A Curvature Tensor Distance for Mesh Visual Quality Assessment

Fakhri Torkhani, Kai Wang, and Jean-Marc Chassery

Gipsa-lab, CNRS UMR 5216, Grenoble, France
firstname.lastname@gipsa-lab.grenoble-inp.fr

Abstract. This paper presents a new objective metric for assessing the visual difference between a reference or 'perfect' mesh and its distorted version. The proposed metric is based on the measurement of a distance between curvature tensors of the two triangle meshes under comparison. Unlike existing methods, our algorithm uses not only eigenvalues but also eigenvectors of the curvature tensor to derive a perceptually-oriented distance. Our metric also accounts for some important properties of the human visual system. Experimental results show good coherence between the proposed objective metric and subjective assessments.

1 Introduction

Three-dimensional (3D) meshes are now used in many multimedia applications such as digital entertainment, medical imaging and computer-aided design. It is common that 3D meshes undergo some lossy operations like simplification, compression and watermarking. Since the end users are often human beings, it is thus important to derive metrics that can faithfully evaluate the perceptual distortions introduced by such operations [1]. Classical metrics of simple geometric distances (e.g. root mean squared error and Hausdorff distance) [2, 3] have been demonstrated not relevant to human visual perception and thus fail to predict the visual difference between a pair of reference and deformed meshes [1].

In order to design an accurate mesh visual distance metric, this paper attempts to reconcile several properties of the human visual system (HVS) with differential geometric quantities. Our contributions are summarized as follows:

- Introduction of an effective approach to the assessment of mesh visual quality (MVQ) based on a novel distance measure between mesh curvature tensors.
- Use of not only curvature values, but also surface principal directions (which have been proven perceptually important) to define the curvature tensor distance. This distance measure seems generic enough to be used in the future in other applications such as mesh segmentation and shape matching.
- Integration of some HVS features in the metric: we introduce a roughness-based weighting of local visual distance to simulate the *visual masking* effect, and a processing step similar to *Divisive Normalization Transform* (DNT) to mimic an important neural mechanism known as *adaptive gain control*.

The proposed metric has the potential to be used, for instance, to benchmark a variety of mesh processing algorithms, or to guide the design of new perceptually-oriented algorithms. The rest of this paper is organized as follows: The relevant research is briefly reviewed in Section 2; Section 3 details the pipeline of the proposed MVQ metric; Section 4 presents some experimental results, including the comparison with state-of-the-art metrics; Finally, we conclude in Section 5.

2 Related Work and Motivation

Mesh visual quality assessment. The first perceptually-oriented MVQ metric was introduced by Karni and Gotsman [4] for the evaluation of their mesh compression algorithm. The authors derived a metric by combining errors in both vertex positions and mesh Laplacian coordinates. Corsini et al. [5] developed two perceptual metrics, named respectively $3DWPM_1$ and $3DWPM_2$, for the visual quality assessment of watermarked meshes. The visual distortion is measured as the roughness difference between the original and watermarked meshes. Bian et al. [6] derived a perceptual measure where the visual difference between a pair of meshes is defined as the amount of strain energy required to induce the deformation between them. Lavoué et al. proposed a metric called mesh structural distortion measure ($MSDM$) [7], which can be considered as an extension of the well-known structure similarity index of 2D images [8] to the case of 3D triangle meshes. $MSDM$ relates the visual degradation to the alteration of local statistics (i.e. mean, variance and covariance) of mesh curvature amplitudes. An improved multiscale version $MSDM2$ [9] has been recently proposed, which also integrates a vertex matching preprocessing step to allow the comparison of two meshes with different vertex connectivities.

Motivation. $MSDM2$ exhibits good correlation with subjective scores [9], though by considering only the modification in mesh curvature amplitudes. We argue that the modification in surface *principal directions* (as defined by the *orthogonal* directions of minimum and maximum curvatures) is also important to MVQ assessment, because intuitively these directions imply structural information of the surface and thus should be visually important. Indeed, when drawing a 3D object, one strategy of caricaturists is to draw strokes on these lines of curvatures. In the digital world, surface principal directions have been successfully used for describing [10] and illustrating [11] complex 3D objects, and for guiding a remeshing algorithm [12]. Motivated by the above observation, we introduce a new MVQ metric, named $TPDM$ for Tensor-based Perceptual Distance Measure, that makes use of full information of the mesh curvature tensor, i.e. both curvature amplitudes and principal directions. In the following, we will briefly present a technique for estimating mesh curvature tensors and explain how to obtain curvature amplitudes and principal directions from the tensor.

