

Comparing Small Visual Differences between Conforming Meshes

Zhe Bian¹, Shi-Min Hu¹ and Ralph Martin²

¹ Department of Computer Science & Technology
Tsinghua University, Beijing, China, 100084

bz05@mails.tsinghua.edu.cn, shimin@tsinghua.edu.cn

² School of Computer Science, Cardiff University, Cardiff, UK
Ralph.Martin@cs.cardiff.ac.uk

Abstract. This paper gives a method of quantifying (small) visual differences between 3D mesh models with conforming topology, based on the theory of strain fields. Our experiments show that our difference estimates are well correlated with human perception of differences. This work has applications in the evaluation of 3D mesh watermarking, 3D mesh compression reconstruction, and 3D mesh filtering.

Key words: 3D Conforming Meshes; Mesh Comparison; Perception; Strain Fields

1 Introduction

3D surface triangle meshes have widespread uses in computer graphics and modelling; such techniques as watermarking, filtering and compression are often applied to such meshes.

Watermarking is used to provide information security and digital rights management, the aim being to hide an ‘invisible’ digital signature into the mesh [1, 2, 7, 11–13]. The basic approach used is to encode the watermark information into small perturbations to the model’s description, e.g. its vertex coordinates, such that the model’s geometry changes by a hopefully imperceptible amount. Relatively little work has been done on methods for evaluating the *quality* of watermarking schemes, which involves the *perceptibility* of the watermark (as well as other considerations such as how difficult it is to remove or destroy the watermark, how much information the watermark can carry, the time taken to process the watermark, etc).

Generally, mesh models constructed from 3D scanner data are often noisy, and must be denoised, or faired, before they are suitable for application purposes [17]. More generally, various filters may be applied to meshes to modify them in some way [18]. Again, it is useful to be able to assess the *visual impact* of filtering algorithms.

Both to economize use of bandwidth, and to save storage, various mesh compression algorithms have been devised [8, 14, 16]. Evaluation of mesh compression algorithms is again an important issue, where *visual differences* between the reconstructed version and the original are again important (as well as other issues such as speed of compression and decompression).

Here, we give a new method for measuring small visual differences between conforming meshes (i.e. meshes with the same number of triangles, connected in the same way), with such applications as the above in mind. Specifically, given an initial mesh, and a modified version of the mesh, we quantify the visible difference between the two models. Such a quality measure should agree with human perception, which is generally considered to be subjective and hard to quantify [19]. However, our objective methodology based on strain fields, which numerically assesses deformation between the meshes, provides a measure which is well correlated with perceptual results provided by human subjects.

A limitation of our method is that it assumes that the processing method only perturbs the vertices of the mesh, and does not change its topology. This is often the case in watermarking and filtering schemes, but perhaps less so in other applications.

2 Related work

We now review previous work on measuring visual errors in various application areas.

Clearly, if a watermark becomes too obvious, or distorts the model too much, a watermarking scheme is unacceptable. We may also wish to *compare* the perceptibility of the same watermarking information added by different schemes when deciding which scheme to adopt. For such assessment, [9] gives a method adapted from image watermarking assessment, while [6] provides two methods based on surface roughness. The shortcomings of these approaches will be discussed shortly. Few objective methodologies have been proposed, and even less attention has been paid to reliable procedures for the tricky problem of subjective evaluation [15]: static 2D images of 3D models are inadequate for assessing the quality of a 3D model. There is a need for 3D quality metrics.

Other papers have considered evaluation of the visual effects of mesh compression and filtering. Some are based on the perceptual metrics used for assessing visual differences between images [9], others consider geometric differences [3, 5, 6], and yet others combine both [21]. However, assessing perceptual degradation of images is a rather different task to assessing perceptual degradation of 3D models. The methodologies in [3, 5] only consider maximum and mean geometric errors based on Hausdorff distance, which have been shown not to correlate well with human opinions of perceptibility [6]. The two methodologies in [6] provide better bases for evaluation. One describes mesh distortion using changes of dihedral angles between faces, but ignores vertex displacements. The other uses a first order equation for displacement in a 2-ring area which can not readily describe gradual changes and sudden changes alike. Furthermore, the time taken is worse than linear in the number of mesh faces.

The most direct way of evaluating the perceptibility of a change made to a 3D mesh is to measure it by means of subjective experiments using human observers. While such methods represent the ‘correct’ answer, provided by the human visual system, there are many difficulties in carrying out such evaluations. Using a pool of ‘typical’ (whatever that means) human observers is time consuming and costly. Human beings do not give consistent and repeatable answers. Different people may give different opinions. They may get tired and lose concentration. It is difficult to carry out tests of this kind

without accidentally introducing bias. Thus, there are various reasons to prefer objective methodologies for assessing visual changes in meshes. This paper provides such a methodology which at the same time produces results correlating well with those subjectively produced by human observers, at least in our limited testing.

