Perceptually Driven Interactive Geometry Remeshing

Lijun Qu* Gary W. Meyer†

Computer Science and Engineering Department and Digital Technology Center
University of Minnesota

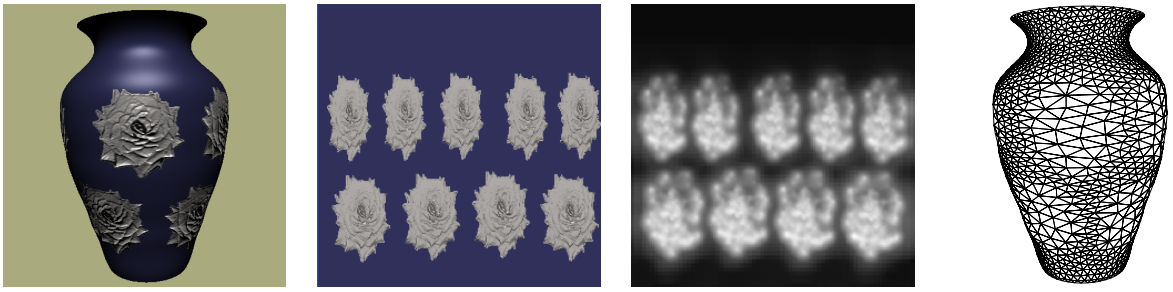

(a) (b) (c) (d)

Figure 1: Bump mapped vase (a) created using a normal map and a vase model. The shading calculation transforms the normal map into a color pattern which is gathered into a color map (b). The perceptual properties of the color map are then evaluated using a visual discrimination metric. The brighter region in the map (c) indicates stronger visual masking. This map is then used to guide the placement of vertex samples (d) in the geometry remeshing stage. The original vase has 7171 vertices, and the remeshed vase has 2000 vertices.

Abstract

Visual patterns on the surface of an object, such as two dimensional texture, are taken into consideration as part of the geometry remeshing process. Given a parameterized mesh and a texture map, the visual perceptual properties of the texture are first computed using a visual discrimination metric. This precomputation is then used to guide the distribution of samples to the surface mesh. The system automatically distributes few samples to texture areas with strong visual masking properties and more samples to texture areas with weaker visual masking properties. In addition, due to contrast considerations, brighter areas receive fewer samples than do darker surface features. Because of the properties of the human visual system, especially visual masking, the artifacts in the rendered mesh are invisible to the human observer. For a fixed number of polygons, this approach also improves the quality of the rendered mesh since the distribution of the samples is guided by the principles of visual perception. The utility of the system is demonstrated by showing that it can also account for other observable patterns on the surface, besides two dimensional texture, such as those produced by bump mapping, lighting variations, surface reflectance, and interreflections.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—surface and object representations

Keywords: Geometry Remeshing, Mesh Simplification, Level of Detail, Visual Masking, Visual Perception

*e-mail:lijun@cs.umn.edu
†e-mail:meyer@cs.umn.edu

1 Introduction

Surface signals, including texture maps, bump maps, surface reflectance, environment maps, can have a dramatic impact on the appearance of a polygon mesh. Today these surface signals are used to produce striking visual effects at little cost by employing the texture mapping and pixel shading hardware available on PC graphics cards. There has been a considerable amount of work in the field of computer graphics on the creation, processing, and usage of these surface signals.

Surface signals have also been used to accelerate global illumination algorithms [Ramasubramanian et al. 1999], to compress the texture map [Balmelli et al. 2002], and to generate a specialized signal parametrization [Sander et al. 2002]. However, very little work in the area of geometric modeling has taken surface signals into consideration.

Some researchers have noticed that surface signals can be useful in the area of geometric modeling. Ferwerda et al. [1997] developed a visual masking model for computer graphics and observed that visual masking can have an impact on the representation of geometric models. The existence of surface signals on a geometric model can raise the geometric error threshold for tessellating the surface. This elevated threshold should be taken into account by algorithms that determine a simplified representation for a geometric model, such as mesh simplification algorithms and surface remeshing algorithms. Figure 2 shows a flat shaded cylinder with and without texture. Faceting artifacts can be seen clearly in the left image, but no faceting artifacts can be seen in the right image. For the cylinder with texture a more coarse geometric representation would suffice. In this paper, we propose an algorithm that samples the geometric mesh taking into account the surface signals on the mesh.

We are particularly interested in the visual perceptual properties of the surface signals. Most surface remeshing algorithms distribute samples on the surface according to the geometric properties of the mesh, such as curvature information. In this paper, the distribution of samples is guided both by the geometric properties of the mesh as well as the perceptual properties of the surface signals.

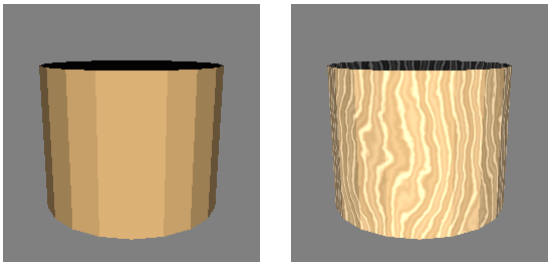

Figure 2: Flat shaded cylinder without (left) and with (right) texture (after Ferwerda et al. 1997).

