



**HAL**  
open science

## SHARP 2020: The 1st Shape Recovery from Partial Textured 3D Scans Challenge Results

Alexandre Saint, Anis Kacem, Kseniya Cherenkova, Konstantinos Papadopoulos, Julian Chibane, Gerard Pons-Moll, Gleb Gusev, David Fofi, Djamila Aouada, Björn Ottersten

► **To cite this version:**

Alexandre Saint, Anis Kacem, Kseniya Cherenkova, Konstantinos Papadopoulos, Julian Chibane, et al.. SHARP 2020: The 1st Shape Recovery from Partial Textured 3D Scans Challenge Results. European Conference on Computer Vision, Aug 2020, Online conference, France. hal-03130337

**HAL Id: hal-03130337**

**<https://hal.science/hal-03130337>**

Submitted on 3 Feb 2021

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# SHARP 2020: The 1st Shape Recovery from Partial Textured 3D Scans Challenge Results

Alexandre Saint<sup>1</sup>, Anis Kacem<sup>1</sup>, Kseniya Cherenkova<sup>1,2</sup>, Konstantinos Papadopoulos<sup>1</sup>, Julian Chibane<sup>3</sup>, Gerard Pons-Moll<sup>3</sup>, Gleb Gusev<sup>2</sup>, David Fofi<sup>4</sup>, Djamila Aouada<sup>1</sup>, and Björn Ottersten<sup>1</sup>

<sup>1</sup> SnT, University of Luxembourg

firstname.lastname@uni.lu

<sup>2</sup> Artec 3D, Luxembourg

{kcherenkova,gleb}@artec-group.com

<sup>3</sup> Max Planck Institute for Informatics

{jchibane,gpons}@mpi-inf.mpg.de

<sup>4</sup> University of Burgundy

david.fofi@u-bourgogne.fr

**Abstract.** The *SHape Recovery from Partial textured 3D scans* challenge, SHARP 2020, is the first edition of a challenge fostering and benchmarking methods for recovering complete textured 3D scans from raw incomplete data. SHARP 2020 is organised as a workshop in conjunction with ECCV 2020. There are two complementary challenges, the first one on 3D human scans, and the second one on generic objects. Challenge 1 is further split into two tracks, focusing, first, on large body and clothing regions, and, second, on fine body details. A novel evaluation metric is proposed to quantify jointly the shape reconstruction, the texture reconstruction and the amount of completed data. Additionally, two unique datasets of 3D scans are proposed, to provide raw ground-truth data for the benchmarks. The datasets are released to the scientific community. Moreover, an accompanying custom library of software routines is also released to the scientific community. It allows for processing 3D scans, generating partial data and performing the evaluation. Results of the competition, analysed in comparison to baselines, show the validity of the proposed evaluation metrics, and highlight the challenging aspects of the task and of the datasets. Details on the SHARP 2020 challenge can be found at <https://cvi2.uni.lu/sharp2020/>.

## 1 Introduction

Representing the physical world in 3D, including shape and colour, is key for industrial and research purposes [1,17,7,10,3,16,4,11]. It includes, for example, areas from virtual reality to heritage conservation, or from medical treatment to fitness, entertainment and fashion. 3D scanning allows to digitise the physical world, *e.g.* objects and humans. Acquired 3D scans vary in quality depending on the scanning system used, the properties of the target and the environment. For example, a high-end photogrammetric scanning system with a fixed camera

array might capture high-quality data at a high frame rate but might be bulky, have a fixed structure, suffer from occlusion and limited in scanning volume. Other high-end systems such as hand-held devices might produce accurate results, while easily transported and easily oriented to limit occlusion, but cannot handle movable targets and are time-consuming. On the other hand, low-end scanning systems might be flexible and easy to manipulate but produce low-quality scans. Limiting factors due to the target or the environment include varying levels of details (*e.g.* finer anatomical parts), occlusion, non-rigidity, movement, and optical properties (*e.g.* fabric, material, hair and reflection). Moreover, for time-consuming acquisition systems or moving targets, it might be desirable or only possible to capture partial data. In this work, defective and/or partial acquisitions, are both viewed as data with missing information that must be completed.

