Computers & Graphics 36 (2012) 808–818

Contents lists available at SciVerse ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

Technical Section

A fast roughness-based approach to the assessment of 3D mesh visual quality

Kai Wang*, Fakhri Torkhani, Annick Montanvert

GIPSA-Lab, CNRS UMR5216, Saint Martin d'Hères F-38402, France

ARTICLE INFO

Article history:
Received 6 February 2012
Received in revised form
22 June 2012
Accepted 22 June 2012
Available online 2 July 2012

Keywords:
Visual quality assessment
Triangle mesh
Human visual system
Objective perceptual metric
Surface roughness

ABSTRACT

We propose in this paper a new objective metric for the visual quality assessment of 3D meshes. The metric can predict the extent of the visual difference between a reference mesh, which is considered to be of perfect quality, and a distorted version. The proposed metric is based on a mesh local roughness measure derived from Gaussian curvature. The perceptual distance between two meshes is computed as the difference between the normalized surface integrals of the local roughness measure. Experimental results from three subjective databases and comparisons with the state of the art demonstrate the efficacy of the proposed metric in terms of the execution time and the correlation with subjective scores. Finally, we show a simple application of the metric in which it is used to automatically determine the optimum quantization level of mesh vertex coordinates.

© 2012 Elsevier Ltd. All rights reserved.

1. Introduction

Three-dimensional (3D) surface models, representing a character or animal, a mechanical object or a human organ, are commonly used in diverse applications such as digital entertainment, computer-aided design and medical imaging. 3D models are commonly represented by polygonal meshes [1], which constitute a piecewise linear approximation of the underlying continuous surface. With the rapid growth of network bandwidth and the increasing capability of personal computers (PCs), it is now common to see 3D meshes transmitted on the Internet and manipulated and visualized on ordinary PCs. The popularity of the mesh model has made it one of the three main formats for object representation in 3DTV that are currently being discussed by international standardization groups [2].

In practical applications, a mesh is usually subject to different kinds of lossy operations, which introduce distortions and modifications to the original model. For example, we may simplify (i.e., reduce the number of vertices) and compress a complex mesh before transmitting it to a cell phone that has limited processing and visualization capacities; a mesh model may be contaminated by random noise during its transmission in a noisy channel; or a digital watermark may be embedded into a 3D mesh for the purpose of copyright protection.

In all such scenarios, it is important to evaluate how much visual distortion has been introduced into the original model by a particular operation and whether this distortion deteriorates the quality of service (QoS). The visual distortion can be measured either *subjectively* (i.e., by a group of human observers) or *objectively* (i.e., by an automatic metric running on a computer). Although the subjective evaluation appears to be more reliable, under most circumstances, it is a solution that is too expensive, laborious and time-consuming to be adopted. It is thus necessary to devise objective and automatic metrics that would correlate well with the subjective assessment.

In the literature, the difference between two meshes is in most cases evaluated by simple geometric measures, such as the Hausdorff distance or the mean squared error (MSE) [3,4]. However, as with the MSE and the PSNR (peak signal-to-noise ratio) for digital images [5], it has been demonstrated that these simple geometric measures do not correlate well with human perception [6,7] (see Fig. 1 for an example). Therefore, the development of a perceptually based quality metric for 3D meshes has recently attracted the attention of many researchers. It is believed that effective perceptually based metrics that can accurately predict the subjective assessment of the visual degradation of a 3D mesh will replace the aforementioned geometric measures in the future in a wide range of geometry processing applications. In general, an objective mesh visual quality (MVQ) metric can be used in the following ways:

- It can be used for the *evaluation* and *benchmarking* of mesh processing algorithms, especially simplification [8,9], watermarking [10,11] and compression [12,13].
- It can be used for the *parameter optimization* of many mesh processing algorithms, which usually aim to achieve an optimum trade-off between a specific performance indicator (e.g., the vertex reduction ratio, watermark robustness or bit rate) and the induced visual degradation.

* Corresponding author. Tel.: +33 476826255; fax: +33 476574790.
E-mail addresses: kai.wang@gipsa-lab.grenoble-inp.fr (K. Wang), fakhri.torkhani@gipsa-lab.grenoble-inp.fr (F. Torkhani), annick.montanvert@gipsa-lab.grenoble-inp.fr (A. Montanvert).

0097-8493/$ - see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.cag.2012.06.004

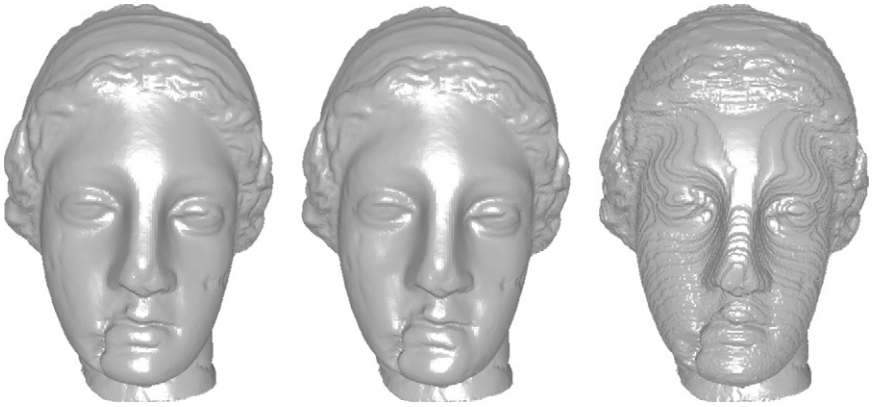

Fig. 1. On the left is the original Venus model, in the middle is the model watermarked using the method of Wang et al. [14], and on the right is the model watermarked using the method of Cho et al. [15]. Compared to the original mesh, the two watermarked models have exactly the same maximum root mean squared error (MRMS, a widely used geometric distance [4]) value of 1.52×10^{-3}, but their visual degradation is obviously very different. The proposed fast mesh perceptual distance (FMPD) metric yields the correct results: the perceptual distances are 0.01 and 0.70 for the two watermarked models. The execution time is less than 3 s on an ordinary laptop for this mesh of 100 K vertices.

- The development of such metrics, together with a better understanding of the behavior of the human visual system (HVS) while observing 3D mesh models, may stimulate the development of fundamentally new *perception-oriented* mesh processing algorithms.

Because the relevant research is still in its early stage, most existing MVQ metrics have limitations, such as being computationally slow, imposing constraints on the input meshes, or not being invariant to similarity transformations (i.e., scaling, translation and rotation). Additionally, as noted by Lavoué and Corsini [6], the distance or similarity values produced by most objective MVQ metrics do not correlate well enough with subjective scores. In this paper, we propose a new and conceptually very simple objective metric for 3D mesh visual quality assessment. The proposed metric, called the fast mesh perceptual distance (FMPD), has the following features:

- The FMPD metric can compare two triangle meshes of different connectivities (i.e., different adjacency relationships between the vertices), and the metric is invariant to mesh similarity transformations.
- The FMPD metric has a very fast execution time (<3 s to compare two meshes of 100 K vertices).
- The FMPD metric achieves performance comparable to (or even slightly better than) the most recent metric described in [16] in terms of the correlation between the objective perceptual distances and the subjective scores on three common databases.

Compared to most existing methods, the proposed FMPD metric is faster and has fewer constraints on the input meshes mainly because of two reasons: FMPD is a global-roughness-based metric, and it is based on geometric quantities that are extremely simple and fast to compute. FMPD has a high correlation with subjective scores because two important HVS phenomena, the visual masking effect and the psychometric saturation effect (which will be described in detail in Section 3), have been taken into account during its design. Moreover, we will describe a simple application of the FMPD in which the metric is used to automatically determine the perceptually optimum quantization level of mesh vertex coordinates. The source code for the proposed metric is shared on the Internet.[1]

[1] http://www.gipsa-lab.inpg.fr/~kai.wang/publications_en.html.

The remainder of this paper is organized as follows. Section 2 briefly reviews the related work. Section 3 describes the pipeline of the proposed metric in detail. Section 4 presents the experimental results, including comparisons with existing metrics and a simple application. We conclude in Section 5.

2. Related work and motivation

Visual quality assessment: With the increasing popularity of various forms of digital multimedia content in our daily life, such as images, audio and video clips, and 3D models, it is important to develop objective visual quality metrics for them. As an example, such metrics would be essential in the QoS control of network-based multimedia services, such as VoD (video-on-demand) systems, online video games, and virtual visits. Many metrics have been proposed during the last two decades, especially for images [17] and video [18]. As presented by Wang and Bovik [17], we can usually classify visual quality metrics using two criteria. First, metrics can be classified according to the amount of information available about a reference source that is assumed to be of perfect quality as *full reference* (i.e., the reference content is completely available), *no reference* (i.e., no information is available about the reference content) or *reduced reference* (i.e., part of the information is available). Second, metrics can be classified according to their design philosophy as either *bottom-up* or *top-down*. Bottom-up approaches attempt to simulate the functionality of different components of the human visual system and then develop a metric based on this component-level simulation. Top-down approaches treat the HVS as a black box and make general system-level assumptions about it, and the objective is to devise a metric that mimics the input–output characteristics of the HVS. However, the distinction between these two types of approaches is not strict, and researchers often combine the methodologies from both bottom-up and top-down approaches to develop hybrid metrics.

