

A deep perceptual metric for 3D point clouds

Maurice Quach[†], Aladine Chetouani^{†*}, Giuseppe Valenzise[†], Frederic Dufaux[†];

[†] Université Paris-Saclay, CNRS, CentraleSupélec, Laboratoire des signaux et systèmes; 91190 Gif-sur-Yvette, France

* Laboratoire PRISME, Université d'Orléans; Orléans, France

Abstract


Point clouds are essential for storage and transmission of 3D content. As they can entail significant volumes of data, point cloud compression is crucial for practical usage. Recently, point cloud geometry compression approaches based on deep neural networks have been explored. In this paper, we evaluate the ability to predict perceptual quality of typical voxel-based loss functions employed to train these networks. We find that the commonly used focal loss and weighted binary cross entropy are poorly correlated with human perception. We thus propose a perceptual loss function for 3D point clouds which outperforms existing loss functions on the ICIP2020 subjective dataset. In addition, we propose a novel truncated distance field voxel grid representation and find that it leads to sparser latent spaces and loss functions that are more correlated with perceived visual quality compared to a binary representation. The source code is available at https://github.com/mauriceqch/2021_pc_perceptual_loss.

Introduction

As 3D capture devices become more accurate and accessible, point clouds are a crucial data structure for storage and transmission of 3D data. Naturally, this comes with significant volumes of data. Thus, Point Cloud Compression (PCC) is an essential research topic to enable practical usage. The Moving Picture Experts Group (MPEG) is working on two PCC standards [1]: Geometry-based PCC (G-PCC) and Video-based PCC (V-PCC). G-PCC uses native 3D data structures to compress point clouds, while V-PCC employs a projection-based approach using video coding technology [2]. These two approaches are complementary as V-PCC specializes on dense point clouds, while G-PCC is a more general approach suited for sparse point clouds. Recently, JPEG Pleno [3] has issued a Call for Evidence on PCC [4].

Deep learning approaches have been employed to compress geometry [5, 6, 7, 8, 9, 10] and attributes [11, 12] of points clouds. Specific approaches have also been developed for sparse LIDAR point clouds [13, 14]. In this work, we focus on lossy point cloud geometry compression for dense point clouds. In existing approaches, different point cloud geometry representations are considered for compression: G-PCC adopts a point representation, V-PCC uses a projection or image-based representation and deep learning approaches commonly employ a voxel grid representation. Point clouds can be represented in different ways in the voxel grid. Indeed, voxel grid representations include binary and Truncated Signed Distance Fields (TSDF) representations [15]. TSDFs rely on the computation of normals; however, in the case of point clouds this computation can be noisy. We then ignore the normal signs and reformulate TSDFs to propose a new Truncated Distance Field (TDF) representation for point clouds

Deep learning approaches for lossy geometry compression



(a) Autoencoder model.

(b) Perceptual loss.

Figure 1. Perceptual loss based on an autoencoder. The grayed out parts do not need to be computed for the perceptual loss.

typically jointly optimize rate and distortion. As a result, an objective quality metric, employed as a loss function, is necessary to define the distortion objective during training. Such metrics should be differentiable, defined on the voxel grid and well correlated with perceived visual quality. In this context, the Weighted Binary Cross Entropy (WBCE) and the focal loss [16] are commonly used loss functions based on a binary voxel grid representation. They aim to alleviate the class imbalance between empty and occupied voxels caused by point cloud sparsity. However, they are poorly correlated with human perception as they only compute a voxel-wise error.

A number of metrics have been proposed for Point Cloud Quality Assessment (PCQA): the point-to-plane (D2) metric [18], PC-MSDM [19], PCQM [20], angular similarity [21], point to distribution metric [22], point cloud similarity metric [23], improved PSNR metrics [24] and a color based metric [25]. These metrics operate directly on the point cloud. However, they are not defined on the voxel grid and hence cannot be used easily as loss functions. Recently, to improve upon existing loss functions such as the WBCE and the focal loss, a neighborhood adaptive loss function [17] was proposed. Still, these loss functions are based on the explicit binary voxel grid representation. We show in this paper that loss functions based on the TDF representation are more correlated with human perception than those based on the binary representation.