The SHARP 2020 challenge for *SHApe Recovery from Partial textured 3D scans* is proposed to foster research and provide a benchmark on 3D shape and texture completion from partial 3D scan data. First, two new unique datasets of 3D textured scans are proposed to serve as reference data. These datasets contain thousands of scans of humans and generic objects with varied identities, clothing, colours, shapes and categories. One challenge is proposed per dataset, challenge 1 focusing on human scans and challenge 2 on object scans. Second, partial scans are generated synthetically, but randomly, to simulate a general pattern of partial data acquisition while still having access to ground truth data. Third, specific evaluation metrics are proposed to quantitatively measure the quality of the shape and texture reconstructions, and the amount of completed data. Fourth, reusable software libraries developed for the challenge are also made available. These contain routines to process 3D scans, to generate partial data, and to evaluate and analyse the submissions on the proposed benchmark. This paper summaries the SHARP challenge with a presentation of the proposed datasets, benchmark and evaluation method, as well as the results of the submitted methods and an analysis of the results. SHARP 2020 is the first edition of the challenge, held in conjunction with the 16th European Conference on Computer Vision (ECCV).

In the following, Section 2 describes the challenge and the proposed datasets. Section 3 describes the proposed evaluation protocol and, in Section 4, an extensive analysis on the results is presented. To conclude, the results and outcomes of the challenge are discussed in Section 5.

## 2 Challenges and Datasets

The SHARP challenge is split into two separate challenges: challenge 1 focuses on human scans, and challenge 2 focuses on generic object scans. Two corresponding datasets are introduced, 3DBodyTex.v2 for human scans, and 3DObjectTex for generic object scans. Table 1 describes the datasets with figures on the number of samples for different subsets, including the splits used in the challenges and categories of body pose and clothing type for 3DBodyTex.v2.

**3DBodyTex.v2** is an extension of 3DBodyTex proposed by Saint et al. [12]. See sample scans in Fig. 1. It contains about 3000 static 3D human scans with high-resolution texture. It features a large variety of poses and clothing types, with about 500 different subjects. Each subject is captured in about 3 poses. Most subjects perform the corresponding poses in both standard close-fitting clothing and arbitrary casual clothing. The faces are anonymised, for privacy reasons, by blurring the shape and the texture.

**3DObjectTex** is a subset of the *viewshape* [2] repository. See sample scans in Fig. 2. It consists of about 1200 textured 3D scans of generic objects with a large variation in categories and in physical dimensions.

Both datasets encode the scans as 3D triangle meshes. The colour is encoded in a texture atlas with an independent and arbitrary UV mapping [18] for each mesh.

dataset	subset		samples				challenge	track
			train	val	test	total		
3DBodyTex.v2	full		2094	454	451	2999	1	1
	clothing	fitness	844	178	182	1204	1	2
		casual	1250	276	269	1795	-	-
	poses	standard	977	219	224	1420	-	-
		other	1117	235	227	1579	-	-
3DObjectTex	-		799	205	205	1209	2	-

Table 1: Contents of the datasets and categorical subsets, along with associated challenge and track. The *standard poses* are the *A* and *U* rest poses. The *other poses* are varied, from a predefined list or arbitrary. The *casual* and *fitness* clothing types are shown in Fig. 1.

## 2.1 Challenge 1: Recovery of Human Body Scans

Challenge 1 covers the reconstruction of human scans, both with loose casual clothing and minimal close-fitting clothing.

