Subjective Visual Quality Assessment of Immersive 3D Media Compressed by Open-Source Static 3D Mesh Codecs

Kyriaki Christaki[1], Emmanouil Christakis[1], Petros Drakoulis[1], Alexandros Doumanoglou[1][0000−0002−4337−1720], Nikolaos Zioulis[1][0000−0002−7898−9344], Dimitrios Zarpalas[1], and Petros Daras[1]

Visual Computing Lab, Information Technologies Institute, Centre for Research and Technology Hellas
{kchristaki, manchr, petros.drakoulis, aldoum, nzioulis, zarpalas, daras}@iti.gr
vcl.iti.gr

Abstract. While studies for objective and subjective evaluation of the visual quality of compressed 3D meshes has been discussed in the literature, those studies were covering the evaluation of 3D-meshes either created by 3D artists or generated by a computationally expensive reconstruction process applied on high quality 3D scans. With the advent of RGB-D sensors operating at high frame-rates and the utilization of fast 3D reconstruction algorithms, humans can be captured and reconstructed into a 3D representation in real-time, enabling new (tele-)immersive experiences. The produced 3D mesh content is structurally different in the two cases. The first type of content is nearly perfect and clean while the second type is much more irregular and noisy. Evaluating compression artifacts on this new type of immersive 3D media, constitutes a yet unexplored scientific area. In this paper, we conduct a survey to subjectively assess the perceived fidelity of 3D meshes subjected to compression using three open-source static 3D mesh codecs compared to the original uncompressed models. The subjective evaluation of the content is conducted in a Virtual Reality setting, using the forced-choice pairwise comparison methodology with existing reference. The results of this study are two-fold; first, the design of an experimental setup that can be used for the subjective evaluation of 3D media, and second, a mapping of the compared conditions to a continuous ranking scale. The latter can be used when selecting codecs and optimizing their compression parameters to achieve optimum balance between bandwidth and perceived quality in tele-immersive platforms.

Keywords: Subjective visual quality study · 3D · Compression · Tele-immersion · Forced Pairwise Comparison · Virtual Reality (VR).

1 Introduction

Nowadays, advances in the fields of 3D capturing, imaging and processing technologies have allowed the advent of new forms of interactive immersive experi-

ences. Mixed reality and tele-immersive platforms [17, 6], are now investigated as applications that increase user engagement and immersion by embedding 3D reconstructed human representations in virtual environments. These 3D representations usually take the form of 3D meshes. A 3D mesh is a collection of vertices and faces (triangles) that defines the surface of an object in three dimensions. In the previous decades, 3D meshes were mostly created in specialized modeling software by 3D expert artists or generated by 3D reconstruction algorithms operating on real world depth data acquired by 3D scanners. In both scenarios, the 3D mesh outcomes are perfect and clean, as in the first case they are manually crafted and in the second, they are produced off-line by computationally expensive algorithms on high precise depth data acquired by the aforementioned 3D scanners. The real-time production of human 3D meshes in modern mixed reality and tele-immersive platforms, conceptually undergo a similar automatic process as mentioned previously, with the fundamental difference of utilizing computationally cheaper reconstruction algorithms operating on depth data of lower precision, acquired at high frame rates. Contrariwise to the above mentioned perfect and clean meshes, the latter are highly irregular and noisy.

To realize the tele-immersion experience, the human 3D meshes are required to be transmitted in real-time to remote parties, often not without compression. However, in that case, an overly aggressive compression scheme can negatively impact the content quality and the viewer's quality of experience. Oftentimes, the assessment of the geometrical similarity between 3D meshes via objective metrics is used to measure the fidelity of a compressed model to the original. However, objective metrics might not correlate well with viewers' perceived visual quality that may also be related to psychophysical factors. While subjective studies on the visual quality of compressed 3D meshes have been extensively discussed in the literature, the overwhelming majority of those works studied compression artifacts on clean and perfect 3D meshes. Subjectively evaluating compression artifacts on immersive 3D media (3D meshes produced in mixed reality and tele-immersive platforms) is a field not yet thoroughly explored.

