CMDM
A Full-Reference Objective Quality Metric for 3D Meshes with Diffuse Colors

November 3, 2020

Authors:
Yana Nehmé
Florent Dupont
Jean-Philippe Farrugia
Patrick Le Callet
Guillaume Lavoué

Overview

This document presents the visual quality metric for 3D meshes with rich attributes, that we proposed in the context of ANR-PISCo project (ANR-17-CE33-0005). The metric we proposed (*CMDM*) is inspired by the MSDM2 frameworks [6] and adapted to address multimodal nature of enriched data (geometry and color information). CMDM is the first metric for quality assessment of 3D meshes with diffuse colors, which works entirely on the mesh domain, at vertex level. It is a full-reference data-driven metric, that incorporates perceptually-relevant curvature-based and color-based features. Our metric demonstrates good results and stability over two datasets.

Moreover, Meynet et al. [12] and our team worked together, so that we successfully designed and implemented a metric for quality assessment of colored 3D point clouds (*PCQM*), that considers the same initial collection of color and geometric features as CMDM.

1. Toward an objective metric for assessment of colored mesh quality

1.1 Introduction

Nowadays, three-dimensional (3D) graphics are widely used in many applications such as digital entertainment, architecture and scientific simulation. These data are increasingly rich and detailed; as a complex 3D scene may contain millions of geometric primitives, enriched with various appearance attributes such as texture maps designed to produce a realistic material appearance. These huge data tend to be visualized on various devices (e.g., smartphone, head mounted display) and possibly via the network. Therefore, to avoid latency or rendering issues, there is a critical need for the compression and simplification of these high quality 3D models. These processing operations are lossy. They operate on both geometry and appearance attributes, which inevitably introduce distortions that impact the perceived quality of the data and thus the quality of user experience (QoE).

Objective quality metrics are thus critically needed to automatically predict the level of annoyance caused by these operations. Most metrics in the literature evaluate only geometric distortions (i.e. they consider meshes without appearance attributes), e.g. [6, 16, 17]. When it comes to meshes with diffuse color information (either in the form of texture or vertex-colors), little work has been published [15] [4]. Actually, for this kind of data, it is still unclear how color and geometry distortions affect quality.

In this report, we address the problem of objective quality assessment of 3D models with diffuse colors. So, we designed an objective quality assessment metric for colored meshes: *CMDM* (Color Mesh Distortion Measure). This is a full-reference data-driven metric that fully operates on the mesh domain, at vertex level. It consists of a linear combination of perceptually-relevant features related to color and geometry. The optimal set of features was selected through logistic regression. To evaluate the performance of *CMDM*, we used 2 subjective ground-truths: the LIRIS Textured Mesh Database [4] and a dataset of animated meshes with vertex colors (not published yet). Our metric demonstrates good results and a better stability than image quality metrics.

To the best of our knowledge, our proposed metric is the first attempt to integrate both geometry and color information for quality assessment of meshes with diffuse colors. The source code of the metric is made publicly available[1] on the MEsh Processing Platform (MEPP).

1.2 CMDM: Color Mesh Distortion Measure

As outlined in the introduction, constructing an objective metric for the quality assessment of 3D content with appearance attributes is no trivial task. The main reasons are: (1) the multimodal nature of the data (geometry and color or texture information) and (2) the complex processing pipeline that constructs the final rendered image from the 3D content (computation of light-material interactions, viewpoint selection, and rasterization). To overcome this problem, we consider a data-driven approach based on the results and data of a subjective study we conducted. Thus, we propose an objective metric for colored mesh quality assessment as a linear combination of accurate geometry and color quality measurements.

1.2.1 Overview of our approach

The metric we propose is a full-reference multiscale metric based on curvature and color statistics computed on local corresponding neighborhoods from the original and distorted models. The metric is largely inspired by the MSDM2 frameworks from which we take the curvature features and the neighborhood correspondence mechanisms [6]. To address the color-related aspects of our metric, we consider the features introduced in the 2D image-difference framework of Lissner et al. [9]. Their color features have recently been used successfully for the quality assessment of colored 3D point clouds [12].

