Eurographics Workshop on 3D Object Retrieval (2020)
T. Schreck and T. Theoharis (Editors)

SHREC'20: Non-rigid Shape Correspondence of Physically-Based Deformations

R. M. Dyke[†1] F. Zhou[†2] Y-K. Lai[†1] P. L. Rosin[†1] D. Guo[3] K. Li[3] R. Marin[4] J. Yang[3]

[1]School of Computer Science & Informatics, Cardiff University, Wales
[2]Beihang University, Bejing, China
[3]Tianjin University, Tianjin, China
[4]Department of Computer Science, University of Verona, Italy

Abstract

Commonly, novel non-rigid shape correspondence techniques focus on particular matching challenges. This can lead to the potential trade-off of poorer performance in other scenarios. An ideal dataset would provide a granular means for degrees of evaluation. In this paper, we propose a novel dataset of real scans that contain challenging non-isometric deformations to evaluate non-rigid point-to-point correspondence and registration algorithms. The deformations included in our dataset cover extreme types of physically-based contortions of a toy rabbit. Furthermore, shape pairs contain incrementally different types and amounts of deformation, this enables performance to be systematically evaluated with respect to the nature of the deformation. A brief investigation into different methods for initialising correspondence was undertaken, and a series of experiments were subsequently conducted to investigate the performance of state-of-the-art methods on the proposed dataset. We find that methods that rely on initial correspondences and local descriptors that are sensitive to local surface changes perform poorly in comparison to other strategies, and that a template-based approach performs the best.

CCS Concepts

• ***Computing methodologies*** → ***Shape analysis;*** *Mesh models;*

1. Introduction

Shape correspondence is a fundamental problem in geometry processing with numerous applications, e.g., texture transfer, shape morphing, and statistical shape analysis. However, accurately identifying correspondences between two or more surfaces continues to be a challenging problem, and scenarios in which contemporary approaches fail persist.

Many potential challenges occur when estimating a correspondence between surfaces: deformations such as non-isometry (i.e., stretching between shapes); ambiguities (e.g., mapping features with little geometric detail); shape incompatibility [MMR*19] (e.g., varying connectivity and vertex density); partial correspondence (e.g., registering two incomplete scans); and semi-compatible shapes (e.g., matching a human and a centaur).

State-of-the-art techniques aim to address challenges such as missing data and non-isometry. However, evaluative datasets fail to provide an appropriate means of identifying precisely where a method's performance degrades. Therefore it may be difficult to determine what degree of robustness a method has for a particular correspondence challenge. This is crucial for understanding what scenarios a method may be useful in.

We seek to address this problem by designing a structured, incrementally challenging dataset. We focus on inducing non-rigid deformations on a real-world object and use texture markers to establish ground-truths. Various internal properties and drastic poses of the shape were selected, causing complex deformations, and varying degrees of local protrusions and indentations on the surface. A subset of deformations were selected—twist, indent, inflate & stretch—from those identified by [SPF19], whose work focuses on how humans interpret deformation of an unknown object.

The organisers briefly investigate the problem of correspondence initialisation in Section 3.1, the results of which are discussed in Section 6.1. Participants were asked to estimate the correspondence between each partial scan in the dataset and one watertight scan of the rabbit. Section 4 details the correspondence methods evaluated using this dataset. Participant results are presented in Section 6.2.

2. Related work

There are many different possible approaches to estimating correspondences between non-rigid shapes. These may be classified as extrinsic approaches (e.g., iterative closest point), intrinsic or

† Track organisers

spectral techniques (e.g., Gaussian kernels [VLR*17], heat kernels [VLB*17] and functional maps [EBC17]), and methods that use a combination of both [DLRT19]. Data-driven methods can perform well in benchmarks [DSL*19], but may fail in scenarios where training data is insufficient or no comparable model is available. Many other methods have been proposed for computing shape correspondence, which are further discussed in the following surveys [vKZHCO11, TCL*13, BCBB16, BBL*17, Sah19].

Many recent works have focused on addressing the problem of non-isometric deformations [VLR*17, VLB*17, EBC17, DLRT19, JYZ*19, EHA*19, EEBC20]— [Sah19] provides a comprehensive overview of many methods. However, the resources available to evaluate the performance of such methods are insufficient.

Previous SHape REtrieval Contests (SHREC) have produced insightful benchmark datasets. Many use synthetic objects [BBC*10, LRB*16, RCB*17, RCL*17]. While these synthetic datasets may have certain advantages, such as easily established ground-truths, they also produce deformations that are not necessarily realistic.

Human body scan datasets [RDP99, ASK*05, BRLB14] capture real-world scans of humans undergoing realistic deformations. This is useful for understanding how well a method may perform on human subjects. However, the range of motion of a human is naturally constrained, thus the degree of non-rigid deformation is relatively limited. Whilst significant non-isometric deformation does occur between persons, features like the head, hands and feet, are particularly distinct, enabling correspondences between distinctive surface regions to be established through smooth interpolations.

3. Dataset

For this paper, 3D scans of a real-world object were captured using a high-precision 3D scanner (Artec3D Space Spider). The scanner is accurate up to 0.05mm. Each scan exhibits one or more types of deformation. Scans were classified into four distinct groups by the type of deformation primarily exhibited by a given pose: *twist, indent, inflate, & stretch*. The challenge of the dataset was further increased by varying the internal properties of the object by filling it with different materials: *couscous, risotto, & chickpea*. Using different internal materials changed the local appearance of the shape's surface, as well as its deformation properties. For each pose, the object was filled, and then scanned, with three different internal materials. This caused incremental changes to the surface and overall deformation exhibited, as illustrated in Fig. 1.