The perceptual loss has previously been proposed as an objective quality metric for images [26]. Indeed, neural networks learn representations of images that are well correlated with perceived visual quality. This enables the definition of the perceptual loss as a distance between latent space representations. For the case of images, the perceptual loss provides competitive performance or even outperforms traditional quality metrics. We hypothesize that a similar phenomenon can be observed for point




Figure 2. Voxel grid representations of a point cloud. The upper bound distance value is 2 for the TDF and the TSDF. Normals are facing up in the TSDF.

Table 1: Objective quality metrics considered in this study.

Domain	Name	Signal type	Block aggregation	Learning based	Description
Points	D1 MSE	Coordinates	✗	✗	Point-to-point MSE
	D2 MSE	Coordinates	✗	✗	Point-to-plane MSE
	D1 PSNR	Coordinates	✗	✗	Point-to-point PSNR
	D2 PSNR	Coordinates	✗	✗	Point-to-plane PSNR
Voxel Grid	Bin BCE	Binary	L1	✗	Binary cross entropy
	Bin naBCE	Binary	L1	✗	Neighborhood adaptive binary cross entropy [17]
	Bin WBCE 0.75	Binary	L2	✗	Weighted binary cross entropy with $w = 0.75$
	Bin PL	Binary	L1	✓	Perceptual loss (explicit) on all feature maps
	Bin PL F1	Binary	L1	✓	Perceptual loss (explicit) on feature map 1
	TDF MSE	Distances	L1	✗	Truncated distance field (TDF) MSE
	TDF PL	Distances	L1	✓	Perceptual loss (implicit) over all feature maps
TDF PL F9	Distances	L1	✓	Perceptual loss (implicit) on feature map 9	

clouds.

Therefore, we propose a differentiable perceptual loss for training deep neural networks aimed at compressing point cloud geometry. We investigate how to build and train such a perceptual loss to improve point cloud compression results. Specifically, we build a differentiable distortion metric suitable for training neural networks to improve PCC approaches based on deep learning. We then validate our approach experimentally on the ICIP2020 [27] subjective dataset. The main contributions of the paper are as follows:

- A novel perceptual loss for 3D point clouds that outperforms existing metrics on the ICIP2020 subjective dataset
- A novel implicit TDF voxel grid representation
- An evaluation of binary (explicit) and TDF (implicit) representations in the context of deep learning approaches for point cloud geometry compression

Voxel grid representations

In this study, we consider different voxel grid representations for point clouds. A commonly used voxel grid representation is the explicit binary occupancy representation where the occupancy

of a voxel (occupied or empty) is represented with a binary value (Figure 2a). In the binary (Bin) representation (Figure 2a), each voxel has a binary occupancy value indicating whether it is occupied or empty. When the i th voxel is occupied, then $x_i = 1$ and otherwise $x_i = 0$.

Another representation is the implicit TSDF representation which has been employed for volume compression [9]. Instead of an occupancy value, the value of a voxel is the distance from this voxel to the nearest point and the sign of this value is determined from the orientation of the normal (Figure 2b). However, this requires reliable normals which may not be available in sparse and/or heavily compressed point clouds.

Hence, we propose an implicit TDF representation which is a variant of the TSDF without signs and therefore does not require normals. In the implicit TDF representation (Figure 2c), the i th voxel value is $x_i = d$, where d is the distance to its nearest occupied voxel. Consequently, $x_i = 0$ when a voxel is occupied and $x_i = d$ with $d > 0$ otherwise. Additionally, we truncate and normalize the distance values into the $[0, 1]$ interval with

$$x_i = \min(d, u)/u, \quad (1)$$

where u is an upper bound value.

In this study, we focus on the explicit binary and implicit TDF representations.

Objective quality metrics

In Table 1, we present the objective quality metrics considered in this study. Specifically, we evaluate metrics that are differentiable on the voxel grid to evaluate their suitability as loss functions for point cloud geometry compression. We include metrics defined on binary and TDF representations and we compare their performance against traditional point set metrics.