[1] https://github.com/MEPP-team/MEPP2

Our framework is as follows: For given distorted M_{dist} and reference M_{ref} meshes, we first establish a correspondence between M_{dist} and M_{ref} (see section 1.2.2). Then for each scale h_i, we define a spherical neighborhood around each vertex v of M_{dist} (see section 1.2.3) and compute a set of local geometry and color based features over the points belonging to the neighborhood of v and their corresponding points on M_{ref} (see section 1.2.4). Local single-scale feature values are pooled into global multiscale features f_j. Finally, CMDM is defined as a linear combination of an optimal subset of features determined through logistic regression (see section 1.2.5).

1.2.2 Correspondence between meshes

The first objective is to establish a correspondence between the meshes being compared (M_{dist} and M_{ref}). Thus, we match each vertex v of the distorted mesh M_{dist} with its nearest 3D point $\hat{v}$ on the surface of the reference mesh M_{ref} using a fast asymmetric projection (as in MSDM2, we consider the AABB tree structure from CGAL [2]). Then, for each projected 3D point ($\hat{v}$), we compute its curvature and color using barycentric interpolation from vertices of the triangle it belongs to. This way, each vertex from M_{dist} has a corresponding point on M_{ref} (with a curvature and a color value).
The correspondence is scale-independent: it takes place once only at the beginning of the process. Nevertheless, the curvature and color values of $\hat{v}$ are updated for each scale h_i.

1.2.3 Neighborhood Computation

As stated above, the features used in our metric are not computed globally on the entire mesh but locally at multiple scales over spherical neighborhoods around each vertex. Thus as in [6], we define, for each scale h, a neighborhood $N(v, h)$ of radius h around each vertex v of M_{dist} as the connected set of vertices belonging to the sphere with center v and radius h. We also add to this neighborhood the intersections between this sphere and the edges of M_{dist}. The curvature and color values of the intersection points are interpolated. Then, we consider for the set of points belonging to $N(v, h)$ their projected 3D points on M_{ref} (corresponding neighborhood of $\hat{v}$). Features are computed by considering curvature and color statistics over $N(v, h) \in M_{dist}$ and $N(\hat{v}, h) \in M_{ref}$. In this paper, we consider the following three scales: $h_i \in \{0.003BB, 0.0045BB, 0.006BB\}$, where BB is the maximum length of the Axis-Aligned Bounding Box (AABB) of the stimulus. The choice of these scales is detailed and justified in the supplementary material.

1.2.4 Perceptually relevant features

For each scale h, the following 8 features are computed over the local corresponding neighborhood of each vertex v of M_{dist}.

A. Geometry-based features

These features are based on mean curvature information defined at multiple scales. To compute curvature, we adopted the method developed by Alliez et al. [3], which evaluates the curvature tensor on a geodesic neighborhood around each vertex. This method is interesting and robust because it avoids the problem of sensitivity to connectivity (M_{dist} and M_{ref} do not necessarily share the same connectivity nor the same level of details). Note that, we used a radius $r = \frac{h}{3}$ for the computation of curvature as a good compromise between small radii which capture tiny details and larger radii which provide strong smoothing effects.
As in [6], we consider the following geometry features:

$$\textit{Curvature comparison} \quad f_1^h(v) = \frac{\left\| \overline{C_v^h} - \overline{C_{\hat{v}}^h} \right\|}{max(\overline{C_v^h}, \overline{C_{\hat{v}}^h}) + k} \tag{1.1}$$

$$\textit{Curvature contrast} \quad f_2^h(v) = \frac{\left\| \sigma_{C_v^h} - \sigma_{C_{\hat{v}}^h} \right\|}{max(\sigma_{C_v^h}, \sigma_{C_{\hat{v}}^h}) + k} \tag{1.2}$$

$$\textit{Curvature structure} \quad f_3^h(v) = \frac{\left\| \sigma_{C_v^h}\sigma_{C_{\hat{v}}^h} - \sigma_{C_v^h C_{\hat{v}}^h} \right\|}{\sigma_{C_v^h}\sigma_{C_{\hat{v}}^h} + k} \tag{1.3}$$

where k is a constant to avoid instability when denominators are close to zero ($k = 1$ as in [6]). $\overline{C_v^h}$ and $\overline{C_{\hat{v}}^h}$ are Gaussian-weighted averages of curvature over the points belonging to the neighborhood $N(v, h)$ and $N(\hat{v}, h)$, respectively. Similarly, $\sigma_{C_v^h}$, $\sigma_{C_{\hat{v}}^h}$ and $\sigma_{C_v^h C_{\hat{v}}^h}$ are Gaussian-weighted standard deviations and covariance of curvature over these neighborhoods.

