Roughness-Adaptive 3D Watermarking of Polygonal Meshes

Kwangtaek Kim[1], Mauro Barni[2], and Hong Z. Tan[1]

[1] Haptic Interface Research Laboratory, Purdue University,
465 Northwestern Avenue, West Lafayette, Indiana 47906 USA
[2] Department of Information Engineering, University of Siena,
via Roma 56, 53100, Siena, Italy

Abstract. We present a general method to improve watermark robustness by exploiting the masking effect of surface roughness on watermark visibility, which, to the best of our knowledge, has not been studied in 3D digital watermarking. Our idea is to adapt watermark strength to local surface roughness based on the knowledge that human eyes are less sensitive to changes on a rougher surface patch than those on a smoother surface. We implemented our idea in a modified version of a well known method proposed by Benedens [3]. As an additional contribution, we modified Benedens's method in two ways to improve its performance. The first improvement led to a blind version of Benedens's method that no longer requires any key that depends on the surface mesh of the cover 3D object. The second improvement concerned the robustness of bit '1' in the watermark. Experimental results showed that our new method permits to improve watermark robustness by 41% to 56% as compared to the original Benedens's method. Further analyses indicated that the average watermark strength by our roughness-adaptive method was larger than that by the original Benedens's method while ensuring watermark imperceptibility. This was the main reason for the improvement in robustness observed in our experiments. We conclude that exploiting the masking property of human vision is a viable way to improve the robustness of 3D watermarks in general, and therefore could be applied to other 3D digital watermarking techniques.

1 Introduction

Over the past few years, 3D display devices and rendering software have become more affordable than before, accelerating the use of 3D meshes in applications such as video games, CAD (computer-aided design), VR (virtual reality) and medical imaging. The rapid growth of 3D contents on the Internet has led to renewed interests in developing 3D watermarking schemes to protect 3D polygonal meshes from illegal reproductions. Compared to 2D digital watermarking, 3D watermarking is more difficult due to the increased complexity associated with 3D objects with arbitrary shapes. In addition, 3D watermarking is more fragile due to the various ways that embedded watermarks can be destroyed

S. Katzenbeisser and A.-R. Sadeghi (Eds.): IH 2009, LNCS 5806, pp. 191–205, 2009.

by simply altering the meshes making up the 3D objects. For this reason, existing 2D watermarking techniques cannot be directly applied to 3D models, necessitating new approaches that are specifically designed for 3D objects. One straight-forward way to embed watermarks into 3D meshes is to modify the locations of vertices making up the 3D surface. The challenge is to design 3D digital watermarks that are *unobtrusive*, *robust*, and *space efficient* [1]. The *unobtrusive* requirement means that the embedded watermarks must not interfere with the intended use of a model, which may imply imperceptibility. *Robustness* refers to the ability for watermarks to survive various intentional and unintentional attacks to the watermarked 3D model. This is a very challenging requirement as no algorithm has been shown to be perfectly robust. However, constant improvements are being made that result in more robust watermarking schemes as compared to previous methods. The last requirement is about having enough space for watermark embedding. To meet all those requirements is not trivial.

This paper exploits visual masking as a general method for improving the robustness of 3D digital watermarks. Masking refers to our decreased ability to perceive a stimulus (e.g., the watermark) in the presence of other signals (e.g., polygonal mesh). In the context of watermarking, we expect that stronger watermarks can be embedded in a rougher surface area without compromising imperceptibility. Therefore, we present a new method that adaptively matches watermark strength to local surface roughness, as opposed to the traditional method of choosing the maximum imperceptible watermark strength for the entire surface of a 3D object. To test our idea, we applied a roughness-adaptive watermarking technique to Benedens's geometry-based 3D watermarking method [3] and evaluated its robustness against simulated additive-noise attacks. We also improved the original method in two ways, by achieving a blind-version of Benedens's method and by improving the robustness of bit '1' in the watermark.

1.1 Literature Review of 3D Digital Watermarking

As one of the earliest studies of watermarking on polygonal meshes, Ohbuchi et al. [1] proposed several relatively robust techniques utilizing a combination of geometry (vertex, edge, facet) and topology (mesh connectivity) properties. For instance, the ratio of lines, tetrahedral volume ratio and density of mesh were modified as the embedding primitive for watermarking. Building on Ohbuchi et al.'s work [1], Cayre and Macq proposed a high-capacity blind watermarking scheme that was an extension of the TSPS (Triangle Strip Peeling Sequence) method by employing quantization indices (0 or 1) during the reflection of a triangular vertex on the list of triangles to be traversed [8]. Watermarks by either method were not robust against re-meshing and mesh simplification that can perturb mesh-connectivity, but were robust against geometrical attacks such as rotation, translation, or uniform scaling transformation. To address the robustness problem, Benedens [3] proposed a more robust method using the distribution of surface normals of the input mesh. The main idea was to group the surface normals into distinct sets called bins. Each bin was formed in a cone shape defined by a unique Bin Center(BC) normal and a bin radius (φ^R) with the

