

TACO: a benchmark for connectivity-invariance in shape correspondence

Simone Pedico¹, Simone Melzi¹, and Filippo Maggioli¹

¹University of Milano-Bicocca, Italy




Figure 1: A visual representation of our processing pipeline. Given a dataset for shape matching applications like TOSCA (on the left), we first apply a non connectivity-preserving remeshing for altering the meshes connectivities, making them different even among shapes of the same class. Then, by exploiting the 3D alignment between the original shape and its remesh, as well as the ground truth correspondence provided by TOSCA, we build a new ground truth correspondence between the remeshed surfaces. As a result, we obtain our new dataset TACO (on the right).

Abstract

In real-world scenarios, a major limitation for shape-matching datasets is represented by having all the meshes of the same subject share their connectivity across different poses. Specifically, similar connectivities could provide a significant bias for shape matching algorithms, simplifying the matching process and potentially leading to correspondences based on the recurring triangle patterns rather than geometric correspondences between mesh parts. As a consequence, the resulting correspondence may be meaningless, and the evaluation of the algorithm may be misled. To overcome this limitation, we introduce TACO, a new dataset where meshes representing the same subject in different poses do not share the same connectivity, and we compute new ground truth correspondences between shapes. We extensively evaluate our dataset to ensure that ground truth isometries are properly preserved. We also use our dataset for validating state-of-the-art shape-matching algorithms, verifying a degradation in performance when the connectivity gets altered.

CCS Concepts

• **Computing methodologies** → *Shape analysis*; • **Theory of computation** → *Computational geometry*;

1. Introduction

Shape matching is a major research topic in the field of computer graphics. Given two meshes, the goal of a shape matching task is to identify a correspondence between geometric elements (*i.e.*, vertices, edges, faces) of the first mesh onto the second. The problem of finding such correspondence is, in its general setting, known to be NP-hard [BLG*21]. Current approaches for solving shape matching tasks are mainly evolving in two directions: either using geometric deep learning techniques and data-driven approaches to learn features for correspondence estimation [VM24; MRMO20], or relying on the functional maps framework to infer the point-wise correspondence from a linear map among functional spaces over the surfaces [OBS*12; MRR*19; MMO*21; MBRM24]. In both these scenarios, shape matching algorithms make massive use of descriptors, a class of scalar functions over the surface with the property of identifying certain features, localizing in certain areas, or providing an almost unique signature of each vertex. While being very useful and effective, shape descriptors are generally derived from discrete differential operators, and consequently they strongly depend on the mesh resolution and connectivity [SOG09; ASC11]. It follows that the performance of the algorithms is influenced by the connectivity and the quality of the shape.

The major downside in this regard is that shape matching techniques can be biased by similar connectivities, recognizing recurring triangle patterns instead of semantically and geometrically meaningful corresponding features. As a consequence, evaluating a method on datasets where shapes of the same class share their connectivity could produce misleading results. To overcome this limitation, we provide a processing pipeline to alter the connectivities of a shape matching dataset and build new robust ground truth correspondences, resulting in a dataset where the connectivity varies across different meshes of the same subject. With this process, we aim to increase the challenge provided by existing datasets for shape matching, reducing the bias introduced by shared connectivities and making the evaluation process for matching algorithms more robust. We apply our procedure to the well-known TOSCA dataset [BBK08], a longstanding benchmark for shape matching techniques, obtaining the new dataset TACO (TOSCA Altered COnectivities). In TACO, shapes representing the same subject in different poses are discretized with significantly different connectivities, and the dataset comes with ground truth correspondences between each pair of shapes belonging to the same class. We evaluate the quality of these correspondences by comparing them with the original ground truth from TOSCA, and we verify the increase in challenge by testing state-of-the-art solutions against both TOSCA and TACO.

In summary, the contribution of this paper is twofold:

- we introduce a novel pipeline to alter existing shape matching dataset and increase the challenge of finding a correspondence by varying the connectivity between shapes, while at the same time preserving the overall geometry of the surface and a robust ground truth correspondence between shape pairs;
- we apply our procedure to the well-established TOSCA dataset, producing the new and more challenging dataset TACO, which we use for evaluating state-of-the-art methods and verify their robustness to varying connectivities.

2. Related work

Due to the relevance of the shape matching task, researchers have designed and developed several datasets for evaluating matching algorithms, each presenting its own features and limitations.