Image-based MVQ assessment: In contrast to the fruitful advances in the fields of image and video quality assessment, there exist relatively few perceptually based metrics designed specifically for 3D meshes. Indeed, in many graphics-related applications, 2D image quality metrics are used to evaluate the distortions introduced in 3D meshes. The evaluation is performed on one or several rendered 2D images of the mesh model. The obtained 2D perceptual distortion, and the related HVS features, has been successfully used to guide the procedures of mesh

simplification [19,20], rendering [21] and level-of-detail control [22,23], and to evaluate the distortion introduced by mesh watermarking algorithms [24]. However, using 2D image quality metrics to assess mesh visual quality in a general context would be difficult and problematic because of two facts: (1) the automatic selection of the optimum viewpoints from which to produce the rendered images is a very difficult problem, and (2) as shown by Rogowitz and Rushmeier [25] through a series of experiments, it appears that 2D image metrics are not completely adequate for assessing the visual quality of 3D models.

Model-based MVQ assessment: Aware of the limitations of the above *image-based* approaches, researchers have attempted to develop quality metrics directly from the 3D shape of meshes. Such methods are often called *model-based* methods [6], and their most significant advantage is that they do not depend on viewpoint selection. Because the metric proposed in this paper follows the model-based principle, we will briefly review existing model-based metrics.

To our knowledge, the first model-based metric was proposed by Karni and Gotsman [26] in 2000 for the evaluation of their spectral mesh compression method. They measured the distance between two meshes as a weighted sum of the vertex root mean squared error and the vertex Laplacian coordinate error. It is argued by the geometry processing community that the Laplacian coordinates are related to the surface normal that is used by most mesh rendering algorithms, and therefore the Laplacian coordinates are relevant to mesh visual quality. Karni and Gotsman's metric was later enhanced by Sorkine et al. [27] by giving a greater weight to the Laplacian coordinate error. Pan et al. [28] experimentally studied the relationship between mesh visual quality and a number of factors including the geometric resolution (i.e., the number of vertices) of the model. Drelie Gelasca, Corsini and colleagues [29,30] made the assumption that the perceived quality is related to the roughness of the mesh surface. They proposed two perceptual metrics using two definitions of the mesh roughness, one definition based on the dihedral angle between neighboring facets and the other definition based on the difference between the input mesh and a carefully smoothed version. Lavoué et al. proposed a metric called the mesh structural distortion measure (MSDM) [31]. Their basic idea was to extend a well-known image quality metric, the SSIM (structural similarity) index [32], to 3D meshes by replacing the pixel intensities in the SSIM index with the mesh mean curvature. Bian et al. [33] proposed a metric based on the theory of surface strain energy. They considered a mesh as an elastic object and related the visual degradation to the amount of strain energy required to produce the corresponding mesh deformation.

Most of the above metrics were evaluated and compared in a recent study [6] using two subjective databases, which showed that the existing metrics do not correlate well with subjective scores. The Pearson correlations of the objective scores produced by the best metric, MSDM, and the mean opinion scores (MOS) provided by human observers were approximately 60%, which is not a satisfactory value. The authors of [6] also noted that some metrics have constraints on the type of input meshes. For example, Karni and Gotsman's [26] and Bian et al.'s [33] metrics, as well as the MSDM [31], assume that the meshes before and after distortion share the same connectivity, and the roughness-based metrics in [29,30] assume a uniform sampling over the mesh surface.

Recently, Lavoué proposed an improved version of the MSDM, called the MSDM2 [16]. There were two main improvements compared to the original version: (1) the MSDM2 can compare triangle meshes with different connectivities using a vertex correspondence preprocessing step, and (2) the visual difference is now evaluated in a multiscale manner so that the MSDM2 correlates better with subjective scores. The MSDM2 produces Pearson correlations of 66.2%, 76.2% and 79.6% with three subjective databases, which is a significant improvement from the MSDM.

Motivation: The MSDM2 measure nevertheless has some disadvantages. First, the execution time is relatively high; for example, it takes approximately 100 s to compare two meshes of 100 K vertices. Second, the correlation values still must be improved so that a perceptually based metric can be reliably used in practical applications. These observations and the relevant state of the art motivated our work. Our objective was to propose an open-source metric with low time complexity, high correlation with subjective scores and no or very few constraints on the input meshes (including the capability to compare triangle meshes with different connectivities and invariance to similarity transformations). Our metric follows a design philosophy that is similar to the metrics of [29,30] in the sense that they are all global-roughness-based metrics. However, we use a quite different definition of the surface roughness that appears to be more relevant to human visual perception, in particular by explicitly taking into account the visual masking effect. Our metric outperforms the metrics in [29,30] by producing much higher correlation values with subjective scores, as will be shown in Section 4.

3. The proposed metric

3.1. Overview of the pipeline

Fig. 2 depicts the block diagram of the proposed metric. First, the local roughness is defined at each vertex of the reference mesh and the deformed mesh based on Gaussian curvature. We then apply a careful modulation of the local roughness to account for the visual masking effect and the psychometric saturation effect, which are both important HVS features. Next, we compute the global roughness using the normalized surface integrals of the local roughness on both meshes. Finally, the perceptual distance between the two meshes is evaluated as the difference of the two surface integrals. In the following, we present the technical details and the motivation for each step of the proposed pipeline.

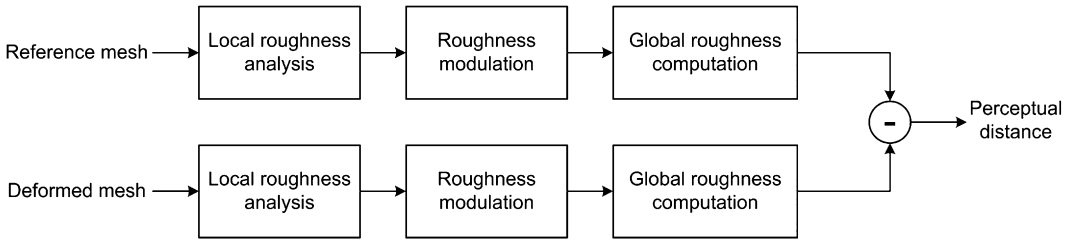

Fig. 2. Block diagram illustrating the pipeline of the proposed metric.

3.2. Local roughness

As in [29,30], we make the assumption that the perceived quality of 3D polygonal meshes is related to the modification of the local and global roughness of the mesh surface. To construct the pipeline shown in Fig. 2, we need to first derive an adequate definition of the mesh local roughness that is consistent with human visual perception. To keep the proposed MVQ metric simple and fast, in this section we identify simple differential geometric quantities that are relevant to the mesh visual quality and use the quantities to define the local roughness.

As mentioned in Section 2, it is believed that there is a link between mesh curvature and the perceived mesh quality. Therefore, curvature has played an important role in many mesh processing and analysis algorithms, such as saliency analysis [34], visual quality assessment [31,16], smoothing and fairing [35], and deformation [36]. Consequently, it is reasonable to define the mesh local roughness based on Gaussian curvature. First, for each vertex v_i in the reference mesh $\mathcal{M}_r$, the discrete Gaussian curvature is defined as follows [37]:

$$GC_i = \left| 2\pi - \sum_{j \in \mathcal{N}_i^{(F)}} \alpha_j \right|, \tag{1}$$

where $\mathcal{N}_i^{(F)}$ is the set of all the neighboring facets of v_i, and α_j is the angle in facet j that is incident to v_i (see Fig. 3 left). Intuitively, discrete Gaussian curvature describes how much the local patch deviates from a planar surface.

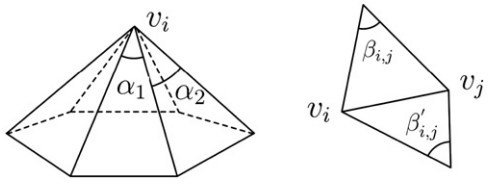

Fig. 3. Geometric quantities used for computing the discrete Gaussian curvature (left) and the discrete Laplacian operator (right).