Curvature tensor estimation. The estimation of mesh curvature tensor is a well-researched problem. So far, the most popular estimation technique is the one from Cohen-Steiner and Morvan [13]. Based on the solid foundation

of normal cycle theory, they derived an elegant per-vertex curvature tensor estimation. Tensors computed on edges are averaged on a geodesic disk window B of user-defined size to obtain the curvature tensor $\mathcal{T}$ on each vertex v:

$$\mathcal{T}(v) = \frac{1}{|B|} \sum_{edges\ e} \beta(e)\ |e \cap B|\ \bar{e}\ \bar{e}^t, \qquad (1)$$

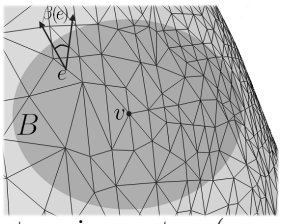

where $|B|$ is the area of the geodesic disk, $\beta(e)$ is the signed angle between the normals of the two triangles incident to edge e, $|e \cap B|$ is the length of the part of e inside B, and $\bar{e}$ is a unit vector in the direction of e (cf. the inset on right). The minimum and maximum curvature amplitudes (κ_{min} and κ_{max}) are two eigenvalues of tensor $\mathcal{T}$, and the principal directions are two eigenvectors (γ_{min} and γ_{max}). The lines of minimum and maximum curvatures define respectively the directions along which surface normals vary the slowest (e.g. along creases) and the fastest (e.g. across creases), which represent structural features of the surface. In the next section, we will derive a perceptually-oriented distance measure between curvature tensors by incorporating the information from both their eigenvalues and eigenvectors, and use this distance to conduct MVQ assessment.

3 Curvature Tensor Distance Based MVQ Assessment

3.1 Overview of the Pipeline

Figure 1 illustrates the pipeline of our metric $TPDM$. First, in order to compare two meshes with potentially different connectivities, we perform a vertex matching step between the two meshes, based on the AABB tree data structure in the CGAL library [14]. The next step is to compute a curvature tensor distance on each local window centered at a vertex. Afterwards, this local tensor distance is weighted by two roughness-based factors, so as to account for the visual masking effect. Finally, we use a surface-weighted Minkowski pooling of the local distances to obtain a global $TPDM$ value.

3.2 Curvature Tensor Distance

After the vertex matching step, each vertex v in the reference mesh $\mathcal{M}_r$ has a corresponding vertex v' on the surface of the distorted mesh $\mathcal{M}_d$. The curvature tensors on the two vertices are denoted respectively by $\mathcal{T}$ and $\mathcal{T}'$. For the comparison between them, we first establish correspondence relationships between the principal directions and curvature amplitudes of the two tensors. More precisely, for γ_{min} of $\mathcal{T}$, we find the principal direction of $\mathcal{T}'$ that has the smallest *angular distance* to it (this direction is denoted by γ'_1), and relate γ_{min} to γ'_1. Accordingly, κ_{min} of $\mathcal{T}$ is related to the curvature amplitude associated to γ'_1 (denoted by κ'_1). Note that κ'_1 and γ'_1 can be the maximum curvature and its direction of $\mathcal{T}'$. Similarly, the following correspondence relationships are