In this paper, we try to investigate the impact of the chosen 3D mesh codec and the value of geometric distortion parameter on the subjective visual quality of compressed immersive 3D media. In particular, we examine the geometric distortion artifacts induced by three open-source static 3D mesh codecs on watertight meshes of reconstructed human models, produced in mixed reality platforms. In order to examine possible effects of the mesh production process on the subject's opinion, we perform two independent subjective experiments by using the same codecs and compression parameters but different datasets. The first dataset is composed of human 3D meshes generated by a real-time 3D reconstruction algorithm used in a mixed-reality platform, while the second one consists of samples taken from an open repository containing high quality 3D reconstructed meshes of precisely 3D scanned real world objects. Tested conditions include various codec and compression parameter combinations. Instead of using the conventional approach of a 2D monitor, we chose to use a virtual reality (VR) headset as our display medium. A VR headset allows for realistic

and natural viewing of the surveyed content as the (human) 3D models can be observed in real-life sizes. In addition, it is also very aligned with the envisaged applications of tele-immersive content that focus on an elevated feeling of presence and natural interactions.

The main contributions of this paper are:

– The setup of a consistent experiment for subjective evaluation of compressed immersive 3D media in a VR setting.
– To provide a concise analysis of the collected subjective data, leading to an overall ranking of the compared conditions, implicitly capturing the subjective preference on the compression algorithm and distortion levels.
– To provide a side-by-side comparison on how the same compression conditions affect subjective preference depending on the source type of the 3D meshes, being either immersive 3D media or meshes generated by 3D reconstructing high quality 3D scans of real objects.

The rest of the paper is organized as follows. Firstly, in Section 2, we present existing works related to objective and subjective quality assessment, where the effect of one or multiple distortion parameters on 3D mesh and point cloud geometry is examined. Section 3 presents this survey's experimental setup. In Section 4 we present and analyze the survey's results and discuss on the relation between an objective metric (Hausdorff distance) and the acquired subjective ratings. Finally, in Section 5 we conclude by highlighting the main findings of this study.

2   Related Work

There exist a number of objective metrics that are used to measure the geometric error between 3D meshes, as there is not a single more appropriate way to measure the difference between 3D geometries [10]. One of the most widely accepted ones, also typically used when comparing the performance of 3D compression methods, is the Hausdorff distance [9]. While objective metrics results might not necessarily agree with the perceived - by the viewers - quality, they are commonly used for evaluating degradation effects on 3D meshes, as subjective evaluation can be time expensive. Nonetheless there are also works that focus on evaluating the quality of 3D meshes and point clouds, affected by distortions, on a subjective basis.

Subjective quality assessment of mesh distortions, other than compression: In [27], the authors perform a subjective evaluation survey to determine how down-sampling or adding coordinate/color noise in 3D point cloud affects the perceived quality using a Mean-Opinion-Score (MOS) methodology. A subjective study of 3D point cloud denoising algorithms with the use of Double-Stimulus-Impairment-Scale (DSIS) methodology is presented in [15] and the correlation with objectives metrics is investigated. However, the focus of this study is to assess the performance of denoising algorithms rather than the quality of

the geometry of the compressed meshes. In a later work [8], a subjective evaluation with the use of DSIS methodology and Mean-Opinion-Scores of 3D point cloud models with different noise distortion levels and geometry resolution in an Augmented Reality (AR) environment, is presented. The study concludes that the geometric complexity of the model affects the evaluation score and that there is low correlation between objective and subjective metrics. While relevant, this study is focused on clean models as opposed to immersive 3D media. Another study [25] in VR assesses the impact of the number of mesh triangles on the perceived quality using a pairwise comparison with forced choice. An interesting conclusion is drawn, which is that by displaying the meshes either using a VR headset or a regular monoscopic display, there are no significant differences in the choices of the subjects. This result, however, cannot be safely extrapolated to the effects of compression.