**Track 1: Recovery of Large Regions.** Track 1 focuses on the reconstruction of large regions of human scans, excluding hands and head. Fig. 1 shows samples of the data, with ground-truth and partial scans. Both shape and texture are considered. These large regions are of relatively high quality in the raw reference data. The fine details (hands and head) are of unreliable quality and thus ignored in this track. The full 3DBodyTex.v2 dataset is used (see Table 1).



Fig. 1: Samples of 3DBodyTex.v2 in casual (top rows) and fitness (bottom row) clothing. For each person, ground truth scan (left), sample synthetic partial scan (right).

**Track 2: Recovery of Fine Details.** Track 2 focuses on fine body details not considered in Track 1, *i.e.*, hands, fingers, and nose. The raw scans are not of reliable quality for these details and the faces are not released due to privacy concerns. Thus, the reference data is generated synthetically from a body model [9] and the scans of 3DBodyTex.v2 in *fitness* clothing (see Table 1), where the fine body details are not occluded, to capture to the real distribution of rough poses and shapes. The reference data is generated in two steps: (1) fit a parametric body model to the scans to obtain the ground-truth data; (2) simulate the scanning process in software to obtain a synthetic scan (with simulated artefacts in the regions of interest, *i.e.* hands, ears...) from which the partial data is generated. The fitting is performed with the approach of Saint et al. [15,12,13]. The texture is not considered in this setting as the raw data does not contain texture of reliable quality (and the parametric body model represents only the shape).

## 2.2 Challenge 2: Recovery of Generic Object Scans

Challenge 2 focuses on the shape and texture completion of generic objects. Fig. 2 shows some example objects.

## 2.3 Partial data

In all challenges, the partial data is generated synthetically by removing surface regions randomly, as shown in Fig. 1 and Fig. 2. This is done with a hole cutting operation following the steps below:

- (1) take as input a central vertex  $v_c$  and the number of vertices to remove  $k$ ;
- (2) select the  $k$  vertices closest to  $v_c$  in Euclidean distance;
- (3) remove the selected vertices and adjacent triangles from the mesh.

The process is repeated 40 times, with  $k$  set to 2% of the points of the mesh. Most



submission to the challenge contains reconstructions for all the samples of the test set. The submission is ranked by averaging the individual scores into a final global score.

**Surface area score.** The surface area score,

$$S_\alpha = 1 - |\bar{A}_X - \bar{A}_Y| \in [0, 1], \quad (2)$$

penalises a reconstruction proportionally to the deviation of its surface area with respect to the ground truth.  $\bar{A}_X = \frac{A_X}{A_X + A_Y}$  is the normalised surface area of  $X$  with respect to  $Y$ . With a surface area lower or greater than the ground truth,  $S_a$  decreases proportionally the overall score (1). As the surface area approaches the ground truth,  $S_a$  approaches 1, less affecting the overall score. The surface area of a mesh is computed in practice by adding the surface areas of all the individual triangles.

**Surface-to-surface distance.** Both shape and texture scores,  $S_s$  and  $S_t$ , are determined by estimating the directed distances,  $d^{XY}$  and  $d^{YX}$ , between the ground truth mesh  $Y$  and the estimated reconstruction  $X$ .

The directed distance between mesh  $A$  and mesh  $B$ ,  $d^{AB}$ , is computed by sampling  $N$  points uniformly on the surface of  $A$ , finding the corresponding closest points on the surface of  $B$ , and averaging the associated distances.

For a point  $p^A \in \mathbb{R}^3$  on  $A$ , the corresponding closest point on  $B$  is determined by computing the smallest point-to-triangle distance between  $p^A$  and all triangles of  $B$ . The point-to-triangle distance,  $d = d_0 + d_1$ , is made of two components. For a specific triangle in  $B$ ,  $d_0$  is the Euclidean distance from  $p^A$  to the closest point  $p_0$  on the plane of the triangle.  $d_1$  is then the Euclidean distance from  $p_0$  to the nearest point  $p_1$  of the triangle (in the plane of the triangle). If the intersection  $p_0$  is inside the triangle, it is denoted as a *hit*, otherwise it is a *miss*. In case of a hit,  $p_0 = p_1$ , thus  $d_1 = 0$  and  $d = d_0$ .