Voxel grid metrics

We partition the point cloud into blocks and compute voxel grid metrics for each block. For each metric, we aggregate metric values over all blocks with either L1 or L2 norm. Specifically, we select the best aggregation experimentally for each metric.

Given two point clouds A and B, we denote the i th voxel value for each point cloud as $x_i^{(A)}$ and $x_i^{(B)}$. Then, we define the WBCE as follows

$$-\frac{1}{N} \sum_i \left(\alpha x_i^{(A)} \log(x_i^{(B)}) + (1 - \alpha)(1 - x_i^{(A)}) \log(1 - x_i^{(B)}) \right), \quad (2)$$

where α is a balancing weight between 0 and 1. The binary cross entropy (BCE) refers to the case $\alpha = 0.5$.

Different from the WBCE, the Focal Loss (FL) amplifies ($\gamma > 1$) or reduces ($\gamma < 1$) errors and is defined as follows

$$-\sum_i \left(\alpha x_i^{(A)} (1 - x_i^{(B)})^\gamma \log(x_i^{(B)}) + (1 - \alpha)(1 - x_i^{(A)}) (x_i^{(B)})^\gamma \log(1 - x_i^{(B)}) \right), \quad (3)$$

where α is a balancing weight and the log arguments are clipped between 0.001 and 0.999.

Compared to the WBCE, the FL adds two factors $(1 - x_i^{(B)})^\gamma$ and $(x_i^{(B)})^\gamma$. However, while in the context of neural training, $x_i^{(B)}$ is an occupancy probability, in the context of quality assessment, $x_i^{(B)}$ is a binary value. As a result, the FL is equivalent to the WBCE since γ has no impact in the latter case. For this reason, we include the WBCE with $\alpha = 0.75$ in our experiments as an evaluation proxy for the FL used in [6].

The neighborhood adaptive BCE (naBCE) [17] was proposed as an alternative to the BCE and FL [16]. It is a variant of the WBCE in which the weight α adapts to the neighborhood of each voxel u resulting in a weight α_u . Given a voxel u , its neighborhood is a window W of size $m \times m \times m$ centered on u . Then, the neighborhood resemblance r_u is the sum of the inverse euclidean distances of neighboring voxels with the same binary occupancy value as u . Finally, the weight α_u is defined as $\alpha_u = \max(1 - r_u / \max(r), 0.001)$ where $\max(r)$ is the maximum of all neighborhood resemblances.

The Mean Squared Error (MSE) on the TDF is expressed as follows

$$\frac{1}{N} \sum_i (x_i^{(A)} - x_i^{(B)})^2. \quad (4)$$

Perceptual Loss

We propose a perceptual loss based on differences between latent space representations learned by a neural network. More precisely, we use an autoencoder as the underlying neural network. The model architecture of autoencoder used and its training procedure are presented in the following.

Model architecture

We adopt an autoencoder architecture based on 3D convolutions and transposed convolutions. Given an input voxel grid x , we perform an analysis transform $f_a(x) = y$ to obtain the latent space y and a synthesis transform $f_s(y) = \tilde{x}$ as seen in Figure 1a. The analysis transform is composed of three convolutions with kernel size 5, stride 2 while the synthesis transform is composed of three transposed convolutions with same kernel size and stride. We use ReLU [28] activations for all layers except for the last layer which uses a sigmoid activation.

Training

Using the previously defined architecture, we train two neural networks: one with explicit representation (binary) and another with implicit representation (TDF).

In the explicit case, we perform training of the perceptual loss with a focal loss function as defined in Eq. (3). In the implicit case, we first define the Kronecker delta δ_i such that $\delta_i = 1$ when $i = 0$ and otherwise $\delta_i = 0$. Then, we define an adaptive MSE loss function for the training of the perceptual loss as follows

$$\frac{1}{N} \sum_i \left(\delta_{1-x_i^{(A)}} w (x_i^{(A)} - x_i^{(B)})^2 + (1 - \delta_{1-x_i^{(A)}}) (1 - w) (x_i^{(A)} - x_i^{(B)})^2 \right), \quad (5)$$

where w is a balancing weight. Specifically, we choose w as the proportion of distances strictly inferior to 1 with

$$w = \min \left(\max \left(\frac{\sum_i 1 - \delta_{1-x_i^{(A)}}}{N}, \beta \right), 1 - \beta \right). \quad (6)$$

where β is a bounding factor such that w is bounded by $[\beta, 1 - \beta]$. This formulation compensates for class imbalance while avoiding extreme weight values.