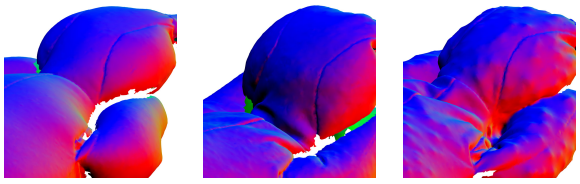

(a) *Couscous* **(b)** *Risotto* **(c)** *Chickpea*

Figure 1: *Illustrations of the surfaces of meshes with different internal materials coloured by the surface normals.*

The dataset consists of a stuffed soft toy rabbit made out of a stretchy jersey material with no type of internal skeleton that could otherwise restrict its movement, see Fig. 2. The rabbit had 590 coloured markers drawn on the surface, which allowed numerous accurate ground-truth correspondences to be established, see Fig. 3. Note that our purpose is to investigate how different types of physically-based deformations affect non-rigid shape correspondence, so a carefully chosen object with different material fillings is sufficient and makes data capture and analysis more manageable.

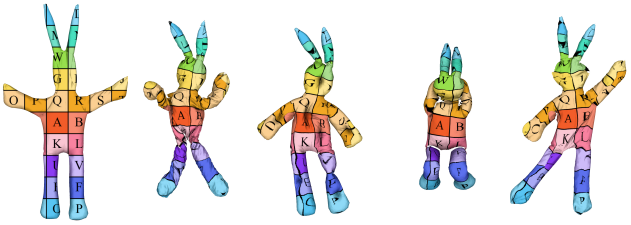

(a) **(b)** **(c)** **(d)** **(e)**

Figure 3: *Examples of texture transfer using the ground-truth correspondences. (a) Target shape. (b-e) Source shapes. Correspondences were transferred and interpolated using a landmark-based correspondence method [EBC17]. The original texture was projected onto the coronal (frontal) plane of the rabbit in (a).*

The poses and materials used were carefully selected to incrementally vary the deformation challenge so that algorithm deficiencies—with respect to these varying properties—may be identified. Some examples of challenging problems are shown in Fig. 4.

Figure 2: *A photo of the model used for the dataset with markers.*

A key point about the proposed dataset is that the exaggerated nature of the deformations, such as twisting, often contradicts the underlying assumptions of state-of-the-art approaches. Due to changes of the internal material and surface creases, geometry is often locally non-isometric, which is problematic for many shape descriptors. The model has intrinsic symmetries and, with the exception of the target scan, all scans are partial.

Information about the data underpinning the results presented here, including how to access them, can be found in the Cardiff University data catalogue (here).

3.1. Initial correspondences

Most correspondence methods require an initial set of sparse or dense correspondences that are subsequently refined. Due to the challenging variations in both the local and global appearance, three sets of initial correspondences using different methods were produced to find an optimal approach. The following methods were used to acquire the initial sets of correspondences: diffusion pruning [TMRL14], region-level correspondence [KO19], and *non-rigid iterative closest point* (N-ICP) [BP13]. All methods took a similar amount of time to compute. Where possible these initial correspondences have been used for fairer comparisons.

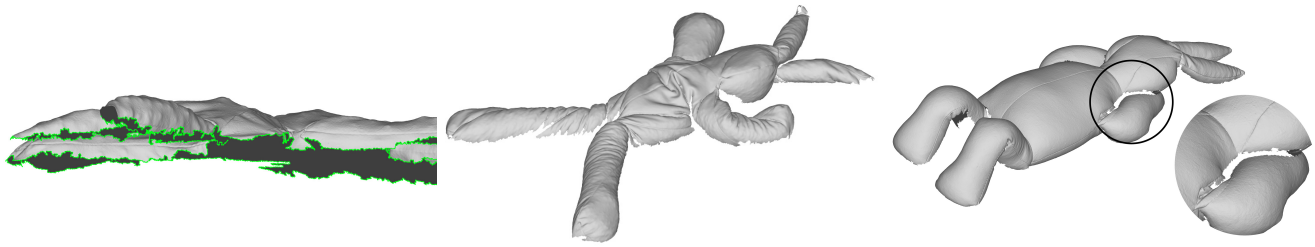

(a) (b) (c)

Figure 4: *Illustrations of some of the challenges in the dataset. (a) Partial scans (green indicates the boundary). (b) Complex deformations. (c) Missing geometry caused by self-occlusion.*

The diffusion pruning method of [TMRL14] was used to produce a set of globally consistent correspondences from an initial set of point-wise descriptors. Due to the significant deformation, local geometry is often non-isometric and thus local descriptors that rely on near-isometries performed poorly. *Signature of histograms of orientations* (SHOT) [TSDS10] signatures were computed with 10 bins at two scales, which covered 2% & 5% of the surface area square rooted. Due to memory limitations and computation time, diffusion pruning was performed with the default parameters, except $K = 5$ and $d = 25\%$.

The region-level correspondence method of [KO19] was used to produce a set of correspondences between segmented shapes. The method is capable of finding region-based matches using a shape graph that describes the connectivity of consistently segmented shapes. The method was run using nearly all default parameters, with `numComponentsRange` $= \{10, 9, 8, 7\}$. A dense point-to-point mapping was subsequently recovered using nearest neighbours of a functional map from a state-of-the-art approach [NO17].

The non-rigid registration method of [BP13] was used to register the shapes together. Point-to-point correspondences were computed using nearest neighbour between the vertices of the two surfaces. The method's parameters were set to `w1` $= 1$ (point-to-plane term), `w2` $= 1$ (point-to-point term), `w3` $= 1$ (global rigidity term), `w4` $=$ 1000 (local rigidity term) and `iter` $= 100$.