**Hit rate.** When computing the surface-to-surface distance from mesh  $A$  to mesh  $B$ , the hit rate,

$$h^{AB} = \frac{H^{AB}}{N} \in [0, 1], \quad (3)$$

is the proportion of the  $N$  points sampled on  $A$  *hitting*  $B$  (see previous paragraph).

**Shape score.** The shape score,

$$S_s = \frac{S_s^{XY} + S_s^{YX}}{2} = \frac{h^{XY} \phi_s(d_s^{XY}) + h^{YX} \phi_s(d_s^{YX})}{2} \quad (4)$$

is a measure of the similarity of the shape of two meshes. The measure is symmetric by averaging the directed measures. The hit rates,  $h^{XY}$  and  $h^{YX}$ , penalise

overcomplete and incomplete reconstructions, respectively. In the directed shape score,

$$S_s^{XY} = h^{XY} \phi_s(d_s^{XY}), \quad (5)$$

the mapping function  $\phi_s : [0, \infty] \mapsto [0, 1]$  converts the computed distance to a score in  $[0, 1]$ . It is defined by a normal distribution function with zero mean,

$$\phi_s(d) = \frac{1}{\sigma_s \sqrt{2\pi}} e^{-\frac{1}{2} \left(\frac{d}{\sigma_s}\right)^2}, \quad (6)$$

where the standard deviation  $\sigma_s$  is estimated from baselines, including the ground-truth data, the input partial data, additional perturbations thereof with local and/or global white Gaussian noise on the shape, and a baseline of shape reconstruction based on a hole-filling algorithm.

**Texture score.** The texture score,

$$S_t = \frac{S_t^{XY} + S_t^{YX}}{2} = \frac{h^{XY} \phi_t(d_t^{XY}) + h^{YX} \phi_t(d_t^{YX})}{2} \quad (7)$$

is similar in principle to the shape score,  $S_s$ , except that the distance  $d_t$  is computed in texture space for all point correspondences obtained by the surface-to-surface distance. Additionally, the parameter  $\sigma_t$  for the mapping function  $\phi_t$  specific to the texture is estimated from the ground-truth data, the input partial data, and additional perturbations with local and/or global white Gaussian noise, in both shape and texture. Below, the texture score is also interchangeably denoted colour score.

## 4 Results and Analysis

This section presents and analyses the submissions to both challenges. The challenge has attracted 32 participants with 9 validated registrations for Challenge 1 and 6 for Challenge 2. Table 2 gives the number of valid submissions, received and accepted, and the number of submitted solutions.

challenge	track	registrations	submissions/participants	
			received	accepted
1	1	9	4/3	3/2
1	2	9	2/2	2/2
2	-	6	4/1	4/1

Table 2: Figures on validated registrations and entries for the challenges of SHARP 2020.



The accepted entries are *Implicit Feature Networks for Texture Completion of 3D Data* [6,5], from RVH (Real Virtual Humans group at Max Planck Institute for Informatics), submitted in several variants to both Challenge 1 and Challenge 2, and *3DBooSTeR: 3D Body Shape and Texture Recovery* [14], from SnT (Interdisciplinary Centre for Security, Reliability and Trust at the University of Luxembourg), submitted to Challenge 1. In the following, the entries are interchangeably abbreviated RVH-IF and SnT-3DB, respectively. Table 3 shows the quantitative results of the submissions for both challenges. The scores are reported in percent in an equivalent way to the scores mapping to  $[0, 1]$  in Section 3. The methods are compared to a baseline consisting of the unmodified partial data. The rest of this section presents and analyses the results.