constraint that no overlapping was allowed among the bins. Watermarks were embedded by modifying the mean normals, the mean angle of normals, or the amount of normals inside each bin predefined by the bin radius (φ^R) from the BC. Due to the advantages associated with the use of surface normals instead of mesh connectivity, the Benedens's method is considered one of the most robust 3D watermarking methods based on geometry property.

In an effort to improve Ohbuchi et al.'s method [1] that uses the ratio of lines and volumes, Wagner [2] proposed a new method that used the relative lengths of vectors of each of two connected points on a polygonal mesh to improve the robustness against affine and similarity transformations. Similarly, Benedens and Busch [4] presented an Affine Invariant Embedding (AIE) technique that was applicable to non-manifolds.

In order to improve the imperceptibility of watermarks, Bors [6,7] employed human detection thresholds to choose the regions where the human eyes are less sensitive to changes. The watermarks were embedded into the selected vertices using two different bounding volumes consisting of two planes and ellipsoids. The method was robust against geometrical transformations such as translation, rotation and scaling as well as cropping and additive noise. Most recently, two blind methods, POA (Principal Object Axis) and SPOA (Sectional Principal Object Axis), were presented by Zafeifiou et al. [12]. For watermark embedding, a conversion to spherical coordinates was required during preprocessing, and the r component, the distance from the origin, was modified to insert bi-directional watermarks along a ray from the origin that passes through a point p determined by a neighborhood operator. The neighborhood operator was formed by computing the value of the mean or the median of the r component of neighbors, or by building a parametric surface using the neighborhood as the control points. POA is robust against rotation, translation, and uniform scaling. SPOA is additionally robust against mesh simplification attacks. Since both schemes are based on principal component analysis, watermarks can be easily destroyed by cropping attacks that perturb the principal object axes.

1.2 Our Contributions

Our work makes two contributions. First, we propose an adaptive method to select watermark strength. A psychophysical experiment on visual masking showed that human detection thresholds for additive noise increase monotonically with local surface roughness. The resulting quantitative model is the basis for roughness-adaptive watermarking algorithm that ensures locally maximal watermark strength while guaranteeing its imperceptibility. Second, we introduce a blind version of Benedens's method with improved robustness of the '1' bit.

The remainder of this paper is organized as follows. In Section 2, we introduce Benedens' method and discuss its limitations. The blind version of Benedens's method is presented in Section 3. Section 4 describes the psychophysical experiment and the roughness-adaptive watermarking algorithm. Evaluation results appear in Section 5. Finally, we conclude the paper in Section 6.

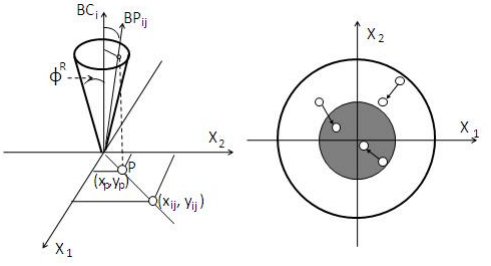

Fig. 1. (left) Transformation of 3D coordinates into 2D coordinates. (right) Embedding a bit '0' by pushing normals into the kernel area. Modified from [3].

2 Benedens's Method and Its Limitations

As mentioned earlier, Benedens's non-blind, geometry-based 3D watermarking method [3] uses the distribution of surface normals on polygonal meshes. The watermark embedding and retrieval procedures of Benedens's methods are briefly described in this section in order to familiarize readers with its basic operations. In Benedens's method, watermarks are embedded by modifying (i) the mean of normals, (ii) the mean angle of normals to a Bin Center (BC) normal, or (iii) the amount of normals in a bin. Our method is based on the third feature because it is most straight-forward to improve blindness and robustness with the amount of normals in a bin. We therefore focus on the third feature in this section. The entire embedding process of Benedens's method is summarized below:

1. Create a unit sphere, and then tessellate the surface of the unit sphere to generate bins defined by a Bin Center (BC) normal and a bin angle (φ^R) (in the following the bin angle will also referred to as bin radius). The same bin radius is used for all the bins. Bins are cone-shaped as illustrated in the left image of Figure 1.
2. Randomly choose a set of bins for embedding the watermark bits and sample surface normals. A surface normal is assigned to a bin if the angle formed between the surface normal (BP) and the Bin Center normal (BC) (Figure 1, left image) is smaller than that formed between the cone's axis (BC) and any line on the surface of the cone that passes through its apex.
3. Compute the ratio of normals (nk_i) inside the bin kernel defined by the kernel angle φ^k for each bin. The 2D projected kernel area is the gray colored inner circle as seen in the right image of Figure 1. The kernel radius (φ^k) is predefined.
4. Transform all the 3D surface normals in each bin into 2D coordinates in the X_1 and X_2 plane (See the left image of Figure 1) and perform the core embedding process as described below.