Our basic approach is based on analyzing shape and size changes of the mesh triangles, in a way which removes any rigid-body motion. We use *strain field energy* to quantify the deformation. Section 3 explains the basic ideas of strain fields, and Section 4 shows how to compute strain in a mesh. Section 5 provides a method to compute a perceptibility distance based on strain field energy. Section 6 and Section 7 describe an experiment used to understand human perception of visual changes in meshes, and experimental tests of our proposed methodology, including comparisons with other simpler candidate measures of perceptual difference. Conclusions are drawn in Section 8.

3 Strain fields

Stress and *strain* are physical quantities used to describe pointwise deformation inside an elastic body [22]. Stress represents forces acting at any point, while strain represents the resulting deformation. Strain field theory makes the following assumptions:

1. The object is solid, and contains no voids.
2. The object is composed of a homogeneous material.
3. The material is isotropic (i.e. behaves the same in all directions).
4. The material reverts to its original shape if external forces are removed; deformation is proportional to the applied external force.
5. No external forces act in the undeformed state, i.e. the original stress is zero.
6. The deformation of the object is small.

Clearly, the first five assumptions are plausible for most mesh models, and so it is reasonable to treat them as (the surfaces of) elastic objects. The final assumption is also satisfied for the kind of visual difference we wish to assess.

We regard the effect of the mesh processing algorithm being evaluated as equivalent to applying a fictitious external force which causes a distortion of the 3D mesh. We use the idea of a strain field to analyze the deformation of the mesh (or more exactly, of the 3D elastic solid whose surface is the mesh).

Changes in position of the mesh vertices represent distortion of the mesh, but may also incorporate a rigid-body motion which leaves the object's shape unchanged. We thus must first remove any such rigid body motion, using the concept of *strain*. However, we first explain *stress*, which causes strain.

When external forces are applied to an object, these create internal forces. Consider two parts of the solid which touch each other—these exert mutual forces on each other. See Figure 1(a): suppose the plane $n-n$ is the interface between the two parts; let M be any point on that plane. Consider the differential area ΔA around M . Let the internal force \mathbf{F}_{in} on ΔA caused by the external force be $\Delta \mathbf{P}$; this will generally have components both within the plane and perpendicular to it. The quantity $\Delta \mathbf{P}/\Delta A$ is called the average pressure. The *stress* \mathbf{F} acting on M is defined to be the limit of the average

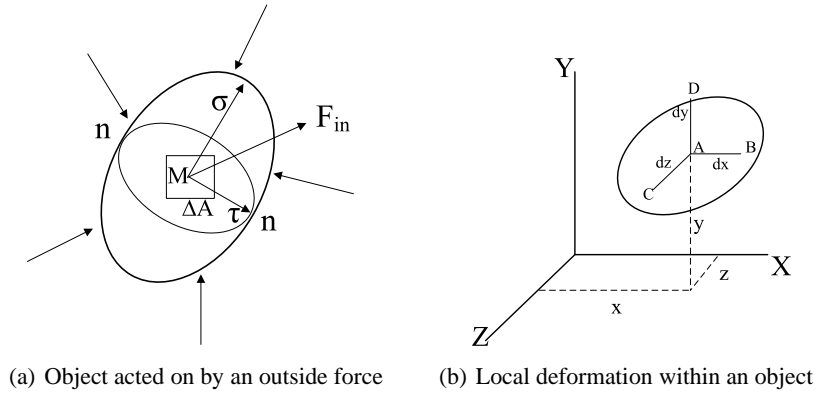


Fig. 1. Object deformation

pressure [20]:

$$\mathbf{F} = \lim_{\Delta A \rightarrow 0} \frac{\Delta \mathbf{P}}{\Delta A} \quad (1)$$

Within this plane, we may divide the *stress* into two components: the *normal stress*, denoted by σ , perpendicular to the section plane, and the *shear stress*, within the section plane, denoted by τ : see Figure 1(a).

We now consider strain: see Figure 1(b). Let A be some arbitrary point of the object. Consider three infinitesimal lines AB , AC , and AD within the object, parallel to the coordinate axes, with respective lengths dx , dy , dz . When the object is acted upon by an external force, the lengths of these lines, and the angles between them, change. The fractional changes in the *lengths* are called the *normal strain* at A , denoted by ξ_x , ξ_y , ξ_z , and the fractional changes in the *angles* are called the *shear strain*, denoted by γ_{xy} , γ_{yz} , γ_{zx} .

When an external force acts on an object, its changes in shape and size can be described in terms of displacements. Given A , at (x, y, z) , let A' be the corresponding point after deformation. We denote by (u, v, w) the components of the *displacement* in the (x, y, z) directions respectively, i.e. $A' = (x + u, y + v, z + w)$.

We may now state the relations between displacement and strain. In 3D these are given by:

$$\begin{aligned} \xi_x &= \partial u / \partial x, & \gamma_{xy} &= \partial v / \partial x + \partial u / \partial y \\ \xi_y &= \partial v / \partial y, & \gamma_{yz} &= \partial w / \partial y + \partial v / \partial z \\ \xi_z &= \partial w / \partial z, & \gamma_{zx} &= \partial u / \partial z + \partial w / \partial x \end{aligned} \quad (2)$$

4 Strain in a mesh

We now consider how stress and strain can be used to provide a simple measure of perceptibility of differences between two conforming meshes.