challenge	track	method	score (%)		
			shape	texture	overall
1	1	baseline	$38.95 \pm 4.67$	$40.29 \pm 4.24$	$39.62 \pm 4.41$
		SnT-3DB	$54.21 \pm 14.28$	$70.55 \pm 7.26$	$62.38 \pm 9.61$
		RVH-IF-1	$85.24 \pm 5.72$	$87.69 \pm 5.96$	$86.47 \pm 5.38$
		<b>RVH-IF-2</b>	<b><math>85.24 \pm 5.72</math></b>	<b><math>88.26 \pm 5.46</math></b>	<b><math>86.75 \pm 5.19</math></b>
1	2	baseline	$41.1 \pm 3.31$	-	-
		SnT-3DB	$60.7 \pm 10.98$	-	-
		<b>RVH-IF</b>	<b><math>83.0 \pm 4.87</math></b>	-	-
2	-	baseline	$42.09 \pm 4.99$	$41.64 \pm 5.24$	$41.87 \pm 4.74$
		RVH-IF-1	$72.68 \pm 23.47$	$76.44 \pm 14.33$	$74.56 \pm 16.89$
		<b>RVH-IF-2</b>	<b><math>73.91 \pm 22.85</math></b>	<b><math>76.93 \pm 14.31</math></b>	<b><math>75.42 \pm 16.68</math></b>
		RVH-IF-3	$72.68 \pm 23.47$	$53.13 \pm 18.21$	$62.91 \pm 16.09$
		RVH-IF-4	$73.91 \pm 22.85$	$53.73 \pm 18.33$	$63.82 \pm 16.01$

Table 3: Reconstruction scores (%) of the baseline and the proposed methods for both challenges. The baseline is the unmodified partial data.

#### 4.1 Challenge 1 - Track 1

In challenge 1, track 1, the RVH-IF [6,5] approach ranks first with an overall reconstruction score of 86.75% (see Table 3). SnT-3DB comes second with 62%. RVH-IF surpasses the baseline unmodified partial data with an overall score of 46% higher and performs significantly better than SnT-3DB [14] with a score increment of 24%. RVH-IF-2 has similar shape and texture scores with differences of 2-3%, while SnT-3DB has a much higher texture score, 16% above the shape score. The RVH-IF-2 variant slightly improves the texture score over RVH-IF-1, with 88.26% instead of 87.69%, but the shape scores are identical.

Fig. 3 shows the frequency distribution of the scores. RVH-IF appears relatively

tightly located around the mean score 85% with a slight tail expanding in the lower scores down to 60%. SnT-3DB appears more spread for the shape score and with a tail reaching very low scores down to 10% below the baseline.

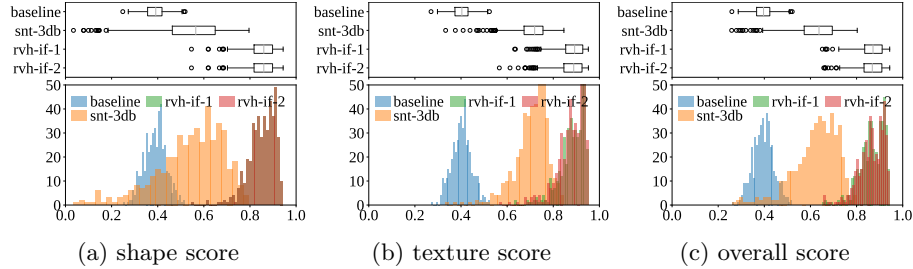


Fig. 3: Challenge 1 - track 1: Boxplots and frequency distributions of the reconstruction scores of the samples of the test set, for the baseline unmodified partial data, and all submissions, SnT-3DB and RVH-IF- $\{1,2\}$ .

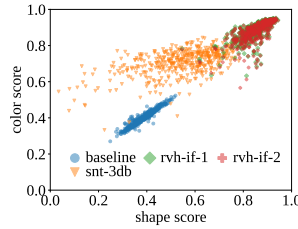


Fig. 4: Challenge 1 - track 1: Correlation between shape and texture scores, for the baseline unmodified partial data, and all submissions, SnT-3DB and RVH-IF- $\{1,2\}$ .