This paper makes contributions in the following areas:

1. We have extended the current state of the art in perceptually based level of detail algorithms to include visual masking. Visual masking requires multiscale and multiorientation decomposition of the image and is difficult to determine in iterative edge collapse mesh simplification algorithms.
2. We propose a new method that can compute the visual perceptual properties of the surface signal based on the Sarnoff visual discrimination metric: a contemporary vision based model that takes advantage of threshold-vs-intensity, contrast sensitivity, and visual masking.
3. Our remeshing algorithm accounts for geometric properties as well as the visual perceptual properties of the surface signals on the mesh. Specifically, our algorithm incorporates the visual perceptual properties of the surface signal into the remeshing framework proposed by Alliez et al. [2002].

The remainder of the paper is organized as follows: Section 2 reviews some of the previous work in surface remeshing and the application of visual perception to geometric modeling. We then introduce our algorithm in section 3. Next we provide a simple but novel approach to pre-compute the visual perceptual properties of the surface signal in section 4. The remeshing process is discussed in section 5. We then improve the remeshing results by considering other major features of the surface signal in section 6. We then discuss in detail the difference between our work and other previous work, and draw some conclusions.

2 Previous Work

In this section, we present some previous work in the areas of perceptually guided level of detail, geometry remeshing, and the relationship between the surface signal and geometry.

2.1 Perceptually Based Level of Detail

Lindstrom and Turk [2000] propose image driven simplification in which the importance of each edge is weighted according to the root mean square image difference, not the geometric difference, it makes when deleted. They demonstrate that image driven simplification can produce results equal to or better than most geometry based mesh simplification algorithms. Lindstrom [2000] employs a perceptually motivated metric in a mesh simplification algorithm.

Luebke and Hallen [2001] propose perceptually driven mesh simplification that controls the simplification using psychophysical models of visual perception. They map an edge collapse operation to the worst contrast grating introduced by the edge in question. They later extended their work to the simplification of lit, textured meshes [Williams et al. 2003]. However, visual masking is not included. Our paper has the same goal as their research. However, we use a contemporary model of the human visual system which includes threshold-vs-intensity, contrast sensitivity and visual masking.

A more recent work by [Lee et al. 2005] introduces the idea of mesh saliency as a measure of regional importance for computer graphic models and integrates this information into a mesh simplification algorithm. It is important to point out that this work only considers mesh saliency for geometry without textures. As was mentioned above, diffuse color textures, normal maps, and lighting patterns can each have a dramatic impact on the appearance of computer graphic models and can thus change the regional saliency of the geometric models. It would be interesting to take these factors into account during mesh simplification.

2.2 Geometry Remeshing

With the advance of model acquisition techniques, there has been a considerable amount of work in the area of surface remeshing. We list publications most significant to our own work. Please see [Alliez et al. 2005] for a recent extended review. Alliez et al. [2002] propose a novel interactive technique that first partitions the model into patches homeomorphic to disks, and then parameterizes each patch over a planar domain. Most of the remeshing operations can then be performed in the 2D parametric domain instead of 3D. In their recent work on anisotropic remeshing [Alliez et al. 2003a], they show that sampling along the principle curvature directions can produce compactly represented meshes. Some researchers have taken another approach to surface remeshing by working directly on the 3D mesh. Turk [1992] designs an elegant algorithm that positions vertices by point repulsion. More recent work [Surazhsky and Gotsman 2003] employ a series of local operations to improve the mesh quality.

2.3 Surface Signals and Geometry

Most meshes come with surface signals. However, the majority of existing work either considers the problem of surface remeshing without taking surface signals into account, or attacks the problem of construction, manipulation and optimization of surface signals without incorporating the geometry. Until recently, there has been almost no work that includes both the geometry and the surface signal. [Sander et al. 2002; Tewari et al. 2004] propose signal-specialized surface parametrization that minimizes the signal stretch instead of the usual geometry stretch and show that the signal-specialized parametrization can improve the image quality due to less texture stretch. Carr and Hart [2004] design an interactive painting system that dynamically adjusts the parametrization of the geometry according to the frequency content of the texture painted on the surface. Their system allocates more texture space to high frequency texture regions, thus preserving the details of the texture during rendering.

3 Algorithm Overview

The input to the remeshing algorithm is a parameterized triangulated mesh and several surface signals that have accumulated on

the mesh during rendering. First, our algorithm generates the composite surface signal from several surface signal sources. Our algorithm then analyzes the perceptual properties of the composite surface signal using the Sarnoff visual discrimination metric. Third, the surface mesh is converted to a map based representation, and the geometry remeshing process is treated as a 2D sampling process based on an importance map. Finally, a Delaunay triangulation operation is performed on these samples. These samples and their connectivity are re-projected back to 3D to form a 3D mesh.

We introduce the algorithm by showing how it can account for a single type of observable surface signal: the color pattern produced by two dimensional texture mapping. Near the end of the paper we will broaden the definition of the surface signal to include the effect of such things as bump mapping, spotlighting, shadow patterns, and interreflections. We will also demonstrate how all of the effects included in this general definition of the surface signal can be accommodated using the same procedures developed to handle two dimensional texture mapping.