B. Color-based features

To compute the color features, we first transform the RGB color values of each vertex of the meshes being compared (M_{dist} and M_{ref}) into the perceptually uniform color space LAB200HL [8]. Lissener et al. recommended working in this color space since there is little cross contamination between the color attributes (lightness, chroma, hue). Each vertex v has of a lightness and two chromatic values (L_v, a_v, b_v). The chroma of the vertex is as follows: $Ch_v = \sqrt{a_v^2 + b_v^2}$.

We transposed for 3D meshes, the 2D image features proposed by [9]. These features take into account not only the luminance but also the chroma and hue components to better assess the chromatic distortions.

$$\textit{Lightness comparison} \qquad f_4^h(v) = \frac{1}{c_1(\overline{L_v^h} - \overline{L_{\hat{v}}^h})^2 + 1} \tag{1.4}$$

$$\textit{Lightness contrast} \qquad f_5^h(v) = \frac{\sigma_{L_v^h}\sigma_{L_{\hat{v}}^h} + c_2}{\sigma_{L_v^h}^2 + \sigma_{L_{\hat{v}}^h}^2 + c_2} \tag{1.5}$$

$$\textit{Lightness structure} \qquad f_6^h(v) = \frac{\sigma_{L_v^h L_{\hat{v}}^h} + c_3}{\sigma_{L_v^h}\sigma_{L_{\hat{v}}^h} + c_3} \tag{1.6}$$

$$\textit{Chroma comparison} \qquad f_7^h(v) = \frac{1}{c_4(\overline{Ch_v^h} - \overline{Ch_{\hat{v}}^h})^2 + 1} \tag{1.7}$$

$$\textit{Hue comparison} \qquad f_8^h(v) = \frac{1}{c_5\overline{\Delta H_{v\hat{v}}^h}^2 + 1} \tag{1.8}$$

where $\overline{L_v^h}$, $\overline{L_{\hat{v}}^h}$, $\overline{Ch_v^h}$ and $\overline{Ch_{\hat{v}}^h}$ denote the Gaussian-weighted averages of Lightness and Chroma computed respectively over the set of points belonging to $N(v,h)$ and $N(\hat{v},h)$. $\sigma_{L_v^h}$, $\sigma_{L_{\hat{v}}^h}$ and $\sigma_{L_v^h L_{\hat{v}}^h}$ are Gaussian-weighted standard deviations and covariance of lightness in the mentioned neighborhood. The term $\overline{\Delta H_{v\hat{v}}^h}$ refers to the Gaussian-weighted average hue difference between $N(v,h)$ and $N(\hat{v},h)$. It is defined as follows: $\Delta H_{v\hat{v}} = \sqrt{(a_v - a_{\hat{v}})^2 + (b_v - b_{\hat{v}})^2 - (Ch_v - Ch_{\hat{v}})^2}$. The constants c_1, c_2, c_3, c_4 and c_5 were set respectively to 0.002, 0.1, 0.1, 0.002 and 0,008 as in [9].

We invert the scaling of the color-based features so that they are consistent with curvature-based features (i.e. each color feature $f_j^h = 1 - f_j^h$). This way, a value of 0 means that there is no local (geometric and color) distortion around vertex v. All features $\in$ [0, 1].

1.2.5 Global perceptual quality score

The set of local geometric and color features, presented in the subsection above, is computed for each vertex of the distorted mesh and for each scale h_i. The local multiscale measure of the features is simply the average of its single-scale values.

$$f_j(v) = \frac{1}{n}\sum_{i=1}^{n} f_j^{h_i}(v) \tag{1.9}$$

where n is the number of scales used. It is defined in section 1.2.3 as well as h_i the scale values used (neighborhood radii) .

We aim to obtain a global score of visual distortion according to each feature (f_j). So, we average the local values of each feature over all the vertices.

$$f_j = \frac{1}{|M_{dist}|}\sum_{v \in M_{dist}} f_j(v) \tag{1.10}$$

where $|M_{dist}|$ is the number of vertices of the distorted mesh. The features f_j are all within the range [0, 1].