Subjective quality assessment of mesh distortions related to compression: In [16], a subjective and objective evaluation of two point cloud compression schemes, Octree-based and Graph-based compression, is attempted. The authors correlate the subjective results obtained using the DSIS methodology, with two different types of objective metrics, namely, Point-to-Point and Point-to-Plane metrics. However, a typical LCD monitor was used as the display medium during the user study. In a similar context [7], performs a study in order to correlate objective and subjective metrics for the evaluation of point clouds subjected to realistic degradations. An interesting finding is that although there is a strong correlation between subjective and objective metrics in noise related distortions, they do not correlate well in the case of compression related distortions. While relevant to our work, their study is focused on clean point clouds and the evaluation is performed on a flat screen. In [20], an objective and subjective evaluation of different compression algorithms applied on reconstructed 3D human meshes in a virtual environment is presented. This work, while in the same application area with ours, is focusing on compression of immersive 3D-media without investigating compression artifacts on clean models. Furthermore, the real-time 3D reconstruction method used to produce the human 3D meshes in [20] is one of the first in the field, while in the current paper we use 3D meshes produced by a more recent and higher quality method [6]. Additionally, while the meshes used in [20] are textured, we choose to exclude texturing from our meshes in an effort to eliminate the unwanted effects of texturing on the subjects' choices, as texture masking can greatly reduce the perception of geometric distortions and degradations [11]. Finally, in the present work we consider more recent and open-sourced 3D mesh codecs than [20].

3   Experiment

In this section we present the implementation details of the conducted survey by describing the evaluation methodology, dataset generation, compared conditions and details of the experiment.
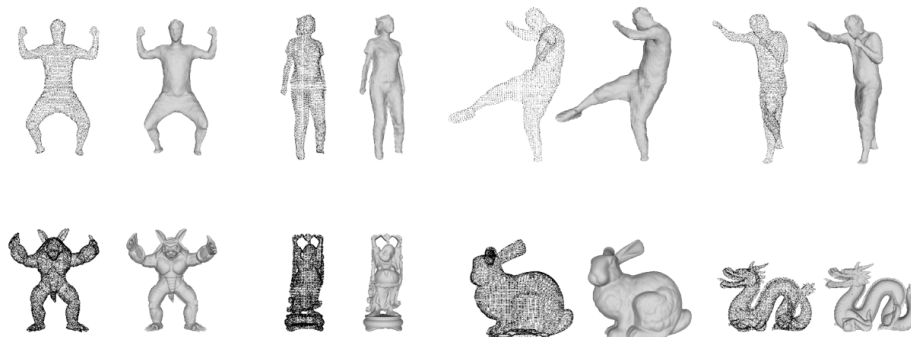
Fig. 1. The 3D models comprising the survey's content. The human 3D reconstructions are presented on top while the scanned objects are presented on the bottom. For each 3D model, a wireframe representation is depicted on the left, followed by a shaded one on the right.

Evaluation Methodology: In both of our experiments we chose a pairwise comparison methodology [21] to evaluate the fidelity of the various conditions to the reference undistorted 3D mesh. More specifically, during the experiment every subject was presented simultaneously with two distorted versions as well as the reference one. Then, each subject was asked to choose which of the two distorted versions bests resembles the reference. The forced choice pairwise comparison method was preferred over standard methodologies recommended by ITU [14],[13] because recent works have shown its superiority in terms of producing statistically robust results with smaller variance [19, 26]. Furthermore, this method is easier to implement and is less mentally demanding for the subjects. However, a significant drawback of this method is that the number of the required comparisons can quickly get prohibitively large if the number of examined conditions is also large. Various ways to mitigate this problem have been investigated [12]. As mentioned in [21], an effective way to deal with this issue is to only consider comparisons between conditions that do not have a big difference in terms of quality.

Datasets: As already discussed in section 1, in order to study the effect of compression distortion in different types of content, mostly related to the mesh production techniques, we use two distinct 3D content groups. The first group ("Model Group #1"), comprises a selection of four 3D reconstructed human meshes that were produced by the tele-immersive platform of [6] during four human performance captures. As in most typical cases of tele-immersive application settings, these models suffer from visual artifacts and the presence of noise is apparent. The second group ("Model Group #2'), comprises four models taken from the widely acknowledged and used Stanford 3D Scanning Repository [5], namely the "Bunny", "Happy Buddha", "Dragon" and "Armadillo" models.