Fig. 4 highlights the correlation of the shape and texture scores. The baseline unmodified partial data is highly correlated, which is expected because the partial data available is exact. RVH-IF methods display a high correlation, with overall a tendency towards better texture scores than shape scores, except for a small proportion of grouped outlier samples with higher shape and lower texture scores. The SnT-3DB method displays the same tendency towards better texture than shape but with more dispersed results. Not a significant group of outliers is observed. The weaker correlation in SnT-3DB could be due to at least two factors: the sequential nature of the method, making the shape and texture reconstructions somewhat independent of each other; the mostly uniform colour of most clothing and of the body shape, resulting in good texture recovery even if the spatial regions are not matched correctly.

Fig. 5 and 6 show the distribution of scores and the correlation between shape and texture scores for subsets of the 3DBodyTex.v2 dataset. The casual clothing subset (Fig. 5a) shows lower and more dispersed scores than the fitness clothing (Fig. 5b). This reflects the diversity in texture and shape in the casual clothing, which is challenging for all proposed methods. Moreover, to a small but noticeable extent, the reconstruction scores are higher and more correlated in shape and texture on standard poses (Fig. 6d) than on non-standard ones (Fig. 6c).




Fig. 5: Challenge 1 - track 1: Boxplots and frequency distributions of the shape (top), texture (middle) and overall (bottom) scores on the test set, for subsets of 3DBodyTex.v2, and for the baseline unmodified partial data and all submissions, SnT-3DB and RVH-IF- $\{1,2\}$ . Subsets (columns): casual and fitness clothing, non-standard and standard (A and U) poses.

As a qualitative assessment, the best and worst reconstructions are shown in Fig. 7 for both approaches SnT-3DB and the RVH-IF. The best reconstructions are for scans in fitness clothing and standard pose, as also observed quantitatively above. The worst reconstructions are for scans in casual clothing and non-standard sitting/crouching poses, containing more shape and texture variation. It could also be that the ground truth for this data is less reliable, biasing the results towards lower scores.

## 4.2 Challenge 1 - Track 2

Challenge 1 - Track 2 targets the reconstruction of fine details on the hands, ears and nose, with only a shape ground truth. The only score evaluated is thus




Fig. 6: Challenge 1 - track 1: Correlation of shape and texture scores on the test set, for subsets of 3DBodyTex.v2, and for the baseline unmodified partial data and all submissions, SnT-3DB and RVH-IF- $\{1,2\}$ . Subsets (columns): casual and fitness clothing, non-standard and standard (A and U) poses.




Fig. 7: Challenge 1 - track 1: Best and worst reconstructions for submissions SnT-3DB and RVH-IF-1. From top to bottom: partial scan, reconstruction, ground truth.

the shape score. As in Track 1, and as visible in Fig. 8, RVH-IF is the highest-performing approach with a score of 83% compared to 60.7% for SnT-3DB and 41.1% for the baseline. The distribution of scores is also more concentrated around higher values for RVH-IF.

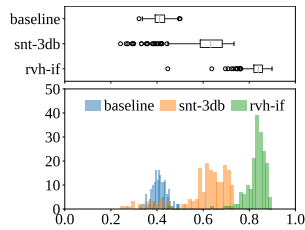


Fig. 8: Challenge 1 - track 2: Boxplots and frequency distributions of shape scores on the test set, for the baseline unmodified partial data and all submissions, SnT-3DB and RVH-IF.

### 4.3 Challenge 2

Challenge 2 contains four submissions by a single participant with the RVH-IF method. As reported in Table 3, RVH-IF-2 achieves the highest score around 75%, which is more than 30% over the baseline. The RVH-IF-1 is the second best-performing approach, a few percent below.

