software.

Acknowledgments

We thank the authors of methods that did not participate but made their code available online [42, 45, 46, 29, 47, 53]. We thank the creators of the 3D models used in this dataset for sharing their work, on Sketchfab and the AIM@SHAPE-VISIONAIR

Shape Repository. We also thank Ira Kemelmacher for giving us permission to incorporate their work [38] into this dataset.

References

- [1] van Kaick O, Zhang H, Hamarneh G, Cohen-Or D. A survey on shape correspondence. Comput Gr Forum 2011;30(6):1681–707. doi:10.1111/j.1467-8659. 2011.01884 x.
- [2] Dyke RM, Lai Y-K, Rosin PL, Tam GKL. Non-rigid registration under anisotropic deformations. Comput Aided Geomet Des 2019a;71:142–56. doi:10.1016/j.cagd. 2019.04.014.
- [3] Jiang T, Yang X, Zhang J, Tian F, Liu S, Xiang N, Qian K. Huber- L_1 -based non-isometric surface registration. Vis Comput 2019;35(6):935–48. doi:10.1007/s00371-019-01670-1.
- [4] Ezuz D, Heeren B, Azencot O, Rumpf M, Ben-Chen M. Elastic correspondence between triangle meshes. Comput Gr Forum 2019;38(2):121–34. doi:10.1111/ cgf 13624
- [5] Melzi S, Marin R, Rodolà E, Castellani U, Ren J, Poulenard A, Wonka P, Ovsjanikov M. Matching humans with different connectivity. In: Proceedings of the Eurographics workshop on 3D object retrieval. The Eurographics Association and John Wiley & Sons, Ltd.; 2019a. p. 121–8. doi:10.2312/3dor.20191070.
- [6] Dyke RM, Stride C, Lai Y-K, Rosin PL, Aubry M, Boyarski A, Bronstein AM, Bronstein MM, Cremers D, Fisher M, Groueix T, Guo D, Kim VG, Kimmel R, Lhner Z, Li K, Litany O, Remez T, Rodolà E, Russell BC, Sahillioğlu Y, Slossberg R, Tam GKL, Vestner M, Wu Z, Yang J. Shape correspondence with isometric and non-isometric deformations. In: Proceedings of the Eurographics workshop on 3D object retrieval. The Eurographics Association and John Wiley & Sons, Ltd.; 2019b. p. 111–19. doi:10.2312/3dor.20191069.
- [7] Ehsani K, Bagherinezhad H, Redmon J, Mottaghi R, Farhadi A. Who let the dogs out? Modeling dog behavior from visual data. In: Proceedings of the computer vision and pattern recognition. IEEE; 2018. p. 4051–60. doi:10.1109/CVPR.2018. 00426.
- [8] Zuffi S, Kanazawa A, Jacobs DW, Black MJ. 3D Menagerie: modeling the 3D shape and pose of animals. In: Proceedings of the computer vision and pattern recognition. IEEE; 2017. p. 5524–32. doi:10.1109/cvpr.2017.586.