These higher quality models, were created via a 3D reconstruction process on higher accuracy 3D scans of real world objects. All the reference models of the two datasets are illustrated in Fig. 1

Compared Conditions: All 3D mesh codecs compress connectivity losslessly and any visual distortions come from geometric loss that is controlled by a quantization parameter expressed in number of bits. Different codecs may use different quantization schemes. While there are four available open source 3D mesh codecs, more specifically, Draco [2, 24], Corto [1, 23, 22], O3DGC [3, 18] and OpenCTM [4], we select the first three, Draco, Corto and O3DGC for this study. From the selected three codecs, Draco and O3DGC use the same quantization strategy leading to visually identical models for the same compression parameters. The only aspect in which those two codecs differ, is in compression performance which is out of this paper's scope. Thus, while we explicitly use Draco in our experiments, the results equivalently apply to O3DGC. The reason why OpenCTM was excluded from our experiments is two-fold. First, and most importantly, including one more codec in the study would result into a much more expensive experiment in terms of time while we would have to take additional measures to counter potential subjects' fatigue. Second, based on an extensive benchmarking we have conducted in our lab, OpenCTM was found to be the least performant codec among the rest in rate-distortion terms.

A forced pairwise comparison methodology comes with a set of restrictions. These are related to the way the number of comparisons scales with the increase in the different levels that are compared. As a result, in order to reduce the burden and stress of the survey on the subjects to obtain higher quality results, three compression levels were selected. These are directly mapped to the quantization bits of each codec, and more specifically, 10, 11 and 12 bits were the target of this study. This choice was based on the empirical assessment that 9 bits produced unpleasing visual results, but even more importantly, it produced a bigger difference in visual quality when compared to the results of 10 bits than the other incremental combinations of the higher quantization bit levels. This would manifest in a clear and universal selection of the higher quality (10 bits), that would in turn create biased or even erroneous results, as reported in [21]. Further, as aforementioned, adding yet another compression bit into the comparisons, would significantly increase the required comparisons for every subject and make it a lot more demanding. Overall, we evaluate 2 codecs in 3 quantization levels, with comparisons being limited among codecs and neighboring quantization levels. This results in a total number of 11 comparisons for each data sample, compared to 15 when comparing between all possible condition combinations. Since we used 4 models in each dataset, the total number of comparisons required for each model group was 44.

Survey: The user study was realized by a VR application, developed in Unity3D[1], that implemented the pairwise methodology for comparing the visual quality of

---

[1] https://unity3d.com

compressed 3D content. We used an HTC Vive head mounted display (HMD) that each subject wore. Within the VR environment, each subject was able to view the undistorted (i.e. uncompressed) 3D model in the middle, and the two distorted (i.e. compressed) on its left and right sides. The models were displayed in life size, effectively simulating realistic tele-immersive scenarios. The content was viewed freely as a combination of natural navigation (i.e. physical movement in the real-world) and user interaction. Exploiting the HMD's tracking system, the users were allowed to freely move into the tracking area and thus, inspect the models in a natural manner. In addition, by using the headset's controllers, they were also able to rotate the models around their origin simultaneously to aid them in inspecting their visual fidelity and compare them to the reference. Overall, the study subjects could view and inspect the models in a free viewpoint fashion within the VR environment.

Inside the VR environment, the subjects would vote for the least distorted mesh by choosing one of the side (left/right) models. Left/right positioning was randomized to avoid any bias in the selection process. The 3D content was rendered un-textured, with flat shading, so as to accentuate differences in lighting, therefore allowing surface normal information to influence the perceived quality.

Moreover, we conducted the study using two different control groups, each one paired with a model group. As a result, half the subjects (Control Group #1) only viewed and assessed 3D content of captured human performances (Model Group #1), while the other half (Control Group #2) focused on the widely used high quality scanned objects models (Model Group #2).

There was no time limit imposed on the subjects, instead they could take their time in inspecting the content in their own pace. The average study session duration was 25 minutes, thereby minimizing fatigue which the subjects are more prone to, due to the use of VR headsets. In total, 40 subjects participated in the study, 8 females and 32 males evenly divided in the 2 Control Groups.

Fig. 2. Screenshot from the VR application showing the reconstructed human meshes. The reference model is in the middle while the two distorted are to its left and right.

Subjective Score vs Compression Parameters

Fig. 3. Survey Results: Subjective ranking score for every compression condition and for both Model Groups. Values are smoothly interpolated in each heat-map. Interpolation used only for visualization purposes. The underlying domain is discrete.