Our metric is then defined as a combination of the features f_j. However, choosing the best combination model is a crucial problem. For prediction of the color-image differences [9], the authors used a factorial combination model, while Meynet et al. considered a linear model for their point cloud quality metric [12]. In our case, we chose to consider a linear model: (1) to make the optimization easier and (2) because we tried nonlinear models such as Minkowski pooling, which did not provide better performance. Thus, the global multiscale distortion (GMD) score is computed as follows:

$$GMD_{M_{dist} \to M_{ref}} = \sum_{j \in S} w_j f_j \tag{1.11}$$

S is the set of feature indexes of our linear model. w_j weights the contribution of each feature to the overall distortion prediction. $GMD_{M_{dist} \to M_{ref}}$ evaluates the distortion of the distorted model regarding the reference model. In order to strengthen the robustness of our method and to obtain a symmetric measure, we also compute $GMD_{M_{ref} \to M_{dist}}$ and we retain the average as the final distortion measure $CMDM$.

$$CMDM = \frac{GMD_{M_{dist} \to M_{ref}} + GMD_{M_{ref} \to M_{dist}}}{2} \tag{1.12}$$

As in [10], the optimal subset of features of $CMDM$ and their corresponding weights are obtained through an optimization computed by logistic regression. The optimization is based on cross-validation (see section 1.3.1).

1.3 Results and evaluation

In this section, we evaluate the performance of our metric and compare it to state-of-the-art approaches, including 2D image metrics. To train and evaluate our metric, we used the ground truth database obtained from one of our subjective studies (not published yet). The database used is composed of 80 stimuli, generated from 5 high-resolution triangle meshes, each having diffuse color information represented by vertex colors (no texture mapping). The source models have been corrupted by 4 types of distortion applied on geometry and color. These selected distortions represent common simplification and compression operations typically used in 3D model modeling and post-processing. Each distortion was applied with 4 different strengths, adjusted manually in order to span the whole range of visual quality from imperceptible levels to high levels of impairment. The databaset was produced from a subjective study based on Double Stimulus Impairment Scale (DSIS) methodology, as recommended by [13]. Each stimulus was rated by at least 24 observers.

1.3.1 Toward an Optimal Combination of features

Our metric contains 8 different features f_j. In this 8 dimensional space, some features are obviously more significant than others. Also, features may be redundant with one another, and if all the features are taken into account, this could potentially lead to an overfitting. Therefore, in the same vein as [10], we conduct two Leave-One Out Cross-Validation tests (LOOCV) on the data obtained from our subjective experiment to select an optimal subset of features. Each cross-validation test divides the database into a training set that serves to optimize feature weights using linear regression and a test used for testing the obtained metric.

1. We split the training and test sets according to the source models. Given that there are 5 sources in our database, we leave 1 source model and its distortions out for testing, while the remaining stimuli (4 models * 16 distorted stimuli) are used for training. Thus after 5 folds, each source model has been used as a test set.

2. Similar to test 1, but we divide the database according to the distortion types (regardless of the model). We train the metric on 3 distortion types out of 4 (5 models * 12 distorted stimuli) and test on the fourth type. After 4 folds, each distortion type has been used once for testing.

These 2 types of LOOCV tests provide a good measure of the robustness of our metric. We exhaustively search through all possible combinations of features (255 combinations), and select the feature-subset that generates the best average performance of $CMDM$ over all the test sets (9 folds) in terms of the mean of Pearson Linear Correlation Coefficient (PLCC) and Spearman Rank Order Correlation Coefficient (SROCC). We obtained that the final model of our metric is composed of only 4 features: Curvature contrast (f_2), Lightness contrast (f_5) and structure (f_6) and chroma comparison (f_7). The results of our metric and comparisons with state-of-the-art approaches are reported in the following sections.

1.3.2 Comparisons of objective metrics

In this section, we present the results of the cross-validation tests, described in the previous subsection. As an ablation study, we compare our metric with two of its versions trained with different subsets of features: *CMDM_Geo* that takes into account only the geometry features and *CMDM_Col* based only on color features. As a baseline, we also include results of a classical color distance D_LAB, which is the average of the color difference (in LAB2000HL) computed symmetrically between the reference and the distorted model. Finally, we compare our metric with 3 state-of-the-art full-reference image quality metrics (IQMs): *SSIM* [18], *HDR-VDP2* [11], *iCID* [14]. To apply these IQMs, we generate for each 3D object in our database, a set of 18 snapshots taken from different viewpoints (camera positions regularly sampled). The global quality score of a stimulus, given by an IQM, is then the average of the objective scores over all its snapshots.