- [9] Robinette KM, Daanen H, Paquet E. The CAESAR project: a 3D surface anthropometry survey. In: Proceedings of the international conference on 3D digital imaging and modeling. IEEE: 1999. p. 380-6. doi:10.1109/IM.1999.805368.
- [10] Sand P, McMillan L, Popović J. Continuous capture of skin deformation. Trans Gr 2003;22(3):578–86. doi:10.1145/882262.882310.
- [11] Anguelov D, Srinivasan P, Koller D, Thrun S, Rodgers J, Davis J. SCAPE: Shape completion and animation of people. Trans Gr 2005;24(3):408–16. doi:10.1145/ 1073204.1073207.
- [12] Vlasic D, Baran I, Matusik W, Popović J. Articulated mesh animation from multi-view silhouettes. Trans Gr 2008;27(3):1–9. doi:10.1145/1360612. 1360696.
- [13] Hasler N, Stoll C, Sunkel M, Rosenhahn B, Seidel HP. A statistical model of human pose and body shape. Comput Gr Forum 2009;28(2):337–46. doi:10. 1111/i.1467-8659.2009.01373.x.
- [14] Bogo F, Romero J, Loper M, Black MJ. FAUST: Dataset and evaluation for 3d mesh registration. In: Proceedings of the computer vision and pattern recognition. IEEE; 2014. p. 3794–801. doi:10.1109/CVPR.2014.491.
- [15] Bogo F, Romero J, Pons-Moll G, Black MJ. Dynamic FAUST: Registering human bodies in motion. In: Proceedings of the computer vision and pattern recognition. IEEE; 2017. p. 5573–82.
- [16] Yang Y, Yu Y, Zhou Y, Du S, Davis J, Yang R. Semantic parametric reshaping of human body models. In: Proceedings of the international conference on 3D vision. In: 3DV '14, 2. IEEE; 2014. p. 41–8. doi:10.1109/3DV.2014.47.
- [17] Loper M, Mahmood N, Romero J, Pons-Moll G, Black MJ. SMPL: a skinned multi-person linear model. Trans Gr 2015;34(6). doi:10.1145/2816795.2818013.
- [18] Pickup D, Sun X, Rosin PL, Martin RR, Cheng Z, Lian Z, Aono M, Hamza AB, Bronstein A, Bronstein M, Bu S, Castellani U, Cheng S, Garro V, Giachetti A, Godil A, Han J, Johan H, Lai L, Li B, Li C, Li H, Litman R, Liu X, Liu Z, Lu Y, Tatsuma A, Ye J. Shape retrieval of non-rigid 3D human models. In: Bustos B, Tabia H, Vandeborre J-P, Veltkamp R, editors. Proceedings of the Eurographics workshop on 3D object retrieval. The Eurographics Association and John Wiley & Sons, Ltd.; 2014. p. 101-10. ISBN 978-3-905674-58-3. doi:10.2312/3dor.2014156
- [19] Groueix T, Fisher M, Kim VG, Russell BC, Aubry M. 3D-CODED: 3D correspondences by deep deformation. In: Proceedings of the European conference on computer vision. Springer Publishing Company; 2018. p. 235–51. doi:10.1007/978-3-030-01216-8_15.
- [20] Bronstein A, Bronstein M, Kimmel R. Numerical geometry of non-rigid shapes. 1. Springer Publishing Company; 2008. ISBN 0387733000.