By starting with real scans of various humans, Angelov *et al.* [ASK*05] developed a data-driven approach for deforming a template shape into the scans, producing SCAPE, a dataset of synthetic human shapes in a one-to-one correspondence. While offering a certain variety, the dataset only contains human shapes with the same connectivity and at low resolution.

Similarly, Bogo *et al.* [BRLB14] produced the FAUST dataset by scanning 10 different subjects in 30 different poses, totalling 300 high-resolution, triangulated, non-watertight, and non-manifold meshes. These meshes are subdivided into a training set, composed of 100 meshes with ground truth available and comprising a set of registrations aligned to a template shape, and a test set, composed of 200 shapes. The FAUST dataset addresses some of the limitations of SCAPE by increasing the resolution of the meshes and the variety of the dataset. Nonetheless, the dataset is still focused on human shapes only, and the low-resolution registration template shares the same connectivity across all the training shapes. Later, the dataset has been expanded with DFAUST (Dynamic FAUST) [BRPB17], which provides dynamic sequences of human body scans captured at 60 fps.

By defining a collection of shapes from other known dataset and registering them to a common template, Melzi *et al.* [MMR*19] produced the SHREC19 dataset for shape correspondences. The dataset contains human shapes at very different resolutions, including both synthetic clean shapes and noisy scans presenting artifacts, holes and clutters. While being a well-established benchmark for shape matching tasks, the SHREC19 dataset is still completely composed of human shapes and more focused on non-isometric deformations (*i.e.*, correspondences between different subjects). Later, Dyke *et al.* [DZL*20] designed the SHREC20 dataset, counting a richer set of subjects in correspondence. Nonetheless, the SHREC20 dataset only provides sparse correspondences and is oriented towards strongly non-isometric matching tasks.

More recently, Li *et al.* [LTT*21] released the Deforming Things 4D (DT4D) dataset. DT4D contains synthetic data for 1972 animation sequences spanning across 31 humanoids and animal subjects. The DT4D dataset is very rich and varied, but it only contains meshes at low resolution. Furthermore, correspondences are provided only for meshes in the same animation sequence, without connectivity variations.

Finally, the TOSCA dataset [BBK08] is a well-established dataset for isometric shape matching, and one of the most widely used for evaluating matching algorithms. It contains a collection of 9 synthetic human and animal subjects at various resolutions, each represented in different poses. All the poses of the same subject are in one-to-one isometric correspondence, but sharing the same connectivity. The goal of our work is to address this main limitation of the TOSCA dataset, providing a new dataset for isometric shape matching between meshes with different connectivity, effectively increasing the challenge posed by the dataset and evaluating the robustness of matching algorithms to variations of connectivity.

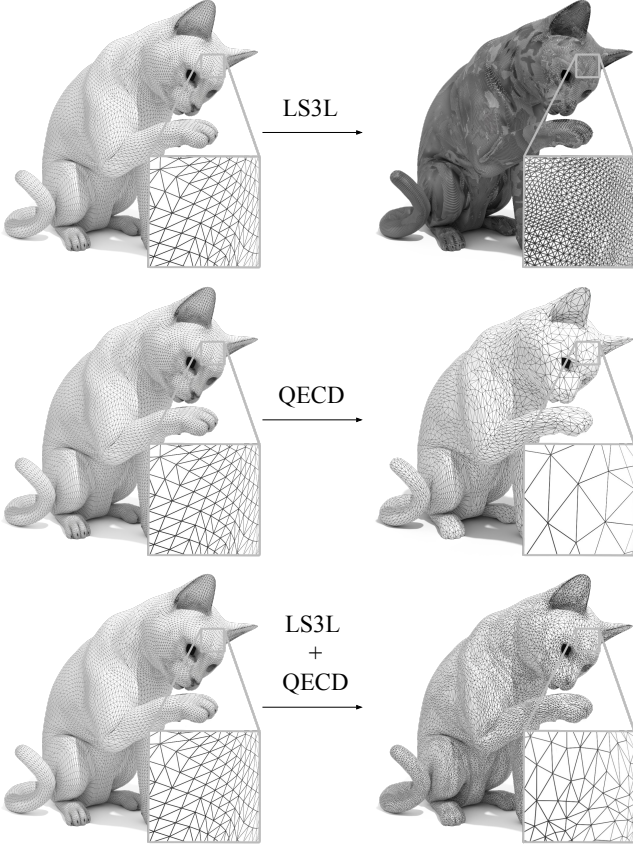


Figure 2: Pipeline of the remeshing procedure. Starting from the original TOSCA mesh (left), we apply a LS3L subdivision (top), followed by a simplification through QECD (middle). As a result, we obtain a mesh representing the same geometry as the original one, but with a different connectivity (bottom).