Fig. 2. Triangular face before and after distortion

The 3D mesh is a shell composed of triangular faces with negligible thickness. We analyze the strain by considering these triangles. See Figure 2(a). If the model is distorted slightly, and we assume that its faces do not bend, to a first approximation the mesh triangles are unchanged in the direction of their normal: only their shapes and locations change [4]. Thus, any elastic deformation only occurs within the plane of each triangle, and the strain for each triangle can be computed in its own plane by ignoring any rigid body motion.

To compute the strain for the whole mesh, using the mesh vertex displacements, we also need the rates of change of displacement. Thus, we interpolate the displacement functions across each triangle. See Figure 2(b). Let D, E, F be the vertices of a triangle in its initial position. We choose a local coordinate system so that this triangle is in the (x, y) plane. After deformation, these vertices go to new positions. We project these final vertices, in the z -direction, back into the (x, y) plane as shown, giving new points D', E', F' . The displacement of each vertex from its initial position is given by $x_{D'} = x_D + u_D, y_{D'} = y_D + v_D$. The finite element method [23] is now used to interpolate the displacement functions u, v at all points inside the triangle. We consider u in detail; v is computed similarly. We approximate u across each triangle by a linear function: $u = a_1 + a_2x + a_3y$ where the a_i are to be determined. We put u, x and y for each vertex of the triangle in turn into this expression and solve the resulting equations to give a_1, a_2, a_3 . The results are:

$$\begin{aligned}
 a_1 &= \frac{1}{2A}((x_E y_F - x_F y_E)u_D + (x_F y_D - x_D y_F)u_E + (x_D y_E - x_E y_D)u_F) \\
 a_2 &= \frac{1}{2A}((y_E - y_F)u_D + (y_F - y_D)u_E + (y_D - y_E)u_F) \\
 a_3 &= \frac{1}{2A}((x_F - x_E)u_D + (x_D - x_F)u_E + (x_E - x_D)u_F)
 \end{aligned} \tag{3}$$

where

$$2A = (y_E - y_F)x_D + (y_F - y_D)x_E + (y_D - y_E)x_F.$$




Fig. 3. Area and distortion deformation

As we assume that each triangle deforms entirely within its own plane, there is no deformation or strain normal to each triangle. Thus, following the approach in [4], Equation (2) may be simplified in this case to give:

$$\begin{aligned}\xi_x &= \frac{\partial u}{\partial x}, & \xi_y &= \frac{\partial v}{\partial y}, & \xi_z &= \frac{\nu}{\nu - 1}(\xi_x + \xi_y), \\ \gamma_{xy} &= \frac{\partial v}{\partial x} + \frac{\partial u}{\partial y}, & \gamma_{yz} &= 0, & \gamma_{zx} &= 0,\end{aligned}\quad (4)$$

where ν is the Poisson ratio (see later). We can now compute ξ_x , ξ_y and γ_{xy} from the displacement functions u, v obtained above.

5 Perceptibility and strain field energy

5.1 Basic Approach

A model of perceptual change for images was given in [7]; we use a similar approach for perceptual differences in meshes. We define a *perceptual distance* $P(m_0, m_p)$ between the original model m_0 and the processed model m_p : the larger the perceptual distance, the more visible the difference. Defining such a function that entirely agrees with human perception is an idealistic goal. However, our experiments show that we can define such a function using stress and strain which agrees well with human perception.

Finding a suitable way to combine ξ_x , ξ_y and γ_{xy} , which have different physical meanings, and different directionality, is a difficult issue which we now consider. As a first step to combining the strain components into a perceptual distance, it helps to combine them into fewer physical quantities with more easily understood meanings and which are more convenient to use. We use the concept of *strain field energy* from elasticity for this purpose.

The *strain field energy density* is defined as

$$D = \frac{1}{2}(\sigma_x \xi_x + \sigma_y \xi_y + \sigma_z \xi_z + \tau_{xy} \gamma_{xy} + \tau_{yz} \gamma_{yz} + \tau_{zx} \gamma_{zx}).$$

Using the relations between stress and strain, the *strain field energy* may be written as [22]:

$$W = DS_{\Delta} = \left(\frac{1}{2}(\lambda + \frac{2}{3}G)\xi_{ii}^2 + \frac{E}{2(1+\nu)}\xi'_{ij}\xi'_{ij}\right)S_{\Delta} \quad (5)$$

where S_{Δ} is the area of each triangle, and

$$\begin{aligned} \xi_{ii}^2 &= \xi_x^2 + \xi_y^2 + \xi_z^2, & \xi'_{ij}\xi'_{ij} &= \xi'_x\xi'_x + \xi'_y\xi'_y + \xi'_z\xi'_z + \frac{1}{2}(\gamma_{xy}^2 + \gamma_{yz}^2 + \gamma_{zx}^2), \\ \xi'_k &= \xi_k - \frac{1}{3}(\xi_x + \xi_y + \xi_z), & k &= x, y, z \end{aligned}$$

E is *Young's modulus* of elasticity, ν is the *Poisson ratio* and $G = E/(2(1+\nu))$ is the *shear modulus*: these physical quantities determine the material's elastic properties. We simply fix them in our methodology: we use $E = 1$ and $\nu = 0$, as these choices seem to give good results in practice.