4 Computing the Visual Masking Map

4.1 Visual Discrimination Metric (VDM)

Visual discrimination metrics have been created to assist in the design and evaluation of imaging systems. These metrics include current knowledge of the human visual system and are designed for physiological plausibility. The computer graphics community is already familiar with these metrics. For example, visual discrimination metrics have been used to place samples adaptively into areas of the image plane that are visually more important [Bolin and Meyer 1998] and to choose a global illumination algorithm from a pool of global illumination algorithms [Volevich et al. 2000]. There has also been work that employs the separate stages of the visual difference metric to speed up rendering algorithms. A recent effort is [Dumont et al. 2003].

In this work, the Sarnoff VDM [Lubin 1995] is used to compute the visual perceptual properties of the texture. The Sarnoff VDM consists of five major components: optics and resampling, bandpass contrast responses, oriented responses, transducer and distance summation. The optics and resampling stage incorporates the optics of the human visual system and models how the rods and cones in the human visual system sample real world images. The bandpass contrast stage models the frequency selectivity of the human visual system, including the decomposition, using image pyramid algorithms, of the original images into seven bandpass images with peak frequencies from 32 through 0.5 cycles/degree. The oriented responses stage models the orientation selectivity of the human visual system. During this stage the images are filtered by a set of filters with different orientations. The transducer stage does the normalization and models the visual masking function of the human visual system. Finally, the distance summation computes the visual difference between the two input images.

Visual discrimination metrics have some limitations that prevent them from being used more widely in computer graphics. First, these metrics are in general slow to compute, which makes it difficult to use them in interactive or realtime computer graphics applications. Second, these metrics were originally designed to take two images at input, but only one image is available in many computer graphics applications. The first limitation has been recently addressed by [Windsheimer and Meyer 2004]. We will address the second limitation in the next section.

4.2 Visual Masking Map Computation

The ability of a base visual stimulus to increase the visibility threshold of a test visual stimulus is called visual masking. The base visual stimulus is sometimes referred to as the masker, and the test visual stimulus is called the signal.

In this paper we want to compute the visual masking properties of a texture and use the results of this computation to guide a surface remeshing algorithm. Since the masking ability of the texture correlates strongly with the spatial frequency, contrast, and orientation of the test stimulus, any visual masking computation is not theoretically correct without considering the test stimulus itself. However, a well designed algorithm based on models of the human visual system can still provide valuable information about the visual masking potential of a texture. There is some previous work in this area. Walter et al. [2002] compute the visual masking properties of a texture using aspects of the JPEG image compression standard. Ramasubramanian et al [1999] propose a novel method to compute the visual masking properties of a texture by handling the luminance dependent processing and spatially dependent processing separately and then combining them in an appropriate manner. We have taken a similar approach as in [Ramasubramanian et al. 1999].

We propose to compute the visual masking properties of a texture using the Sarnoff VDM. This allows us to take advantage of the accumulated experience and robustness that is built into this metric. Since the Sarnoff VDM takes two images as input, we need to have a second comparison image to feed in as input along with the original texture. Some researchers have tried novel ways to derive the second image or both images. Bolin and Meyer [1998] determine two candidate images while ray tracing by using current estimates of the mean value and variance at each pixel. Volevich et al. [2000] employ two intermediate global illumination solutions as input to the visual discrimination metric.

According to Fourier theory, a texture can be decomposed into multiple frequencies. Since any nonzero signal frequency potentially causes visual masking, we can remove all nonzero frequencies from the original texture and compare the resulting image (basically, just the DC component) to the original texture. Since the Sarnoff VDM employs contemporary models of the human visual system, given the original texture and the DC component of the texture, the Sarnoff VDM will pick out visual differences for any nonzero frequencies in the original texture.

This approach would work if the original texture had similar intensity values across the texture. However, in general this is not true for real world textures which have very nonuniform intensities. Regions with different intensities will be averaged together and they can not be handled well by this approach. To solve this problem, some low frequencies in the original texture are allowed in the second comparison image to preserve the local average of the texture. Allowing some low frequencies in the second comparison image doesn't cause significant error in the final visual masking map because frequencies close to zero have relatively weaker visual masking compared to higher frequencies. Figure 3 shows the visual masking caused by frequencies without considering the threshold-vs-intensity function of the human visual system. Notice that the right image correctly shows the visual masking caused by the step function in the original texture, while the middle image incorrectly shows visual masking occurring across almost the entire texture.

This can be implemented efficiently by low pass filtering the original texture. In our implementation, we have used a Gaussian filter to remove most of the high frequencies. It is important to choose the right filter kernel size to filter the texture. If the filter only removes

(a) (b) (c)

Figure 3: Image (a) is the original texture. Image (b) is the visual masking map computed by using only the DC component of the original texture as the second comparison image. Image (c) is the visual masking map using low pass filtering of the original texture as the second comparison image.

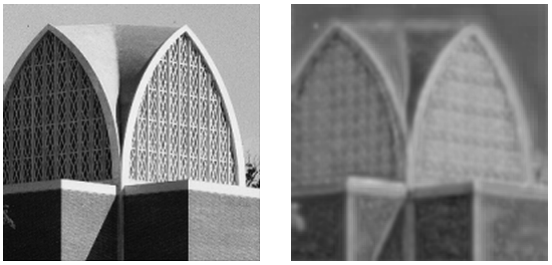

Figure 4: Image on the left is the original chapel image, image on the right is the visual masking map. Brighter region of the visual masking map indicates stronger visual masking.

a small portion of the high frequencies in the original image, the visual masking caused by those frequencies left out in the second image will not show up in the final visual masking map. Therefore, the visual masking caused by those frequencies cannot be utilized in the remeshing algorithm. On the other hand, if the filter removes too many of the frequencies, the problem shown in the middle image of Figure 3 will occur. The above mentioned problem has more chance to happen if the image contains irregular intensity regions and these regions have sharp boundaries (see Figure 3(a)). Hence, the optimal kernel size is a function of image content. In our implementation we have experimentally selected a filter kernel of size 15 (one degree of vision for our viewing distance and display dot pitch), and it works well for all the examples shown in this paper.