In that way, the loss function adapts the contributions from the voxels that are far from occupied voxels ($x_i^{(A)} = 1$) and voxels that are near occupied voxels ($x_i^{(A)} < 1$). We train the network with the Adam [29] optimizer.

Metric

As seen in Figure 1b, in order to compare two point clouds $x^{(A)}$ and $x^{(B)}$, we compute their respective latent spaces $f_a(x^{(A)}) = y^{(A)}$ and $f_a(x^{(B)}) = y^{(B)}$ using the previously trained analysis transform. These latent spaces each have F feature maps of size $W \times D \times H$ (width, depth, height). Then, we define the MSE between latent spaces as follows

$$\frac{1}{N} \sum_i (y_i^{(A)} - y_i^{(B)})^2. \quad (7)$$

We compute this MSE either over all F feature maps or on single feature maps.

Table 2: Statistical analysis of objective quality metrics.

Method	PCC	SROCC	RMSE	OR
TDF PL F9	0.951	0.947	0.094	0.375
D2 MSE	0.946	0.943	0.100	0.469
TDF MSE	0.940	0.940	0.103	0.385
D1 MSE	0.938	0.933	0.109	0.479
TDF PL	0.935	0.933	0.110	0.490
Bin PL F1	0.922	0.916	0.115	0.406
D2 PSNR	0.900	0.898	0.129	0.500
Bin WBCE 0.75	0.875	0.859	0.144	0.531
Bin PL	0.863	0.867	0.151	0.552
D1 PSNR	0.850	0.867	0.158	0.448
Bin naBCE	0.740	0.719	0.201	0.573
Bin BCE	0.713	0.721	0.207	0.635

Point set metrics

Point-to-point (D1) and point-to-plane (D2)

The point-to-point distance (D1) [30] measures the average error between each point in A and their nearest neighbor in B:

$$e_{A,B}^{D1} = \frac{1}{N_A} \sum_{\forall a_i \in A} \|a_i - b_j\|_2^2 \quad (8)$$

where b_j is the nearest neighbor of a_i in B.

In contrast to D1, the point-to-plane distance (D2) [18] projects the error vector along the normal and is expressed as follows

$$e_{A,B}^{D2} = \frac{1}{N_A} \sum_{\forall a_i \in A} ((a_i - b_j) \cdot n_i)^2 \quad (9)$$

where b_j is the nearest neighbor of a_i in B and n_i is the normal vector at a_i .

The normals for original and distorted point clouds are computed with local quadric fittings using 9 nearest neighbors.

The D1 and D2 MSEs are the maximum of $e_{A,B}$ and $e_{B,A}$ and their Peak Signal-to-Noise Ratio (PSNR) is then computed with a peak error corresponding to three times the point cloud resolution as defined in [30].

Experiments

Experimental setup

We evaluate the metrics defined above on the ICIP2020 [27] subjective dataset. It contains 6 point clouds [31, 32] compressed using G-PCC Octree, G-PCC Trisoup and V-PCC with 5 different rates yielding a total of 96 stimuli (with 6 references) and their associated subjective scores.

For each metric, we compute the Pearson Correlation Coefficient (PCC), the Spearman Rank Order Correlation Coefficient (SROCC), the Root Mean Square Error (RMSE) and the Outlier Ratio (OR). We evaluate the statistical significance of the differences between PCCs using the method in [33]. These metrics are computed after logistic fittings with cross-validation splits. Each split contains stimuli for one point cloud (i.e. reference point cloud and its distorted versions) as a test set and stimuli of all

other point clouds as a training set. The metrics are then computed after concatenating results for the test set of each split. They are summarized in Table 1 and the values before and after logistic fitting are shown in Figure 3.