3. Proposed Method

Our pipeline, summarized in Figure 1, is divided into two parts. Firstly, we use a sequence of MeshLab [CCC*08] filters to obtain a remeshing of all the meshes of TOSCA, changing their connectivity and creating a new collection of shapes. Then, by exploiting the 3D alignment between the meshes of TOSCA and our collection, we obtain a point-wise correspondence between the original surfaces and the remeshed shapes. By combining these correspondences with the ground truths in TOSCA, we are able to compute robust point-wise correspondences between pairs in our collection of shapes, completing our dataset TACO.

3.1. Remeshing

Our remeshing pipeline relies on the MeshLab implementation of two well-known algorithms: LS3 Loop (LS3L) [BGS10] and Quadric Edge Collapse Decimation (QECD) [GH97]. By applying these filters in sequence, we first obtain an high-resolution smoothed version of the shape with LS3L, and then we simplify the

geometry with QECD obtaining a mesh with a similar vertex count to the original. Both the algorithms are designed to preserve the original surface as much as possible, hence resulting in an almost perfect 3D alignment between the source shape and its remesh. At the same time, the result of both LS3L and QECD do not only depend on the intrinsic geometry, but also on the 3D embedding, which guarantees a change in connectivity among different poses of the same subject. Figure 2 shows the effect of applying the two algorithms individually and in sequence.

By conducting extensive tests on different configurations of parameters, we find that by using 3 iterations of LS3L and fixing the face quality [FB99] threshold for QECD to 0.3 we achieve a significant variation of connectivity. We evaluate the connectivity variation by comparing the number of vertices and triangles among shapes of the same class, as a different number of geometric elements is more likely to correspond to a change in connectivity. The evaluation is also verified through a visual inspection.

3.2. Matching

For completing our shape matching dataset, we exploit the 3D alignment between the original TOSCA meshes and our remeshed shapes, as well as the original one-to-one TOSCA ground truth, to retrieve a ground truth correspondence between the meshes in our dataset. This process is carried out in two steps, which involve finding first the correspondences between the meshes from TOSCA and TACO, and then extending these correspondences to pairs of meshes in our dataset.

3.2.1. Mesh-Remesh Correspondence

In this phase, given a mesh $M = (V, E, T)$ from TOSCA and its remesh $\hat{M} = (\hat{V}, \hat{E}, \hat{T})$, we calculate the vertex-vertex correspondence $\pi_{\hat{M}M} : \hat{V} \rightarrow V$ between them. To accomplish this task, we exploit the fact that our remeshing procedure does not alter the 3D embedding of the original mesh, and thus M and \hat{M} are aligned in the 3D space. We perform a nearest neighbor search between the vertices of M and \hat{M} , and for sake of efficiency we use the FLANN algorithm [ML09]. In a similar fashion, we also compute the inverse correspondence $\pi_{M\hat{M}} : V \rightarrow \hat{V}$.

3.2.2. Remesh-Remesh Correspondence

The objective of the second step is to calculate the correspondences between two meshes from TACO belonging to the same class, for which we will take advantage of the existing ground truth correspondences provided by TOSCA. Let $M = (V_M, E_M, T_M)$ and $N = (V_N, E_N, T_N)$ be two meshes from TOSCA, both belonging to the same class, and let $\hat{M} = (\hat{V}_M, \hat{E}_M, \hat{T}_M)$ and $\hat{N} = (\hat{V}_N, \hat{E}_N, \hat{T}_N)$ be their remeshed counterparts in TACO, respectively. Since the TOSCA dataset provides the one-to-one ground truth correspondence $\pi_{MN} : V_M \rightarrow V_N$ (with inverse $\pi_{NM} = \pi_{MN}^{-1}$), we can compute the ground truth correspondences between \hat{M} and \hat{N} in both directions by function composition. Namely, we have $\pi_{\hat{M}\hat{N}} = \pi_{\hat{N}\hat{N}} \circ \pi_{MN} \circ \pi_{\hat{M}M}$ and $\pi_{\hat{N}\hat{M}} = \pi_{\hat{M}\hat{M}} \circ \pi_{NM} \circ \pi_{\hat{N}N}$. This composition of functions for retrieving ground truth correspondences in TACO is depicted in Figure 1.