Elastic deformation of a planar triangle can be decomposed to into two parts, *area* change and *distortion* deformation: see Figure 3. The former describes change in *area* of the triangle; the latter describes change in *shape* of the triangle. As these two deformations may have different visual effects when looking at a deformed mesh, we may divide the *strain field energy* W into two corresponding parts: *area strain field energy*, W_{Area} and *distortion strain field energy*, $W_{\text{Distortion}}$ [22]. This allows us to combine them with different weights according to their respective visual importance. To calculate these energies, we decompose the stress and strain into independent parts corresponding to area change and distortion deformation. The *strain field energy* model for a single triangle can now be written as $W = W_{\text{Area}} + W_{\text{Distortion}}$, where

$$W_{\text{Area}} = \frac{1}{2}(\lambda + \frac{2}{3}G)\xi_{ij}^2 S_{\Delta}, \quad W_{\text{Distortion}} = \frac{1}{2}(2G)\xi'_{ij}\xi'_{ij} S_{\Delta}.$$

5.2 Candidates for improved methods

Distortions in 3D meshes may be considered to be of varying kinds, with varying perceptibility. We therefore considered how we might improve upon the basic idea of using strain energy to assess visibility of mesh cvhanges. In particular, we considered two issues, and candidate improvements:

- **Projection** Many applications of 3D meshes render them. The most important element determining the appearance of a triangle is its normal vector. Strain field energy does not directly capture this idea. For example, when vertices are displaced within the local tangent plane of the mesh, the strain field can be large although a rendered image remains almost the same. Consider several adjacent triangles which lie almost in the same plane; we call such triangles *nearly tangential triangular faces*: see Figures 4(a) and 4(b). These triangular faces have the same normals. We may consider two types of distortion involving these triangles: ones in which the center vertex still lies in the plane (Figure 4(c)), and others in which the center vertex moves perpendicular to the tangent plane (Figure 4(d)). The former have little




Fig. 4. Nearly tangential triangular faces and their distortions

visual effect. We should de-emphasise the strain field energy of such distortions which have little or no visual effect.

More generally, we can represent any local distortion as a combination of within-tangent-plane distortion and normal distortion. To reduce the contribution of the in-tangent-plane distortion, we may project the vertex after distortion into the original tangent plane as shown in Figure 4(d), and *adjust* the triangles connected to O accordingly. We then calculate the strain field energy with respect to these *adjusted original triangles*.

- **Edge triangles** Triangular faces in different locations generally have differing visual impact. In particular, small distortions in triangular faces near sharp edges of the model are likely to be more noticeable than faces in the interior of smooth surfaces. It is plausible that we should thus give edge triangles a greater weighting. This can be done by applying a weight w_i to each triangle’s strain field energy according to the angle between it and neighbouring faces. A simple way to choose the weights is to set $w_i = \pi - \alpha_i$, if α_i is less than $\pi/2$ and $w_i = 1$ otherwise, where α_i is the smallest dihedral angle between face i and its neighbours.

As we will see later, these two proposed improvements actually have little useful effect on our results.

5.3 Perceptual distance

We now consider how to use strain field energy as a model for *perceptual distance*. We must take into account that when the same geometric change is produced in models of different sizes, visual differences will be smaller in the bigger model. Our method should also ideally be independent of the number of triangles in the model. We thus define the *perceptual distance* $P(m_0, m_p)$ to be the weighted average strain field energy (ASFE) over all triangles (processing tangent triangular faces), normalized by S , the total area of the triangular faces:

$$P(m_0, m_p) = \frac{1}{S} \sum w_i W_i. \quad (6)$$

We may go further, however, and define two separate components $P_{\text{Area}}(m_0, m_p)$ and $P_{\text{Distortion}}(m_0, m_p)$ by separately summing the area and distortion components of W_i . The simplest way to combine them is to just sum them, giving $P(m_0, m_p) = P_{\text{Area}}(m_0, m_p) + P_{\text{Distortion}}(m_0, m_p)$ as before. Alternatively, we might wish to ensure that *both* types of distortion are independently less than some threshold for a difference

to be considered imperceptible, or that *both* quantities are less for a given mesh change than some other mesh change.

Alternatively, we could try to determine the relative importance of these two strain field energies to perceptual mesh differences, and produce an overall perceptual distance based on a projected or weighted sum of these quantities. However, in practice, this does not seem to produce an obvious improvement.

5.4 Algorithm

We now summarize the above ideas as an algorithm for computing the visual change in a mesh.

1. Input the original mesh and the mesh under comparison.
2. Find and process the tangent triangular faces; (optionally) compute the weight for each triangle.
3. Compute and interpolate the displacement function over each triangle.
4. Compute the components of strain for each triangle.
5. Sum the area and distortion strain field energies over the mesh and combine using Equation (6).

6 Evaluation

We now demonstrate testing of our methodology using three experiments, and show that our measure of perceptibility of small changes to a mesh produces results that correlate well with those assessed by human subjects.