Figure 1.1 compares the overall performance of the tested metrics for the 2 cross-validation scenarios presented in 1.3.1. AUC_{DS} and AUC_{BW} are 2 measures (Area Under the Curve values), proposed by Krasula et al. [5], that determining the classification abilities of the metrics according to two scenarios: (1) AUC_{DS} assesses how well can the metric distinguish between significantly different and similar pairs of stimuli, and (2) AUC_{BW} evaluates how well the metric is able to detect the stimulus of better quality in a significantly different pair of stimuli. These measures take into account the uncertainty of the subjective scores.

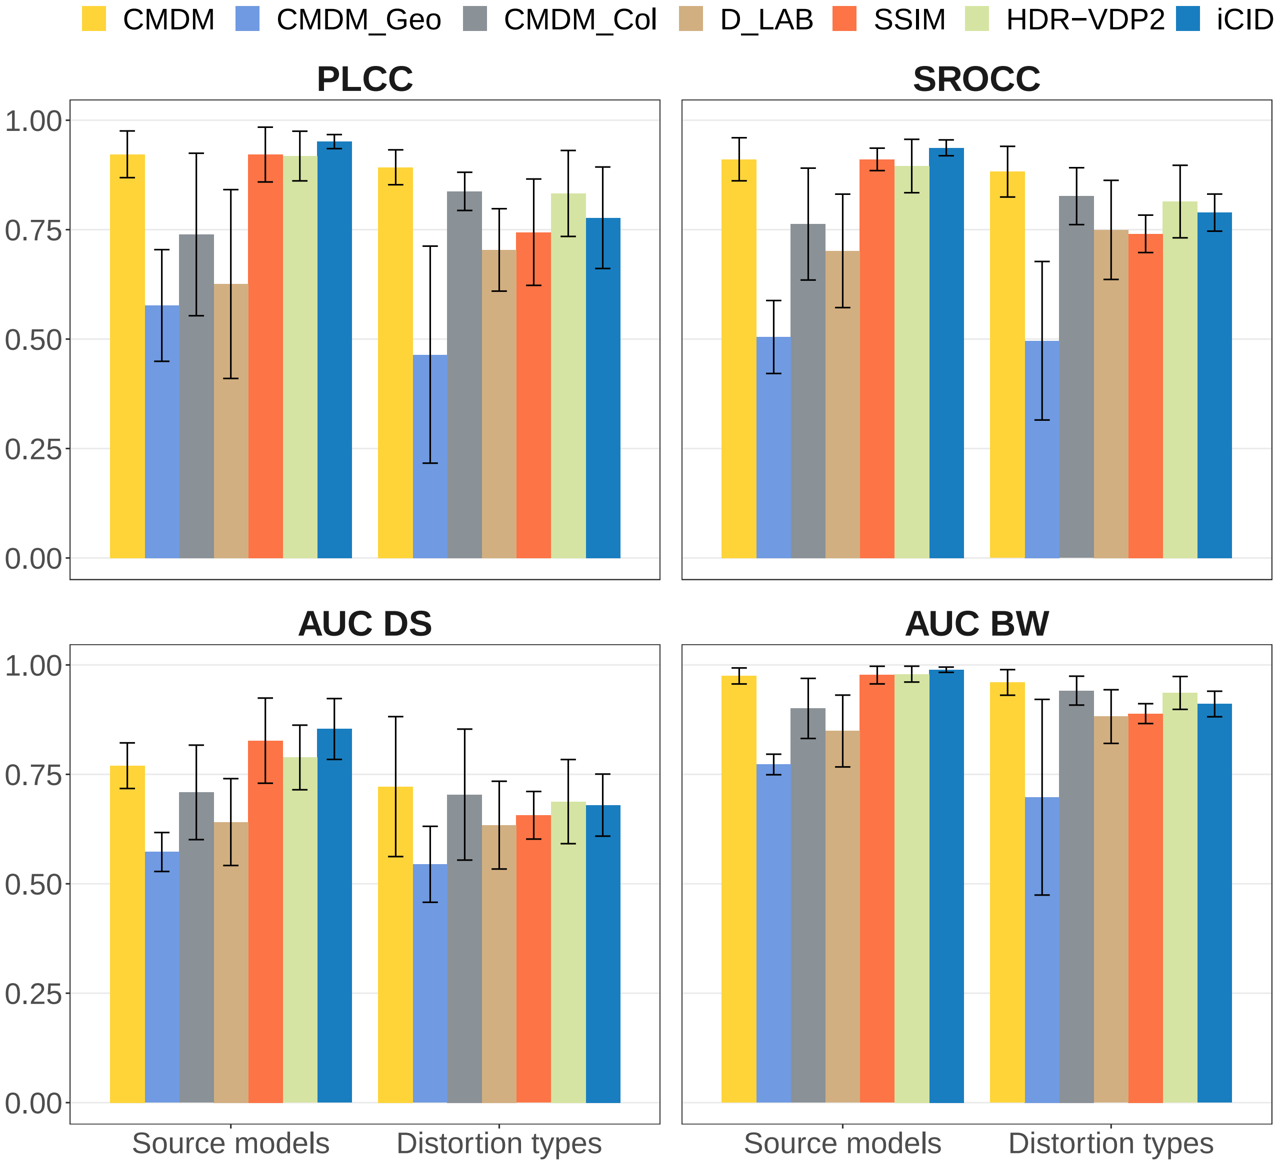

Figure 1.1 – Performance comparison of several metrics on two cross-validation tests. Mean performance evaluation measures are reported. Error bars indicate the standard deviation over the test sets.

For the LOOCV test according to the source models, Figure 1.1 demonstrates that *CMDM* outperforms other model-based metrics. It shows almost the same performance as IQMs in terms of correlations and detection of better quality stimuli (AUC_{BW}). IQMs provide better results in identifying the significantly different pairs of stimuli (AUC_{DS}). We believe this is primarily related to the advantage of IQMs over our metric and other model-based metrics regarding their natural incorporation/knowledge of the entire rendering pipeline. Indeed, IQMs operate on snapshots that consider the same rendering, apparent brightness and lighting conditions as those seen by participants. On the contrary, our metric only considers 3D data, without any knowledge of the rendering conditions. Considering the LOOCV test among the distortions, we notice that our metric performs better than the others, including IQMs. The color-based version of our metric (*CMDM_ Col*) also produces good results. IQMs show a significant decrease in performance, compared to the LOOCV based on source models. These observations corroborate previous results by Lavoué et al. [7]: image-based metrics perform very well when evaluating the quality of different versions of a single source, yet they are less accurate when differentiating/ranking distortions applied on different sources.

1.3.3 Recommended weights