4. Results

To evaluate our dataset, we compare the results of state-of-the-art shape matching algorithms on both TOSCA and TACO, and we verify the decrease in performance. In order to ensure a meaningful comparison, we also evaluate the similarity between the original mesh and the remeshed shapes, both intrinsically and extrinsically. Furthermore, we ensure that the ground truth correspondences that we provide have comparable quality with the ground truth correspondences in TOSCA.

4.1. Quality of meshes

For each original mesh M and its remesh \hat{M} , we evaluate the quality of the remeshing process through a wide set of measurements. We notice that both the original surfaces and our remeshed shapes are all 2-manifold, made of a single connected component, and have the same number of boundary components. The `gorilla` mesh constitutes the only exception: in the original TOSCA dataset it presents non-manifold vertices and various disconnected components, but after our remeshing process it becomes 2-manifold and most of the smallest disconnected components are removed.

Our first step is to verify that the remeshed shape $\hat{M} = (\hat{V}, \hat{E}, \hat{T})$ represents the same surface as the original mesh $M = (V, E, T)$. For this task, we employ both the Hausdorff distance d_H and the Chamfer distance d_C for triangular meshes. They are defined as

$$d_H(M, \hat{M}) = \max \left(\max_{\hat{v} \in \hat{V}} \min_{t \in T} d(\hat{v}, t), \max_{v \in V} \min_{\hat{t} \in \hat{T}} d(v, \hat{t}) \right), \quad (1)$$

$$d_C(M, \hat{M}) = \frac{1}{|\hat{V}|} \sum_{\hat{v} \in \hat{V}} \min_{t \in T} d^2(\hat{v}, t) + \frac{1}{|V|} \sum_{v \in V} \min_{\hat{t} \in \hat{T}} d^2(v, \hat{t}), \quad (2)$$

where $d(v, t)$ is the usual point-triangle distance. Before computing the distances, we translate all the shapes so that their center of mass is the origin of the axes, and we scale them inside the cube $[-1, 1]^3$. The preprocessing step ensures that all the shapes are scaled uniformly and the comparison is meaningful. Figure 3 shows the results of our measurement. We group the values between the different classes that constitute the dataset and compute the average. For all the classes, we register very small distances, ensuring that the remeshed shapes are effectively a close approximation of the original surface. Due to the large amount of small disconnected components in the original mesh, the `gorilla` class registers a larger distance. In the case of the `wolf` class, instead, we notice that the low resolution negatively affects the computation of the Chamfer distance. Indeed, a smaller vertex count prevents some vertices to be placed close to the original surface, and a larger error on these outliers cannot be compensated by averaging across the small number of vertices. Nonetheless, for both classes the distance is small enough to be considered negligible, as shown in the example from Figure 4. Here we apply a procedural pattern to two shapes from the `gorilla` and `wolf` class. The pattern is very sensitive to coordinate variation [MBMR22], demonstrating the similarity between the original meshes and the remeshed shapes, despite the larger Hausdorff and Chamfer distances.

As an additional measure for the similarity between the original mesh M and its remesh \hat{M} , we also evaluate their spectral similarity. The classical result from Reuter *et al.* [RWPO6] showed that the




Figure 3: Chamfer distance and Hausdorff distance between the original meshes from TOSCA and the remeshed shapes from TACO. The distances are grouped by class and averaged. The Chamfer distance is shown in logarithmic scale.




Figure 4: Application of a very coordinate-sensitive procedural pattern to the `gorilla` and `wolf` meshes from TACO (left) and TOSCA (right).

Laplacian spectrum is a good measure of similarity between isometric shapes, and thus, if two meshes represents the same surface, we expect their Laplacian eigenvalues to be very similar. However, instead of directly aligning the spectra, we adopt the technique introduced by Moschella *et al.* [MMC*22], which uses the vector of the first 100 spectral offsets $\text{off}(\lambda_i) = \lambda_i - \lambda_{i-1}$ to compensate for the Weyl's law [Wey11]. The results, grouped by class and averaged, are shown in Figure 5. We notice that `centaur` and `wolf` register a larger error. This is expected, since they are the shapes with lowest resolution in the dataset, and hence moving around their vertices has a larger impact on the spectrum.