4. Correspondence methods

This section presents the methods that are examined in this work. Eight methods were evaluated using the benchmark; namely: non-rigid registration with re-weighted sparsities [LYLG19], non-rigid partial functional maps [RCB*17], anisotropic non-rigid registration [DLRT19], non-rigid partial functional maps [VLB*17], precise recovery of non-isometric functional maps [EBC17], continuous and orientation-preserving functional maps [RPWO18], data-driven non-rigid registration *FARM+*, and a commercial non-rigid registration tool *R3DS Wrap 3.4*.

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

To cope with non-rigid deformation with large motion, [LYLG19] proposes a robust non-rigid registration method using re-weighted sparsities on position and transformation to estimate deformations between 3D shapes. The energy function is formulated with position and transformation sparsity on both the data and smoothness terms, and a smoothness constraint is defined using local rigidity. The double sparsity-based non-rigid registration model is enhanced with a re-weighting scheme, and solved by transferring the model into four alternately-optimised sub-problems which have exact solutions and guaranteed convergence. Experimental results on both public datasets and real scanned datasets show that the method outperforms other state-of-the-art approaches [LYLG19,DSL*19], and is more robust to noise and outliers than conventional non-rigid registration methods.

The set of pruned and N-ICP correspondences were provided, and was run using parameters optimised by the authors.

4.2. Partial Functional Correspondence

[RCB*17] extends the functional mapping framework to address the problem of partial input data. The authors discovered that it is possible to estimate a partial correspondence with functional maps. Computation of a partial functional map $\boldsymbol{C}$ is formulated as an alternating minimisation problem with two steps. In the first step, correspondences are regularised based on assumptions about the overlap of the source mesh. The slope, orthogonality, and rank of $\boldsymbol{C}^\top \boldsymbol{C}$ are incorporated to help regularise this step. In the second step, the mapping is regularised by the part. Solutions with a dissimilar area and long boundaries are penalised. This method assumes the deformation to be near-isometric. While this is not necessarily the case, being able to effectively handle partial data is critical for this dataset, and is typically sufficient. SHOT was used to compute dense point-wise descriptors.

For the experiments, shapes were re-scaled—as described in the method's code. All parameters remained default, with the exception of `n_eigen` $= 100$.

4.3. Non-rigid Registration under Anisotropic Deformations

[DLRT19] proposes a non-rigid registration pipeline that estimates a correspondence between two surfaces in two steps. First, diffusion pruning is used to compute an initial sparse set of globally consistent correspondences. An implementation of N-ICP with an as-rigid-as-possible regularisation term that is generalised to support r-ring neighbourhoods is then used to find a dense set of correspondences. Finally, the dense correspondence is used to estimate

local anisotropy on the source surface. For subsequent iterations, the original geodesics used for diffusion pruning are replaced with anisotropic geodesics.

The described pipeline is fully automatic. However, for comparative purposes, our pre-computed correspondences are used for initialisation. All other parameters remain as described in the original paper.

4.4. Efficient Deformable Shape Correspondence via Kernel Matching

[VLB*17] proposes an algorithm to compute vertex-to-vertex correspondence between non-isometric shapes. The work extends [VLR*17], primarily incorporating the use of heat kernels and a method to handle partial matching cases. A *linear assignment problem* (LAP) is solved using a multi-scale approach for cases where objects have a high number of vertices to relieve the quadratic memory requirements of solving the LAP. In the original paper SHOT and *heat kernel signatures* (HKS) [SOG09] were used as point-wise descriptors.

Based on the discussions and experimental results in the original paper, the parameters α and t were set to 10^{-7} and $\{500, 323, 209, 135, 88, 57, 37, 24, 15, 10\}$ respectively. Only SHOT is used to compute point-wise descriptors, as HKS is sensitive to changes in topology. Optimal results were achieved by scaling input shapes with respect to the unit area of the source shape. Since all shapes were captured using the same device, the original scale was the same.

4.5. Deblurring and Denoising of Maps between Shapes

[EBC17] observes that functional maps typically lack specificity, moreover recovering vertex-to-vertex correspondences from a functional map leads to undesirable noise. The authors propose a method to refine a functional map, which is capable of recovering vertex-to-point correspondences. This is desirable when transporting high frequency information between surfaces—such as textures. The authors incorporate a correspondence regulariser that favours smooth maps. The method removes noise by first blurring the map, and then applying their proposed smoothness term. Point-wise correspondences are subsequently recovered from the mapping, projected to a facet.

This method was run in three configurations, using each respective set of pre-computed initial correspondences. Due to memory limitations caused by computing the wave kernel map, all configurations used a subset of 200 correspondences—selected using *farthest point sampling* (FPS) based on geodesic distances. The following parameters were modified $k_1 = 120$ & $k_2 = 120$.

4.6. Continuous and Orientation-preserving Correspondences via Functional Maps

[RPWO18] proposes a method that attempts to address the problem of symmetric ambiguity for functional mapping methods. The method also attempts to ensure that the mapping remains bijective and continuous when recovering point-wise correspondences from a given functional map. The *wave kernel signature* (WKS) is used as a point-wise descriptor. WKS was shown to be robust to some topological changes in the seminal paper [ASC11], although the descriptor is known to be strictly invariant only under isometric deformation. Due to the symmetrical appearance of the shape used in this dataset, distinguishing such ambiguity is beneficial. However, the authors remark that the method is not designed to directly handle partial cases.

The method was run both with and without our sets of initial correspondences. Similarly to [EBC17], due to memory limitations of computing the wave kernel map, 200 correspondences were selected using geodesic-based FPS. Parameters were configured to the default settings used in the paper, with `num_iters` $= 10$. The parameters for computing WKS were `numTimes` $= 100$ and `skipSize` $= 10$.