The local roughness at v_i is measured as the Laplacian of the discrete Gaussian curvature. To this end, we first compute the mesh Laplacian matrix as

$$\begin{cases} D_{i,j} = \dfrac{\cot(\beta_{i,j}) + \cot(\beta'_{i,j})}{2} & \text{for } j \in \mathcal{N}_i^{(V)}, \\ D_{i,i} = -\sum\limits_j D_{i,j}, \end{cases} \tag{2}$$

where $\mathcal{N}_i^{(V)}$ is the set of all the neighboring vertices of v_i, and $\beta_{i,j}$ and $\beta'_{i,j}$ are the two angles opposite to the edge that connects v_i and v_j (see Fig. 3 right). The above sparse matrix, representing the discrete Laplacian operator on 2-manifold triangle meshes, can be derived using either finite element modeling theory [38] or discrete exterior calculus [39]. This matrix yields better performance than other mesh Laplacians [40] in applications such as smoothing [39] and watermarking [41,42]. Next, the local roughness LR_i at v_i is defined as

$$LR_i = \left| GC_i - \frac{\sum_{j \in \mathcal{N}_i^{(V)}} D_{i,j} \cdot GC_j}{\sum_{j \in \mathcal{N}_i^{(V)}} D_{i,j}} \right| = \left| GC_i + \frac{\sum_{j \in \mathcal{N}_i^{(V)}} D_{i,j} . GC_j}{D_{i,i}} \right|. \tag{3}$$

The roughness is evaluated as a weighted difference between GC_i and the Gaussian curvatures of the neighbors, where the weights are determined according to the entries in the Laplacian matrix of Eq. (2). The same computation is performed at each vertex v'_i on the deformed mesh $\mathcal{M}_d$. Fig. 4 illustrates the roughness maps of several meshes that we will use in our experiments in Section 4. In general, this simple roughness measure is quite consistent with human perception.

The first motivation for using the Laplacian of the Gaussian curvature, instead of the Gaussian curvature itself, is that computing the Laplacian allows the roughness evaluation to be performed in a more contextual manner; i.e., more vertices and facets are involved in the evaluation. Moreover, the variation of the curvature is conjectured to be related to surface fairness [43], i.e., the aesthetic measure of "well-shapedness," so it is reasonable to assume that the Laplacian of

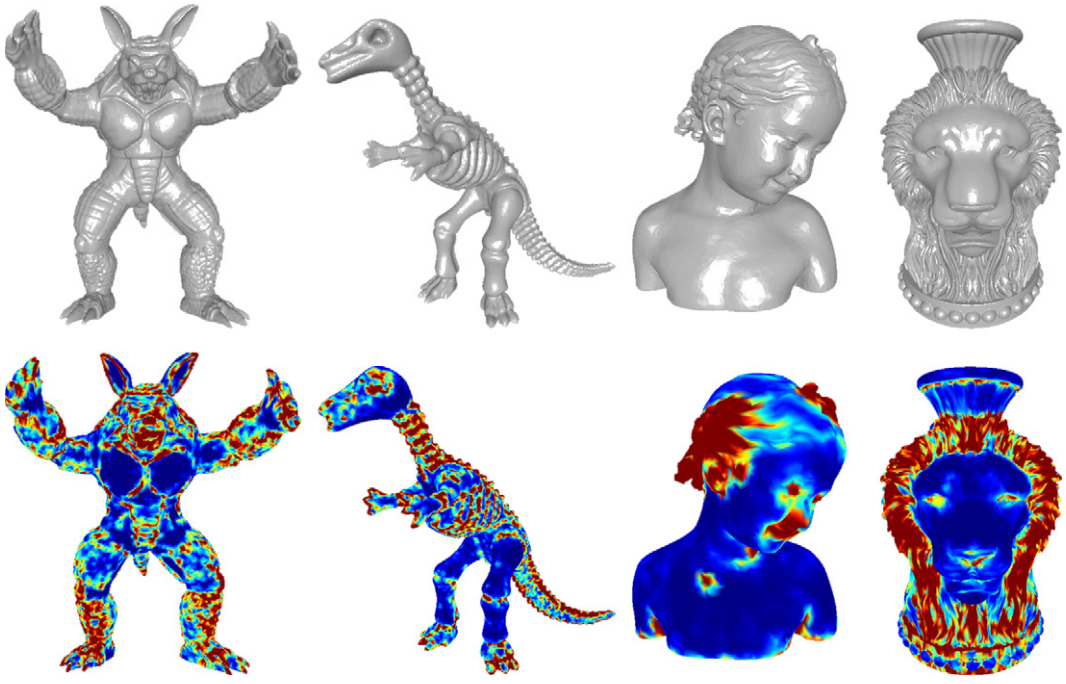

Fig. 4. Roughness maps of sample meshes. In the top row are flat-shaded renderings; in the bottom row are roughness maps in which warmer colors represent larger values.

the Gaussian curvature has a strong link to the surface visual quality. Finally, in general using the variation of the curvature instead of the curvature itself we can avoid mistakenly classifying a high-curvature smooth surface as a rough region. We choose Gaussian curvature, rather than mean curvature, in our algorithm mainly for two reasons. First, compared to the mesh mean curvature, the discrete Gaussian curvature is extremely simple to compute, and the Gaussian curvature is intrinsically invariant to scaling. Second, in our experiments, we have used Gaussian curvature, the Laplacian of the Gaussian curvature, the mean curvature and the Laplacian of the mean curvature as local roughness descriptors, while keeping the rest of the pipeline of the MVQ metric unchanged. We have found that the Laplacian of the Gaussian curvature yields the best results (cf. Section 4.4). In the future, we plan to design and conduct psychovisual experiments to study the relationship between these differential geometric quantities and human visual perception.

3.3. Local roughness modulation

In this step, the local roughness obtained in the previous step is modulated to account for the *visual masking* effect and the so-called *psychometric saturation* effect. The visual masking effect [44] means that the existence of one visual signal may hide or reduce the visibility of another signals. In the case of MVQ assessment, the visual masking effect mainly implies that a local surface modification, e.g., due to the addition of noise or to vertex coordinate quantization, is more visible in a smooth region than in a rough region. This effect is illustrated in Fig. 5. For example, the Bimba model in Fig. 5(b) is contaminated by a uniform random noise whose amplitude is the same throughout the surface; however, the noise is much more visible in smooth regions, such as the chest, than in rough regions, such as the hair. The psychometric saturation effect means that when asked to assess the intensity of a stimulus, humans tend to provide a constant response at extreme quantities that are beyond or below a certain threshold. Typically, human observers do not distinguish between very small (or large) stimuli of slightly different intensities, and observers will assign the same subjective score for slightly different but very bad (or good) qualities (e.g., as mentioned in [45,46] in the context of image quality assessment).

To incorporate and model these two effects in our metric, our solution is to carefully modulate the local roughness. First, the range of the roughness is limited both above and below so that the roughness lies in the interval $[Th_l, Th_h]$. The main objective of this thresholding is to account for the psychometric saturation effect. The thresholding also makes the results more robust by avoiding instability near zero. We then modulate the thresholded roughness using a power function as follows:

$$LRM_i = f(LR_i) = (LR_i)^a - (Th_l)^a, \tag{4}$$

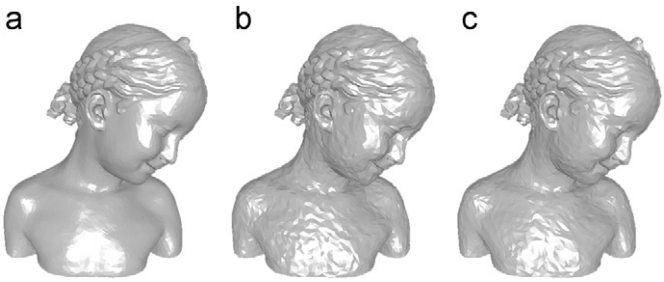

Fig. 5. This figure illustrates the visual masking effect on 3D meshes: (a) the original Bimba model; (b) the modified model after the addition of uniform random noise; (c) the model after uniform vertex coordinate quantization. The visual degradation is much more visible in smooth regions than in rough regions.