Figure 5: Average relative difference between the first 100 spectral offset of a mesh M and its remesh \hat{M} . Results are grouped by class and averaged.

4.2. Preserving the isometries

Preserving the underlying surface is not enough to guarantee the quality of a shape matching dataset. With our dataset, we also provide a ground truth correspondence between the remeshed shapes. Thus, we perform a set of experiments to validate both the quality of the correspondence and the steps through which it is computed.

Since our ground truth correspondence is obtained via a nearest neighbor correspondence $\pi_{M\hat{M}}$ between an original mesh M and its remeshed counterpart \hat{M} , we have to ensure that this correspondence is semantically correct. To accomplish this task, we measure the geodesic distortion induced by our correspondence. By definition, if $\pi_{M\hat{M}}$ is an isometry it must preserve the distances; namely

$$d_M(x, y) = d_{\hat{M}}(\pi_{M\hat{M}}(x), \pi_{M\hat{M}}(y)), \quad (3)$$

where $d_S(p, q)$ represents the geodesic distance over the surface S between points p and q .

We compute the matrix of geodesic distances among the vertices V of M , and we do the same for the vertices \hat{V} of \hat{M} . Then, we use the correspondence $\pi_{M\hat{M}}$ to evaluate the geodesic distortion between pairs of vertices and compute the average. For a uniform comparison, we normalize the difference by the shape diameter $\delta_M = \max_{v_1, v_2 \in V} d_M(v_1, v_2)$, obtaining

$$\frac{1}{|V|^2} \sum_{v_1, v_2 \in V} \text{abs} \left(\frac{d_M(v_1, v_2) - d_{\hat{M}}(\pi_{M\hat{M}}(v_1), \pi_{M\hat{M}}(v_2))}{\delta_M} \right). \quad (4)$$

Computing the matrices of all the geodesic distance pairs and comparing them would be computationally unfeasible, due to the density of the meshes. Thus, we approximate the computation with 5000 vertices selected with a geodesic farthest point sampling, to ensure uniform coverage of the surface. The results for all the 80 shapes in the dataset are summarized in Figure 6. We see that for most of the dataset the average geodesic distortion is under 1.5% of the shape diameter, and for all the shapes it is always less than 2.2% of δ_M , proving that the correspondence we compute between the mesh M and its remesh \hat{M} is very close to an isometry.

For evaluating the quality of the correspondences between two shapes M_1, M_2 of the same class in different poses, we rely on the




Figure 6: Distribution of the average geodesic distortion across the shapes in the dataset.




Figure 7: Distribution of the Laplacian commutativity and the map orthogonality across all the dataset pairs for TACO (blue) and TOSCA (orange).

metrics discussed by Lescoat *et al.* [LLT*20]: the Laplacian commutativity $\|\cdot\|_L$ and the map orthogonality $\|\cdot\|_D$. These metrics are defined as

$$\|\mathbf{C}\|_L = \frac{\|\mathbf{C}\mathbf{A}_1 - \mathbf{A}_2\mathbf{C}\|_2}{\|\mathbf{C}\|_2}, \quad (5)$$

$$\|\mathbf{C}\|_D = \frac{\|\mathbf{C}^\top \mathbf{C} - \mathbf{I}_k\|_2}{\sqrt{k}}, \quad (6)$$

where \mathbf{A}_i is the diagonal matrix that has the Laplacian eigenvalues of M_i on its diagonal and \mathbf{I}_k is the $k \times k$ identity matrix.

For each of the 420 possible pairs of original shapes M_1 and M_2 in the same class, we compute their ground truth functional map \mathbf{C} . Then, given the remeshed shapes \hat{M}_1 and \hat{M}_2 , we compute the functional map $\hat{\mathbf{C}}$ inferred by our computed correspondence. For both \mathbf{C} and $\hat{\mathbf{C}}$, we compute the Laplacian commutativity and map orthogonality. Figure 7 shows that the measures have similar distribution and the results are comparable, meaning that the correspondence between the remeshed shapes is about the same quality as the correspondence between the original shape.