To provide the recommended model of our metric, we averaged the weights obtained for each training subset of the two LOOCV tests. CMDM is thus defined, for the three selected scales ($h_i \in \{0.003BB, 0.0045BB, 0.006BB\}$), as follows:

$$CMDM_{rec} = 0.091f_2 + 0.22f_5 + 0.032f_6 + 0.656f_7 \tag{1.13}$$

In order to reveal the relative importance of each of the 4 features, we scaled the weights presented in the equation above with the standard deviation of the features. Scaled weights are 0.333, 0.46, 0.07 and 0.136, respectively, for f_2, f_5, f_6 and f_7. The curvature and lightness contrast features (f_2 and f_5) have the highest overall importance. It would seem that users are particularly sensitive to artifacts that harm the contrast (both geometric and color contrasts).

We evaluate the performance of the tested metrics, including $CMDM_{rec}$, on the whole dataset (80 stimuli). The results are reported in Table 1.1.

Table 1.1 – Performance comparison of different metrics on the whole dataset.

	$PLCC$	$SROCC$	AUC_{DS}	AUC_{BW}
$\mathrm{CMDM_{rec}}$	**0.913**	**0.9**	**0.782**	**0.968**
CMDM_Geo	0.501	0.437	0.604	0.749
CMDM_Col	0.745	0.746	0.732	0.893
D_LAB	0.55	0.603	0.651	0.805
SSIM	0.797	0.799	0.716	0.912
HDR-VDP2	0.853	0.84	0.703	0.944
iCID	0.825	0.83	0.747	0.924

CMDM performs notably better than the others in terms of correlations. Moreover, the AUC values reflect its good classification abilities in both Different vs. Similar and Better vs. Worse analyses. This shows the good robustness of our metric: it is able to differentiate and rank stimuli from different sources and different distortions.

1.3.4 Validation on a dataset of textured 3D meshes

To evaluate the robustness of our recommended metric (eq. 1.13) and to verify that it did not just learn the distortions that are specific to our dataset, we tested $CMDM_{rec}$ on a new dataset: the LIRIS Textured Mesh Database [4]. We selected a subset containing a source model (a dwarf statue) corrupted by 36 mixed distortions (combination of geometry and texture distortions). Before applying our metric, we transferred the texture color information into vertex colors. The results are summarized in Table 1.2. We include results of the IQMs presented previously, as well as the results obtained by Guo et al. [4] for different metrics either applied on rendered videos of the stimuli or directly applied on textured meshes.