We use an upper bound value $u = 5$ when computing the TDF in Eq. (1) and a block size of 64 when block partitioning point clouds. The naBCE window size is $m = 5$ as in the original paper. The perceptual loss is trained with a learning rate of 0.001, $\beta_1 = 0.9$ and $\beta_2 = 0.999$ on the ModelNet dataset [34] after block partitioning using Python 3.6.9 and TensorFlow [35] 1.15.0.

Comparison of perceptual loss feature maps

In our experiments, we first considered the perceptual loss computed over all feature maps. However, we observed that some feature maps are more perceptually relevant than others. Consequently, we include the best feature maps for each voxel grid representation in our results. This corresponds to feature map 9 (TDF PL F9) for TDF PL and 1 (Bin PL F1) for Bin PL.

Moreover, We observe that some feature maps are unused by the neural network (constant). Therefore, they exhibit high RMSE values (all equal to 0.812) as their perceptual loss MSE are equal to 0. Specifically, we observe that TDF PL has 6 unused feature maps, while Bin PL has a single unused feature map. This suggests that the perceptual loss learns a sparser latent space representation when using TDF compared to binary. Thus, implicit representations may improve compression performance compared to explicit representations as fewer feature maps may be needed.

Comparison of objective quality metrics

In Table 2, we observe that the TDF PL F9 is the best method overall. In particular, identifying the most perceptually relevant feature map and computing the MSE on this feature map provides a significant improvement. Specifically, the difference between the PCCs of TDF PL F9 and TDF PL is statistically significant with a confidence of 95%.

For voxel grid metrics, we observe that TDF metrics perform better than binary metrics. In particular, the RMSEs of the former are noticeably lower for point clouds compressed with G-PCC Octree compared to the RMSE of the latter as can be seen in Table 3. This suggests that implicit representations may be better at dealing with density differences between point clouds in the context of point cloud quality assessment.

Conclusion

We proposed a novel perceptual loss that outperforms existing objective quality metrics and is differentiable in the voxel grid. As a result, it can be used as a loss function in deep neural networks for point cloud compression and it is more correlated with perceived visual quality compared to traditional loss functions such as the BCE and the focal loss. Overall, metrics on the proposed implicit TDF representation performed better than explicit binary representation metrics. Additionally, we observed that the TDF representation yields sparser latent space representations compared to the binary representations. This suggests that switching from binary to the TDF representation may improve compression performance in addition to enabling the use of better loss functions.