The threshold-vs-intensity function gives the error detection threshold corresponding to a given luminance background. To compute the error threshold described by the threshold-vs-intensity function, we have used the piecewise approximation described by [Ferwerda et al. 1996]. To get the final visual masking map, we use a linear combination of these two maps. Note that we combine the results differently from [Ramasubramanian et al. 1999] since they compute an elevation map in the second step (which in our case is the just noticeable difference map, a kind of error threshold). Figure 4 shows a chapel image and the final visual masking map generated by our algorithm. Notice that the window of the chapel shows stronger potential for visual masking while the background shows less possibility of visual masking. In addition, the right window shows visual masking is more likely than the left window because it has a higher base luminance level.

Since the Sarnoff visual discrimination metric has been designed for physiological plausibility and has been verified by a number of applications, our approach is simple but has a strong underlying foundation.

4.3 Computing the Reflected Surface Signal

We propose to compute the reflected surface signal in the parametric domain. This parametric domain approach has proven to be useful in many applications [Gu et al. 2002; Nguyen et al. 2005; Yuan et al. 2005; Wang et al. 2005]. The advantage to working in the parametric domain is that a 3D problem is converted into a 2D problem and thus a simpler problem is solved. In addition the tools for image processing are available for the problem at hand. For example, in our case we can use the Sarnoff VDM which is a 2D algorithm. Without mapping into the parametric domain it would be difficult, if not impossible, to adapt this tool.

As mentioned previously, what matters is the reflected surface signal that reaches the observer's eye. We need to compute the reflected surface signal. While the reflected surface signal is a three dimensional entity and is view-dependent, our modified visual discrimination metric only handles two dimensional surface signals. We therefore need to unwrap the 3D surface signal into a 2D map. Conceptually this can be done by pointing an array of cameras at the model, collecting all the reflected signal, and unwrapping the signal into a plane. Since we assume we have a parameterized model, we can unwrap the surface signal by rendering the model into a 2D map using texture coordinates instead of the original vertices as vertex positions. Moreover, we can unwrap several surface signals into one composite surface signal, analyze the visual perceptual properties of the composite surface signal, and use the result to guide the remeshing process.

We take the following procedures into account for the reflected surface. First, we render the model into a 2D map by using texture coordinates as vertex positions. The model can be textured using traditional 2D textures, projected textures, spotlight textures and environment maps. During the rendering, specular highlights are not computed. Once we have the composite 2D surface signal, the 2D map can be treated just like a traditional texture map where visual masking properties can be computed using the visual discrimination metric and its visual masking properties can be exploited during the remeshing process.

We point out that this approach introduces distortions into the computed surface signal which, in turn, introduces distortions into the computed visual importance map. The general solution for this problem is that all the applications, including [Nguyen et al. 2005; Yuan et al. 2005; Wang et al. 2005], compute another map called the area distortion map to compensate for this distortion.

5 Surface Remeshing

5.1 Map Based Representation

Once we compute the visual discrimination map for the texture we can take advantage of this information to perform geometry remeshing. In this paper we have adapted the remeshing approach developed by [Alliez et al. 2002]. This method computes a set of 2D maps to represent the geometry properties of the model. The advantage of this technique is that most of the remeshing and filtering operations can be easily done in the 2D parametric domain.

To represent the geometric properties of the model we have computed the following 2D maps: an area distortion map, a curvature map, and a regular sampling of the 2d parametric domain as was done in [Alliez et al. 2002]. When combined with the previously computed visual discrimination map, we can perform perceptually based geometry remeshing.

5.2 Importance Sampling based on Centroidal Voronoi Tessellation

A density map is computed using the maps determined previously. Ideally, high curvature areas and low visual masking texture areas require denser sampling while low curvature and strong visual masking areas require less sampling. We have used two parameters to guide the generation of the density map. The curvature gamma adjusts the relative importance of the curvature map. The visual perceptual gamma adjusts the relative importance of the visual perceptual map.

Once the density map is computed, we need to discretize the density map to a set of samples. Alliez et al. [2002] use error diffusion to generate the samples, then switched to centroidal Voronoi tessellation [Alliez et al. 2003b]. In this work, we take the second approach because it generates highly regular samples and thus makes post-processing unnecessary.

Given a region A and a density function ρ defined over this region, the mass of centroid c of A is defined by

$$c = \frac{\int_A x\rho(x)dx}{\int_A \rho(x)dx} \quad (1)$$

One way to compute the weighted centroidal Voronoi tessellation is to use Lloyd's relaxation [Lloyd 1982]. Lloyd's relaxation can be considered as a fixed point iteration. Given a density map and an initial set of n sites, it consists of the following three steps:

1. Build a Voronoi diagram of the n sites.
2. Compute the centroid of each site and move the n sites to their respective centroid.
3. Repeat step 1 and 2 until a satisfactory solution is reached.