During the core watermark embedding process, watermark bits are inserted by changing the number of normals inside the kernel area in each bin. For instance, to embed a bit '0', all the normals outside the kernel are moved inside it as depicted in the right image of Figure 1. It means that nk_i (the ratio of normals inside the kernel) becomes 1.0, which is the maximum for any bin. Conversely, a bit '1' is embedded by taking all the normals inside the kernel out of the kernel so that nk_i goes to 0.0, which is the minimum for any bin. In order to search for the new locations of vertices that move the normals to the new directions, an optimization algorithm, the Downhill simplex method, is used to relocate the vertices. Optimization is performed according the two cost functions defined in Eqn. 1, respectively for bit 1 and 0.

$$costs_{f;v\to v'} \begin{cases} cos\left[\frac{\pi}{2}\left(\frac{cos^{-1}(BP'_{ij} * BC_i)}{\varphi_i^R}\right)\right] & S_i = 1 \\ \frac{cos^{-1}(BP'_{ij} * BC_i)}{\varphi_i^R} & S_i = 0 \end{cases} \tag{1}$$

In order to minimize distortions on the surface of the input model, the following constraints are imposed during the watermark embedding process:

- The normal of a face adjacent to a vertex v in the bin is not allowed to change by an angle that is $\geq\alpha$;
- The normal of a face adjacent to a vertex v that is not in the bin is not allowed to change by an angle that is $\geq \beta$;
- No normal is allowed to leave its bin.

For retrieval of the embedded watermarks, the information about bins (bin radius, kernel radius, the ratio of normals in each bin, and the chosen bins used in the embedding process) need to be delivered to the extraction stage. With the watermarked polygonal mesh, repeat the same steps (1 to 6) of the embedding process. Then the ratio of normals nk_i in each bin is compared with the original value of nk_i. If nk_i of the watermarked mesh is larger than the nk_i value of the original model, then the embedded watermark is a bit '0'. Otherwise, it's '1'.

As mentioned earlier, watermarks embedded by Benedens's method are especially robust against mesh-simplification and vertex randomization, because the distribution of surface normals is invariant to these two types of modifications to the polygonal mesh. Benedens's method therefore addressed two major drawbacks of most of the earlier geometry-based watermarking methods.

However, no evaluation against additive noise was conducted in Benedens's paper. Our tests indicated that about one tenth of the embedded bits were destroyed by Gaussian additive noise when its magnitude was set to cause visible, but not annoying, perturbation on the surface of the input mesh. This means that Benedens's method is not robust enough against additive noise. One trivial way to increase the robustness of Benedens's method is simply to increase Bin radius (φ^R) but that will lead to decrease watermark capacity. The other drawback of Benedens's method is its non-blindness because the original nk_is of all bins have to be carried to the extraction stage. This is not acceptable in many practical

applications. In the following two sections, we present our approaches to address these major drawbacks.

3 A Blind Version of Benedens's Method

Retrieval of the watermarks embedded by Benedens's method requires the availability of a priori knowledge including bin radius, number of bins, and the original ratio of normals (nk) in the kernel of each bin, to the extractor. This is the secret key needed for retrieval of the watermarks later on. Since the original values of nk depend on the polygonal mesh of the 3D object, Benedens's method is a non-blind watermarking technique. Our work first focused on eliminating the need to carry the original nk values to the extracting stage by using the probability distribution of normals in the kernel area of each bin. The key idea is to choose a kernel radius such that the ratio of surface normals inside the kernel is a fixed value with the assumption that normals are uniformly distributed over the object's surface and hence probabilities correspond to area ratios. Then, the retrieved watermark bit is 0 if the nk value of the watermarked mesh is greater than the fixed ratio at the extractor. Otherwise, it is 1. Therefore, all we have to do is to compute the exact kernel radius φ^k that satisfies the constraint that the ratio of normals inside the kernel area is fixed. A ratio of 0.5 was chosen so that there were equal number of surface normals inside and outside the kernel.