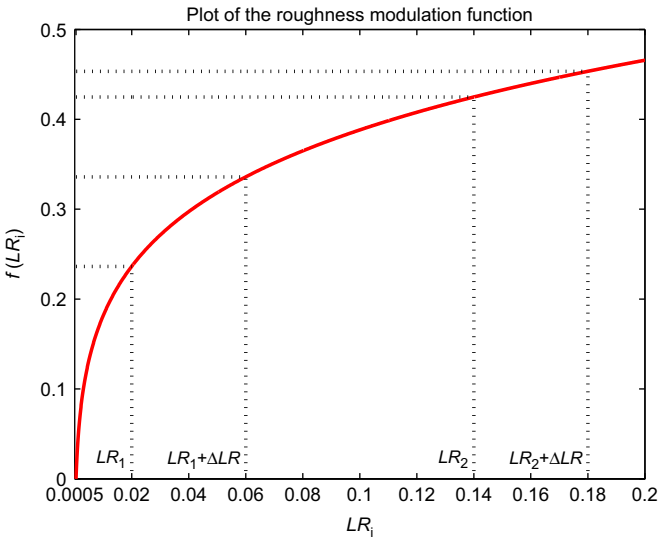

Fig. 6. Modulation of the local roughness by a power function.

where $0 < a < 1$ is a parameter that controls the shape of the power function and is fixed as a constant for all meshes. This power modulation function is illustrated in Fig. 6. It can be observed that $f(LR_i)$ is a monotonically increasing function with a monotonically decreasing first derivative. We adopt this function to capture the visual masking effect. The basic idea is to induce a large visual difference result where a smooth region becomes a rough region (or vice versa) after modification but to induce a small difference in situations where an originally rough region still remains rough after modification. For example, suppose that in the reference mesh $\mathcal{M}_r$ there are two vertices v_1 and v_2 with local roughness values LR_1 and LR_2, respectively, where $LR_1 < LR_2$. Assume that after a deformation, the two vertices have undergone the same roughness increase ΔLR; i.e., the roughnesses become $LR'_1 = LR_1 + \Delta LR$ and $LR'_2 = LR_2 + \Delta LR$ for v_1 and v_2, respectively. Using the power function in Eq. (4), it can be ensured that after applying the modulation we have $f(LR'_1) - f(LR_1) > f(LR'_2) - f(LR_2)$ (see Fig. 6). Therefore, we can guarantee that the same roughness modification ΔLR will induce a greater change in the modulated roughness (and thus a greater visual difference) in a smooth region than in a rough region. Experimentally, other functions (e.g., the logarithmic function and the right part of the cumulative Gaussian function) could also be used for the roughness modulation, and similar overall performance could be attained. We choose the power function because of its simplicity and its common use in visual perception research [47,48].

The above modulation accounts for the visual masking effect in a general fashion that is common for all meshes. However, in practice this modulation alone is not sufficient to capture the masking effect well, and it is necessary to carry out a second modulation that is adapted to the characteristics of the two meshes under comparison. Experimentally, we have to further reduce the influence of the values that are greater than the mesh's average roughness. For this purpose, first, the average roughness (before any modulation) $\overline{LR}$ of the reference mesh $\mathcal{M}_r$ is defined as

$$\overline{LR} = \frac{\sum_i LR_i \cdot s_i}{\sum_i s_i}, \tag{5}$$

where s_i is one third of the total area of the incident facets of v_i, and $\sum_i s_i$ is the total area of the triangular manifold mesh $\mathcal{M}_r$. The average roughness $\overline{LR}$ is actually a surface-weighted average of the mesh's local roughness. The same computation is performed on the deformed mesh $\mathcal{M}_d$ to obtain its average roughness $\overline{LR}'$. $\overline{LR}$ and $\overline{LR}'$

are then modulated by the function $f(\cdot)$ defined in Eq. (4), and the modulated values are denoted by $\overline{LRM}=f(\overline{LR})$ and $\overline{LRM}'=f(\overline{LR}')$. We then set a threshold $Th_{LRM}=\min\{\overline{LRM},\overline{LRM}'\}$ and further reduce any modulated roughness that is greater than this threshold according to the following equation:

$$LRF_i = Th_{LRM} + b(LRM_i - Th_{LRM}) \quad \text{for } LRM_i > Th_{LRM}, \tag{6}$$

where LRF_i is the final local roughness at v_i and $0<b<1$ is a parameter that controls the magnitude of the reduction. After this second modulation, the local modifications in rough regions of the model (i.e., where $LRM_i > Th_{LRM}$) will induce an even smaller visual difference, reflecting the masking effect.

The same modulation is performed on the deformed mesh $\mathcal{M}_d$, and the final local roughness of vertex v_i' is denoted by LRF_i'. In our experiments, the parameter values are chosen as follows: $Th_l=5.0\times10^{-4}$, $Th_h=\max\{0.20,5Th_{LRM}\}$, $a=0.15$ and $b=0.50$. These settings were used for all the models that we tested, and they yielded consistently good results, as will be shown in Section 4.

3.4. Global roughness and perceptual distance

The global roughness GR of the reference mesh $\mathcal{M}_r$ is evaluated as a normalized surface integral of the local roughness

$$GR = \frac{\sum_i LRF_i \cdot s_i}{\sum_i s_i}, \tag{7}$$

where s_i is the same as in Eq. (5). This surface-based weighting of the local roughness ensures that the metric is to some extent robust to the variation of the vertex sampling density over the mesh surface.

Finally, the perceptual distance $\mathrm{FMPD}_{\mathcal{M}_r,\mathcal{M}_d}$ between the reference and the deformed meshes is defined as follows:

$$\mathrm{FMPD}_{\mathcal{M}_r,\mathcal{M}_d} = c\left|GR - GR'\right|, \tag{8}$$

where GR and GR' are the global roughness of $\mathcal{M}_r$ and $\mathcal{M}_d$, respectively, and $c=8.0$ is a scaling factor that brings the perceptual distance into the $[0,1]$ interval (values greater than 1 are simply thresholded to be 1).

The proposed FMPD metric is symmetric (i.e., the distance from $\mathcal{M}_r$ to $\mathcal{M}_d$ is the same as the distance from $\mathcal{M}_d$ to $\mathcal{M}_r$), invariant to similarity transformations and capable of comparing triangle meshes of different connectivities. The time complexity of the metric is a linear function of the total number of vertices in $\mathcal{M}_r$ and $\mathcal{M}_d$.

4. Experimental results

The proposed metric has been implemented in Matlab, and the source code is freely available on the Internet. In the following, we present the experimental results from the application of our metric to three subjective databases, a comparison of the metric with state-of-the-art methods, and a simple practical application of the metric to vertex coordinate quantization.

4.1. Tests and comparisons on subjective databases

The standard index for evaluating the performance of an objective MVQ metric is the correlation between the perceptual distances or similarities produced by the metric and subjective mean opinion scores. There are two commonly used types of correlation: the Pearson linear correlation coefficient (PLCC) and the Spearman rank-order correlation coefficient (SROCC). The PLCC measures the linear dependence between the objective and subjective scores and is generally considered a more effective and more important index than the SROCC. The SROCC measures how well the relationship between the objective and subjective scores can be described by a monotonic function. Only the ranks of the scores are used in the computation of the SROCC, not the actual score values.

As noted in [16], there exist only three publicly available subject-rated databases for the evaluation of MVQ metrics:

- The LIRIS/EPFL general-purpose database[2] [31].
- The LIRIS masking database[3] [49].
- The IEETA simplification database[4] [50].

All three databases contain several reference and deformed meshes with a series of subjective mean opinion scores (MOSs), each of which reflects the extent of the visual difference between a pair of reference and deformed models. The MOS values were obtained in three steps, according to the relevant ITU recommendations [51,52]: (1) the collection of raw data from multiple human observers, (2) the screening of observers to remove outliers, and (3) the normalization of the scores to reduce the deviation among observers.

The LIRIS/EPFL general-purpose database comprises 88 models: four reference meshes and 21 deformed models for each reference mesh. The deformations in the database include the addition of noise and smoothing of different intensities in different spatial regions (i.e., smooth, rough and intermediate regions and the whole surface). It is believed that the included deformations simulate a wide range of common processing techniques on polygonal meshes [31]. Fig. 7 shows several models from this repository with the output of our FMPD metric and the corresponding MOS values (where MOS $\in[0,1]$ reflects the extent of the visual difference: the higher the MOS, the greater the visual difference).

Before computing the PLCC between the FMPD and the MOS, it is recommended that a psychometric curve fitting between the two measures [30] be conducted to partially remove the non-linearity between the two groups of values. In our experiments, we choose to use the cumulative Gaussian psychometric function [53] to perform this fitting

$$g(m,n,R) = \frac{1}{\sqrt{2\pi}}\int_{m+nR}^{+\infty} e^{-t^2/2}\,dt, \tag{9}$$

where m and n are the two parameters that we need to estimate, and R is the objective distance. We derive the values of m and n using the FMPD and MOS values of one group of mesh models from the general-purpose database, the reference Dinosaur model and its 21 deformed versions. After a non-linear least squares fitting based on the Levenberg–Marquardt algorithm, we obtain $m=0.4123$ and $n=-1.981$. Note that the values of m and n for the fitting between the FMPD and MOS data are fixed as constant in all our experiments, even for models from the other two databases. The consistently good results obtained, as presented in the following, demonstrate the efficacy of this curve fitting as well as the stability of the proposed metric. Fig. 8(a) illustrates this curve fitting technique using the Dinosaur models, with each circle symbol representing a pair of FMPD and MOS values. Fig. 8(b) shows the same curve plotted with all the pairs of FMPD and MOS values from the general-purpose database.

Table 1 presents the PLCC and SROCC values of the FMPD from the general-purpose database as well as the values from five other metrics: the Hausdorff distance (HD) [3,4], the root mean squared

[2] http://liris.cnrs.fr/guillaume.lavoue/data/datasets.html.
[3] http://liris.cnrs.fr/guillaume.lavoue/data/datasets.html.
[4] http://www.ieeta.pt/~sss/repository/.