Efficiently computing the centroid of each site is not a trivial problem. Determining the centroid requires evaluation of Equation 1 for each site. Inspired by the work of [Hoff et al. 1999], we resort to the use of computer graphics hardware to compute the centroid of each site. A fragment program is used to perform the integrations in Equation 1 using the vector reduce operation together with summation [Krueger and Westermann 2003]. One major limitation of computing the Voronoi diagram using graphics hardware is the frame buffer resolution issue. This is especially true for our case since there can be millions of samples for large models. Instead of computing the centroids of all sites at the same time, we compute one centroid at a time. This allows us to get around the resolution issue.

5.3 Results

The left image in figure 5 demonstrates the result of rendering a textured model of the head of Venus (see Figure 7(e)) into a 2D map. The visual discrimination map that corresponds to this 2D map is shown as the right image.

Perceptually based surface remeshing of the texture mapped head of Venus depicted in Figure 7 begins with the generation of samples from the visual discrimination map shown in Figure 5. The result of applying Lloyd's relaxation on the map for 20 iterations is illustrated in Figure 6. Finally, the samples are reprojected to 3D to generate the 3D mesh.

Figure 7 shows the remeshing of the Venus model both with and without using the visual perceptual properties of the surface signal.

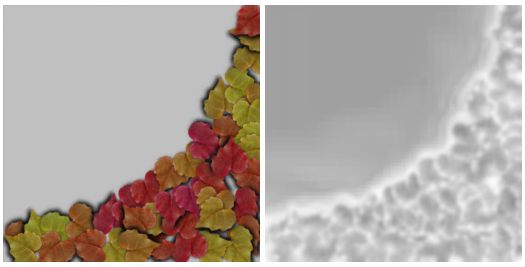

Figure 5: Image on the left is the original texture, image on the right is the visual discrimination map indicating visual masking properties of the texture.

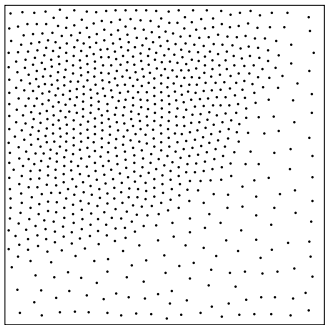

Figure 6: Samples generated by centroidal Voronoi tessellation based the above computed visual masking map.

The two pairs (a) and (b), (c) and (d) are generated with different gamma values for the curvature and perceptual components. The original mesh contains 5000 vertices. Image (a) shows a uniformly remeshed model (curvature gamma is 0) with 2000 vertices. Image (b) is produced with curvature gamma 0 and perceptual gamma 1.0. Image (c) is produced with curvature gamma 1.2, image (d) is produced with the same curvature gamma as (c) and perceptual gamma 2.0.

Notice that the geometric details on the top part of the original mesh are further removed as shown in image (b) and (d) compared to (a) and (c) respectively since it is covered by texture. This further reduction of polygon count in textured areas will not be noticeable due to the visual masking properties of the texture. The triangles saved in the textured areas are used in other parts of the model. As can be seen in the figure, the eyes, nose, and mouth of the model have denser samples than the model without the perceptual component, thus more details are preserved in these areas.

On a Xeon 1.8Ghz, 1G memory machine, it takes less than 1s to compute the surface signal and convert the geometry into a map based representation. The evaluation of the Sarnoff VDM takes about 4s for an image pair of size 512x512. The centrical Voronoi tessellation using 10 Lloyd's iterations for Figure 7 takes about 20s. This is the most expensive part of the algorithm. Fortunately, very few iterations are required to generate good samples. Furthermore, generating samples using image halftoning techniques can be used in the design phase to create the initial samples.

6 Other Types of Surface Signals

Two dimensional texture mapping has been used in this paper to demonstrate how the perceptual properties of the texture, such as