Experiment I gives human opinions concerning the perceptibility of changes caused by making changes to models used in Experiments II and III. Experiment II considers two simple measures of perceptibility of changes, based on triangle areas, and triangle normals: clearly we hope to do better than using such simple measures. Experiment III demonstrates the correlation between perceptual distance computed by our approaches with and without the suggested improvements, and the human perception results.

Experiment I: Human perception of differences versus strain field measure

To obtain ground truth, i.e. some subjective human results on the perceptibility of various changes in the mesh models [19], we followed the approach in [6], used to evaluate and multimedia content quality. We embedded data into mesh models of a chess king and a horse, using the *Triangle Similarity Quadruple* watermarking method [12], and a noise embedding method.

To measure the subjective degree of mesh deformations perceived by human observers, we produced meshes with different amounts of deformation. We prepared variants of the horse, and of the chess king, resulting in 15 meshes for each: the undeformed model, 7 with watermarks embedded in different parts of the mesh, incorporating different amounts of data, and 7 with different intensity noise embedded at different places.

Subjects were asked to rate the differences between the original model and the processed model on a scale from 0 to 10, where 0 meant identical and 10 meant very dissimilar, to give a *opinion score* (OS).

In order to help the subjects evaluate the 3D mesh models, and to assist subjects' determination of differences, we paid careful attention to rendering conditions, as suggested by [6]:

- *Color.* We used black for the background to help models stand out. Models were coloured grey, which makes edges more visible and deformations easier to see.
- *Light source.* All models used the same single white point light source: multiple lights can confuse observers.
- *Lighting.* Although a local illumination model can produce more realistic effects for textured models, it can hide parts with high distortion. Thus we also used a global illumination model.
- *Texture.* Models were untextured, as textures can hide any distortion.
- *Test subjects.* 30 test subjects (20 male, 10 female) were drawn from a pool of computer science students aged 22–25. For impartiality, some of the chosen test subjects had knowledge of computer graphics, and others did not.
- *Screen and Model Resolution.* The models were displayed on an 17-in LCD monitor, with resolution 1280×1024 . The watermarked and original model were displayed together so as to fill the screen. The chess king model had 12170 triangles and horse model had 10024 triangles, allowing clear observation of detail. The screen was viewed from a distance of approximately 0.6m.
- *Interaction.* We allowed the subjects to rotate and zoom the models: [15] suggests that evaluation of alterations to 3D objects should permit interaction.

The experiment comprised three steps:

1. *Oral instructions and training.* First, we told the subjects about 3D mesh models, watermarking, compression and filters. We then gave examples of an unaltered mesh, to be scored as 0, and a worst-case altered mesh, to be scored as 5.
2. *Practice with a sample model.* Next, the subjects were allowed to interact with various processed models to familiarize themselves with the experiment.
3. *Experimental trials.* In this step, the subjects were asked to score the differences between the original models and altered models.

While human observers were somewhat variable in their opinions as to perceptibility of differences, there was general correlation between (average) strain field energy and opinion scores, as will be discussed in more detail in the next Section. Figures 5(a) and 5(b) show plots of the individual opinion scores (OS) against strain field energy. Each circle in Figure 5(a) and cross in Figure 5(b) corresponds to one model assessed by one subject.

We note that as might be expected, human subjects vary considerably in their opinions of the perceptibility of differences. Thus, in subsequent experiments we used the *means* of these human opinion scores (MOS) for each altered model as being representative of the amount of visual differences perceived by human subjects.

Experiment II: Strain field measure versus other simpler measures

We next investigated the relationship between mean opinion score values (MOS) and two other simple perceptual distance measures which might plausibly be used for assessing mesh distortion: the *fractional change in the total area of the triangles*




Fig. 5. Unimproved ASFE perceptual distance versus subjective opinion score (OS)

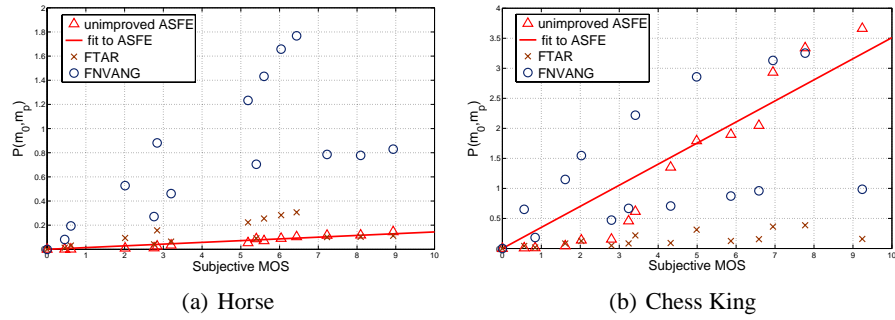


Fig. 6. Various perceptual difference measures versus mean subjective opinion score (MOS)

(P_{FTAR}) and the *fractional change in angle between normal vectors of adjacent faces* (P_{FNVANG}). The first of these measures is defined as follows:

$$P_{\text{FTAR}}(m_0, m_p) = \sum_{i=1}^n |\Delta S_i| / \sum_{i=1}^n S_i,$$

where n is the number of faces of the mesh, ΔS_i is the change in area of face i and S_i is its area. The second is defined as

$$P_{\text{FNVANG}}(m_0, m_p) = \sum_{i=1}^m |\Delta \alpha_i| / \sum_{i=1}^m \alpha_i,$$

where m is the number of the edges of the mesh, $\Delta \alpha_i$ is change between in angle of normal vectors between edge i and α_i is the angle of normal vectors between edge i .