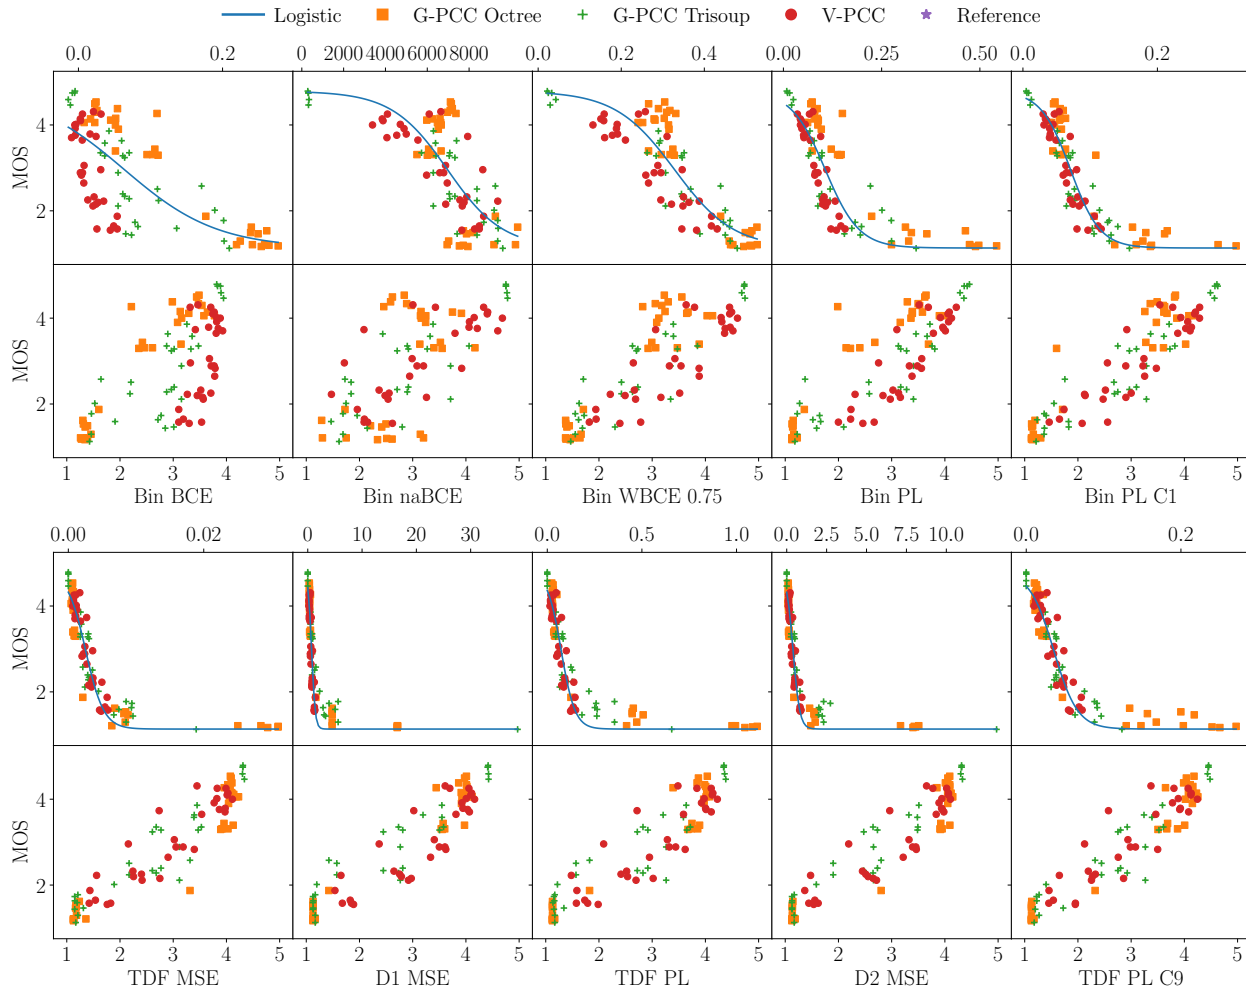


Figure 3. Scatter plots between the objective quality metrics and the MOS values. The plots before and after logistic fitting are shown.

Table 3: Statistical analysis of objective quality metrics by compression method.

Method	G-PCC Octree			G-PCC Trisoup			V-PCC		
	PCC	SROCC	RMSE	PCC	SROCC	RMSE	PCC	SROCC	RMSE
TDF PL F9	0.975	0.859	0.078	0.936	0.910	0.101	0.897	0.850	0.106
D2 MSE	0.962	0.829	0.094	0.954	0.924	0.103	0.903	0.860	0.103
TDF MSE	0.952	0.839	0.106	0.933	0.917	0.106	0.912	0.867	0.098
D1 MSE	0.976	0.851	0.082	0.937	0.918	0.126	0.876	0.844	0.119
TDF PL	0.970	0.840	0.087	0.918	0.900	0.127	0.876	0.837	0.115
Bin PL F1	0.941	0.786	0.138	0.927	0.907	0.107	0.898	0.865	0.109
D2 PSNR	0.943	0.890	0.110	0.926	0.895	0.108	0.738	0.723	0.166
Bin WBCE 0.75	0.923	0.747	0.163	0.918	0.886	0.112	0.850	0.786	0.164
Bin PL	0.931	0.852	0.186	0.892	0.886	0.130	0.880	0.852	0.142
D1 PSNR	0.903	0.859	0.156	0.910	0.895	0.117	0.599	0.689	0.202
Bin naBCE	0.552	0.357	0.277	0.846	0.786	0.154	0.748	0.692	0.170
Bin BCE	0.946	0.841	0.188	0.776	0.800	0.177	0.574	0.500	0.250

Acknowledgments

We would like to thank the authors of [17] for providing their implementation of naBCE. This work was funded by the ANR ReVeRy national fund (REVERY ANR-17-CE23-0020).