4.7. FARM+: Functional Automatic Registration Method for 3D Human Bodies

FARM+ a variant of [MMRC20] is used. This method relies on the registration of a morphable model exploiting the functional map framework [OBCS*12]. In the original paper, automatic landmarks are detected on protrusions relying only on geometrical information of the *discrete-time evolution process* (DEP) [MOR*18] descriptor. These landmarks are used to initialise a dense correspondence in the functional domain, and also to retrieve finer local correspondence over hands and heads. Then, a learned deformable template of humans bodies is optimised to fit the target model. Finally, a local refinement is applied to align the template to the target using an as-rigid-as-possible regularisation. The correspondence between the template and target is achieved using nearest-neighbour between vertices. This method has been refined by [MMRC19], removing some iterative steps and using ZoomOut [MRR*19] refinement for functional maps. In this challenge, we start from [MMRC19]. Since the provided template is not capable of deformation, it is animated using Mixamo [Ado20] and some deformation basis is defined to inflate or shrink the template along the direction of the surface normals. To compute landmarks, the minimum and maximum of the first Laplacian eigenfunctions are used, and six landmarks are classified over extremal points of the rabbit's ears, arms, and legs. Another main change of the original methods of [MMRC19] and [MMRC20] is that no local correspondence is used. All parameters were left unchanged from the original method, which were tuned for the specific domain of human bodies.

4.8. R3DS Wrap 3.4

[Rus19] develops a commercial software product with functionality to non-rigidly register surfaces called *Wrap 3*. The method uses landmark points to compute an optimal alignment between two non-rigid surfaces. The method uses a multiple-step sub-division surface approach with N-ICP in a coarse-to-fine manner to establish an approximate non-rigid alignment, which is subsequently refined in further steps.

Wrap 3 was used to register the scans both with and without our initial pruned correspondences. Better results were achieved by set-

ting the final weight of control points (the initial correspondences) to zero, all other settings remained as per default.

5. Evaluation

The quality of correspondences is measured following the protocol described by [KLF11] using the sparse texture markers as ground-truths. For convenience, we describe the protocol for computing the geodesic error of correspondences here. For an estimated correspondence $(x,y) \in X \times Y$ and the respective ground-truth correspondence $(x,y^*) \in X \times Y$. The geodesic distance between the corresponding points on Y is $d_Y(y,y^*)$

$$\varepsilon(x) = \frac{d_Y(y,y^*)}{\text{area}(Y)^{1/2}}. \quad (1)$$

6. Results

In this section we discuss the performance achieved by correspondence methods on our dataset. We examine how different initialisation techniques impact the performance of a selection of methods. Then the affects of variations in pose and internal material are scrutinised. It should be noted that the density of the correspondences acquired is not wholly reported because all methods report a dense mapping of $\approx 100\%$ in most cases.

6.1. Initial correspondences

In Fig. 5 the performance of the initial correspondence methods described in Section 3.1 is reported. The methods are abbreviated as follows: *N-ICP* [BP13] refers to [BP13], *Pruning* [TMRL14] refers to the diffusion pruning method of [TMRL14], and *SEG* [KO19] refers to the region-level correspondence method of [KO19]. N-ICP [BP13] performed the best. This may be due to the shapes having a reasonable initial rigid alignment, which is necessary when using N-ICP. All deformations are largely locally non-isometric, this causes point-wise descriptors that methods use on unreliable.

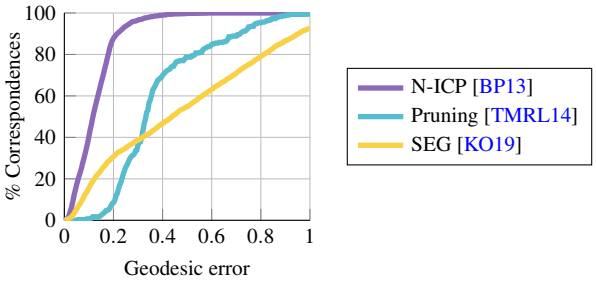

Figure 5: *Error curves for the methods used to establish an initial correspondence. Note, Pruning [TMRL14] only estimated* $\approx 7.4250\%$ *of correspondences.*

These initial correspondences were subsequently used to initialise the relevant shape correspondence and registration methods. In Fig. 6, a comparison of the performance of a selection of methods using different initial correspondences is reported. The results show that using the initial correspondences of N-ICP [BP13] greatly reduces the resulting error of the subsequent methods.

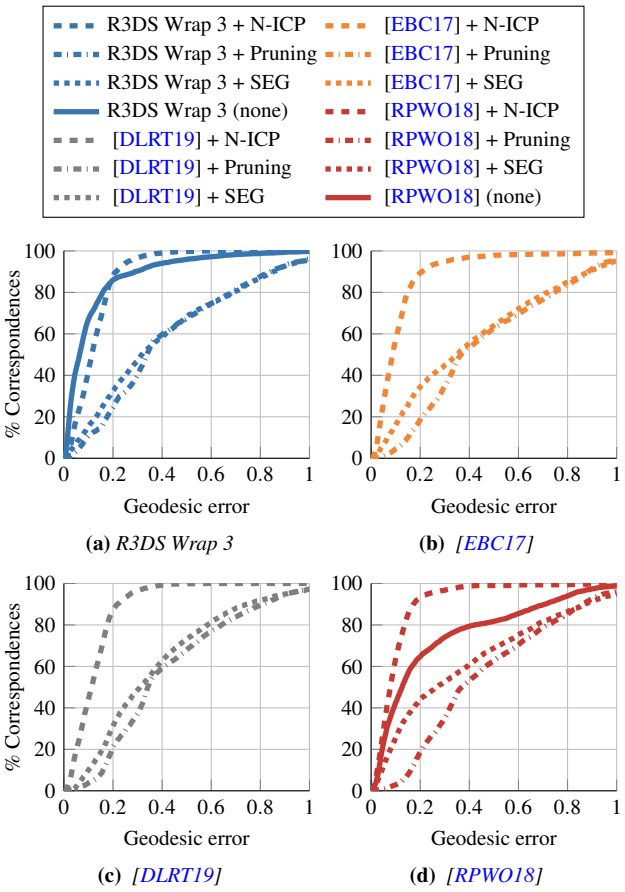

Figure 6: *Comparison of the error of methods initialised with different initial correspondence techniques (N-ICP [BP13], Pruning [TMRL14], SEG [KO19], or none).*