Fig. 9 shows the distribution of reconstruction scores. All approaches achieve high shape scores. For the texture, RVH-IF- $\{3,4\}$  perform less well and overlap the baseline, not improving much on the unmodified partial data. Overall, the scores are widely spread, even for the baseline. This reflects the diversity of object classes in 3DObjectTex, making both the reconstruction and the evaluation tasks challenging. This is confirmed in the shape-texture correlation in Fig. 10 where the correlation is less than in Challenge 1.

As Fig. 11 shows, the best reconstructions are of meshes with a closed surface, and with both a smooth shape and a relatively uniform texture. The worst reconstructions are of meshes with complex shapes, *i.e.* non-convex and with holes, creases, sharp edges, steps, etc.

## 5 Conclusion

This paper presented the SHARP 2020 challenge for *SHApe Recovery from Partial textured 3D scans* held in its first edition in conjunction with the 16th European Conference on Computer Vision. The SHARP challenge proposes to foster research and provide a benchmark on 3D shape and texture completion from partial 3D scan data. With two aspects, the recovery of human scans and of object scans, two unique datasets of reference high-quality textured 3D scans

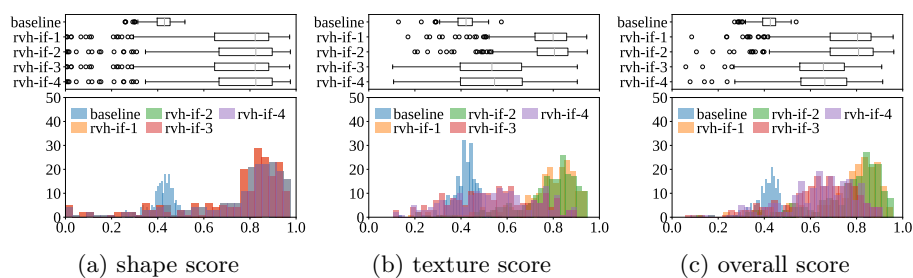


Fig. 9: Challenge 2: Boxplots and frequency distributions of scores on the test set for the baseline unmodified partial data, and all submissions RVH-IF- $\{1,4\}$ .




Fig. 10: Challenge 2: Correlation between shape and texture on the test set, for the baseline unmodified partial data, and all submissions RVH-IF- $\{1,4\}$ .




Fig. 11: Challenge 2: Best and worst reconstruction results for all submissions (RVH-IF- $\{1,4\}$ ). From top to bottom: partial scan, reconstruction, ground truth.

were proposed and released to the scientific community. Moreover, new specific evaluation metrics were proposed to measure simultaneously the quality of shape and texture reconstruction, and the amount of completion. The results of the participants show the validity of the proposed metrics, the challenging nature of the datasets and highlight the difficulties of the task. The RVH-IF submission obtained the highest scores in both challenges with the *Implicit Feature Networks for Texture Completion of 3D Data*. The variation in clothing seems a major difficulty in Challenge 1, and to a lesser extent, the pose. The texture seems more easily reconstructed than the shape, probably due to it being mostly uniform on clothing. The reconstruction of fine details was more demanding than the reconstruction of the full body. Challenge 2 shows that the proposed dataset of generic objects contains a lot of variation that makes both the shape and the texture challenging to recover. As in Challenge 1, smooth objects were better reconstructed.

The SHARP 2020 challenge has thus promoted the development of new methods for 3D shape and texture completion. The addition of the texture aspect on top of the shape aspect makes the contributions stand out from the current scientific literature. As the proposed metrics show, there is still room for improvement. This will be the object of the next editions of SHARP challenges.