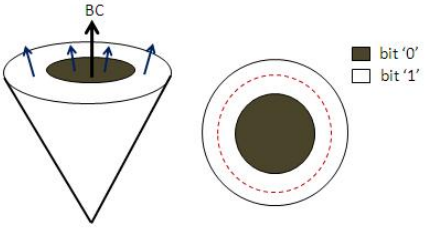

Fig. 2. Two views of a bin with sampled normals. The dark area (inner circle) is the kernel area defined by φ^k, which is used for embedding '0'; the bin excluding the dark area is used for embedding '1'. The dashed circle represents the new zone for embedding '1'.

Given a spherical cap defined by a sphere of radius R and a height H from the top of the spherical cap to the bottom of the base circle, its surface area can be calculated as $A_{spherical\,cap} = 2\pi RH$.
Since $H = R\text{-}Rcos\varphi^R$ for a bin defined by φ^R, the surface area of the bin becomes

$$A_{spherical\,cap\,of\,bin} = 2\pi R^2(1 - cos\varphi^R). \tag{2}$$

The surface area of the kernel defined by φ^k can be calculated similarly as

$$A_{spherical\,cap\,of\,kernel} = 2\pi R^2(1 - cos\varphi^k). \tag{3}$$

We require that

$$\frac{A_{spherical\,cap\,of\,kernel}}{A_{spherical\,cap\,of\,bin}} = 0.5. \tag{4}$$

Therefore the size of the kernel, φ^k, can be computed from Eqns. 2, 3, and 4.

$$\varphi^k = cos^{-1}(1 - \frac{1}{2}(1 - cos\varphi^R)). \tag{5}$$

This way, the original nk values no longer need to be carried to the extraction stage, thus achieving a blind version of Benedens's method.

An additional modification we have made to Benedens's method was to improve the robustness of '1's bits. During the embedding process of the original method, the normals are moved in two opposite directions. When embedding a bit '0', all normals in the bin are moved inside the kernel area (the dark inner circle shown in Figure 2). For better robustness, the normals should be enforced to be as close to the BC line (bin center normal) as possible. When embedding a bit '1', however, the normals are moved towards the border curve and are pushed as closed to the rim of the bin as possible. There are therefore two imaginary embedding zones: one around the BC and the other around the rim of the bin. Ideally, all normals should be placed at either the BC (for bit '0') and one on the rim of the bin (for bit '1'). The problem, however, is that the normals located on the rim of the bin can be easily pushed out of the bin. This is not the case with the normals located at the BC. Therefore, bit 1 is much less robust than bit 0, which is an undesirable feature.

To verify the difference in robustness of bit '0' and bit '1', an experiment with the Bunny model was conducted. We first embedded a sequence of '0' and '1' (20 bits in total), then additive noise following a Gaussian distribution N(0, σ=0.005) was added, and the robustness of bits '0' and '1' were as compared. This experiment was repeated ten times and the results averaged. As expected, the error of the bit '1' (15 %) was higher than that of bit '0' (11 %).

To improve the robustness of bit '1', the ideal embedding zone was moved away from the bin rim, as shown by the dashed circle on the right of Figure 2. The new embedding zone defined by the red dashed circle was defined by a new radius φ^{k1} such that the surface area of the spherical cap is 3/4 of that of the bin. By using Equations 3 to 4, we have:

$$\varphi^{k1} = cos^{-1}(1 - \frac{3}{4}(1 - cos\varphi^R)). \tag{6}$$

With this new embedding zone for bit '1', the error of bit '1' was reduced to 10%. In general, the robustness of bit '0' is expected to be superior to that of bit '1', especially under the condition of strong attacks, because the new embedding zone for bit '1' is still closer to the rim of the bin as compared to the BC which is the embedding zone of bit '0'.

4 Roughness-Adaptive 3D Watermarking

In this section, we present a novel way of adaptively selecting watermark strength that optimizes the conflicting requirements of robustness and imperceptibility. It is well known that stronger watermarks improve robustness while weaker watermarks ensure imperceptibility. The key to optimizing watermark strength is to choose the strongest watermark with the constraint that it remains invisible. The optimization can be performed globally or locally based on certain properties of the polygonal mesh. We propose an adaptive algorithm that selects watermark strengths based on local surface roughness measures. Our algorithm takes advantage of the fact that the human eyes are more sensitive to distortions of smooth surface patches than to those of rough surface patches. The watermark strength is therefore set to the human detection threshold; i.e., the largest watermark that remains invisible, for a given surface roughness. In the remainder of this section, we first describe the psychophysical experiment that estimated the quantitative relation between human watermark detection threshold and local surface roughness. We then present our roughness-adaptive watermarking algorithm.