Table 1 summarises the same performance measurements numerically by reporting the *area under the curve* AUC of each method for each set of initial correspondences. All methods perform best using N-ICP [BP13] for initialisation, with the exception of R3DS Wrap 3, which achieved its best results using none. This slight improvement may be a consequence of the coarse-to-fine strategy being able to converge to a more optimal solution. [RPWO18] achieve better results using no correspondence initialisation compared to Pruning [TCL*13] or SEG [KO19], or just using WKS [ASC11] descriptors. In all cases, using SEG [KO19] produced better results than using Pruning [TMRL14].

6.2. Track results

Hereafter, methods initialised with N-ICP [BP13] correspondences are qualified by an asterisk (*). Comprehensive results—measured by AUC—of running each method on each scan—grouped by pose or by internal material—are reported in Table 2. Error curves that complement Table 2 are given in Figs. 7 & 8. Scans that are members of a given column were examined collectively. FARM+ achieved the highest overall accuracy of any method using a semi-automatic approach. Of the fully-automatic methods, [RPWO18]

Table 1: *(left) The AUCs are calculated from the cumulative error curves of the initial correspondence methods reported in Fig. 5. (right) The resulting AUCs of methods using each set of initial correspondences—reported in Fig. 6. Note, [TMRL14] reports an average of 7.4250% vertex correspondences.*

Initialisation method	AUC
N-ICP [BP13]	0.8736
Pruning [TMRL14]	0.6151
SEG [KO19]	0.5323

	Initial correspondences			
Method	N-ICP [BP13]	Pruning [TMRL14]	SEG [KO19]	None
R3DS Wrap 3	0.8763	0.5869	0.6050	0.8837
[EBC17]	0.8829	0.5430	0.5891	N/A
[DLRT19]	0.8771	0.5812	0.6290	N/A
[RPWO18]	0.9015	0.5480	0.6347	0.7609

performed the best. The pose that all methods performed the best on, on average, was the indented pose. The sporadic non-isometry, limited deformation, and slight topological change may have helped. Conversely, twist was the most challenging pose. The most challenging internal material was the risotto. This is likely to be due to the subtle variations induced on local geometry and how the risotto grains affected the way the rabbit bent in the distinct poses. However, excluding the results of [RCB*17] and [VLB*17], all methods achieve a particularly high accuracy on the stretch pose. Although there is a large amount of non-rigid motion, the topology remains consistent, and the non-isometric deformation is relatively simple.

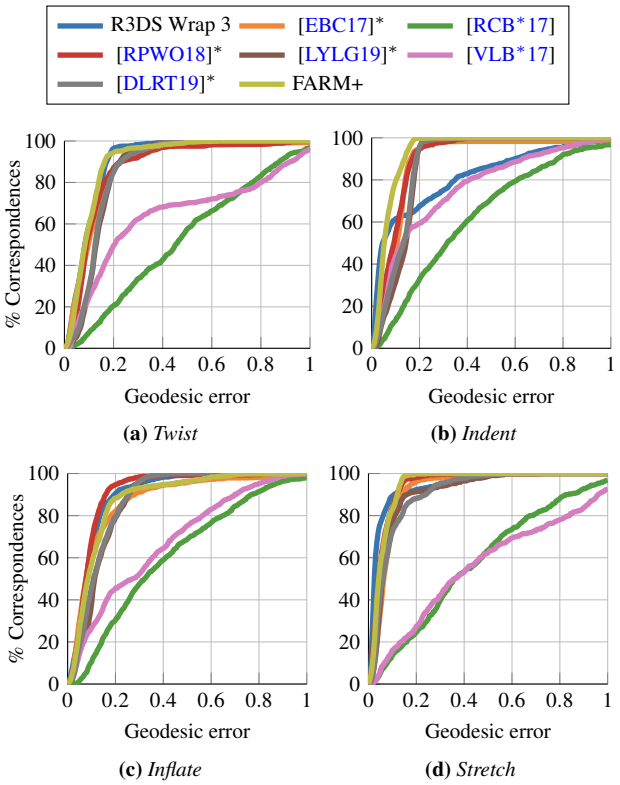

Figure 7: *Cumulative error curves with scans grouped by the pose exhibited.*

R3DS Wrap 3 performed unexpectedly poorly on the indent pose, with the worst performance of any method on scan no. 5. [RCB*17] performed poorly in most scenarios, especially in cases