4   Results

After collecting the responses from all the subjects in our pairwise comparison survey, we processed these data using the toolbox provided by Perez-Ortiz and Mantiuk [21]. This results in an assignment of a rating in a continuous scale to each of the compared conditions. By default, a zero rating is given to the worst, as perceived by the subjects, condition. No voting outliers were detected in our subjects for either of the control groups, as reported by the use of the toolbox [21]. The final ratings for both groups are illustrated in Fig 3 as heat-maps. Moreover, Fig. 4 presents the rankings and their corresponding 95% confidence intervals for both control groups (*left and middle* for Control Group #1 & #2 respectively). The scores can be interpreted as such: a rating difference of a single unit between two conditions means that in a hypothetical comparison between these conditions, the probability of a subject to choose the one with the higher rating is 75%. A detailed mapping between probabilities and quality scores are depicted in Fig 4 (Figure re-printed with permission from [21]).

During the experiment, the majority of the subjects commented on the difficulty in distinguishing the differences between the compared conditions and the reference mesh. Despite this fact, it is obvious from the results that they had a clear preference towards meshes compressed with a higher number of quantization bits across codecs and model contents. This difficulty in making a confident vote, interestingly aligns with a similar outcome presented in [25], where subjects were asked to subjectively evaluate the visual quality of simplified meshes (meshes with less number of triangles compared to the original). In [25], despite the subjects reporting a guessing behavior in their votes, their choices actually aligned with the true quality of the meshes (i.e. meshes with higher number of triangles were consistently voted preferable).

In our case, the transition from 10 to 11 bits had a higher impact on subjective visual quality than the one from 11 to 12, especially for the immersive

Fig. 4. a,b) Survey Results: Quality scores of the compared conditions and their confidence intervals. c) Mapping of the difference in the rating between two conditions into the estimated probability of a random subject to select the condition with the higher rating. Reprinted from "A practical guide and software for analysing pairwise comparison experiments." by Perez-Ortiz, M., Mantiuk, R.K, ArXiv e-prints (Dec 2017). Reprinted with permission.

3D media content (Model Group #1). Further, as illustrated in Fig. 4, there is a greater confidence in assessing the drop in perceived quality when switching to 10 bits. Instead, there is a high overlap between 11 and 12 bits, showcasing the difficulty in selecting between them. Nonetheless, there was a perceived difference in the upper quantization bits, as indicated by the rankings, albeit with lower confidence. One one hand, this is reasonable as increasing the quantization bits should progressively lead to less perceived differences between the encoded contents. This is expected as the distortion would be getting smaller given that the gains are diminishing. In other words, the high overlap of confident intervals is an indicator that we approach the perception threshold. On the other hand, understanding where the first jump in perceived quality happens is a very important indicator when choosing the quantization level.

For the "noisy" models (Model Group #1), Draco was preferred for the lowest bits (10) but Corto was scored as more visually pleasing for the 11 and 12 quantization bits. However, for the "clean" models (Model Group #2), the compressed representation produced by Draco was always preferable to the one produced by Corto for the same number of quantization bits. For the scanned objects (Model Group #2), the difference in the relative score between the lowest and the highest rated condition was smaller than in the case of the human meshes. Effectively, there was a larger distribution of the scores, and by extension a wider ranking, for Control Group #1. This points to an increased sensitivity to distortions for the Model Group #1 that can be attributed to the content itself being actual persons' 3D reconstructions and the fact that we are highly attuned to the human body and face forms. Consequently, such content requires higher presentation quality as it is more susceptible to perceived distortions.

Codec-Parameter	3D Reco. Human Meshes		Scanned Objects	
	Hausdorff distance	RSD	Hausdorff distance	RSD
Corto-10	0.0012025	10%	0.00111425	15%
Corto-11	0.00060775	9%	0.00062375	30%
Corto-12	0.00030325	11%	0.00027675	13%
Draco-10	0.000594	11%	0.00073425	56%
Draco-11	0.00030525	10%	0.000279	17%
Draco-12	0.00015	8%	0.0001365	13%

Table 1. Average distortion across models (measured in Hausdorff distance with respect to the bounding box) and its Relative Standard Deviation (RSD) for the two model groups for all compression parameters.