4.1 Psychophysical Experiment on Human Detection Thresholds

The psychophysical experiment was designed to estimate the relation between visual watermark detection threshold (in terms of watermark strength δ used to restrict the search space in the core embedding process) and roughness of spherical surfaces. Note that the original Beneden's method restricted the search space in the Downhill Simplex optimization method based on the minimum of four parameters including δ. For simplicity and without loss of generality, we used only the δ parameter in our experiments. In the following, δ is expressed as a ratio over the bounding box diameter of the 3D object. Four participants (three males and one female) took part in the experiment. None of the participants reported any visual deficiencies.

Stimuli. The visual display consisted of a spherical surface rendered with 3752 vertices and 7500 faces. The image of the sphere occupied a visual angle of roughly 30 degrees. Six spherical surfaces with different roughness levels were created by introducing varying amount of additive noise to the vertices. The surface roughness level was controlled by the magnitude of the additive noise generated by a Gaussian probability distribution function $N(0,\sigma)$. The direction of the additive noise was randomly chosen between 0 and 360 degrees. Figure 3 shows the six reference stimuli with increasing surface roughness. The left most sphere has a smooth surface with no additive noise. The roughness level of each spherical surface was estimated with a 1-ring roughness measure based on the mutli-scale roughness estimation method proposed by Corsini et al. in [13]. They were 0.000027, 0.000209, 0.001704, 0.003139, 0.010937, and 0.025304, respectively.

Fig. 3. The six spherical surfaces with increasing roughness used in the present study. The estimated roughness values were 0.000027, 0.000209, 0.001704, 0.003139, 0.010937, and 0.025304, respectively, from left to right. The six spherical surfaces contained the same number of vertices (3752) and faces (7500).

In order to make our experiment as general as possible, we decided not to use any particular watermarking technique. Instead, we measured the visibility of a general noise-like visual stimulus on a surface with controlled roughness.

Additive noise was used to randomly alter chosen vertices of the sphere in the direction of surface normal, as specified below.

$$\overrightarrow{v}_{new}(i) = \overrightarrow{v}(i) + \delta * \overrightarrow{n}(i)$$

where $\overrightarrow{v}_{new}(i)$ is the modified coordinate vector of the i-th vertex $\overrightarrow{v}(i)$, δ denotes stimulus intensity that varied according to the participant's responses, and $\overrightarrow{n}(i)$ is the surface normal vector[1]. In our experiment, the amount of modified vertices was fixed to 3% so that images of spheres with embedded noise could be rendered in a reasonable amount of time from trial to trial.

Procedures. A two-interval forced choice (2IFC) one-up three-down adaptive procedure [15] was used to measure the (watermark) detection thresholds as a function of surface roughness. The threshold so obtained corresponded to the 79.4 percentile point on the psychometric function. On each trial, the participant looked at two spherical surfaces, a reference surface (no embedded noise) and a target surface (reference plus noise), presented on the left and right sides of the monitor. The location of the surface with embedded noise was randomly chosen to be on the left or right of the monitor on each trial. The participant's task was to indicate which spherical surface contained noise. According to the one-up three-down adaptive rule, the stimulus intensity (δ) was increased after a single incorrect response and decreased after three successive correct responses. Otherwise, the value of δ remained the same on the next trial. The initial δ value was chosen to be large enough so that the stimulus was clearly perceptible to the participant. The value of δ then decreased or increased by a fixed step size (0.0004), depending on the participant's responses. After the initial three reversals (a reversal occurred when the value of δ decreased after one more increase, or vice versa), the value of δ changed by a smaller step size (0.0002). The initial larger change in δ was necessary for faster convergence of the δ

[1] Though in the experiments the stimulus has a form that does not directly correspond to the watermark distortion, we used the same symbol δ, since the results of the experiments will be used to set the watermark strength.

values, whereas the later smaller change in δ improved the resolution of threshold estimates. The adaptive series was terminated after 12 reversals at the smaller step size. The detection threshold was computed by taking the average of the δ values from the last 12 reversals. Each participant was tested once per surface roughness, resulting in a total of six adaptive series per participant and a total of 24 series for all participants.