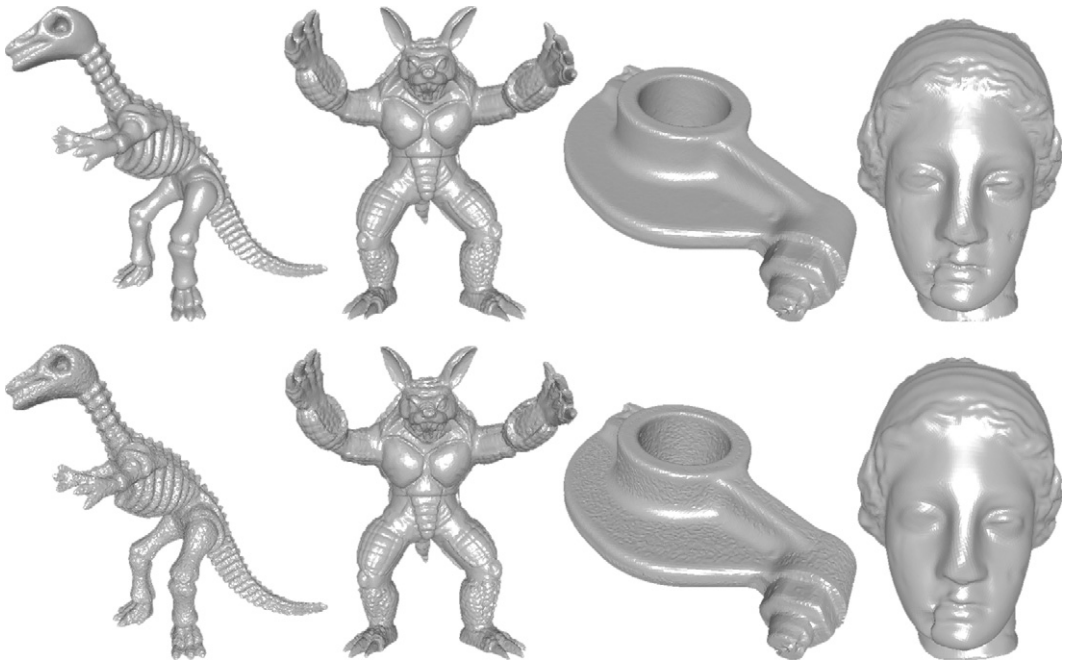

Fig. 7. Some models from the LIRIS/EPFL general-purpose database. In the top row are the four reference models. In the bottom row are four deformed models, from left to right are respectively: Dinosaur with global noise (FMPD=0.82, MOS=0.89), Armadillo with noise in rough regions (FMPD=0.23, MOS=0.34), Rockerarm with noise in smooth regions (FMPD=0.71, MOS=0.74), and Venus after global smoothing (FMPD=0.25, MOS=0.40).

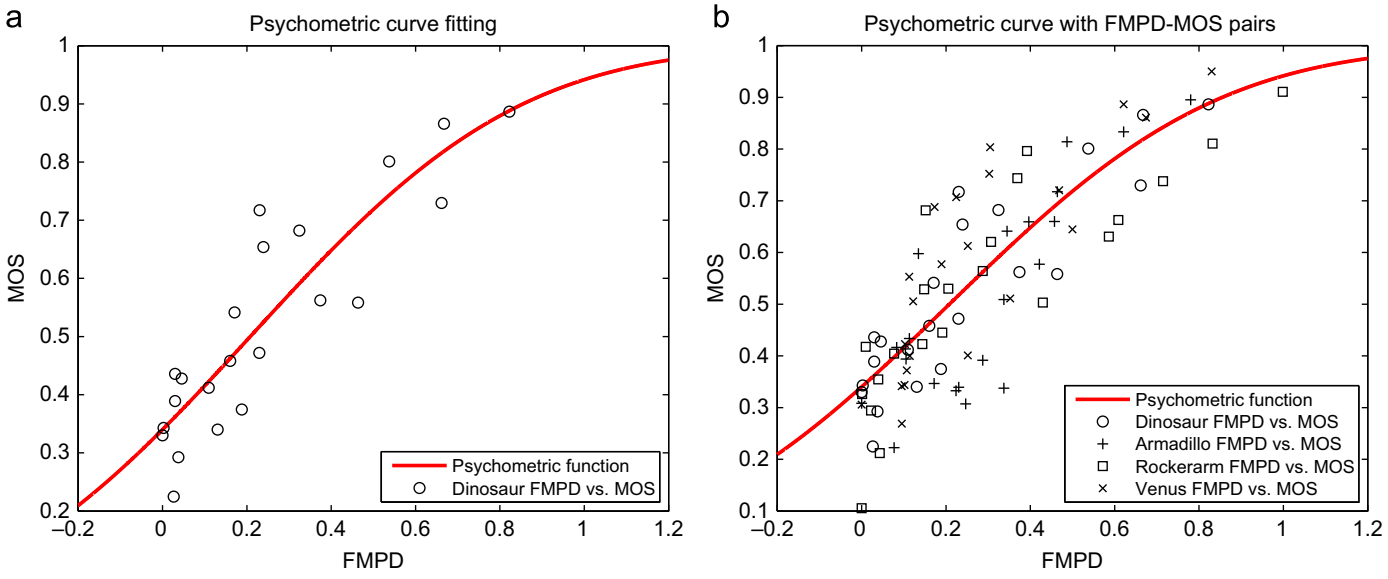

Fig. 8. Psychometric curve fitting between the FMPD and MOS values: (a) the fitting performed using the FMPD and MOS values from the Dinosaur models; (b) the psychometric curve plotted with the FMPD–MOS pairs from all the models in the LIRIS/EPFL general-purpose database.

Table 1
PLCC (r_p) and SROCC (r_s) values (%) of different MVQ metrics on the LIRIS/EPFL general-purpose database.

Metric	Armadillo		Dinosaur		Rockerarm		Venus		All models	
	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s
HD [3,4]	30.2	69.5	22.6	30.9	5.5	18.1	0.8	1.6	1.3	13.8
RMS [3,4]	32.2	62.7	0.0	0.3	3.0	7.3	77.3	**90.1**	7.9	26.8
3DWPM$_1$ [30]	35.7	65.8	35.7	62.7	53.2	87.5	46.6	71.6	38.3	69.3
3DWPM$_2$ [30]	43.1	74.1	19.9	52.4	29.9	37.8	16.4	34.8	24.6	49.0
MSDM2 [16]	72.8	**81.6**	73.5	85.9	76.1	**89.6**	76.5	89.3	66.2	80.4
FMPD	**83.2**	75.4	**88.9**	**89.6**	**84.7**	88.8	**83.9**	87.5	**83.5**	**81.9**

error (RMS) [3,4], the two roughness-based metrics $3DWPM_1$ and $3DWPM_2$ of Corsini et al. [30], and MSDM2 [16]. The values of the state-of-the-art metrics are obtained from [16]. In general, our FMPD metric produced high correlations, in terms of both the PLCC and the SROCC, on this database. In particular, FMPD produces the best PLCC for every model in this repository. Additionally, the PLCC and the SROCC of the FMPD for the whole database (shown in the last two columns of Table 1) are the highest among all the metrics, including MSDM2, the best metric proposed so far (the PLCC values are 83.5% for FMPD vs. 66.2% for MSDM2, and the SROCC values are 81.9% for FMPD vs. 80.4% for MSDM2). The good performance of the FMPD metric can also be confirmed in Fig. 8(b), where most of the FMPD-MOS points in the plot are very close to the psychometric curve.

The LIRIS masking database was designed specifically to test the capacity of an objective MVQ metric to capture the visual masking effect. The database contains 28 models: four reference meshes and six deformed models for each reference mesh. The deformed models were obtained by adding noise of different intensities to either the rough or the smooth regions of the reference model. Fig. 9 shows some models from this database, with the corresponding FMPD and MOS values. Normally, noise of the same magnitude induces much higher visual degradation in smooth regions than in rough regions, as reflected by the FMPD and MOS values. Table 2 provides the results of the different metrics on this database. As on the general-purpose database, the FMPD has the best individual PLCC values on all four models in the masking database, as well as the highest PLCC for the whole repository. However, the overall SROCC for the FMPD is not as good as the SROCC of the MSDM2 metric (80.2% for FMPD vs. 89.6% for MSDM2) because the FMPD values of the four models are not exactly in the same range. The understanding and improvement of this limitation constitutes one aspect of our future work. Nevertheless, in general our metric captures the visual masking effect well, as proven by the high individual and overall PLCC and SROCC values (all $>80\%$) on the masking database.