- [21] Rodolà E, Bulò S, Windheuser T, Vestner M, Cremers D. Dense non-rigid shape correspondence using random forests. In: Proceedings of the computer vision and pattern recognition. IEEE; 2014. p. 4177–84. doi:10.1109/CVPR.2014.532.
- [22] Kim VG, Lipman Y, Funkhouser T. Blended intrinsic maps. Trans Gr 2011;30(4):1–12. doi:10.1145/2010324.1964974.
- [23] Lian Z, Godil A, Bustos B, Daoudi M, Hermans J, Kawamura S, Kurita Y, Lavoué G, Nguyen HV, Ohbuchi R, Ohkita Y, Ohishi Y, Porikli F, Reuter M, Sipiran I, Smeets D, Suetens P, Tabia H, Vandermeulen D. SHREC'11 track: Shape retrieval on non-rigid 3d watertight meshes. In: Proceedings of the Eurographics workshop on 3D object retrieval. The Eurographics Association and John Wiley & Sons, Ltd.; 2011. p. 79–88. doi:10.2312/3DOR/3DOR11/079-088.
- [24] Sumner RW, Popović J. Deformation transfer for triangle meshes. Trans Gr 2004;23(3):399–405. doi:10.1145/1015706.1015736.
- [25] Siddiqi K, Zhang J, Macrini D, Shokoufandeh A, Bouix S, Dickinson S. Retrieving articulated 3D models using medial surfaces. Mach Vis Appl 2008;19(4):261– 75. doi:10.1007/s00138-007-0097-8.
- [26] Bronstein AM, Bronstein MM, Castellani U, Dubrovina A, Guibas LJ, Horaud RP, Kimmel R, Knossow D, Lavante Ev, Mateus D, Ovsjanikov M, Sharma A. SHREC'10 track: Correspondence finding. In: Proceedings of the Eurographics workshop on 3D object retrieval. The Eurographics Association and John Wiley & Sons, Ltd.; 2010a. p. 87–91. doi:10.2312/3DOR/3DOR10/087-091.
- [27] Bronstein AM, Bronstein MM, Castellani U, Falcidieno B, Fusiello A, Godil A, Guibas LJ, Kokkinos I, Lian Z, Ovsjanikov M, Patan G, Spagnuolo M, Toldo R. SHREC'10 track: Robust shape retrieval. In: Proceedings of the Eurographics workshop on 3D object retrieval. The Eurographics Association and John Wiley & Sons, Ltd.; 2010b. p. 71–8. doi:10.2312/3DOR/3DOR10/071-078.
- [28] Boyer E, Bronstein AM, Bronstein MM, Bustos B, Darom T, Horaud R, Hotz I, Keller Y, Keustermans J, Kovnatsky A, Litmany R, Reininghaus J, Sipiran I, Smeets D, Suetens P, Vandermeulen D, Zaharescu A, Zobel V. SHREC '11: Robust feature detection and description benchmark. In: Proceedings of the Eurographics workshop on 3D object retrieval. The Eurographics Association and John Wiley & Sons, Ltd.; 2011. p. 71–8. doi:10.2312/3DOR/3DOR11/071-078.
- [29] Rodolà E, Cosmo L, Bronstein MM, Torsello A, Cremers D. Partial functional correspondence. Computer Graphics Forum 2017;36(1):222–36. doi:10.1111/cgf. 12797.
- [30] Cosmo L, Rodolà E, Bronstein MM, Torsello A, Cremers D, Sahillioğlu Y. Partial matching of deformable shapes. In: Proceedings of the Eurographics workshop