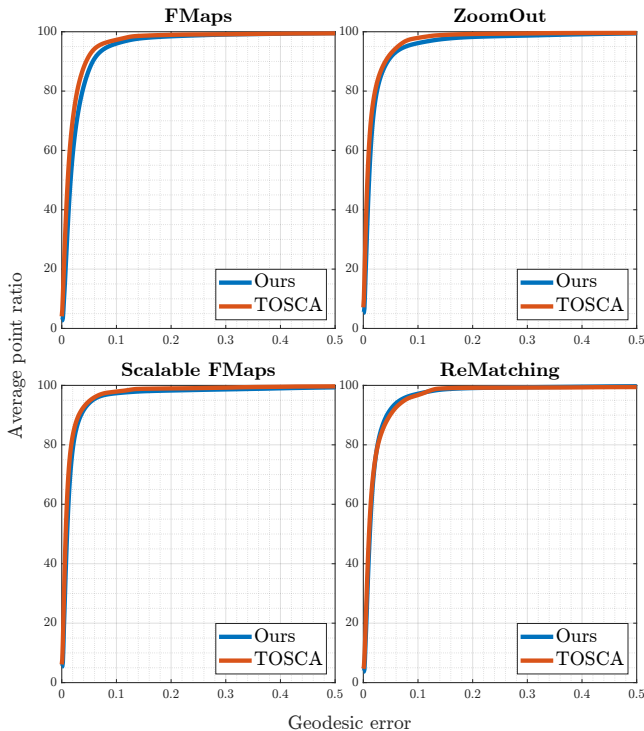


Figure 8: Accuracy curves for all the tested methods on both TOSCA (orange) and TACO (ours, in blue). The curves are averaged across all the 420 possible pairs of shape from the same class.

Method	Dataset	AGE($\cdot 10^{-2}$) ↓	AUC ↑
FMaps	TACO	3.01	47.06
FMaps	TOSCA	2.33	47.77
ZoomOut	TACO	2.50	47.58
ZoomOut	TOSCA	1.85	48.22
Scalable FMaps	TACO	2.24	47.85
Scalable FMaps	TOSCA	1.72	48.34
ReMatching	TACO	2.25	47.79
ReMatching	TOSCA	2.34	47.77

Table 1: Average geodesic error (lower is better) and area under the accuracy curve (higher is better) for each combination of tested methods and datasets.

4.3. Shape matching challenge

To assess the impact of a varying connectivity in shape matching applications, we evaluate state-of-the-art methods for computing isometric correspondences to both TOSCA and TACO. For this task, we used the well-established approaches like functional maps with product preservation (FMaps) [NO17] and ZoomOut [MRR*19], as well as more recent techniques like Scalable Functional Maps [MO23] and ReMatching [MBRM24].

The results of our experiments are summarized in Figure 8, where we show the accuracy curves of each method tested on both TOSCA and TACO. We also provide the area under the ac-

curacy curve (AUC) and the average geodesic error (AGE) in Table 1. The results show that using our dataset reduces the quality of the correspondence, increasing the AGE by 29% (FMaps) to 35% (ZoomOut). Notably, ReMatching achieves better results on TACO than on TOSCA. We guess this being due to the remeshing step applied by the algorithm prior to the functional map computation, which effectively invalidates the effects of changing the original connectivity.

5. Limitations and future directions

While altering the connectivity between the source and target shape results in a more challenging shape matching task, some methods are already overcoming this additional difficulty. On top of that, our dataset specifically targets correspondence tasks on 3D triangular meshes, and does not take into account other representations. To overcome these limitations, we plan to extend our dataset with several additional challenges:

- adding shapes with a larger variation in vertex count can substantially hinder the accuracy of many methods, as it would forbid a bijectivity (even approximate) in the correspondence;
- by introducing noisy and partial shapes, as well as strong deformations, we would be able to target other applications like non-isometric and partial shape matching;
- finding a correspondence between 3D shapes is not limited to triangular mesh, and thus we plan to extend our dataset by including other types of data and representations, such as point clouds and implicit surfaces.

With these future developments in mind, we believe that TACO has the potential for becoming a unified benchmark for 3D matching algorithms.