Fig. 9. Some models from the LIRIS masking database. From left to right are, respectively, the reference Lion-vase model, the deformed model after noise addition only in rough regions (FMPD=0.35, MOS=0.42), and the deformed model after noise addition only in smooth regions (FMPD=0.69, MOS=0.84). The added noise is of the same intensity in both models.

The IEETA simplification database comprises 35 models: five reference meshes and six simplified models for each reference mesh. The simplified models were obtained using three simplification algorithms with two different vertex reduction ratios. Fig. 10 shows some models from this database, and Table 3 presents the results of the MVQ metrics. The FMPD has quite good overall performance on this database, and it outperforms the HD and MSDM2 metrics in terms of both the overall PLCC and the overall SROCC. In particular, the overall PLCC of the FMPD is much better than MSDM2 (89.3% for FMPD vs. 79.6% for MSDM2). However, the FMPD does not always provide very good results for each individual model, especially the Head model because sometimes the FMPD metric has difficulty in capturing the slight and subtle quality differences between simplified models with the same vertex reduction ratio that have been generated by different simplification algorithms. However, the MSDM2 metric has a similar problem; for instance, its relatively low SROCC on the Lung model is due to the same problem.

From the above experimental results and comparisons, we can conclude that the FMPD is quite efficient in predicting the results of the subjective assessment of MVQ. Compared to the existing

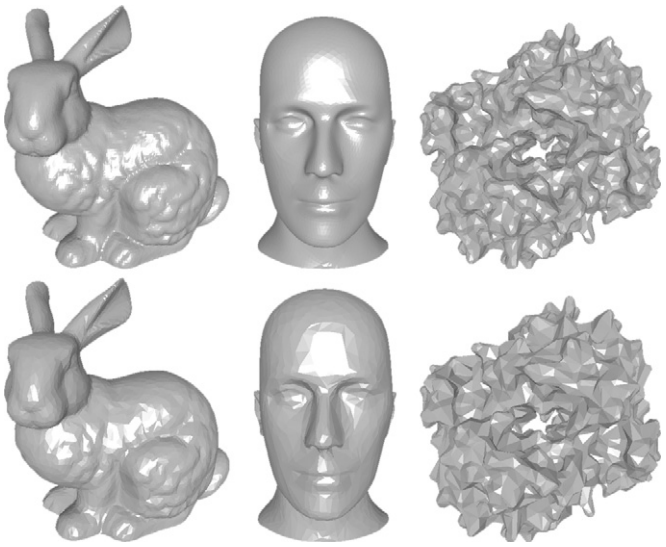

Fig. 10. Some models from the IEETA simplification database. In the top row are three reference models. In the bottom row, from left to right are, respectively, the Bunny model simplified with a reduction ratio of 80% by the QSlim method [8] (FMPD=0.45, MOS=0.59), the Head model simplified with a reduction ratio of 80% by the error-quadratics-based method in OpenMesh [54] (FMPD=0.60, MOS=0.74), and the Strange model simplified with a reduction ratio of 50% by the normal-flipping-criterion-based method in OpenMesh [54] (FMPD=0.30, MOS=0.46).

Table 2
PLCC (r_p) and SROCC (r_s) values (%) of different MVQ metrics on the LIRIS masking database.

Metric	Armadillo		Bimba		Dinosaur		Lion-vase		All models	
	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s
HD [3,4]	37.7	48.6	7.5	25.7	31.1	48.6	25.1	71.4	4.1	26.6
RMS [3,4]	44.6	65.7	21.8	71.4	50.3	71.4	23.8	71.4	17.0	48.8
$3DWPM_1$ [30]	41.8	58.0	8.4	20.0	45.3	66.7	9.7	20.0	10.2	29.4
$3DWPM_2$ [30]	37.9	48.6	14.4	37.1	50.1	71.4	22.0	38.3	18.2	37.4
MSDM2 [16]	65.8	**88.6**	93.7	**100**	91.5	**100**	87.5	**94.3**	76.2	**89.6**
FMPD	**94.2**	**88.6**	**98.9**	**100**	**96.9**	94.3	**93.5**	**94.3**	**80.8**	80.2

Table 3
PLCC (r_p) and SROCC (r_s) values (%) of different MVQ metrics on the IEETA simplification database.

Metric	Bones		Bunny		Head		Lung		Strange		All models	
	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s
HD [3,4]	84.8	**94.3**	14.3	39.5	53.0	**88.6**	64.9	**88.6**	27.4	37.1	25.5	49.4
MSDM2 [16]	**96.7**	77.1	96.3	**94.3**	**79.0**	**88.6**	85.3	65.7	**98.1**	**100**	79.6	86.7
FMPD	96.0	88.6	**98.0**	**94.3**	70.4	65.7	**95.5**	**88.6**	96.0	65.7	**89.3**	**87.2**

Table 4
Comparison of the execution times (in seconds) of the FMPD and MSDM2 [16] metrics and the Metro tool [3], on a laptop equipped with a 2.27 GHz Intel i5 CPU and 4 GB memory. The experiments were conducted on Venus models (Fig. 1, left) of different complexities.

# vertices $\mathcal{M}_r$	# vertices $\mathcal{M}_d$	FMPD	MSDM2	Metro
100 K	100 K	2.86	104.23	11.14
100 K	50 K	2.11	69.49	10.73
100 K	10 K	1.59	52.87	8.91
100 K	1 K	1.46	50.04	8.63
50 K	50 K	1.33	35.40	5.09
50 K	10 K	0.79	19.45	4.27
50 K	1 K	0.69	16.97	3.88
10 K	10 K	0.24	3.35	0.78
10 K	1 K	0.13	1.66	0.70
1 K	1 K	0.03	0.19	0.13

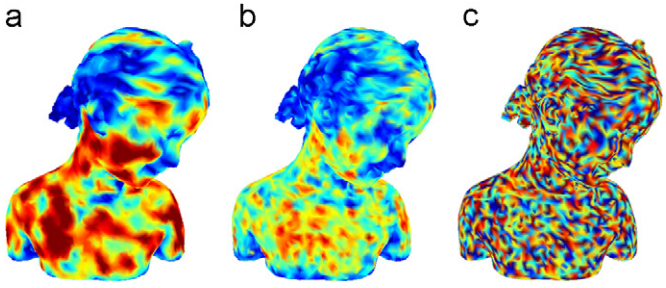

Fig. 11. The distance maps between the Bimba models shown in Fig. 5(a) and (b), produced by (left to right) the (a) FMPD, (b) MSDM2 and (c) RMS metrics. Warmer colors represent higher values. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this article.)

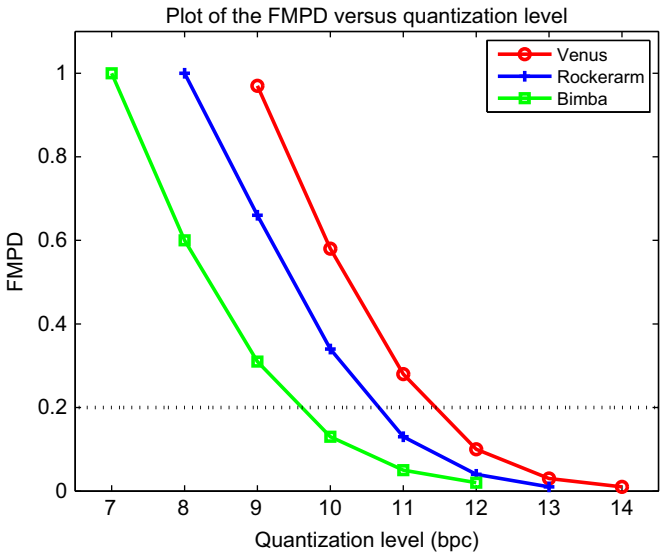

Fig. 12. Plot of FMPD versus quantization level (in bits per coordinate, bpc) of three meshes of different complexities.

metrics, the FMPD has the highest overall PLCC on all three subjective databases that we tested. The overall PLCC and SROCC of the FMPD are greater than 80% on the three databases, which demonstrates the efficacy and stability of the proposed metric. The FMPD outperforms the global-roughness-based metrics, $3DWPM_1$ and $3DWPM_2$, with much higher correlation values on the general-purpose and masking databases. This better performance is mainly due to the use of an efficient local roughness descriptor and the integration of the visual masking effect in the design of the FMPD.