Table 1.2 – Performance comparison of different metrics on a new dataset. For metrics marked with a *, the values are reprinted from [4].

	PLCC	SROCC
$\mathrm{CMDM_{rec}}$	**0.862**	**0.872**
SSIM	0.624	0.657
HDR-VDP2	0.83	0.844
iCID	0.502	0.552
Video-DCT*	0.32	0.50
Video-PSNR*	0.33	0.58
Video-MS-SSIM*	0.67	0.66
FQM*	0.64	0.66
CM_1*	0.74	0.77
CM_2*	0.80	0.85

Our metric provides the best results, although it was trained on a different dataset presenting different sources and different distortions and even a different color representation.It outperforms CM_2, which represents the state-of-the-art of textured mesh quality assessment, and which was learned on similar data. This metric is a global combination of mesh and texture distortion measures (*MSDM2* and *MS-SSIM*, respectively). This tends to validate the fact that operating fully on the mesh domain (like our metric) ensures a better performance than combining errors computed on different domains (i.e., mesh and texture image). These results also confirm the great robustness of our metric compared to IQMs.

1.4 Conclusion and Future work

We developed a perceptually-validated full-reference metric *CMDM* for evaluating the quality of colored 3D meshes. To achieve this, we adapted a set of perceptually-relevant curvature-based and color-based features. We further show how to select an optimal subset of features and use them to train the metric (LOOCV tests using a ground truth dataset). Extensive evaluation shows that *CMDM* provides good results and good stability in terms of correlations and classification abilities. It also demonstrates a good robustness: *CMDM* is able to differentiate and rank stimuli from different sources and different distortions, unlike IQMs which perform very well when assessing the quality of different versions of a single source, but are less accurate when ranking distortions applied on different sources. Last but not least, we demonstrate that our metric can also be used for textured meshes.
The metric code will be made publicly available online[2] on the MEsh Processing Platform (MEPP).

As future work, we would also like to produce a huge subject-rated database of 3D models, in order to be able to envisage the creation of end-to-end deep-learning approaches.

Acknowledgments

This work was supported by French National Research Agency as part of ANR-PISCo project (ANR-17-CE33-0005).

[2] https://github.com/MEPP-team/MEPP2

User Guide

We present here the user manual for CMDM, an objective visual quality metric for 3D meshes with diffuse color, publicly available on the MEsh Processing Platform (MEPP2)[3] .

CMDM is the first metric for quality assessment of meshes with diffuse color information represented by vertex colors (no texture mapping), which works entirely on the mesh domain, at vertex level. It is a full-reference metric, based on an optimally-weighted linear combination of geometry-based and color-based features.

For more detail about this work, please refer to our paper: Y. Nehmé, F. Dupont, J.P. Farrugia, P. Le Callet, and G. Lavoué, "Visual Quality of 3D Meshes with Diffuse Colors in Virtual Reality: Subjective and Objective Evaluation", IEEE Transactions on Visualization ans Computer Graphics, 2020.

Here are the steps to follow to use CMDM.
- Install MEPP2 using the following documentation: `https://projet.liris.cnrs.fr/mepp/doc/nightly/_install_wrapper_page.html`

- To use the filters implemented in MEPP2 (CMDM, MSDM2, JND ...), there are 2 options:

1. **Using the GUI**: After completing the installation of MEPP2, run the GUI. Open the reference and the distorted models in the "space" mode (the 2 models are loaded into the same scene). Choose CGAL Surface Mesh as a data structure of the models.
   Afterward, go to the "Plugin filters" menu and select the CMDM plugin. A small window appears containing different configuration options:
   (a) *1 to 2*: $CMDM_{M_{dist} \to M_{ref}}$ evaluates the distortion of the distorted model regarding the reference model.
   (b) *2 to 1*: $CMDM_{M_{ref} \to M_{dist}}$ evaluates the distortion of the refrence model regarding the distorted model.
   (c) *Symmetric*: $\frac{CMDM_{M_{dist} \to M_{ref}} + CMDM_{M_{ref} \to M_{dist}}}{2}$ used to strengthen the robustness of CMDM and to obtain a symmetric measure.
   (d) *Scales*: determine the number of scales used to compute CMDM (since CMDM is a multiscale metric). The smallest scale is $0.003BB$, where BB is the maximum length of the bounding box of the stimulus. Each scale is greater than the previous/smaller scale of 0.0015 ($scale_2 = sacle_1 + 0.0015$).
   (e) *Use CMDM as color map*: display local distortions (at vertex level) as a color map.