References

- [1] S. Schwarz *et al.*, “Emerging MPEG Standards for Point Cloud Compression,” *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, pp. 1–1, 2018.
- [2] “ITU-T Recommendation H.265: High efficiency video coding (HEVC),” Nov. 2019.
- [3] T. Ebrahimi *et al.*, “JPEG Pleno: Toward an Efficient Representation of Visual Reality,” *IEEE MultiMedia*, vol. 23, no. 4, pp. 14–20, Oct. 2016.
- [4] “Final Call for Evidence on JPEG Pleno Point Cloud Coding,” in *ISO/IEC JTC1/SC29/WG1 JPEG output document N88014*, Jul. 2020.
- [5] M. Quach, G. Valenzise, and F. Dufaux, “Learning Convolutional Transforms for Lossy Point Cloud Geometry Compression,” in *2019 IEEE Int. Conf. on Image Process. (ICIP)*, Sep. 2019, pp. 4320–4324.
- [6] —, “Improved Deep Point Cloud Geometry Compression,” in *2020 IEEE Int. Workshop on Multimedia Signal Process. (MMSP)*, Oct. 2020.
- [7] J. Wang *et al.*, “Learned Point Cloud Geometry Compression,” *arXiv:1909.12037 [cs, eess]*, Sep. 2019.
- [8] —, “Multiscale Point Cloud Geometry Compression,” *arXiv:2011.03799 [cs, eess]*, Nov. 2020.
- [9] D. Tang *et al.*, “Deep Implicit Volume Compression,” in *2020 IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 1290–1300.
- [10] S. Milani, “A Syndrome-Based Autoencoder For Point Cloud Geometry Compression,” in *2020 IEEE Int. Conf. on Image Process. (ICIP)*, Oct. 2020, pp. 2686–2690.
- [11] M. Quach, G. Valenzise, and F. Dufaux, “Folding-Based Compression Of Point Cloud Attributes,” in *2020 IEEE Int. Conf. on Image Process. (ICIP)*, Oct. 2020, pp. 3309–3313.
- [12] E. Alexiou, K. Tung, and T. Ebrahimi, “Towards neural network approaches for point cloud compression,” in *Applications of Digit. Image Process. XLIII*, vol. 11510. Int. Society for Optics and Photonics, Aug. 2020, p. 1151008.
- [13] L. Huang *et al.*, “OctSqueeze: Octree-Structured Entropy Model for LiDAR Compression,” in *2020 IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 1310–1320.
- [14] S. Biswas *et al.*, “MuSCL: Multi Sweep Compression of LiDAR using Deep Entropy Models,” *Adv. Neural Inf. Process. Syst.*, vol. 33, 2020.
- [15] B. Curless and M. Levoy, “A volumetric method for building complex models from range images,” in *Proc. of the 23rd annual conf. on Computer graphics and interactive techniques - SIGGRAPH '96*. ACM Press, 1996, pp. 303–312.
- [16] T.-Y. Lin *et al.*, “Focal Loss for Dense Object Detection,” in *2017 IEEE Int. Conf. on Computer Vision (ICCV)*, Oct. 2017, pp. 2999–3007.
- [17] A. Guarda, N. Rodrigues, and F. Pereira, “Neighborhood Adaptive Loss Function for Deep Learning-based Point Cloud Coding with Implicit and Explicit Quantization,” *IEEE MultiMedia*, pp. 1–1, 2020.
- [18] D. Tian *et al.*, “Geometric distortion metrics for point cloud compression,” in *2017 IEEE Int. Conf. on Image Process. (ICIP)*. Beijing: IEEE, Sep. 2017, pp. 3460–3464.
- [19] G. Meynet, J. Digne, and G. Lavoué, “PC-MSDM: A quality metric for 3D point clouds,” in *2019 Eleventh Int. Conf. on Quality of Multimedia Experience (QoMEX)*, Jun. 2019, pp. 1–3.
- [20] G. Meynet *et al.*, “PCQM: A Full-Reference Quality Metric for Colored 3D Point Clouds,” in *12th Int. Conf. on Quality of Multimedia Experience (QoMEX 2020)*, Athlone, Ireland, May 2020.