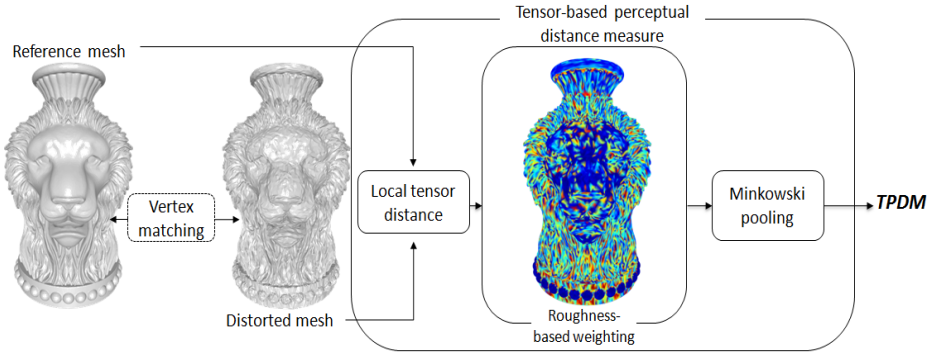

Fig. 1: Block diagram of the pipeline of $TPDM$ (Tensor-based Perceptual Distance Measure). In the roughness map, warmer colors represent larger values.

established: $\kappa_{max} \rightarrow \kappa'_2$ and $\gamma_{max} \rightarrow \gamma'_2$. We find that the above correspondence based on the minimum angular distance criterion yields better results of MVQ assessment than the straightforward min→min/max→max correspondence. In particular, it enhances the stability of $TPDM$ under the situations where the principal directions are severely disturbed after strong deformations and where the values of κ_{min} and κ_{max} are very close to each other.

A local tensor distance (LTD) is computed for each pair of v and v' as

$$LTD = \theta_{min}\delta_{\kappa_{min}} + \theta_{max}\delta_{\kappa_{max}}, \tag{2}$$

where θ_{min} is the angle between γ_{min} and γ'_1 (similarly for θ_{max}), and $\delta_{\kappa_{min}}$ is a Michelson-like contrast of the curvature amplitudes κ_{min} and κ'_1, i.e. $\delta_{\kappa_{min}} = \left|\frac{\kappa_{min}-\kappa'_1}{\kappa_{min}+\kappa'_1+\varepsilon}\right|$ with ε a stabilization constant fixed as 5% of the avarage mean curvature of $\mathcal{M}_r$ (similarly for $\delta_{K_{max}}$). Both the differences in curvature amplitudes and in principal directions are involved in the derivation of LTD.

3.3 Roughness-Based Weighting of Local Tensor Distance

For the development of an effective MVQ metric, we should take into account some HVS features, in particular the visual masking effect. In the context of MVQ assessment, this effect mainly means that a same distortion is less visible in rough regions of the mesh surface than in smooth regions. In order to account for the visual masking effect, our solution is to modulate the values of $LTD_{i,i=1,2,\ldots,N}$ (evaluated at each vertex v_i of $\mathcal{M}_r$) by two roughness-based weights (the rougher the local surface is, the smaller the weights are). The local tensor-based perceptual distance measure $LTPDM_i$ is computed as:

$$LTPDM_i = RW_i^{(\gamma)}.RW_i^{(\kappa)}.LTD_i, \tag{3}$$

where $RW_i^{(\gamma)}, RW_i^{(\kappa)} \in [0.1, 1.0]$ are respectively the weights derived from principal directions and curvature amplitudes in 1-ring neighborhood of v_i. For $RW_i^{(\gamma)}$,

we first project all the principal directions at the 1-ring neighbors on the tangent plane of v_i, and take the sum of the two angular standard deviations of the projected minimum and maximum curvature directions as the local roughness value. This value is then mapped to $[0.1, 1.0]$ to obtain $RW_i^{(\gamma)}$. Similarly, to get $RW_i^{(\kappa)}$, we compute the ratio of the Laplacian of mean curvature in the 1-ring neighborhood and the mean curvature on v_i as the local roughness and map it to $[0.1, 1.0]$. It is worth mentioning that the derivation of the roughness weight $RW_i^{(\kappa)}$ includes a divisive normalization similar to that in the neural mechanism of HVS that partially explains the visual masking effect [15]. Also note that the vertices in isotropic areas, i.e. where κ_{min} and κ_{max} are close to each other, are treated differently. For these vertices, we set $RW_i^{(\gamma)}$ close to 1, and therefore the final weight is dominated by the value of $RW_i^{(\kappa)}$. The reason is that in isotropic areas, principal directions are not well-defined and their estimation is not reliable. A roughness map that combines both weights is shown in Fig. 1.