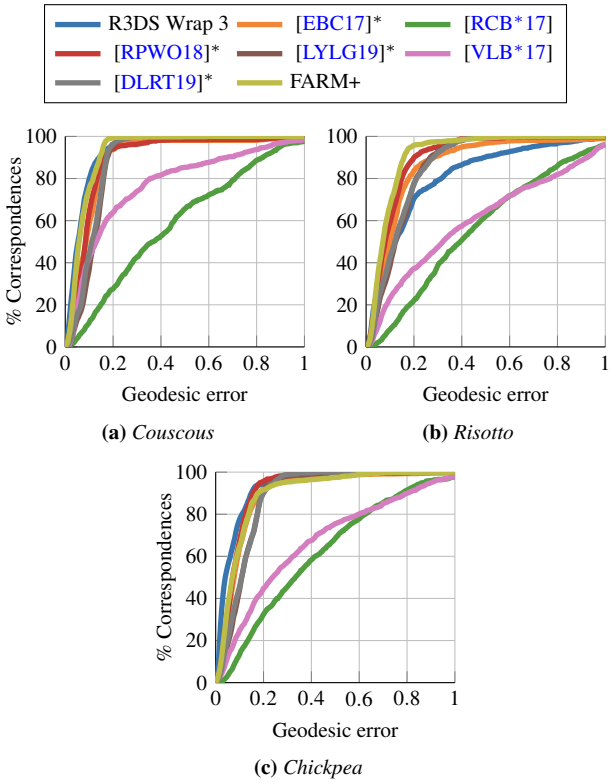

Figure 8: *Cumulative error curves with scans grouped by the internal material selected.*

where a pose exhibited greater non-isometry. For all poses, except inflate, FARM+ converges to 100% correspondence within the smallest geodesic error. [RPWO18] performs the worst on the twist pose, this is understandable as the fused geometry caused by twisting contradicts the assumption of continuity.

As shown in Fig. 9, overall, FARM+ performed the best of all methods. [RPWO18] performed the best of all the fully-automatic approaches. The performance of [DLRT19] is comparable to [LYLG19], with only a slight improvement over the initial N-ICP [BP13] correspondences.

7. Conclusions

We consider the problem of establishing correspondence between shapes with different internal materials in challenging poses. A new dataset was created with high quality texture-based ground-truths.

Table 2: *The total area under the curve of scans grouped by the type of pose exhibited, scans grouped by material, and the overall performance of each method over all scans is reported. The method that achieved the best results in each configuration is emphasised in* ***bold****, and the second best in* ***bold italics****. In the final two rows the mean and standard deviation of each column is reported.*

	Pose				Internal material			
Method	Twist	Indent	Inflate	Stretch	Couscous	Risotto	Chickpea	Overall
R3DS Wrap 3	**0.9025**	0.8129	***0.9005***	0.9418	***0.9265***	0.8071	**0.9277**	0.8837
[EBC17]*	0.8730	0.8830	0.8689	0.9195	0.8944	0.8591	0.8979	0.8829
[RCB*17]	0.5249	0.6247	0.6136	0.5778	0.5840	0.5570	0.6170	0.5863
[RPWO18]*	0.8708	***0.9004***	**0.9070**	***0.9420***	0.8956	***0.8950***	***0.9122***	***0.9015***
[BP13]	0.8620	0.8713	0.8646	0.9087	0.8842	0.8579	0.8813	0.8736
[LYLG19]*	0.8598	0.8719	0.8651	0.9088	0.8822	0.8581	0.8816	0.8733
[VLB*17]	0.6369	0.7610	0.6827	0.5521	0.7660	0.5966	0.6695	0.6692
[DLRT19]*	0.8616	0.8796	0.8712	0.9058	0.8886	0.8604	0.8850	0.8771
FARM+	***0.8999***	**0.9338**	0.8780	**0.9438**	**0.9281**	**0.9111**	0.8992	**0.9113**
Mean	0.7636	0.8023	0.7967	0.7932	0.8126	0.7623	0.8008	
Std.	0.1698	0.1318	0.1307	0.2040	0.1484	0.1581	0.1551	

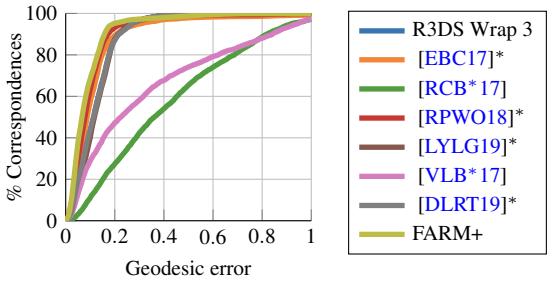

Figure 9: *Overall performance of methods.*

The resulting accuracy achieved by using different correspondence initialisation techniques was investigated. We discover that in this scenario, N-ICP [BP13] performed the best for correspondence initialisation. Of all the methods, FARM+ achieved the greatest accuracy using a semi-automatic approach, while [RPWO18] performed the best of the fully-automatic approaches.

Further work would include an extended investigation into the performance of different shape descriptors and initial correspondence techniques on this dataset.

For many of the methods evaluated, it is unclear how to best optimise the parameters. It was possible to incrementally change the parameters to achieve better results, however this is a time consuming strategy and not necessarily possible in cases where ground-truths are not available. Further investigation into optimal parameters on a range of benchmark datasets is required to give a greater overview of what parameters should be selected in a given scenario.

Acknowledgements

We thank the authors of methods that did not participate but made their code available online [VLB*17, EBC17, RCB*17, RPWO18, Rus19]. We also thank Seana Dykes for helping to fabricate the toy rabbit. This work has been supported by the Cardiff University EPSRC Doctoral Training Partnership [grant ref. EP/N509449/1].