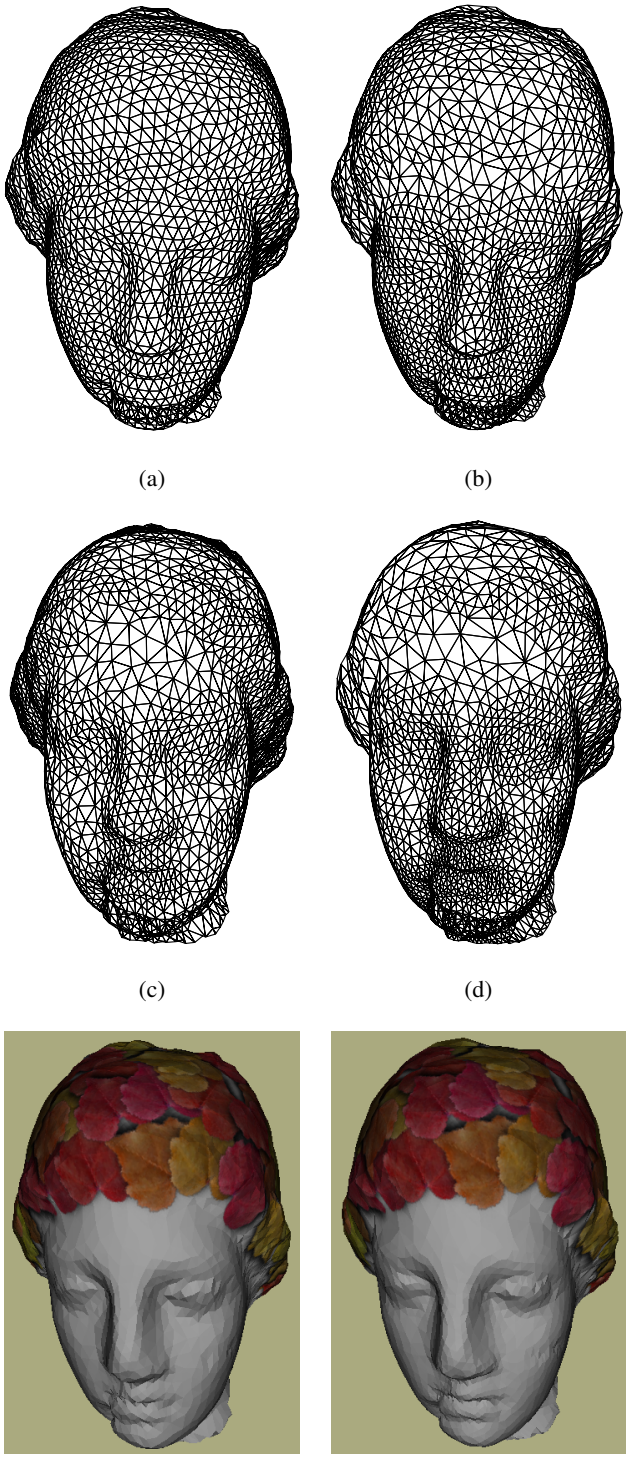

(a) (b) (c) (d) (e) (f)

Figure 7: (a) is a uniformly remeshed model. (b) is a uniformly remeshed model with perceptual component. Notice that the details on the top part of the original mesh are further removed since it is covered by texture. The same for pairs (c) and (d) but with different gamma values. (e) and (f) are rendered images of models shown in (c) and (d), respectively. Notice that more details can be seen in the eyes, nose, and mouth area in image (f).

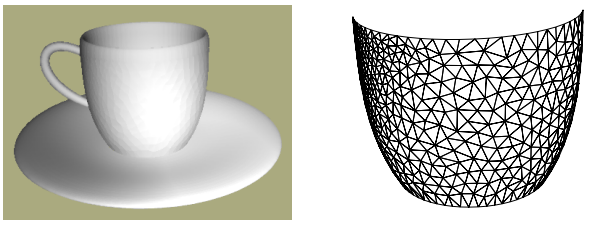

Figure 8: Spotlighted region raises the visual threshold and decreases the number of polygons required.

masking, can be used to guide the remeshing of the geometry to which the texture has been applied. However, two dimensional texture mapping is only one example of several processes that combine to produce the final color pattern that is seen on the surface of the object. We call the variation of lightness and color that is seen by a viewer looking at the object, and that is generated by mechanisms independent of the underlying geometry, the surface signal. To achieve the most dramatic reduction in polygons the complete surface signal should be used in the remeshing process.

In this section of the paper we enumerate the processes by which the surface signal can be altered. In each case we demonstrate how our approach makes use of a single framework to exploit the resulting surface signal and decrease the number of polygons in the underlying geometric mesh. We note that some of the methods by which the surface signal is altered are viewpoint independent and could be taken into account once for a static background like those found in most animations and video games. In other viewpoint dependent cases one would need to page in different mesh representations or remesh on the fly as the observer's position was changed.

6.1 Viewpoint Independent Surface Signals

Bump mapping is another means by which the surface signal can be altered without manipulating the underlying geometry. When only simple diffuse shading is used to perform the bump mapping the result will be viewpoint independent. Figure 1 demonstrates that the illusion of an embossed pattern on the surface due to bump mapping can have a masking effect similar to that produced by two dimensional texture mapping. The area beneath the embossing requires fewer polygons than the homogeneous surfaces adjacent to the embossed area.

Variations in the intensity of a light source across a surface can be another component of the surface signal. The most straightforward way for this to happen is when the light source is focused into a spotlight. This can produce a bright spot on the surface and raise the visual threshold within that pool of light. An example of this is given in Figure 8. Here we see that fewer polygons are required within the bright region produced by the spotlight. Alternatively, obstructions in front of a light source can produce intensity variations that have a similar impact. Masking effects are even possible, as shown in Figure 10, when the pattern of shadows has the necessary frequency content. Here the required number of polygons is reduced in the shadowed areas.

6.2 Viewpoint Dependent Surface Signals

Evaluation of a surface reflection model is an obvious way to alter the surface signal. Implicit in the viewpoint independent surface signals described above is a diffuse shading calculation. Here we

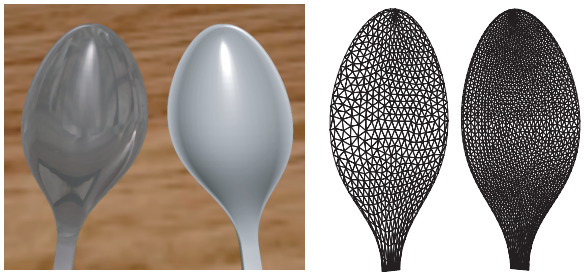

Figure 9: Reflections produce a masking pattern and reduce the number of polygons required in the mesh. The shiny teaspoon on the left has only 1027 vertices, and the diffuse teaspoon on the right has 2761 vertices.

consider the effect of adding a strong specular term to the reflection model that is employed. The result can be a bright highlight on the surface of the object. In a manner similar to the spotlight discussed above, the number of polygons required beneath the specular highlight is reduced because the visual threshold has been elevated in this region. Other more complex BRDFs may produce surface signal variations that can also be exploited to reduce the number of polygons in a mesh.