- on 3D object retrieval. The Eurographics Association and John Wiley & Sons, Ltd.; 2016, p. 61-7. doi:10.2312/3dor.20161089.
- [31] Corman E, Ovsjanikov M, Chambolle A. Supervised descriptor learning for non-rigid shape matching. In: Proceedings of the computer vision. Springer Publishing Company; 2015. p. 283–98. ISBN 978-3-319-16220-1. doi:10.1007/ 978-3-319-16220-1_20.
- [32] Cignoni P, Rocchini C, Scopigno R. Metro: Measuring error on simplified surfaces. Comput Gr Forum 1998;17(2):167–74. doi:10.1111/1467-8659.00236.
- [33] Aspert N, Santa-Cruz D, Ebrahimi T. Mesh: measuring errors between surfaces using the Hausdorff distance. In: Proceedings of the international conference on multimedia and expo. In: ICME '02, 1. IEEE; 2002. p. 705–8. doi:10.1109/ ICME.2002.1035879.
- [34] Edelstein M, Ezuz D, Ben-Chen M. ENIGMA: evolutionary non-isometric geometry matching. Trans Gr 2020;39(4). doi:10.1145/3386569.3392447.
- [35] Bronstein AM, Bronstein MM, Kimmel R, Mahmoudi M, Sapiro G. A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. Int J Comput Vis 2010c;89(2):266–86. doi:10.1007/s11263-009-0301-6.
- [36] Azencot O, Corman E, Ben-Chen M, Ovsjanikov M. Consistent functional cross field design for mesh quadrangulation. Trans Gr 2017;36(4). doi:10.1145/ 3072959.3073696.
- [37] Nogneng D, Melzi S, Rodolà E, Castellani U, Bronstein MM, Ovsjanikov M. Improved functional mappings via product preservation. Comput Gr Forum 2018;37:179–90. doi:10.1111/cgf.13352.
- [38] Kemelmacher I, Basri R. Indexing with unknown illumination and pose. In: Proceedings of the computer vision and pattern recognition. In: CVPR '05, 1. IEEE; 2005. p. 909–16. doi:10.1109/CVPR.2005.186.
- [39] Tombari F, Salti S, Di Stefano L. Unique signatures of histograms for local surface description. In: Proceedings of the European conference on computer vision. Springer Publishing Company; 2010. p. 356–69. doi:10.1007/ 978-3-642-15558-1_26.
- [40] Limberger FA, Wilson RC. Feature encoding of spectral signatures for 3D nonrigid shape retrieval. In: Proceedings of the British machine vision conference. BMVA Press; 2015. p. 56.1–56.13. doi:10.5244/C.29.56.
- [41] Tam GKL, Martin RR, Rosin PL, Lai Y-K. Diffusion pruning for rapidly and robustly selecting global correspondences using local isometry. Trans Gr 2014;33(1):4:1–4:17. doi:10.1145/2517967.
- [42] Bouaziz S, Pauly M. Dynamic 2D/3D registration for the Kinect. In: Proceedings of the SIGGRAPH courses. Association for Computing Machinery; 2013. p. 21:1–21:14. doi:10.1145/2504435.2504456.
- [43] Sorkine O, Alexa M. As-rigid-as-possible surface modeling. In: Proceedings of the symposium on geometry processing. The Eurographics Association and John Wiley & Sons, Ltd.; 2007. p. 109–16. doi:10.2312/SGP/SGP07/109-116.
- [44] Li K, Yang J, Lai Y, Guo D. Robust non-rigid registration with reweighted position and transformation sparsity. Trans Vis Comput Gr 2019;25(6):2255–69.
- [45] Vestner M, Lhner Z, Boyarski A, Litany O, Slossberg R, Remez T, Rodolà E, Bronstein A, Bronstein M, Kimmel R, Cremers D. Efficient deformable shape correspondence via kernel matching. In: Proceedings of the international conference on 3D vision; 2017. p. 517–26. doi:10.1109/3DV.2017.00065.
- [46] Ezuz D, Ben-Chen M. Deblurring and denoising of maps between shapes. Comput Gr Forum 2017;36(5):165–74. doi:10.1111/cgf.13254.
- [47] Ren J, Poulenard A, Wonka P, Ovsjanikov M. Continuous and orientation-preserving correspondences via functional maps. Trans Gr 2018;37:248:1-248:16.
- [48] Kleiman Y, Ovsjanikov M. Robust structure-based shape correspondence. Comput Gr Forum 2019;38(1):7–20. doi:10.1111/cgf.13389.
- [49] Nogneng D, Ovsjanikov M. Informative descriptor preservation via commutativity for shape matching. Comput Gr Forum 2017;36(2):259-67. doi:10.1111/cof.13124
- [50] Marin R, Melzi S, Musoni P, Bardon F, Tarini M, Castellani U. CMH: Coordinates manifold harmonics for functional remeshing. In: Proceedings of the Eurographics workshop on 3D object retrieval. The Eurographics Association and John Wiley & Sons, Ltd.; 2019. p. 63–70. doi:10.2312/3dor.20191063.
- [51] Melzi S, Marin R, Musoni P, Bardon F, Tarini M, Castellani U. Intrinsic/extrinsic embedding for functional remeshing of 3D shapes. Comput Gr 2020;88:1–12. doi:10.1016/j.cag.2020.02.002.
- [52] Melzi S, Ren J, Rodolà E, Wonka P, Ovsjanikov M. ZoomOut: Spectral upsampling for efficient shape correspondence. Trans Gr 2019b;38:155:1–155:14. doi:10.1145/3355089.3356524.
- [53] Russian3DScanner. Wrap 3.4. https://www.russian3dscanner.com/; 2019.
- [54] Sahillioğlu Y. Recent advances in shape correspondence. Vis Comput 2019:1– 17. doi:10.1007/s00371-019-01760-0.
- [55] Kimmel R, Sethian JA. Computing geodesic paths on manifolds. Proc Natl Acad Sci 1998;95(15):8431–5. doi:10.1073/pnas.95.15.8431.