Correlation with objective metrics: Interestingly, the correlation between objective metrics and subjective score is not consistent in the two model sets. The average Hausdorff distances across all models for the same quantization parameters and their Relative Standard Deviations (RSDs) for both Model Groups are presented in Table 1. In the scanned objects group, we notice a strong correlation between Hausdorff distance and the subjective score. Draco codec has smaller Hausdorff distance compared with the Corto one, for the same quantization bits. Therefore, the preference of Draco over Corto among user study subjects, can be easily explained. The human meshes experiment does not seem to follow this pattern. There, users seems to prefer Corto over Draco in all quantization levels, apart from the lowest one. The fact that Corto compressed meshes receive higher subjective score, despite having a greater Hausdorff distance from the original mesh, implies that the human perceived quality does not correlate well with the objective metrics for that Model Group. The high presence of noise in the human meshes may influence the subjects' preference towards Corto compressed meshes, as in that case Corto's quantization scheme may produce subjectively more pleasant forms. This may also mean that codec choice matters for subjective visual quality, depending on the production process of the 3D meshes.

5   Conclusion

In this work, we performed a survey to evaluate the effects of compression on the subjectively perceived quality of 3D meshes generated by two different processes: a) human meshes produced by a real-time 3D reconstruction algorithm used in a tele-immersive platform and b) meshes produced by computationally expensive 3D reconstruction algorithms applied on high quality 3D scans of real objects. The evaluation of three open-source static 3D mesh codecs was conducted with two of them producing an exact visual output and only differing in compression performance. The reference meshes were compressed by all codecs in three different distortion levels, controlled by a quantization parameter. The experimental results showed that the quantization scheme applied by each individual codec

matters to the subjective visual quality of the compressed mesh and the pre-ferred codec generally depends on the generation process that produced the 3D mesh. While the distortion levels on the output meshes that were induced by the quantization parameters were distinguished from the subjects, for higher values of the quantization parameter the differences are less apparent. The study was conducted in Virtual Reality to better emulate tele-immersive experience and followed the forced choice pairwise methodology with full reference.

Based on the findings of this work, future studies for tele-immersive applications may discuss on making codec and distortion level choices based on rate-distortion performance of the codecs and the subjectively perceived visual quality of the 3D meshes.

References

1. Corto, https://github.com/cnr-isti-vclab/corto, accessed: 2018-06-07
2. Google Draco, https://github.com/google/draco, accessed: 2018-06-07
3. Open 3D Graphics Compression (O3DGC), https://github.com/amd/rest3d/tree/master/server/o3dgc, accessed: 2018-06-07
4. OpenCTM, http://openctm.sourceforge.net/, accessed: 2018-06-07
5. The Stanford 3D Scanning Repository, http://graphics.stanford.edu/data/3Dscanrep/, accessed: 2018-06-07
6. Alexiadis, D.S., Chatzitofis, A., Zioulis, N., Zoidi, O., Louizis, G., Zarpalas, D., Daras, P.: An integrated platform for live 3d human reconstruction and motion capturing. IEEE Transactions on Circuits and Systems for Video Technology 27(4), 798–813 (2017)
7. Alexiou, E., Ebrahimi, T.: On subjective and objective quality evaluation of point cloud geometry. In: 2017 Ninth International Conference on Quality of Multimedia Experience (QoMEX). pp. 1–3 (May 2017). https://doi.org/10.1109/QoMEX.2017.7965681
8. Alexiou, E., Upenik, E., Ebrahimi, T.: Towards subjective quality assessment of point cloud imaging in augmented reality. In: 2017 IEEE 19th International Workshop on Multimedia Signal Processing (MMSP). pp. 1–6 (Oct 2017). https://doi.org/10.1109/MMSP.2017.8122237
9. Aspert, N., Santa-Cruz, D., Ebrahimi, T.: Mesh: measuring errors between surfaces using the hausdorff distance. In: Proceedings. IEEE International Conference on Multimedia and Expo. vol. 1, pp. 705–708 vol.1 (2002). https://doi.org/10.1109/ICME.2002.1035879
10. Berjn, D., Morn, F., Manjunatha, S.: Objective and subjective evaluation of static 3d mesh compression. Signal Processing: Image Communication 28(2), 181 – 195 (2013). https://doi.org/https://doi.org/10.1016/j.image.2012.10.013, http://www.sciencedirect.com/science/article/pii/S0923596512002019, mPEG-V
11. Bulbul, A., Capin, T., Lavou, G., Preda, M.: Assessing visual quality of 3-d polygonal models. IEEE Signal Processing Magazine 28(6), 80–90 (Nov 2011). https://doi.org/10.1109/MSP.2011.942466