2. **Using the CMDM executable**: There is an executable of CMDM, which considers a symmetrical measure of the distortion and three scales ($h_i \in \{0.003BB, 0.0045BB, 0.006BB\}$, where BB is the maximum length of the bounding box of the stimulus).
   example_ CMDM.exe path_ to_ the_ reference_ mesh path_ to_ the_ degraded_ mesh

[3] https://github.com/MEPP-team/MEPP2

Bibliography

[1] P. Alliez, S. Tayeb, and C. Wormser. Aabb tree.cgal 3.5 edition. 2009.

[2] Pierre Alliez, David Cohen-Steiner, Olivier Devillers, Bruno Lévy, and Mathieu Desbrun. Anisotropic polygonal remeshing. *ACM Trans. Graph.*, 22(3):485–493, July 2003.

[3] Jinjiang Guo, Vincent Vidal, Irene Cheng, Anup Basu, Atilla Baskurt, and Guillaume Lavoue. Subjective and objective visual quality assessment of textured 3D meshes. *ACM Transactions on Applied Perception*, 14(2), 2016.

[4] L. Krasula, K. Fliegel, P. Le Callet, and M. Klíma. On the accuracy of objective image and video quality models: New methodology for performance evaluation. In *2016 Eighth International Conference on Quality of Multimedia Experience (QoMEX)*, pages 1–6, 2016.

[5] Guillaume Lavoué. A Multiscale Metric for 3D Mesh Visual Quality Assessment. *Computer Graphics Forum*, 30(5):1427–1437, 2011.

[6] Guillaume Lavoue, Mohamed Chaker Larabi, and Libor Vasa. On the Efficiency of Image Metrics for Evaluating the Visual Quality of 3D Models. *IEEE Transactions on Visualization and Computer Graphics*, 22(8):1987–1999, 2016.

[7] I. Lissner and P. Urban. Toward a unified color space for perception-based image processing. *IEEE Transactions on Image Processing*, 21(3):1153–1168, 2012.

[8] Ingmar Lissner, Jens Preiss, Philipp Urban, Matthias Scheller Lichtenauer, and Peter Zolliker. Image-difference prediction: From grayscale to color. *IEEE Transactions on Image Processing*, 22(2):435–446, 2013.

[9] Yiming Liu, Jue Wang, Sunghyun Cho, Adam Finkelstein, and Szymon Rusinkiewicz. A no-reference metric for evaluating the quality of motion deblurring. *ACM Transactions on Graphics*, 32(6):171–175, 2013.

[10] Rafał Mantiuk, Kil Joong Kim, Allan G. Rempel, and Wolfgang Heidrich. Hdr-vdp-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM Trans. Graph.*, 30(4), July 2011.

[11] Gabriel Meynet, Yana Nehmé, Julie Digne, and Guillaume Lavoué. PCQM: A Full-Reference Quality Metric for Colored 3D Point Clouds. In *International Conference on Quality of Multimedia Experience*, 2020.

[12] Yana Nehmé, Jean Philippe Farrugia, Florent Dupont, Patrick Le Callet, and Guillaume Lavoué. Comparison of subjective methods, with and without explicit reference, for quality assessment of 3D graphics. In *ACM Conference on Applied Perception*, 2019.

[13] Jens Preiss, Felipe Fernandes, and Philipp Urban. Color-image quality assessment: From prediction to optimization. *IEEE Transactions on Image Processing*, 23(3):1366–1378, 2014.

[14] Dihong Tian and G. AlRegib. Batex3: Bit allocation for progressive transmission of textured 3-d models. *IEEE Transactions on Circuits and Systems for Video Technology*, 18(1):23–35, 2008.

[15] Libor Váša and Jan Rus. Dihedral Angle Mesh Error: a fast perception correlated distortion measure for fixed connectivity triangle meshes. *Computer Graphics Forum*, 31(5), 2012.

[16] Kai Wang, Fakhri Torkhani, and Annick Montanvert. A Fast Roughness-Based Approach to the Assessment of 3D Mesh Visual Quality. *Computers & Graphics*, 2012.

[17] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.