References

[Ado20] ADOBE INC.: Mixamo. https://www.mixamo.com/, July 2008–2020. 4

[ASC11] AUBRY M., SCHLICKEWEI U., CREMERS D.: The Wave Kernel Signature: A quantum mechanical approach to shape analysis. In *International Conference on Computer Vision Workshops* (Nov. 2011), ICCV '11, IEEE, pp. 1626–1633. doi:10.1109/ICCVW.2011.6130444. 4, 5

[ASK*05] ANGUELOV D., SRINIVASAN P., KOLLER D., THRUN S., RODGERS J., DAVIS J.: SCAPE: Shape completion and animation of people. *ACM Transactions on Graphics 24*, 3 (July 2005), 408–416. doi:10.1145/1073204.1073207. 2

[BBC*10] BRONSTEIN A. M., BRONSTEIN M. M., CASTELLANI U., DUBROVINA A., GUIBAS L. J., HORAUD R. P., KIMMEL R., KNOSSOW D., LAVANTE E. V., MATEUS D., OVSJANIKOV M., SHARMA A.: SHREC'10 track: Correspondence finding. In *Eurographics Workshop on 3D Object Retrieval* (Oct. 2010), 3DOR '10, The Eurographics Association, pp. 87–91. doi:10.2312/3DOR/3DOR10/087-091. 2

[BBL*17] BRONSTEIN M. M., BRUNA J., LECUN Y., SZLAM A., VANDERGHEYNST P.: Geometric deep learning: Going beyond euclidean data. *Signal Processing Magazine 34*, 4 (July 2017), 18–42. doi:10.1109/MSP.2017.2693418. 2

[BCBB16] BIASOTTI S., CERRI A., BRONSTEIN A., BRONSTEIN M.: Recent trends, applications, and perspectives in 3D shape similarity assessment. *Computer Graphics Forum 35*, 6 (Oct. 2016), 87–119. doi:10.1111/cgf.12734. 2

[BP13] BOUAZIZ S., PAULY M.: Dynamic 2D/3D registration for the Kinect. In *ACM SIGGRAPH 2013 Courses* (July 2013), SIGGRAPH '13, Association for Computing Machinery, pp. 21:1–21:14. doi:10.1145/2504435.2504456. 2, 3, 5, 6, 7

[BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: FAUST: Dataset and evaluation for 3D mesh registration. In *Conference on Computer Vision and Pattern Recognition* (June 2014), CVPR '14, IEEE, pp. 3794–3801. doi:10.1109/CVPR.2014.491. 2

[DLRT19] DYKE R. M., LAI Y.-K., ROSIN P. L., TAM G. K.: Non-rigid registration under anisotropic deformations. *Computer Aided Geometric Design 71* (May 2019), 142–156. doi:10.1016/j.cagd.2019.04.014. 2, 3, 5, 6, 7

[DSL*19] DYKE R. M., STRIDE C., LAI Y.-K., ROSIN P. L., AUBRY M., BOYARSKI A., BRONSTEIN A. M., BRONSTEIN M. M., CREMERS D., FISHER M., GROUEIX T., GUO D., KIM V. G., KIMMEL R., LÄHNER Z., LI K., LITANY O., REMEZ T., RODOLÀ E., RUSSELL

B. C., SAHILLIOĞLU Y., SLOSSBERG R., TAM G. K. L., VESTNER M., WU Z., YANG J.: Shape correspondence with isometric and non-isometric deformations. In *Eurographics Workshop on 3D Object Retrieval* (May 2019), 3DOR '19, The Eurographics Association, pp. 111–119. doi:10.2312/3dor.20191069. 2, 3

[EBC17] EZUZ D., BEN-CHEN M.: Deblurring and denoising of maps between shapes. *Computer Graphics Forum 36*, 5 (Aug. 2017), 165–174. doi:10.1111/cgf.13254. 2, 3, 4, 5, 6, 7

[EEBC20] EDELSTEIN M., EZUZ D., BEN-CHEN M.: ENIGMA: Evolutionary Non-Isometric Geometry MAtching. *ACM Transactions on Graphics 39*, 4 (July 2020). doi:10.1145/3386569.3392447. 2

[EHA*19] EZUZ D., HEEREN B., AZENCOT O., RUMPF M., BEN-CHEN M.: Elastic correspondence between triangle meshes. *Computer Graphics Forum 38*, 2 (June 2019), 121–134. doi:10.1111/cgf.13624. 2

[JYZ*19] JIANG T., YANG X., ZHANG J., TIAN F., LIU S., XIANG N., QIAN K.: Huber-L_1-based non-isometric surface registration. *The Visual Computer 35*, 6 (May 2019), 935–948. doi:10.1007/s00371-019-01670-1. 2

[KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. *ACM Transactions on Graphics 30*, 4 (July 2011). doi:10.1145/2010324.1964974. 5

[KO19] KLEIMAN Y., OVSJANIKOV M.: Robust structure-based shape correspondence. *Computer Graphics Forum 38*, 1 (Apr. 2019), 7–20. doi:10.1111/cgf.13389. 2, 3, 5, 6

[LRB*16] LÄHNER Z., RODOLÀ E., BRONSTEIN M. M., CREMERS D., BURGHARD O., COSMO L., DIECKMANN A., KLEIN R., SAHILLIOĞLU Y.: Matching of deformable shapes with topological noise. In *Eurographics Workshop on 3D Object Retrieval* (May 2016), 3DOR '16, The Eurographics Association, pp. 55–60. doi:10.2312/3dor.20161088. 2

[LYLG19] LI K., YANG J., LAI Y. K., GUO D.: Robust non-rigid registration with reweighted position and transformation sparsity. *Transactions on Visualization and Computer Graphics 25*, 6 (June 2019), 2255–2269. doi:10.1109/TVCG.2018.2832136. 3, 6, 7

[MMR*19] MELZI S., MARIN R., RODOLÀ E., CASTELLANI U., REN J., POULENARD A., WONKA P., OVSJANIKOV M.: Matching humans with different connectivity. In *Eurographics Workshop on 3D Object Retrieval* (May 2019), 3DOR '19, The Eurographics Association, pp. 121–128. doi:10.2312/3dor.20191070. 1

[MMRC19] MARIN R., MELZI S., RODOLÀ E., CASTELLANI U.: High-resolution augmentation for automatic template-based matching of human models. In *International Conference on 3D Vision* (Sept. 2019), 3DV '19, IEEE, pp. 230–239. doi:10.1109/3DV.2019.00034. 4