4.2. Execution time and distance map

Time complexity is an important performance index for MVQ metrics. A low time complexity is desired in many practical applications, particularly in network-based QoS control and in the parameter optimization of mesh processing algorithms. In QoS control, it is essential to use a very fast MVQ metric to realize real-time QoS control, so as not to disturb the clients. In parameter optimization, it is often necessary to iteratively evoke the MVQ routine to find the optimum parameter value; therefore, a low-cost metric is required. Table 4 presents the execution time of the FMPD and MSDM2 [16] metrics and the Metro tool (which is a highly efficient implementation of the HD and RMS metrics) [3]. The comparisons were conducted on Venus models (Fig. 1, left) of different complexities (with the number of vertices ranging from 1 K to 100 K). All the data in Table 4 were obtained on a laptop equipped with a 2.27 GHz Intel i5 CPU and 4 GB memory. Our metric largely outperforms the MSDM2 metric and is also faster than the widely used Metro tool. For example, the FMPD metric takes less than 3 s to compare two meshes of 100 K vertices, whereas the MSDM2 metric takes approximately 100 s and the Metro tool takes approximately 11 s. Therefore, we believe that due to its high speed and simplicity, the FMPD metric has the potential to be used in many mesh applications that have strict requirements on time complexity.

A distance map is not a standard output of the proposed metric because the FMPD is a global-roughness-based metric. However, when correspondence information between the reference and deformed meshes is available, the FMPD can produce a distance map based on this information. This is the case for all the models in the general-purpose and masking databases because the models before and after deformation share the same connectivity. The local distance on each vertex is simply computed as the difference between its local roughness and that of its counterpart in the other mesh. Fig. 11(a) shows the distance map from the FMPD between the Bimba models shown in Fig. 5(a) and (b). This map is consistent with human perception; i.e., the perceived degradation is higher in smooth regions than in rough regions. A similar map produced by the MSDM2 metric is shown in Fig. 11(b), but the map from the local RMS metric in Fig. 11(c) is not relevant to human perception.

4.3. Application to vertex coordinate quantization

Vertex coordinate quantization is an important and almost mandatory preprocessing step of many mesh processing algorithms, especially compression [12,55]. The intensity of this quantization is usually measured in bits per coordinate (bpc): the lower the bpc, the greater the quantization. It is desired to find a proper and adaptive quantization level for each mesh, which in general is not an easy task. Here we define the optimum quantization level as the greatest intensity that does not introduce visible and disturbing distortion to the original mesh. A weaker quantization (i.e., with a higher bpc) than the optimum will increase the number of bits necessary for representing each coordinate; a greater quantization (i.e., with a lower bpc) than the optimum will introduce visible distortion. Meshes of different complexities or geometric details may have different optimum quantization levels, and a universal bpc setting is not appropriate. Currently, we often rely on human observation to find the optimum level, which is not fully automatic.

Our metric provides potentially a simple way to automatically determine the optimum quantization level for each triangle mesh. After a series of experiments, we have found that the optimum bpc can be found by simply fixing a threshold on the FMPD between the original and quantized meshes. More precisely, we consider the optimum bpc to be the lowest value that produces an FMPD of less than 0.20. Fig. 12 plots the FMPD versus the bpc for three meshes of different complexities: the Venus model with 100 K vertices shown in Fig. 1 left, the Rockerarm model with 40 K vertices shown in the third column of Fig. 7 and the Bimba model with 8.8 K vertices shown in Fig. 5(a). The evolution of the FMPD versus the bpc is different from mesh to mesh, but the three curves all exhibit a rapid increase once the FMPD is greater than 0.20. According to the simple rule that we propose, the optimum quantization level is 12, 11 and 10 bpc for the Venus, Rockerarm and Bimba models, respectively. These results are quite consistent with human observation, as shown in Fig. 13. Of course, users can adjust the threshold to adapt to the specific requirements of an application. It is worth mentioning that it would be nearly impossible to use simple geometric measures, such as the HD and RMS metrics, to automatically determine the optimum bpc because the same quantization level would produce nearly the same HD and RMS errors (relative to the size of the object) on all meshes.

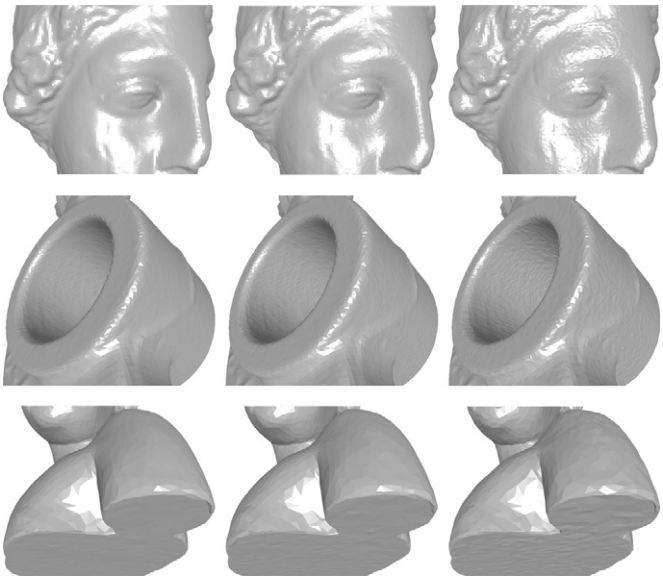

Fig. 13. Close-ups of quantized meshes with different bpc values. Close-ups in the middle correspond to the optimum quantization level (respectively, 12, 11 and 10 bpc for Venus, Rockerarm and Bimba, from top to bottom), whereas the close-ups on the left and right correspond, respectively, to a bpc that is one bit higher and one bit lower than the optimum level.

4.4. Discussion

One limitation of our method is that it cannot compare the visual difference between a pair of meshes that have exactly the same global roughness. This could happen in some extreme cases, but it has not been encountered in our experiments. Despite this limitation, an advantage of the FMPD is that it is possible to use the FMPD as a reduced-reference MVQ metric because only the global roughness of the reference mesh, along with some additional parameters, is required to assess the visual quality of a distorted mesh at the receiver side. Nevertheless, to overcome this global roughness ambiguity problem, it would be possible to modify the FMPD so that the modified metric would be based on the comparison of the local roughness between corresponding vertices of the two meshes. One side effect of this solution would be that we would need to perform a mesh correspondence preprocessing step, which could be difficult and computationally expensive in the general case.

As mentioned in Section 3.2, we tested the performance of our MVQ assessment pipeline, shown in Fig. 2, with different local roughness descriptors and found that the Laplacian of Gaussian curvature yielded the best results. For example, the overall SROCC value for the general-purpose database is reduced to 62.1%, 61.8% and 80.3% when using the mean curvature, the Laplacian of the mean curvature and the Gaussian curvature, respectively (compared to 81.9% for the Laplacian of Gaussian curvature). These results also demonstrate that the Gaussian curvature appears to be experimentally more relevant to human visual perception than the mean curvature. We also tested the pipeline with higher-order derivatives of the Gaussian curvature as the local roughness descriptor; the SROCC values for the general-purpose database are, respectively, 80.8% and 80.1% for the bi-Laplacian and the tri-Laplacian of the Gaussian curvature. The increased order of the Laplacian operator does not result in improvement in the metric performance. This may be because the higher-order geometric quantities are more sensitive to noisy data and thus may have lower numerical stability. At present, the Laplacian of the Gaussian curvature appears to be a good choice to obtain a high correlation with subjective scores while maintaining sufficient computational stability.

5. Conclusions and future work

We have proposed in this paper a fast and efficient objective metric, called FMPD, for the assessment of mesh visual quality. The perceptual comparison between a pair of reference and deformed models is based on a local roughness definition of the mesh surface. Careful modulation of the local roughness is applied to take into account the visual masking effect and the psychometric saturation effect. In particular, we show that using existing differential geometric quantities, combined with simple mathematical operations, we can derive an effective MVQ metric that is conceptually simple and easy to implement. Through a series of experiments and comparisons with existing perceptually based metrics, it is shown that the FMPD metric has a high correlation with subjective scores and a very low processing time. The simple application of the proposed metric to vertex coordinate quantization demonstrates the potential of this metric for its use in practical mesh processing applications. The source code for the FMPD implementation is shared on the Internet.

Our future work will mainly comprise three research directions: (1) conducting psychovisual experiments to study the relationship between the different differential geometric quantities and human visual perception in an attempt to build scientific and theoretical foundations for future research on MVQ assessment, (2) applying the FMPD metric to other mesh processing algorithms, especially compression and watermarking, and (3) extending the proposed metric for the visual quality assessment of point clouds and dynamic meshes.