Results. The average detection thresholds for the four participants are shown in Figure 4. The thresholds follow a monotonically increasing trend that is well modeled by the following power function,

$$\delta = 0.0291 * S^{0.3907}, \tag{7}$$

where δ is the detection threshold, and S denotes surface roughness.[2]

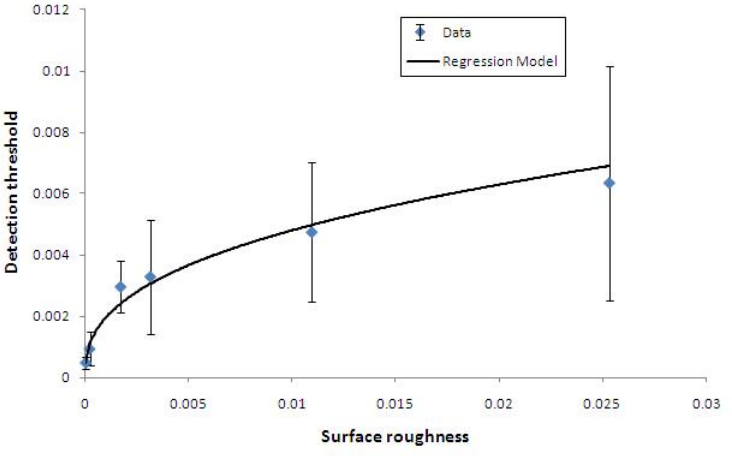

Fig. 4. Human watermark detection threshold data and a power regression model (solid line). See texts for details.

4.2 Selection of Adaptive Watermark Strengths

The results of the psychophysical experiment indicated that a stronger watermark can be hidden into a bumpier surface area with higher roughness. Specifically, the human detection threshold for watermarks is related to the local surface roughness as shown in Eqn.7. So our goal was to use an adaptive watermark strength determined by the local roughness of a surface instead of the constant watermark strength used in Benedens's method. Specifically, we first

[2] The r^2 value, which indicates how well a regression model approximates real data points, was 0.9851, where r^2=1.0 means a perfect fit.

estimated the roughness per vertex, and then chose the maximum imperceptible watermark strength as

$$\begin{cases} \delta & = 0.0291 * S^{0.3907} - constant, \;\; if \; S \geq 0.000027 \\ \delta & = 0.0004, \;\; otherwise \end{cases} \tag{8}$$

where the *constant* (0.00004) was used to uniformly reduce the watermark strength so that the embedded watermark strength was not detectable. When the surface roughness is less than 0.000027 which corresponds to the smooth spherical surface, the watermark strength was fixed to 0.0004.

To estimate the local surface roughness around the vertex to be modified for embedding watermarks, we estimated the roughness of all adjacent faces around the vertex using the 1-ring roughness estimation method described in Corsini et al. [13]. Then the value of δ was determined by Eqn. 8 during the watermark embedding process. Note that watermark strength is directly related to δ, since δ defines the radius of the spherical region wherein the new watermarked vertex location is looked for, during the Downhill simplex optimization in Benedens's method. The resulting distance between the old vertex and the new optimized vertex becomes the watermark strength which is linearly proportional to δ. In the present study, bin radius φ^R was fixed at 11 degrees as a compromise between robustness and capacity.

The key steps in our roughness-adaptive watermarking algorithm can be summarized as follows: (1) Measure the local surface roughness per vertex by Corsini et al.'s method [13]; (2) Compute the maximum δ from Eqn. 8; (3) Apply the chosen δ value to in the watermark embedding process; and (4) Repeat steps 1 to 3 for all the vertices to be watermarked.

5 Evaluations

The experiments that we carried out were aimed at evaluating the improvements brought by the modifications we made to the original Benedens's watermarking scheme, with particular reference to roughness-based watermark strength adaptation. We focused our assessment by comparing the watermark strengths between the original and the watermarked models. Note that by increasing

Table 1. A comparison of error to additive-noise attacks. Shown under the two methods are the averages and standard deviations over 10 tests. See texts for details.

Model	Benedens' method		Our method		Improved (%)
-	Error (%)	Std. Dev.	Error (%)	Std. Dev.	-
M1	11	4.95	9.5	4.38	13.6
M2	9	5.68	7	6.32	22.2
M3	8	3.5	7.5	6.77	6.25

Table 2. Key parameters of the 3D models used in the evaluation experiments

Model	# of vertices	# of faces	Avg. roughness	Std. Dev.
Bunny (M1)	5050	9999	1.888E-07	1.754E-07
Happy Buddha (M2)	4952	9932	7.519E-07	4.476E-07
Dragon (M3)	3512	6999	1.3302E-06	8.842E-07

the magnitude of the watermarks in order to improve robustness, the watermarked model will necessarily have higher MSE (mean-square error), under the same condition of watermark invisiability. To simulate random attacks, we used additive Gaussian noise to alter vertex locations. Note that Benedens did not evaluate his original method against additive noise, so the data reported in this section provide additional information on the robustness of Benedens's method. Specifically, we compared the performance of the original Benedens's method with the performance obtained after our proposed improvements. To do so, we implemented the original Benedens's method and our improvements with Visual C++ with CGAL and OpenGL libraries on a Windows machine with a 2GHz CPU processor. Three 3D models, "Bunny", "Happy Buddha" and "Dragon" were used for the evaluation experiments. The key characteristics of the three models are summarized in Table 2. In what follows, we present two experiments. The first experiment was an evaluation of the proposed improvement to bit '1' robustness. The second experiment measured the additional improvements due to the roughness-adaptive watermarking algorithm.