Figure 6 show the mean opinion scores from Experiment I are not well correlated with the proposed P_{FTAR} or P_{FNVANG} perceptual distances, either for the horse or the chess king—we can see these two simpler perceptual distance measures produce much more scattered results than our strain energy measure. They do not adequately predict human opinion of mesh differences—whereas the strain field perceptual distance measure produces results which lie much closer to a straight line.

Table 1. Correlation between various perceptual distance measures and human opinion

Perceptual distance based on	Correlation coefficient	
	Horse	Chess King
FTAR	0.56	0.71
FNVANG	0.67	0.54
<i>unimproved</i> ASFE	0.97	0.97

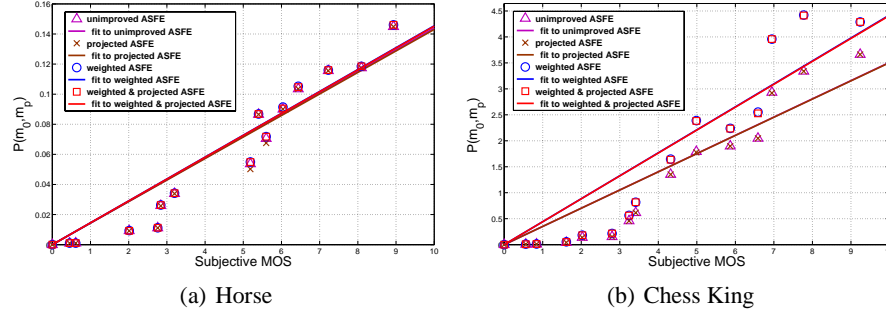


Fig. 7. Kinds of ASFE perceptual distance and mean opinion score

To more precisely analyze this observation, we calculated the correlation coefficient between each perceptual distance measure and the human mean observation scores: see Table 1. As, expected, the simpler FTAR and FNVANG perceptual distance have much lower correlation than our *unimproved* ASFE strain energy perceptual distance, and hence are of less value for measuring human opinion of distortion. We further conclude, given the very high correlation coefficients observed, that the (unimproved) perceptual distance based on strain field energy is a useful replacement for subjective mean human opinions of mesh differences.

Experiment III: Variants of strain field energy measure

We next investigated the relationship between strain field energy and mean opinion score values (MOS) when using the suggested improvements based on projection and edge weights (see Section 5.2). We compared compared 4 variants of our method: *unimproved* perceptual distance (without projection and face weights), perceptual distance using projection, perceptual distance using edge weights, and perceptual distances using both projection and edge weights.

Figure 7 shows the results of comparing the four different perceptual measures with mean human opinion scores, as in Experiment II. In each case a similar close-to-linear relationship can be seen. close-to-linear fashion for both the horse and chess king. It can be seen that the projection idea has almost no effect on perceptual distances. Using edge weighting produces a bigger difference in the perceptual distance measure, although a visual inspection of the results does not show any obvious immediate improvement (see Figure 7) in terms of linearity of relationship. Again correlation coefficients were computed and are given in Table 2. The edge weights *do* make a small improvements for the horse model, but projection does the opposite, while for the chess king model,

Table 2. Correlation between the four perceptual distance measures and human opinion

Perceptual distance based on	Correlation coefficient Horse	Correlation coefficient Chess King
<i>unimproved</i> ASFE	0.9745	0.9660
<i>projected</i> ASFE	0.9715	0.9661
<i>weighted</i> ASFE	0.9751	0.9560
<i>projected & weighted</i> ASFE	0.9752	0.9557

Table 3. Perceptual distance based on strain energy $P \times 10^{-5}$ for Buddha models with varying deformations

Watermarked models	FTAR	P	Models with noise	FTAR	P
(a)	0	0	(g)	0	0
(b)	661	4	(h)	115	4
(c)	3335	114	(i)	568	114
(d)	6644	455	(j)	1074	523
(e)	10020	1042	(k)	1702	1778
(f)	12342	1533	(l)	2313	3988

the opposite is true, projection making a small improvement but edge weighting doing the opposite. Overall, our experiments show no benefit to using either suggested improvement: using projection and/or edge weights makes almost no difference to the correlation between the computed measure and human opinion. We thus recommend the *unimproved* ASFE based perceptual distance which is simpler to compute, and the other methods increase the computation time, without any benefit.

7 Further demonstrations

We now provide two further demonstrations of our approach. The first is intended to give the reader a visual impression of how our measure works in practice, while the second demonstrates an application of our methodology by comparing the relative perceptibility deformations produced by two different mesh deformation methods.