3.4 Global Perceptual Distance

The global tensor-based perceptual distance measure $TPDM$ from $\mathcal{M}_r$ to $\mathcal{M}_d$ is computed as a *weighted Minkowski sum* of the local distances $LTPDM_i$:

$$TPDM = \left(\sum_{i=1}^{N} w_i \left| LTPDM_i \right|^p \right)^{\frac{1}{p}}, \tag{4}$$

where $w_i = s_i / \sum_{i=1}^{N} s_i$ with s_i one third of the total area of all the incident faces of v_i, and we set $p = 2.5$. The surface-based weighting can, to some extent, enhance the stability of the metric to the variation of vertex sampling density over the mesh surface. Compared to the standard mean-squared error in which $p = 2.0$, the choice of $p = 2.5$ can increase the importance of the local distances of high amplitude. This is perceptually relevant since the part of mesh with high-amplitude distortion has experimentally more impact on the result of subjective assessment. Finally, a cumulative Gaussian psychometric function [16] is applied to bring the $TPDM$ value to the $[0, 1]$ interval. More details on the psychometric function will be provided in the next section.

4 Experimental Results

In order to verify its efficacy, the proposed metric $TPDM$ has been extensively tested and compared with existing metrics on three subject-rated databases:

- The LIRIS/EPFL general-purpose database[1] [7]: It contains 4 reference meshes and in total 84 deformed models. The distortion types include noise addition and smoothing, applied either locally or globally on the reference mesh. Subjective evaluations were made by 12 observers.

[1] `http://liris.cnrs.fr/guillaume.lavoue/data/datasets.html`

- The LIRIS masking database[2] [17]: It contains 4 reference meshes and in total 24 deformed models. The local noise addition distortion included in this database was designed specifically for testing the capability of MVQ metrics in capturing the visual masking effect. 11 observers participated in the subjective tests.
- The IEETA simplification database[3] [18]: It contains 5 reference meshes and in total 30 simplified models. 65 observers participated in this study.

$TPDM$ has been compared with five state-of-the-art metrics, i.e. the Hausdorff distance HD [2, 3], the root mean squared error RMS [2, 3], the two roughness-based metrics $3DWPM_1$ and $3DWPM_2$ from Corsini et al. [5], and $MSDM2$ [9]. The coherence between the distance values produced by the objective metrics and the subjective mean opinion scores (MOS) is measured by using two different kinds of correlation: the Pearson linear correlation coefficient ($PLCC$) that measures the prediction accuracy of the objective metrics, and the Spearman rank-order correlation coefficient ($SROCC$) that measures the prediction monotonicity. Before computing $PLCC$, it is recommended to conduct a psychometric fitting between the objective measures and the MOS values, in order to partially remove the non-linearity between them. In our tests, we apply a cumulative Gaussian psychometric function [16]:

$$g(a, b, R) = \frac{1}{\sqrt{2\pi}} \int_{a+bR}^{\infty} e^{-(t^2/2)} dt, \tag{5}$$

where R is the raw $TPDM$ value, and the two parameters a and b are obtained through non-linear fitting using the raw $TPDM$ and the corresponding MOS values of the group of Dinosaur models in the general-purpose database. As shown in Fig. 2.(b) and (c), the same psychometric function has been used for the models in the masking and simplification databases.

Tables 1, 2 and 3 present the results on the general-purpose, masking and simplification databases, respectively. In general, $TPDM$ exhibits quite good performance on all the three databases, reflected by its high correlation with subjective scores on most individual models and on the whole repositories. In particular, $TPDM$ has always the highest overall $PLCC$ value (the second last column in the tables), thus the highest prediction accuracy on all the three databases. On the general-purpose database (cf. Table 1), $TPDM$ has the highest $PLCC$ and $SROCC$ for almost every individual model as well as for the whole repository, and there is much improvement in the overall $PLCC$ compared to $MSDM2$, the best metric proposed so far. On the masking database (cf. Table 2), although the overall $SROCC$ of $TPDM$ is not as high as that of $MSDM2$, we can still conclude that $TPDM$ well captures the visual masking effect, as reflected by the high individual and overall correlation values (all $> 80\%$). $TPDM$ has slightly better overall performance than $MSDM2$ on the simplification database (cf. Table 3). However for the Head model the correlation