5.1 Improvements to Bit '1' Robustness

The purpose of the first experiment was to evaluate the improvement to robustness when the new embedding zone (φ^{k1}=9.53 degrees, φ^{R} = 11degrees) computed by Eqn. 6 for '1' bit was introduced. Watermark strength (δ value) was set to a constant value during the embedding process; e.g., 0.0064 for M1 (see Table 3) . In order to evaluate robustness, a 20-bit watermark was embedded. Then attacks simulated with additive noises generated by a Gaussian distribution with zero mean and standard deviation σ=0.0005 was applied to the watermarked model. Error was calculated as the percentage of bits that failed to be retrieved. The error was computed for both the original Benedens's method and our robust bit '1' method. For each test condition the experiment was repeated ten times and the results were averaged. As shown in Table 1, our method which improved bit '1' robustness outperformed the original Benedens's method for all three models tested. The average improvement, calculated as the ratio of reduction in error over the error for Benedens's method, was about 14 %.

5.2 Additional Improvements Due to Roughness-Adaptive Watermarking

In the second experiment, we implemented both the improvement to bit '1' robustness and roughness-based strength adaptation. Attacks were simulated by additive noises generated with a Gaussian distribution (N(0, σ=0.0005)). The performance of our blind and improved method and Benedens's original method were compared in terms of their robustness against additive noises for a 20-bit watermark. For the experiment on the original Benedens's method, the parameters were set as follows: φ^R = 11 degrees, φ^k= 9.53 calculated by Eqn. 5 and δ=0.0064, 0.007, and 0.0059 for M1, M2, and M3, respectively (i.e., constant watermark strength). The constant δ values were chosen by finding the maximum values ensuring invisibility. For the experiment on our method, the same φ^R and φ^k values were used, but adaptive δ values were selected based on the estimated local surface roughness. The experiment with each method was repeated ten times. The resulting images for three models are shown in Figure 5.

Table 3. A comparison of errors to additive-noise attacks for Beneden's method and our new method. See texts for details.

Model	Benedens' method			Our method			Improved (%)
-	Error (%)	Std. Dev.	δ	Error (%)	Std. Dev.	δ	-
M1	11	4.95	0.0064	6.5	3.37	0.00753	40.9
M2	9	5.68	0.007	5	3.33	0.0129	44.4
M3	8	3.5	0.0059	3.5	4.12	0.0116	56.3

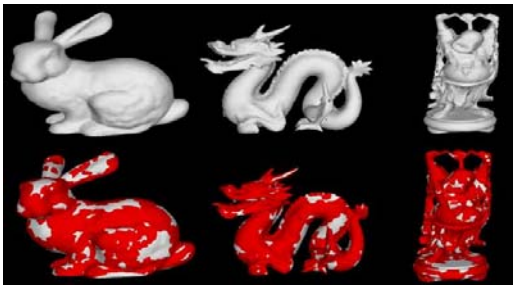

Fig. 5. Resulting watermarked models. Original models (upper row), and the watermarked models (lower row) indicating the modified triangles (red colored triangles). Bunny, Dragon, and Happy Buddha models from the left.

Table 3 shows the results with the original Benedens's method using a constant δ and our new method utilizing adaptive δ values. It is clear that watermarks embedded by our adaptive δ method is more robust against additive-noise attacks for all three models (M1, M2, and M3). The robustness improvement rate

achieved by our method was 56.3% with the Dragon model, 44.4% with the Happy Buddha model, and 40.9% with the Bunny model. The largest improvement observed with the Dragon model was to be expected because the standard deviation of the estimated surface roughness for the Dragon model was the largest among the three models tested (see Table 2). It was also found that the average δ size was larger with our method than with the constant used in Benedens's method, more so with the Dragon model than with the Bunny model (see Table 3). Therefore, the model with the largest variation in surface roughness (the Dragon model) benefited most with our method.