When the specular reflection becomes even stronger and interreflections are calculated, the surface signal will include the reflected image of other objects in the environment. These mirror reflections will produce a pattern on the surface that can be exploited to reduce the number of polygons in the object mesh. An example of this is shown in Figure 9. Here the shiny teaspoon that reflects the surrounding environment requires fewer polygons than the diffuse teaspoon. This illustration was produced using an environment mapping technique to simulate the interreflections. It is interesting to note that an environment map that might not produce a masking effect as a two dimensional surface texture can create a surface signal that will mask the underlying polygons when it is distorted by reflection onto a surface.

6.3 Complex Scenes

As a final example we present a complex scene composed of multiple objects where several different types of surface signals occur in combination. In contrast to the modest gains that can be achieved in the above examples where only a single surface signal is exploited, many opportunities present themselves to recover polygons in a complicated environment where texture mapping is used extensively, surface reflection properties vary widely, and there are several different light sources. A background scene for a computer video game is an example of an environment where there are numerous surface signals, where the objects do not move, and where the viewpoint remains relatively constant. In this case the polygons that are collected from the background objects can be used to create a more detailed moving foreground object that is the center of the viewer's attention.

In Figure 10 we demonstrate how the algorithm presented in this paper can be used to simplify the meshes in a complicated scene that has multiple surface signals. All of the objects in this picture were decimated using our perceptually based remeshing algorithm. Examples of the mesh reduction achieved for the cup and the teapot are shown in this figure.

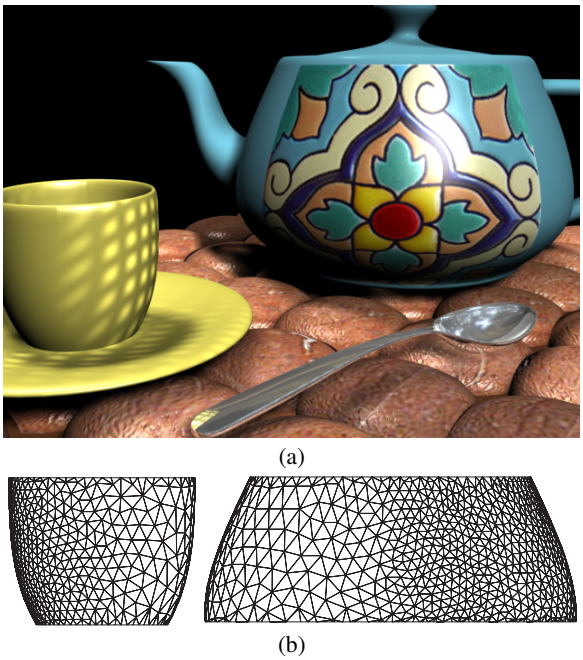

(a)

(b)

Figure 10: (a) Image in which all surfaces are modeled using quadric and cubic surfaces. (b) Example of the remeshing achieved for the cup and the teapot (rotated to show the texture boundary).

7 Discussion

There are major differences between our approach and mesh simplification algorithms that preserve appearance properties such as colors, positions and normals of the geometric models [Cohen et al. 1998; Garland and Heckbert 1998]. Appearance preserving simplification algorithms compute new vertex positions and texture coordinates so as to minimize geometric error and texture stretch. Our approach focuses on the reduction of the polygon count by taking advantage of the perceptual properties of the surface signals. Both approaches take surface signals into account, but with different goals. We do not claim our approach performs better than the above mentioned approaches [Cohen et al. 1998; Garland and Heckbert 1998]. However, we do believe that the above mentioned approaches can be further improved if the perceptual properties of the texture are taken into account.

The biggest difference of our work from previous work in perceptually driven mesh simplification [Luebke and Hallen 2001; Williams et al. 2003] is that we have taken visual masking into account. Visual masking is a very strong phenomenon and can be used to further decrease the sampling rate of the geometry in certain regions (refer to Figure 2). It also requires a much more expensive computation than the other aspects of the visual system such as the contrast sensitivity function and the threshold-vs-intensity function. State of the art visual masking algorithms [Lubin 1995; Ferwerda et al. 1997] require multiscale, multiorientation decomposition of the images. Considering the amount of computation required to evaluate a full perceptual metric including visual masking, it is almost impossible with currently available hardware to evaluate this metric for each edge collapse operation. This explains why we have taken the remeshing approach instead of the mesh simplification approach. The expensive visual discrimination metric is evaluated only once which makes our algorithm tractable.

8 Conclusions

Our system automatically distributes samples uniformly over the polygon mesh by taking the visual perceptual properties of the surface signal into account during the remeshing process. Due to the properties of the human visual system, especially visual masking, the artifacts in the final rendered mesh are invisible to the human observer. This approach also improves the quality of the images in a budget based system since the distribution of polygons across all of the objects in a scene is guided by the principles of visual perception.