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

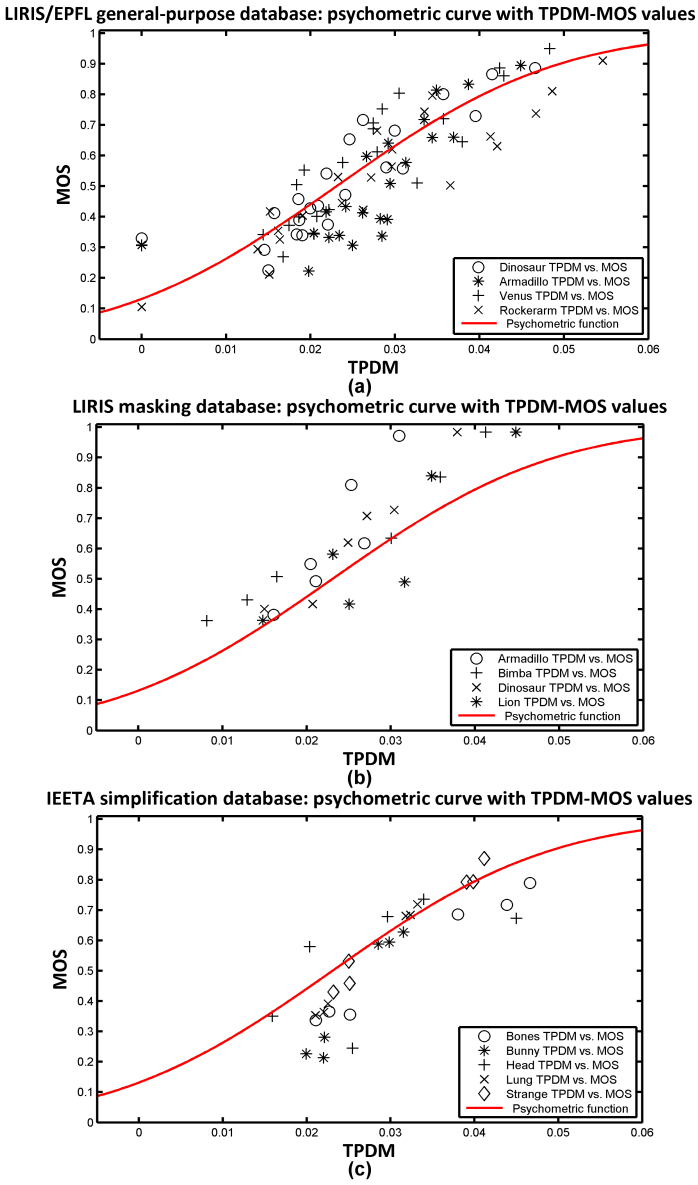

Fig. 2: Psychometric function curve plotted with $TPDM$-MOS values of all the reference/distorted models in: (a) the LIRIS/EPFL general-purpose database; (b) the LIRIS masking database; and (c) the IEETA simplification database.

is rather low, and the reason is that $TPDM$ has difficulties in distinguishing the quality of the simplified Heads generated by different mesh simplification algorithms but with the same vertex reduction ratio.

Table 1: $PLCC$ (r_p) and $SROCC$ (r_s) (%) of the different objective 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 [2, 3]	30.2	69.5	22.6	30.9	5.5	18.1	0.8	1.6	1.3	13.8
RMS [2, 3]	32.2	62.7	0.0	0.3	3.0	7.3	77.3	**90.1**	7.9	26.8
$3DWPM_1$ [5]	35.7	65.8	35.7	62.7	53.2	87.5	46.6	71.6	38.3	69.3
$3DWPM_2$ [5]	43.1	74.1	19.9	52.4	29.9	37.8	16.4	34.8	24.6	49.0
$MSDM2$ [9]	72.8	81.6	73.5	85.9	76.1	89.6	76.5	89.3	66.2	80.4
$TPDM$	**79.3**	**85.4**	**89.4**	**92.2**	**91.4**	**90.6**	**87.6**	89.9	**84.9**	**85.2**

Table 2: $PLCC$ (r_p) and $SROCC$ (r_s) (%) of the different objective metrics on the LIRIS masking database.