[MMRC20] MARIN R., MELZI S., RODOLÀ E., CASTELLANI U.: FARM: Functional automatic registration method for 3D human bodies. *Computer Graphics Forum 39*, 1 (Feb. 2020), 160–173. doi:10.1111/cgf.13751. 4

[MOR*18] MELZI S., OVSJANIKOV M., ROFFO G., CRISTANI M., CASTELLANI U.: Discrete time evolution process descriptor for shape analysis and matching. *ACM Transactions on Graphics 37*, 1 (Jan. 2018), 4:1–4:18. doi:10.1145/3144454. 4

[MRR*19] MELZI S., REN J., RODOLÀ E., SHARMA A., WONKA P., OVSJANIKOV M.: ZoomOut: Spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics 38*, 6 (Nov. 2019), 155:1–155:14. doi:10.1145/3355089.3356524. 4

[NO17] NOGNENG D., OVSJANIKOV M.: Informative descriptor preservation via commutativity for shape matching. *Computer Graphics Forum 36*, 2 (May 2017), 259–267. doi:10.1111/cgf.13124. 3

[OBCS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: A flexible representation of maps between shapes. *ACM Transactions on Graphics 31*, 4 (July 2012), 30:1–30:11. doi:10.1145/2185520.2185526. 4

[RCB*17] RODOLÀ E., COSMO L., BRONSTEIN M. M., TORSELLO A., CREMERS D.: Partial functional correspondence. *Computer Graphics Forum 36*, 1 (Jan. 2017), 222–236. doi:10.1111/cgf.12797. 2, 3, 6, 7

[RCL*17] RODOLÀ E., COSMO L., LITANY O., BRONSTEIN M. M., BRONSTEIN A. M., AUDEBERT N., HAMZA A. B., BOULCH A., CASTELLANI U., DO M. N., DUONG A.-D., FURUYA T., GASPARETTO A., HONG Y., KIM J., SAUX B. L., LITMAN R., MASOUMI M., MINELLO G., NGUYEN H.-D., NGUYEN V.-T., OHBUCHI R., PHAM V.-K., PHAN T. V., REZAEI M., TORSELLO A., TRAN M.-T., TRAN Q.-T., TRUONG B., WAN L., ZOU C.: Deformable shape retrieval with missing parts. In *Eurographics Workshop on 3D Object Retrieval* (Apr. 2017), 3DOR '17, The Eurographics Association, pp. 85–94. doi:10.2312/3dor.20171057. 2

[RDP99] ROBINETTE K. M., DAANEN H., PAQUET E.: The CAESAR project: A 3D surface anthropometry survey. In *International Conference on 3D Digital Imaging and Modeling* (Aug. 1999), 3DIM '99, IEEE, pp. 380–386. doi:10.1109/IM.1999.805368. 2

[RPWO18] REN J., POULENARD A., WONKA P., OVSJANIKOV M.: Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics 37* (Dec. 2018), 248:1–248:16. doi:10.1145/3272127.3275040. 3, 4, 5, 6, 7

[Rus19] RUSSIAN3DSCANNER: Wrap 3.4. https://www.russian3dscanner.com/, June 2019. 4, 7

[Sah19] SAHILLIOĞLU Y.: Recent advances in shape correspondence. *The Visual Computer* (Sept. 2019). doi:10.1007/s00371-019-01760-0. 2

[SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A concise and provably informative multi-scale signature based on heat diffusion. In *Proceedings of the Symposium on Geometry Processing* (Aug. 2009), vol. 28 of *SGP '09*, Eurographics Association, pp. 1383–1392. doi:10.1111/j.1467-8659.2009.01515.x. 4

[SPF19] SCHMIDT F., PHILLIPS F., FLEMING R. W.: Visual perception of shape-transforming processes: 'shape scission'. *Cognition 189* (Aug. 2019), 167–180. doi:10.1016/j.cognition.2019.04.006. 1

[TCL*13] TAM G. K. L., CHENG Z., LAI Y., LANGBEIN F. C., LIU Y., MARSHALL D., MARTIN R. R., SUN X., ROSIN P. L.: Registration of 3D point clouds and meshes: A survey from rigid to nonrigid. *Transactions on Visualization and Computer Graphics 19*, 7 (July 2013), 1199–1217. doi:10.1109/TVCG.2012.310. 2, 5

[TMRL14] TAM G. K. L., MARTIN R. R., ROSIN P. L., LAI Y.-K.: Diffusion pruning for rapidly and robustly selecting global correspondences using local isometry. *ACM Transactions on Graphics 33*, 1 (Feb. 2014), 4:1–4:17. doi:10.1145/2517967. 2, 3, 5, 6

[TSDS10] TOMBARI F., SALTI S., DI STEFANO L.: Unique signatures of histograms for local surface description. In *European Conference on Computer Vision* (Sept. 2010), ECCV '10, Springer, pp. 356–369. doi:10.1007/978-3-642-15558-1_26. 3

[vKZHCO11] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. *Computer Graphics Forum 30*, 6 (July 2011), 1681–1707. doi:10.1111/j.1467-8659.2011.01884.x. 2

[VLB*17] VESTNER M., LÄHNER Z., BOYARSKI A., LITANY O., SLOSSBERG R., REMEZ T., RODOLÀ E., BRONSTEIN A., BRONSTEIN M., KIMMEL R., CREMERS D.: Efficient deformable shape correspondence via kernel matching. In *International Conference on 3D Vision* (Oct. 2017), 3DV '17, IEEE, pp. 517–526. doi:10.1109/3DV.2017.00065. 2, 3, 4, 6, 7

[VLR*17] 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 *Conference on Computer Vision and Pattern Recognition* (July 2017), CVPR '17, IEEE, pp. 6681–6690. doi:10.1109/CVPR.2017.707. 2, 4