## Acknowledgements

We thank Artec3D for sponsoring this challenge with cash prizes and releasing the data for the 3DObjectTex dataset. This work and the data collection of 3D human scans for 3DBodyTex.v2 were partly supported by the Luxembourg National Research Fund (FNR) (11806282 and 11643091). We also gratefully acknowledge the participation, at different times, of all members of the Computer Vision, Imaging and Machine Intelligence (CVI<sup>2</sup>) Research Group at the SnT, University of Luxembourg, including the moderation of the workshop event by Renato Baptista and the support of Pavel Chernakov in the development of the evaluation software. Finally, we express our appreciation to all the reviewers of the workshop submissions.

## References

1. Anguelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., Davis, J.: Scape: shape completion and animation of people. In: ACM SIGGRAPH 2005 Papers, pp. 408–416 (2005)
2. Artec3D: viewshape: online repository of 3d scans. <https://viewshape.com/> (2020), online; accessed 10/09/2020
3. Bogo, F., Romero, J., Loper, M., Black, M.J.: Faust: Dataset and evaluation for 3d mesh registration. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3794–3801 (2014)
4. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Numerical geometry of non-rigid shapes. Springer Science & Business Media (2008)
5. Chibane, J., Alldieck, T., Pons-Moll, G.: Implicit feature networks for texture completion of 3d data. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops (2020)
6. Chibane, J., Alldieck, T., Pons-Moll, G.: Implicit functions in feature space for 3d shape reconstruction and completion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6970–6981 (2020)
7. Hasler, N., Stoll, C., Sunkel, M., Rosenhahn, B., Seidel, H.P.: A statistical model of human pose and body shape. In: Computer graphics forum. vol. 28, pp. 337–346. Wiley Online Library (2009)
8. Jensen, R., Dahl, A., Vogiatzis, G., Tola, E., Aanaes, H.: Large scale multi-view stereopsis evaluation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 406–413 (2014)
9. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3d hands, face, and body from a single image. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 10975–10985 (2019)
10. Pishchulin, L., Wuhler, S., Helten, T., Theobalt, C., Schiele, B.: Building statistical shape spaces for 3d human modeling. *Pattern Recognition* **67**, 276–286 (2017)
11. Robinette, K.M., Daanen, H., Paquet, E.: The caesar project: a 3-d surface anthropometry survey. In: Second International Conference on 3-D Digital Imaging and Modeling (Cat. No. PR00062). pp. 380–386. IEEE (1999)
12. Saint, A., Ahmed, E., Cherenkova, K., Gusev, G., Aouada, D., Ottersten, B., et al.: 3dbodytex: Textured 3d body dataset. In: 2018 International Conference on 3D Vision (3DV). pp. 495–504. IEEE (2018)
13. Saint, A., Cherenkova, K., Gusev, G., Aouada, D., Ottersten, B., et al.: Bodyfitr: Robust automatic 3d human body fitting. In: 2019 IEEE International Conference on Image Processing (ICIP). pp. 484–488. IEEE (2019)
14. Saint, A., Kacem, A., Cherenkova, K., Aouada, D.: 3dbooster: 3d body shape and texture recovery. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops (2020)
15. Saint, A., Shabayek, A.E.R., Aouada, D., Ottersten, B., Cherenkova, K., Gusev, G.: Towards automatic human body model fitting to a 3d scan. In: Proceedings of 3DBODY. TECH 2017-8th International Conference and Exhibition on 3D Body Scanning and Processing Technologies, Montreal QC, Canada, 11-12 Oct. 2017. pp. 274–280. Hometrica Consulting (2017)
16. Xu, Z., Zhang, Q., Cheng, S.: Multilevel active registration for kinect human body scans: from low quality to high quality. *Multimedia Systems* **24**(3), 257–270 (2018)



17. Yang, Y., Yu, Y., Zhou, Y., Du, S., Davis, J., Yang, R.: Semantic parametric reshaping of human body models. In: 2014 2nd International Conference on 3D Vision. vol. 2, pp. 41–48. IEEE (2014)
18. Zayer, R., Lévy, B., Seidel, H.P.: Linear angle based parameterization (2007)