Metric	Armadillo		Bimba		Dinosaur		Lion		All models	
	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s	r_p	r_s
HD [2, 3]	37.7	48.6	7.5	25.7	31.1	48.6	25.1	71.4	4.1	26.6
RMS [2, 3]	44.6	65.7	21.8	71.4	50.3	71.4	23.8	71.4	17.0	48.8
$3DWPM_1$ [5]	41.8	58.0	8.4	20.0	45.3	66.7	9.7	20.0	10.2	29.4
$3DWPM_2$ [5]	37.9	48.6	14.4	37.1	50.1	71.4	22.0	38.3	18.2	37.4
$MSDM2$ [9]	65.8	**88.6**	93.7	**100**	91.5	**100**	**87.5**	**94.3**	76.2	**89.6**
$TPDM$	**91.3**	**88.6**	**97.1**	**100**	**97.1**	**100**	86.7	82.9	**87.1**	87.3

Table 3: $PLCC$ (r_p) and $SROCC$ (r_s) (%) of the different objective 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 [2, 3]	84.8	**94.3**	14.3	39.5	53.0	**88.6**	64.9	88.6	27.4	37.1	25.5	49.4
$MSDM2$ [9]	96.7	77.1	96.3	**94.3**	**79.0**	**88.6**	85.3	65.7	98.1	**100**	79.6	**86.7**
$TPDM$	**98.9**	**94.3**	**97.9**	**94.3**	63.1	65.7	**99.9**	**100**	**98.8**	94.3	**86.4**	**86.7**

Figure 3 illustrates the distance maps produced by $TPDM$ and RMS for a noised Bimba model. The map of $TPDM$ is quite consistent with human perception (i.e. the perceived distortion is higher in smooth regions than in rough regions), while the map of RMS is purely geometric. Figure 4 shows an application of our metric in the visual quality assessment of watermarked meshes. The two watermarked models have exactly the same geometric maximum root mean squared error ($MRMS$) [3] compared to the original mesh, but their visual

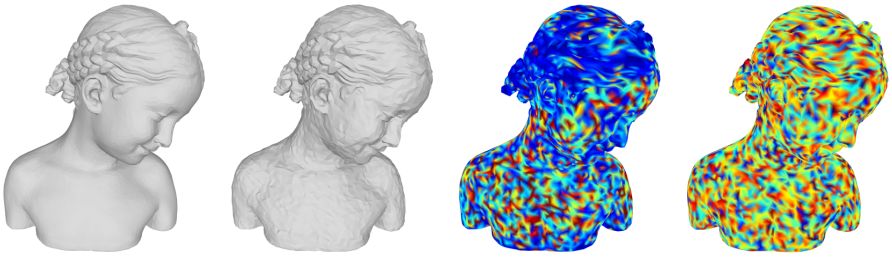

Fig. 3: From left to right: the original Bimba model, the deformed model after uniform noise addition, the distance map of $TPDM$, and the distance map of RMS. In the maps, warmer colors represent higher values.

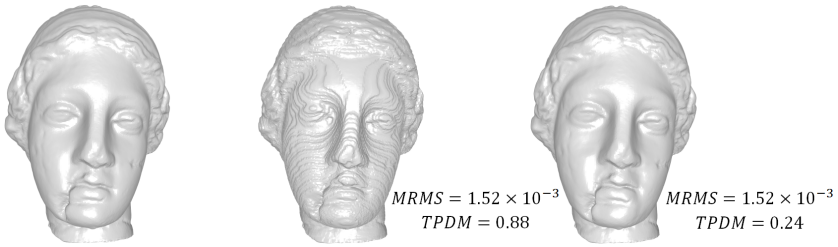

Fig. 4: From left to right: the original Venus model, the model watermarked by the method in [19], and the model watermarked by the method in [20].

quality is quite different. $TPDM$ provides correct MVQ evaluation results that are consistent with a subjective assessment.

5 Conclusions and Future Work

A new curvature-tensor-based approach to the objective evaluation of mesh visual quality has been proposed. We show that it is beneficial to use the full information of the curvature tensor for MVQ assessment. The local curvature distance and the local roughness measures that we propose may be found useful in other mesh applications. Experimental results show that our metric $TPDM$ has high correlation with subjective scores and that it slightly outperforms existing metrics. Future work mainly consists of the integration of more HVS features in the metric (e.g. the contrast sensitivity function), the improvement of the roughness measure which at present appears a little noisy, and the development of a curvature-tensor-based visual quality metric for dynamic meshes.