6. Conclusions

We introduced the new TACO dataset for shape matching applications. Our dataset is built on top of the well-known TOSCA dataset, and it is obtained through a non-connectivity preserving remeshing of the original shapes. By exploiting the 3D alignment of the original and remeshed surfaces, as well as the original ground truth correspondence in TOSCA, we also produced a new ground truth for TACO. Through a set of experiments, we validated the quality of the correspondences, and proved that they are comparable with the original ground truth in TOSCA.

By evaluating state-of-the-art shape matching algorithms on both TOSCA and TACO, we proved that finding a correspondence between meshes with different connectivities is a more challenging task, and thus that our dataset can find an application in making shape matching algorithm more robust.

Acknowledgments

This work was partially supported by the MUR for REGAINS, the Department of Excellence DISCo at the University of Milano-Bicocca, and by the PRIN project GEOPRIDE Prot. 2022-NAZ-0115, CUP H53D23003400001.

References

- [ASC11] AUBRY, MATHIEU, SCHLICKWEI, ULRICH, and CREMERS, DANIEL. “The wave kernel signature: A quantum mechanical approach to shape analysis”. *2011 IEEE international conference on computer vision workshops (ICCV workshops)*. IEEE. 2011, 1626–1633 2.
- [ASK*05] ANGUELOV, DRAGOMIR, SRINIVASAN, PRAVEEN, KOLLER, DAPHNE, et al. “Scape: shape completion and animation of people”. *ACM SIGGRAPH 2005 Papers*. 2005, 408–416 2.
- [BBK08] BRONSTEIN, ALEX, BRONSTEIN, MICHAEL, and KIMMEL, RON. *Numerical Geometry of Non-Rigid Shapes*. New York, NY: Springer, 2008 2.
- [BGS10] BOYÉ, SIMON, GUENNEBAUD, GAEL, and SCHLICK, CHRISTOPHE. “Least squares subdivision surfaces”. *Computer Graphics Forum*. Vol. 29. 7. Wiley Online Library. 2010, 2021–2028 3.
- [BLG*21] BENKNER, MARCEL SEELBACH, LÄHNER, ZORAH, GOLYANIK, VLADISLAV, et al. “Q-match: Iterative shape matching via quantum annealing”. *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, 7586–7596 2.
- [BRLB14] BOGO, FEDERICA, ROMERO, JAVIER, LOPER, MATTHEW, and BLACK, MICHAEL J. “FAUST: Dataset and evaluation for 3D mesh registration”. *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. Piscataway, NJ, USA: IEEE, June 2014 2.
- [BRPB17] BOGO, FEDERICA, ROMERO, JAVIER, PONS-MOLL, GERARD, and BLACK, MICHAEL J. “Dynamic FAUST: Registering human bodies in motion”. *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017, 6233–6242 2.
- [CCC*08] CIGNONI, PAOLO, CALLIERI, MARCO, CORSINI, MASSIMILIANO, et al. “Meshlab: an open-source mesh processing tool.” *Eurographics Italian chapter conference*. Vol. 2008. Salerno, Italy. 2008, 129–136 3.
- [DZL*20] DYKE, ROBERTO M., ZHOU, FENG, LAI, YU-KUN, et al. “SHREC 2020 Track: Non-rigid Shape Correspondence of Physically-Based Deformations”. *Eurographics Workshop on 3D Object Retrieval*. Ed. by SCHRECK, TOBIAS, THEOHARIS, THEOHARIS, PRATIKAKIS, IOANNIS, et al. The Eurographics Association, 2020. ISBN: 978-3-03868-126-7. DOI: [10.2312/3dor.20201161](https://doi.org/10.2312/3dor.20201161) 2.
- [FB99] FREY, PASCAL J and BOROUCAKI, HOUMAN. “Surface mesh quality evaluation”. *International journal for numerical methods in engineering* 45.1 (1999), 101–118 3.
- [GH97] GARLAND, MICHAEL and HECKBERT, PAUL S. “Surface simplification using quadric error metrics”. *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. 1997, 209–216 3.