- [21] E. Alexiou and T. Ebrahimi, “Point Cloud Quality Assessment Metric Based on Angular Similarity,” in *2018 IEEE Int. Conf. on Multimedia and Expo (ICME)*, Jul. 2018, pp. 1–6.
- [22] A. Javaheri *et al.*, “Mahalanobis Based Point to Distribution Metric for Point Cloud Geometry Quality Evaluation,” *IEEE Signal Process. Lett.*, vol. 27, pp. 1350–1354, 2020.
- [23] E. Alexiou and T. Ebrahimi, “Towards a Point Cloud Structural Similarity Metric,” in *2020 IEEE Int. Conf. on Multimedia Expo Workshops (ICMEW)*, Jul. 2020, pp. 1–6.
- [24] A. Javaheri *et al.*, “Improving PSNR-based Quality Metrics Performance For Point Cloud Geometry,” in *2020 IEEE Int. Conf. on Image Process. (ICIP)*, Oct. 2020, pp. 3438–3442.
- [25] I. Viola, S. Subramanyam, and P. Cesar, “A Color-Based Objective Quality Metric for Point Cloud Contents,” in *2020 Twelfth Int. Conf. on Quality of Multimedia Experience (QoMEX)*, May 2020, pp. 1–6.
- [26] R. Zhang *et al.*, “The Unreasonable Effectiveness of Deep Features as a Perceptual Metric,” in *2018 IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*. IEEE, Jun. 2018, pp. 586–595.
- [27] S. Perry *et al.*, “Quality Evaluation Of Static Point Clouds Encoded Using MPEG Codecs,” in *2020 IEEE Int. Conf. on Image Process. (ICIP)*, Oct. 2020, pp. 3428–3432.
- [28] V. Nair and G. E. Hinton, “Rectified Linear Units Improve Restricted Boltzmann Machines,” in *Proc. of the 27th Int. Conf. on Mach. Learn. (ICML)*, 2010, pp. 807–814.
- [29] D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” in *2015 3rd Int. Conf. on Learn. Representations*, Dec. 2014.
- [30] “Common test conditions for point cloud compression,” in *ISO/IEC JTC1/SC29/WG11 MPEG output document N19084*, Feb. 2020.
- [31] C. Loop *et al.*, “Microsoft voxelized upper bodies - a voxelized point cloud dataset,” in *ISO/IEC JTC1/SC29 Joint WG11/WG1 (MPEG/JPEG) input document m38673/M72012*, May 2016.
- [32] E. d’Eon *et al.*, “8i Voxelized Full Bodies - A Voxelized Point Cloud Dataset,” in *ISO/IEC JTC1/SC29 Joint WG11/WG1 (MPEG/JPEG) input document WG11M40059/WG1M74006*, Geneva, Jan. 2017.
- [33] G. Y. Zou, “Toward using confidence intervals to compare correlations,” *Psychol. Methods*, vol. 12, no. 4, pp. 399–413, Dec. 2007.
- [34] N. Sedaghat *et al.*, “Orientation-boosted Voxel Nets for 3D Object Recognition,” *arXiv:1604.03351 [cs]*, Apr. 2016.
- [35] M. Abadi *et al.*, “TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems,” *arXiv:1603.04467 [cs]*, Mar. 2016.

Author Biography

Maurice Quach received the Computer Science Engineer Diploma from University of Technology of Compiègne in 2018. He is currently studying for a PhD on Point Cloud Compression and Quality Assessment under the supervision of Frederic Dufaux and Giuseppe Valenzise at Université Paris-Saclay, CNRS, CentraleSupélec, Laboratoire des signaux et systèmes, France.


JOIN US AT THE NEXT EI!

IS&T International Symposium on

Electronic Imaging

SCIENCE AND TECHNOLOGY

Imaging across applications . . . Where industry and academia meet!



- **SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •**
- **INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •**

www.electronicimaging.org