Demonstration I: Distortion in a Buddha model

We now present a series of Buddha models, with 62224 faces, which have been processed by watermarking (Figures 8(a–f)) and noise addition (Figures 8(g–l)) algorithms, making changes controlled by successively increasing the FTAR measure as stated in Table 3.

These figures and numbers allow the reader to gain some impression of their own concerning our (basic) perceptual distance measure based on strain energy for these models.

Demonstration II: Application to comparison of mesh processing methods

We next demonstrate a simple application of our methodology. We compared meshes processed in two different ways, to demonstrate that our method provides the same opin-




Fig. 8. Watermarked Buddhas (a–f), Buddhas with added noise (g–l)

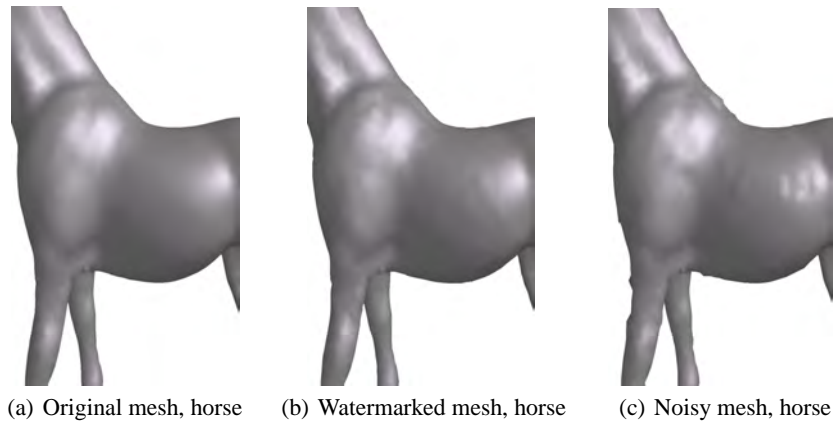


Fig. 9. Original and processed models, horse

ion as human observers as to which produced less perceptible results. Again we used the horse model with 10024 faces and the chess king with 12170 faces. The latter is smoother than the horse and its edges are more obvious than those of the horse, and so we expected differences to be more visually apparent for this model. The meshes were distorted using the same two processing methods as in Experiment I—watermarking, and adding noise. We embedded an average 2 bits distortion into each triangular face of the meshes using each method. Figure 9(a) shows the original horse model, Figure 9(b) shows the watermarked horse model, and Figure 9(c) shows the noisy horse model. Differences in the distortions are clearly perceptible in these close up views. Figures 10(a)–10(c) are analogous figures for the chess king model.

For both models, the reader will probably agree with other human observers that the original and watermarked meshes show quite small differences, whereas the differences between the original and noisy meshes are more obvious. The watermarks are embedded into the all parts of the mesh fairly evenly, whereas more noise is embedded into the mesh at some places than others. As a result distortions in the latter case is more obvious, and correspondingly our perceptual distance measure is also bigger.

Furthermore, as expected, because the king is smoother than the horse, the traces of the processing in the king are more obvious and the deformations are more noticeable, as predicted in Table 4 by our strain energy perceptual distance; here we have used the projected and edge weighted measure. For both models, our measure of the difference is lower for the watermarked mesh than the noisy mesh, which corresponds with our visual inspection of these models. Our own visual subjective results, and the measure are in agreement that the watermarking method produces less visible differences than adding noise, both for the smoother king model and for the more textured horse model.

Note, however, that both of the simpler area-based and normal-based perceptual distance measures predict the *opposite*: that the watermarking changes should be more visible. This is further evidence that these particular, simple, perceptual distance measures are not satisfactory.

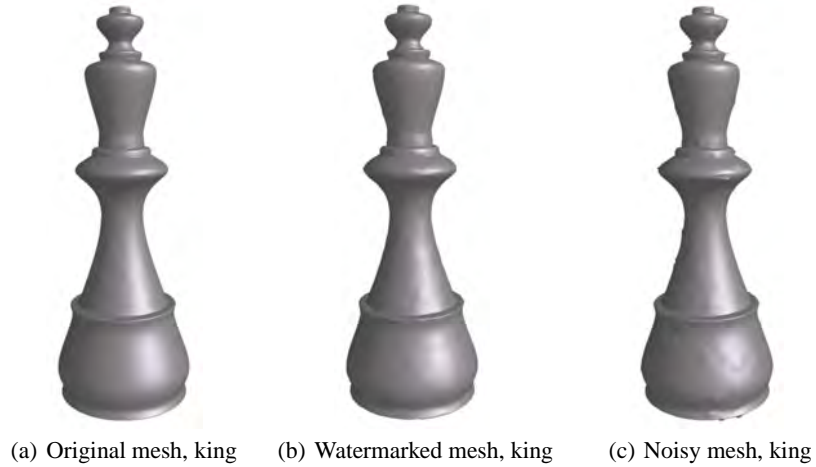


Fig. 10. Original and processed models, king

Table 4. Perceptibility distances for three measures and two models

Models	Projected & weighted ASFE	FTAR	FNVANG
Watermarked Horse	0.0002	0.0623	0.3596
Noisy Horse	0.0017	0.0414	0.2712
Watermarked King	0.0140	0.0835	1.1492
Noisy King	0.0591	0.0499	0.4720

8 Conclusions and future work

This paper has proposed a new methodology for comparing small visual differences between conforming models. Our methodology is based on the use of strain field theory to describe the deformation arising due to application of some algorithm. Our experiments show that this objective method can produce results which are strongly correlated with the average of subjective human opinions. The same is *not* true of simpler measures based on changes in triangle areas, or surface normals.