- [LLT*20] LESCOAT, THIBAUT, LIU, HSUEH-TI DEREK, THIERY, JEAN-MARC, et al. “Spectral mesh simplification”. *Computer Graphics Forum*. Vol. 39. 2. Wiley Online Library. 2020, 315–324 5.
- [LTT*21] LI, YANG, TAKEHARA, HIKARI, TAKETOMI, TAKAFUMI, et al. “4dcomplete: Non-rigid motion estimation beyond the observable surface”. *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, 12706–12716 2.
- [MBMR22] MAGGIOLI, FILIPPO, BAIERI, DANIELE, MELZI, SIMONE, and RODOLÀ, EMANUELE. “Newton’s Fractals on Surfaces via Bicomplex Algebra”. *ACM SIGGRAPH 2022 Posters*. SIGGRAPH ’22. Vancouver, BC, Canada: Association for Computing Machinery, 2022. ISBN: 9781450393614. DOI: [10.1145/3532719.3543211](https://doi.org/10.1145/3532719.3543211) 4.
- [MBRM24] MAGGIOLI, FILIPPO, BAIERI, DANIELE, RODOLÀ, EMANUELE, and MELZI, SIMONE. “ReMatching: Low-Resolution Representations for Scalable Shape Correspondence”. *European Conference on Computer Vision*. Springer. 2024 2, 6.
- [ML09] MUJA, MARIUS and LOWE, DAVID G. “Fast approximate nearest neighbors with automatic algorithm configuration.” *VISAPP (1)* 2.331–340 (2009), 2 3.
- [MMC*22] MOSCHELLA, LUCA, MELZI, SIMONE, COSMO, LUCA, et al. “Learning spectral unions of partial deformable 3D shapes”. *Computer Graphics Forum*. Vol. 41. 2. Wiley Online Library. 2022, 407–417 4.
- [MMO*21] MAGGIOLI, FILIPPO, MELZI, SIMONE, OVSJANIKOV, MAK-SIM, et al. “Orthogonalized fourier polynomials for signal approximation and transfer”. *Computer Graphics Forum*. Vol. 40. 2. Wiley Online Library. 2021, 435–447 2.
- [MMR*19] MELZI, SIMONE, MARIN, RICCARDO, RODOLÀ, EMANUELE, et al. “Shrec 2019: Matching humans with different connectivity”. *Eurographics Workshop on 3D Object Retrieval*. Vol. 7. The Eurographics Association. 2019, 3 2.
- [MO23] MAGNET, ROBIN and OVSJANIKOV, MAK-S. “Scalable and efficient functional map computations on dense meshes”. *Computer Graphics Forum*. Vol. 42. 2. Wiley Online Library. 2023, 89–101 6.
- [MRMO20] MARIN, RICCARDO, RAKOTOSAONA, MARIE-JULIE, MELZI, SIMONE, and OVSJANIKOV, MAK-S. “Correspondence learning via linearly-invariant embedding”. *Advances in Neural Information Processing Systems* 33 (2020), 1608–1620 2.
- [MRR*19] MELZI, SIMONE, REN, JING, RODOLÀ, EMANUELE, et al. “ZoomOut: spectral upsampling for efficient shape correspondence”. *ACM Transactions on Graphics (TOG)* 38.6 (2019), 1–14 2, 6.
- [NO17] NOGNGENG, DORIAN and OVSJANIKOV, MAK-S. “Informative descriptor preservation via commutativity for shape matching”. *Computer Graphics Forum*. Vol. 36. 2. Wiley Online Library. 2017, 259–267 6.
- [OBS*12] OVSJANIKOV, MAK-S, BEN-CHEN, MIRELA, SOLOMON, JUSTIN, et al. “Functional maps: a flexible representation of maps between shapes”. *ACM Transactions on Graphics (ToG)* 31.4 (2012), 1–11 2.
- [RWP06] REUTER, MARTIN, WOLTER, FRANZ-ERICH, and PEINECKE, NIKLAS. “Laplace–Beltrami spectra as ‘Shape-DNA’ of surfaces and solids”. *Computer-Aided Design* 38.4 (2006), 342–366 4.
- [SOG09] SUN, JIAN, OVSJANIKOV, MAK-S, and GUIBAS, LEONIDAS. “A concise and provably informative multi-scale signature based on heat diffusion”. *Computer graphics forum*. Vol. 28. 5. Wiley Online Library. 2009, 1383–1392 2.
- [VM24] VIGANÒ, GIULIO and MELZI, SIMONE. “Bijective upsampling and learned embedding for point clouds correspondences”. *Computers & Graphics* 122 (2024), 103985 2.
- [Wey11] WEYL, HERMANN. “Über die asymptotische Verteilung der Eigenwerte”. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* 1911 (1911), 110–117 4.