6 Concluding Remarks

Developing robust 3D digital watermarking techniques is an ongoing challenging research topic in the field of information hiding. In this paper, we have presented a general way to improve watermark robustness by exploiting visual masking. Our method is based on measured human sensitivity to additive noise as a function of surface roughness of input meshes. Similar approaches have been previously adopted by Bors et al. [6,7] and by Kanai et al. [14]. Bors et al. embedded watermarks into a region where the human eyes are less sensitive to changes. Kanai et al. embedded watermarks by modulating the high frequency signals in an input model where changes are imperceptible to human vision. While both Bors et al. and Kanai et al. focused mainly on watermark imperceptibility, we aimed to improve watermark robustness while maintaining imperceptibility. In addition, these researches used a constant watermark strength in surface areas where humans are less sensitive to watermarks, and they did not embed watermarks in surface areas where humans can detect watermarks more easily. In comparison, our method allows varying watermark strengths per vertex to maximize robustness while maintaining imperceptibility. The watermark strengths were adapted based on local surface roughness, thereby utilizing the whole object surface for watermark embedding. The evaluation experiments of the new scheme with Benedens's method confirmed that the overall watermark robustness was improved greatly by employing roughness-adaptive watermark strengths. We achieved further improvement to robustness by introducing a more robust embedding zone for bit '1', and we introduced a blind version of Benedens's method by choosing kernel radii that resulted in a constant ratio of surface normals inside the kernels. As expected, the experiments demonstrated that the roughness-adaptive watermarking technique brings more benefits for data models with a larger standard deviation of roughness. We showed that, on average, stronger watermarks can be embedded with roughness-adaptive watermark strengths than could be achieved with a constant watermark strength as used by most watermarking methods.

Our approach suggests promising new directions for improving the performance of 3D digital watermarking schemes. In the future, we will test our scheme with other existing geometry-based techniques to verify its generalizability and to compare the results (e.g., improvements in robustness and MSE).

Acknowledgments

The first author (KK) was partially supported by the US National Science Foundation under Grant no. 0836664. The second author (MB) was partially supported by the Italian Ministry of Research and Education under FIRB project no. RBIN04AC9W.

References

1. Ohbuchi, R., Masuda, H., Aono, M.: Watermarking Three-Dimensional Polygonal Models. In: Proceedings of the ACM International Multimedia Conference and Exhibition, pp. 261–272 (1997)
2. Wagner, M.G.: Robust Watermarking of Polygonal Meshes. In: Proceedings of Geometric Modeling and Processing, pp. 201–208 (2000)
3. Benedens, O.: Geometry-Based Watermarking of 3D models. IEEE Computer Graphics and Applications 19(1), 46–55 (1999)
4. Benedens, O., Busch, C.: Towards Blind Detection of Robust Watermarks in Polygonal Models. Computer Graphics Forum 19(3), 199–208 (2000)
5. Yeo, B., Yeung, M.M.: Watermarking 3D Objects for Verification. IEEE Computer Graphics and Applications 19(1), 36–45 (1999)
6. Bors, A.G.: Blind Watermarking of 3D Shapes using localized constraints. In: Proceedings of 3D Data Processing, Visualization and Transmission 2004, 3DPVT, pp. 242–249 (2004)
7. Bors, A.G.: Watermarking Mesh-Based Representations of 3-D Objects Using Local Moments. IEEE Transactions on Image processing 15(3), 687–701 (2006)
8. Cayre, F., Macq, B.: Data Hiding on 3-D Triangle Meshes. IEEE Transactions on Signal Processing 51(4), 939–949 (2003)
9. Cayre, F., Rondao-Alface, P., Schmitt, F., Macq, B., Maitre, H.: Application of spectral decomposition to compression and watermarking of 3D triangle mesh geometry. Signal Processing 18(4), 309–319 (2003)
10. Ohbuchi, R., Takahashi, S., Miyazawa, T., Mukaiyama, A.: Watermarking 3D Polygonal Meshes in the Mesh Spectral Domain. In: Proceedings of Graphics interface, pp. 9–17 (2001)
11. Lin, H.-Y., Liao, H.-Y., Lu, C.-S., Lin, J.-C.: Fragile watermarking for authenticating 3-D polygonal meshes. IEEE Trans. Multimedia 7(6), 997–1006 (2005)
12. Zafeiriou, S., Tefas, A., Pitas, I.: Blind Robust Watermarking Schemes for Copyright Protection of 3D Mesh Objects. IEEE Trans. Visualization and Computer Graphics 11(5), 596–607 (2005)
13. Corsini, M., Gelasca, E.D., Ebrahimi, T., Barni, M.: Watermarked 3D Mesh Quality Assessment. IEEE Transactions on Multimedia (2007)
14. Kanai, S., Date, H., Kishinami, T.: Digital Watermarking for 3D Polygons Using Multiresolution Wavelet Decomposition. In: Proceeding of IFIP WG, pp. 296–307 (1998)
15. Levitt, H.: Transformed Up-Down Methods in Psychoacoustics. The Journal of the Acoustical Society of America 49, 467–477 (1971)