This work has applications to the comparison of 3D mesh processing algorithms such as watermarking, compression and other filtering in terms of the perceptible effects they produce on models. For example, in watermarking we could use this approach to decide which part of a model is the most suitable place for embedding a watermark, to decide how much data can be hidden in a mesh model, or to choose between watermarking schemes.

One area we wish to explore in future is to try to statistically determine thresholds of perceptibility through subjective experiments, which can then be used with our method to decide whether a visual difference is perceptible.

References

1. P. Agarwal, K. Adi and B. Prabhakaran: Robust Blind Watermarking Mechanism for Motion Data Streams. Proc. 8th Workshop on Multimedia and security, Geneva, Switzerland, 230–235. (2006)
2. M. Arnold, D. Wolthusen and M. Schmucker: Techniques and Applications of Digital Watermarking and Content Protection. Artech House Publishers. (2003)
3. N. Aspert, D. Santa-Cruz and T. Ebrahimi: Mesh: Measuring Error Between Surfaces Using the Hausdorff Distance. Proc. IEEE Int. Conf. Multimedia and Expo 2002 (ICME), 1, 705–708. (2002)
4. W. F. Chen and A. F. Salipu: Elasticity and Plasticity. China Architecture & Building Press. (2005)
5. R. S. P. Cignoni and C. Rocchini.: Metro: Measuring Error on Simplified Surfaces. Comput. Graph. Forum, 17 (2), 167–174. (1998)
6. M. Corsini, E. D. Gelasca, T. Ebrahimi and M. Barni: Watermarked 3-D Mesh Quality Assessment. IEEE Trans. Multimedia, 9 (2), 247–256. (2007)
7. I. Cox, M. Miller, J. Bloom and M. Miller: Digital Watermarking: Principles & Practice. Morgan Kaufman. (2002)
8. P. Heckbert and M. Garland: Surface Simplification Using Quadric Error Metrics. Proc. SIGGRAPH'97 209–216. (1997)
9. P. Lindstrom and G. Turk: Image-Driven Simplification. ACM Trans. Graph., 19 (3), 204–241. (2000)
10. Methodology for Subjective Assessment of the Quality of Television Pictures Recommendation BT. 500–11 Geneva, Switzerland. (2002)
11. R. Ohbuchi, H. Masuda and M. Aono: Watermarking Three-Dimensional Polygonal Models. Proc. ACM International Conference on Multimedia'97, Seattle, 261–272. (1997)
12. R. Ohbuchi, H. Masuda and M. Aono: Watermarking Three-Dimensional Polygonal Models through Geometric and Topological Modifications. IEEE J. Selected Areas in Communication, 16 (4), 551–560. (1998)
13. R. Ohbuchi, S. Takahashi, T. Miyazawa and A. Mukaiyama: Watermarking 3D Polygonal Meshes in the Mesh Spectral Domain. Proc. Graphics Interface 2001, 9–17. (2001)
14. E. Puppo and R. Scopigno.: Simplification, LOD and Multiresolution-Principles and Applications. Eurographics'97 Tutorial Notes PS97 TN4. (1997)
15. B. Rogowitz and H. Rushmeier: Are Image Quality Metrics Adequate to Evaluate the Quality of Geometric Objects. Proc. SPIE Human Vis. Electron. Imag. VI, B. E. Rogowitz and T. N. Pappas, Eds. Volume 4299, 340–348. (2001)
16. J. Rossignac and A. Szymczak: Wrap & Zip Decompression of the Connectivity of Triangle Meshes Compressed with Edgebreaker. Computational Geometry, Theory and Applications. 14 (1/3), 119–135. (1999).
17. X. Sun, P. L. Rosin, R. R. Martin and F. C. Langbein: Fast and Effective Feature-Preserving Mesh Denoising. IEEE Trans. Visualization and Computer Graphics 13 (5), 925–938. (2007)
18. G. Taubin. A Signal Processing Approach to Fair Surface Design. In SIGGRAPH'95 351–358. (1995)
19. B. Watson, A. Friedman and A. McGaffey: Measuring and Predicting Visual Fidelity. In SIGGRAPH'01 213–220. (2001)
20. F. Wilhelm: Tensor Analysis and Continuum Mechanics. Berlin: Springer-Verlag. (1972)
21. N. Williams, D. Luebke, J. D. Cohen, M. Kelley and B. Schubert: Perceptually Guided Simplification of Lit, Textured Meshes. In Proc. 2003 Symp. Interactive 3D Graphics. Monterey, CA. 113–121. (2003)
22. Z. Xun: Elasticity Mechanics. Peoples Education Press. (1979)
23. O. C. Zienkiewicz: The Finite Element Method. McGraw-Hill. (2000)