Acknowledgments

We would like to thank the anonymous reviewers for their helpful and constructive comments. This work has been in part supported by the MOOV3D project of the Minalogic competitive cluster.

References

1. Corsini, M., Larabi, M.C., Lavoué, G., Petřík, O., Váša, L., Wang, K.: Perceptual metrics for static and dynamic triangle meshes. In: Proc. of Eurographics State-of-the-Art Rep. (2012) 135–157
2. Cignoni, P., Rocchini, C., Scopigno, R.: Metro: measuring error on simplified surfaces. Comput. Graphics Forum **17**(2) (1998) 167–174
3. Aspert, N., Santa-Cruz, D., Ebrahimi, T.: MESH: measuring errors between surfaces using the Hausdorff distance. In: Proc. of IEEE Int. Conf. on Multimedia & Expo. (2002) 705–708
4. Karni, Z., Gotsman, C.: Spectral compression of mesh geometry. In: Proc. of ACM Siggraph. (2000) 279–286
5. Corsini, M., Drelie Gelasca, E., Ebrahimi, T., Barni, M.: Watermarked 3-D mesh quality assessment. IEEE Trans. on Multimedia **9**(2) (2007) 247–256
6. Bian, Z., Hu, S.M., Martin, R.R.: Evaluation for small visual difference between conforming meshes on strain field. Journal of Comput. Sci. and Technol. **24**(1) (2009) 65–75
7. Lavoué, G., Drelie Gelasca, E., Dupont, F., Baskurt, A., Ebrahimi, T.: Perceptually driven 3D distance metrics with application to watermarking. In: Proc. of SPIE Electronic Imaging. (2006) 63120L.1–63120L.12
8. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: From error visibility to structural similarity. IEEE Trans. on Image Process. **13**(4) (2004) 600–612
9. Lavoué, G.: A multiscale metric for 3D mesh visual quality assessment. Comput. Graphics Forum **30**(5) (2011) 1427–1437
10. Brady, M., Ponce, J., Yuille, A.L., Asada, H.: Describing surfaces. Comput. Vision, Graphics, and Image Process. **32**(1) (1985) 1–28
11. Hertzmann, A., Zorin, D.: Illustrating smooth surfaces. In: Proc. of ACM Siggraph. (2000) 517–526
12. Alliez, P., Cohen-Steiner, D., Devillers, O., Lévy, B., Desbrun, M.: Anisotropic polygonal remeshing. ACM Trans. on Graphics **22**(3) (2003) 485–493
13. Cohen-Steiner, D., Morvan, J.M.: Restricted delaunay triangulations and normal cycle. In: Symp. on Computational Geometry. (2003) 312–321
14. Alliez, P., Tayeb, S., Wormser, C.: 3D fast intersection and distance computation (AABB tree). In: CGAL User and Reference Manual. (2012)
15. Li, Q., Wang, Z.: Reduced-reference image quality assessment using divisive normalization-based image representation. IEEE J. Sel. Topics Signal Process. **3**(2) (2009) 202–211
16. Engeldrum, P.G.: Psychometric Scaling: A Toolkit for Imaging Systems Development. Imcotek Press (2000)
17. Lavoué, G.: A local roughness measure for 3D meshes and its application to visual masking. ACM Trans. on Appl. Perception **5**(4) (2009) 21:1–21:23
18. Silva, S., Santos, B.S., Ferreira, C., Madeira, J.: A perceptual data repository for polygonal meshes. In: Proc. of Int. Conf. in Visualization. (2009) 207–212
19. Cho, J.W., Prost, R., Jung, H.Y.: An oblivious watermarking for 3-D polygonal meshes using distribution of vertex norms. IEEE Trans. on Signal Process. **55**(1) (2007) 142–155
20. Wang, K., Lavoué, G., Denis, F., Baskurt, A.: Robust and blind mesh watermarking based on volume moments. Comput. & Graphics **35**(1) (2011) 1–19