References

[1] Botsch M, Kobbelt L, Pauly M, Alliez P, Lévy B. Polygon mesh processing. AK Peters; 2010.
[2] Smolic A, Muller K, Stefanoski N, Ostermann J, Gotchev A, Akar GB, Triantafyllidis GA, Koz A. Coding algorithms for 3DTV—a survey. IEEE Trans Circuits Syst Video Technol 2007;17(11):1606–21.
[3] Cignoni P, Rocchini C, Scorpigno R. Metro: measuring error on simplified surfaces. Comput Graph Forum 1998;17(2):167–74.
[4] Aspert N, Santa-Cruz D, Ebrahimi T. MESH: measuring error between surfaces using the Hausdorff distance. In: Proceedings of the IEEE international conference on multimedia & expo; 2002. p. 705–8.
[5] Wang Z, Bovik AC. Mean squared error: love it or leave it?—a new look at signal fidelity measures IEEE Signal Process Mag 2009;26(1):98–117.
[6] Lavoué G, Corsini M. A comparison of perceptually-based metrics for objective evaluation of geometry processing. IEEE Trans Multimed 2010;12(7):636–49.
[7] Bulbul A, Capin T, Lavoué G, Preda M. Measuring visual quality of 3-D polygonal models. IEEE Signal Process Mag 2011;28(6):80–90.
[8] Garland M, Heckbert PS. Surface simplification using quadric error metrics. In: Proceedings of the ACM siggraph; 1997. p. 209–16.
[9] Luebke DP. A developer's survey of polygonal simplification algorithms. IEEE Comput Graph Appl 2001;21(3):24–35.
[10] Wang K, Lavoué G, Denis F, Baskurt A. A comprehensive survey on three-dimensional mesh watermarking. IEEE Trans Multimed 2008;10(8):1513–27.
[11] Wang Y-P, Hu S-M. A new watermarking method for 3D models based on integral invariants. IEEE Trans Visual Comput Graph 2009;15(2):285–94.
[12] Alliez P, Gotsman C. Recent advances in compression of 3D meshes. In: Dodgson NA, Floater MS, Sabin MA, editors. Advances in multiresolution for geometric modelling. Springer-Verlag; 2003. p. 3–26.
[13] Lee H, Dikici C, Lavoué G, Dupont F. Joint reversible watermarking and progressive compression of 3D meshes. Visual Comput 2011;27(6–8):781–92.
[14] Wang K, Lavoué G, Denis F, Baskurt A. Robust and blind mesh watermarking based on volume moments. Comput Graph 2011;35(1):1–19.
[15] Cho JW, Prost R, Jung HY. An oblivious watermarking for 3-D polygonal meshes using distribution of vertex norms. IEEE Trans Signal Process 2007;55(1):142–55.
[16] Lavoué G. A multiscale metric for 3D mesh visual quality assessment. Comput Graph Forum (Proc Symp Geometry Process) 2011;30(5):1427–37.
[17] Wang Z, Bovik AC. Modern image quality assessment. Morgan & Claypool; 2006.
[18] Seshadrinathan K, Bovik AC. Video quality assessment. In: Bovik AC, editor. The essential guide to video processing. Academic Press; 2009. p. 417–36.
[19] Lindstrom P, Turk G. Image-driven simplification. ACM Trans Graph 2000;19(3):204–41.
[20] Qu L, Meyer GW. Perceptually guided polygon reduction. IEEE Trans Visual Comput Graph 2008;14(5):1015–29.
[21] Bolin MR, Meyer GW. A perceptually based adaptive sampling algorithm. In: Proceedings of the ACM siggraph; 1998. p. 299–309.
[22] Reddy M. Perceptually optimized 3D graphics. IEEE Comput Graph Appl 2001;21(5):68–75.
[23] Zhu Q, Zhao J, Du Z, Zhang Y. Quantitative analysis of discrete 3D geometrical detail levels based on perceptual metric. Comput Graph 2010;34(1):55–65.
[24] Rondao-Alface P, Craene MD, Macq B. Three-dimensional image quality measurement for the benchmarking of 3D watermarking schemes. In: Proceedings of the SPIE electronic imaging, vol. 5681; 2005. p. 230–40.
[25] Rogowitz BE, Rushmeier HE. Are image quality metrics adequate to evaluate the quality of geometric objects. In: Proceedings of the SPIE human vision and electronic imaging; 2001. p. 340–8.
[26] Karni Z, Gotsman C. Spectral compression of mesh geometry. In: Proceedings of the ACM siggraph; 2000. p. 279–86.
[27] Sorkine O, Cohen-Or D, Toledo S. High-pass quantization for mesh encoding. In: Proceedings of the symposium on geometry processing; 2003. p. 42–51.
[28] Pan Y, Cheng LI, Basu A. Quality metric for approximating subjective evaluation of 3-D objects. IEEE Trans Multimed 2005;7(2):269–79.
[29] Drelie Gelasca E, Ebrahimi T, Corsini M, Barni M. Objective evaluation of the perceptual quality of 3D watermarking. In: Proceedings of the IEEE international conference on image processing; 2005. p. 241–4.
[30] Corsini M, Drelie Gelasca E, Ebrahimi T, Barni M. Watermarked 3-D mesh quality assessment. IEEE Trans Multimed 2007;9(2):247–55.
[31] Lavoué G, Drelie Gelasca E, Dupont F, Baskurt A, Ebrahimi T. Perceptually driven 3D distance metrics with application to watermarking. In: Proceedings of the SPIE electronic imaging; 2006. p. 63120L.1–63120L.12.
[32] Wang Z, Bovik AC, Sheikh HR, Simoncelli E. Image quality assessment: from error visibility to structural similarity. IEEE Trans Image Process 2004;13(4):1–14.
[33] Bian Z, Hu S-M, Martin RR. Evaluation for small visual difference between conforming meshes on strain field. J Comput Sci Technol 2009;24(1):65–75.
[34] Lee CH, Varshney A, Jacobs DW. Mesh saliency. ACM Trans Graph (Proc ACM Siggraph) 2005;24(3):659–66.
[35] Desbrun M, Meyer M, Schröder P, Barr AH. Implicit fairing of irregular meshes using diffusion and curvature flow. In: Proceedings of the ACM siggraph; 1999. p. 317–24.
[36] Botsch M, Sorkine O. On linear variational surface deformation methods. IEEE Trans Visual Comput Graph 2008;14(1):213–30.
[37] Meyer M, Desbrun M, Schröder P, Barr AH. Discrete differential-geometry operators for triangulated 2-manifolds. In: Hege H-C, Polthier K, editors. Visualization and mathematics III. Springer-Verlag; 2003. p. 35–57.
[38] Reuter M, Wolter F-E, Peinecke N. Laplace–Beltrami spectra as 'Shape-DNA' of surfaces and solids. Comput Aided Des 2006;38(4):342–66.
[39] Vallet B, Lévy B. Spectral geometry processing with manifold harmonics. Comput Graph Forum (Proc Eurograph) 2008;27(2):251–60.
[40] Sorkine O. Differential representations for mesh processing. Comput Graph Forum 2006;25(4):789–807.
[41] Liu Y, Prabhakaran B, Guo X. A robust spectral approach for blind watermarking of manifold surfaces. In: Proceedings of the ACM workshop on multimedia and security; 2008. p. 43–52.
[42] Wang K, Luo M, Bors AG, Denis F. Blind and robust mesh watermarking using manifold harmonics. In: Proceedings of the IEEE international conference on image processing; 2009. p. 3657–60.
[43] Burchard HG, Ayers JA, Frey WH, Sapidis NS. Approximation with aesthetic constraints. In: Sapidis NS, editor. Designing fair curves and surfaces: shape quality in geometric modeling and computer-aided design. SIAM; 1994. p. 3–28.
[44] Breitmeyer BG. Visual masking: past accomplishments, present status, future developments. Adv Cogn Psychol 2007;3(1–2):9–20.
[45] Teo PC, Heeger DJ. Perceptual image distortion. In: Proceedings of the IEEE international conference on image processing; 1994. p. 982–6.
[46] Sheikh HR, Sabir MF, Bovik AC. A statistical evaluation of recent full reference image quality assessment algorithms. IEEE Trans Image Process 2006;15(11):3440–51.
[47] Legge GD. A power law for contrast discrimination. Vis Res 1981;21(4):457–67.
[48] Daly S. The visible differences predictor: an algorithm for the assessment of image fidelity. In: Watson AB, editor. Digital images and human vision. MIT Press; 1993. p. 179–206.
[49] Lavoué G. A local roughness measure for 3D meshes and its application to visual masking. ACM Trans Appl Percept 2009;5(4):1–23.
[50] Silva S, Santos BS, Ferreira C, Madeira J. A perceptual data repository for polygonal meshes. In: Proceedings of the international conference in visualization; 2009. p. 207–12.
[51] International Telecommunication Union. Methodology for the subjective assessment of video quality in multimedia applications. Recommendation BT.1788; 2007.
[52] International Telecommunication Union. Subjective video quality assessment methods for multimedia applications. Recommendation P.910; 2008.
[53] Engeldrum PG. Psychometric scaling: a tool for imaging systems development. Imcotek Press; 2000.
[54] Botsch M, Steinberg S, Bischoff S, Kobbelt L. OpenMesh — a generic and efficient polygon mesh data structure. In: Proceedings of the OpenSG symposium; 2002. p. 1–5.
[55] Peng J, Kim C-S, Kuo C-CJ. Technologies for 3D mesh compression: a survey. J Vis Commun Image Represent 2005;16(6):688–733.