12. D. Amnon Silverstein, J.E.F.: Efficient method for paired comparison. Journal of Electronic Imaging 10, 10 – 10 – 5 (2001). https://doi.org/10.1117/1.1344187, https://doi.org/10.1117/1.1344187
13. International Telecommunication Union: Recommendation ITU-T P.910: Subjective video quality assessment methods for multimedia applications (2008)
14. International Telecommunication Union: Recommendation ITU-R BT.500: Methodology for the subjective assessment of the quality of television pictures (2012)
15. Javaheri, A., Brites, C., Pereira, F., Ascenso, J.: Subjective and objective quality evaluation of 3d point cloud denoising algorithms. In: 2017 IEEE International Conference on Multimedia Expo Workshops (ICMEW). pp. 1–6 (July 2017). https://doi.org/10.1109/ICMEW.2017.8026263
16. Javaheri, A., Brites, C., Pereira, F., Ascenso, J.: Subjective and objective quality evaluation of compressed point clouds. In: 2017 IEEE 19th International Workshop on Multimedia Signal Processing (MMSP). pp. 1–6 (Oct 2017). https://doi.org/10.1109/MMSP.2017.8122239
17. Karakottas, A., Papachristou, A., Doumanoglou, A., Zioulis, N., Zarpalas, D., Daras, P.: Augmented VR, IEEE Virtual Reality, Mar 18 - 22, 2018, https://www.youtube.com/watch?v=7O_TrhtmP5Q
18. Mamou, K., Zaharia, T., Prêteux, F.: Tfan: A low complexity 3d mesh compression algorithm. Comput. Animat. Virtual Worlds 20, 343–354 (Jun 2009). https://doi.org/10.1002/cav.v20:2/3, http://dx.doi.org/10.1002/cav.v20:2/3
19. Mantiuk, R., Tomaszewska, A., Mantiuk, R.: Comparison of four subjective methods for image quality assessment 31 (11 2012)
20. Mekuria, R., Cesar, P., Doumanis, I., Frisiello, A.: Objective and subjective quality assessment of geometry compression of reconstructed 3d humans in a 3d virtual room. In: Applications of Digital Image Processing XXXVIII. Proceedings of the SPIE, vol. 9599, p. 95991M (Sep 2015). https://doi.org/10.1117/12.2203312
21. Perez-Ortiz, M., Mantiuk, R.K.: A practical guide and software for analysing pairwise comparison experiments. ArXiv e-prints (Dec 2017)
22. Ponchio, F., Dellepiane, M.: Fast decompression for web-based view-dependent 3d rendering. In: Proceedings of the 20th International Conference on 3D Web Technology. pp. 199–207. Web3D '15, ACM, New York, NY, USA (2015). https://doi.org/10.1145/2775292.2775308, http://doi.acm.org/10.1145/2775292.2775308
23. Ponchio, F., Dellepiane, M.: Multiresolution and fast decompression for optimal web-based rendering. Graphical Models 88, 1 – 11 (2016). https://doi.org/10.1016/j.gmod.2016.09.002, http://www.sciencedirect.com/science/article/pii/S1524070316300285
24. Rossignac, J.: Edgebreaker: Connectivity compression for triangle meshes. IEEE Trans. Vis. Comput. Graph. 5, 47–61 (1999)
25. Thorn, J., Pizarro, R., Spanlang, B., Bermell-Garcia, P., González-Franco, M.: Assessing 3d scan quality through paired-comparisons psychophysics test. CoRR abs/1602.00238 (2016), http://arxiv.org/abs/1602.00238
26. Zerman, E., Hulusic, V., Valenzise, G., Mantiuk, R., Dufaux, F.: The relation between MOS and pairwise comparisons and the importance of cross-content comparisons. In: Human Vision and Electronic Imaging Conference, IS&T International Symposium on Electronic Imaging (EI 2018). Burlingame, United States (Jan 2018), https://hal.archives-ouvertes.fr/hal-01654133
27. Zhang, J., Huang, W., Zhu, X., Hwang, J.N.: A subjective quality evaluation for 3d point cloud models pp. 827–831 (01 2015)