We have also demonstrated that there are many opportunities, besides simple two dimensional texture mapping, to exploit the masking properties of the surface signal and redistribute the polygons available to render a scene. Among the additional mechanisms that contribute to the surface signal are bump mapping, spot lighting, shadow patterns, surface reflectance, and interreflections. In a video game or an animated film where many of the objects and much of the lighting in the scene remains static, a large number of the polygons allocated for these background objects can be recovered and used to render principal characters or objects in the foreground. This can reduce rendering times and improve the overall quality of the final animated sequence.

9 Acknowledgments

We would like to thank all anonymous reviewers for their helpful comments and suggestions, which have significantly improved our paper. Funding was provided by the National Science Foundation Grant number CCR-0242757.

References

ALLIEZ, P., MEYER, M., AND DESBRUN, M. 2002. Interactive geometry remeshing. In *ACM SIGGRAPH Conference Proceedings*, 347–354.

ALLIEZ, P., COHEN-STEINER, D., DEVILLERS, O., LEVY, B., AND DESBRUN, M. 2003. Anisotropic polygonal remeshing. In *ACM SIGGRAPH Conference Proceedings*, 485–493.

ALLIEZ, P., DE VERDIÈRE, É. C., DEVILLERS, O., AND ISENBURG, M. 2003. Isotropic surface remeshing. In *Proceedings of Shape Modeling International*, 49–58.

ALLIEZ, P., UCELLI, G., GOTSMAN, C., AND ATTENE, M. 2005. Recent advances in remeshing of surfaces. In *the state-of-the-art report of the AIM@SHAPE EU network*.

BALMELLI, L., TAUBIN, G., AND BERNARDINI, F. 2002. Space-optimized texture maps. In *Proceedings of Eurographics 2002*, 411–420.

BOLIN, M. R., AND MEYER, G. W. 1998. A perceptually based adaptive sampling algorithm. In *ACM SIGGRAPH Conference Proceedings*, 299–309.

CARR, N. A., AND HART, J. C. 2004. Painting detail. In *ACM SIGGRAPH Conference Proceedings*, 845–852.

COHEN, J., OLANO, M., AND MANOCHA, D. 1998. Appearance-perserving simplification. In *ACM SIGGRAPH Conference Proceedings*, 115–122.

DUMONT, R., PELLACINI, F., AND FERWERDA, J. A. 2003. Perceptually-driven decision theory for interactive realistic rendering. *ACM Trans. Graph. 22*, 2, 152–181.

FERWERDA, J. A., PATTANAIK, S. N., SHIRLEY, P., AND GREENBERG, D. P. 1996. A model of visual adaptation for realistic image synthesis. In *ACM SIGGRAPH Conference Proceedings*, 249–258.

FERWERDA, J. A., SHIRLEY, P., PATTANAIK, S. N., AND GREENBERG, D. P. 1997. A model of visual masking for computer graphics. In *ACM SIGGRAPH Conference Proceedings*, 143–152.

GARLAND, M., AND HECKBERT, P. S. 1998. Simplifying surfaces with color and texture using quadric error metrics. In *IEEE Visualization '98*, 263–269.

GU, X., GORTLER, S. J., AND HOPPE, H. 2002. Geometry images. In *ACM SIGGRAPH Conference Proceedings*, 355–361.

HOFF, K. E., KEYSER, J., LIN, M., MANOCHA, D., AND CULVER, T. 1999. Fast computation of generalized voronoi diagrams using graphics hardware. In *ACM SIGGRAPH Conference Proceedings*, 277–286.

KRUEGER, J., AND WESTERMANN, R. 2003. Linear algebra operators for gpu implementation of numerical algorithms. In *ACM SIGGRAPH Conference Proceedings*, 908–916.

LEE, C. H., VARSHNEY, A., AND JACOBS, D. W. 2005. Mesh saliency. In *ACM SIGGRAPH Conference Proceedings*, 659–666.

LINDSTROM, P., AND TURK, G. 2000. Image-driven simplification. *ACM Trans. Graph. 19*, 3, 204–241.

LINDSTROM, P. 2000. *Model Simplification using Image and Geometry-Based Metrics*. PhD thesis, Georgia Institute of Technology.

LLOYD, S. P. 1982. Least squares quantization in PCM. *IEEE Transactions on Information Theory IT-28*, 2 (Mar.), 129–137.

LUBIN, J. 1995. A visual discrimination model for imaging system design and evaluation. In *Vision Models for Target Detection and Recognition*. World Scientific, 245–283.

LUEBKE, D. P., AND HALLEN, B. 2001. Perceptually-driven simplification for interactive rendering. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, 223–234.

MITCHELL, D. P. 1987. Generating antialiased images at low sampling densities. In *ACM SIGGRAPH Conference Proceedings*, 65–72.

NGUYEN, M. X., YUAN, X., AND CHEN, B. 2005. Geometry completion and detail generation by texture synthesis. *The Visual Computer 21*, 8-10, 669–678.

RAMASUBRAMANIAN, M., PATTANAIK, S. N., AND GREENBERG, D. P. 1999. A perceptually based physical error metric for realistic image synthesis. In *ACM SIGGRAPH Conference Proceedings*, 73–82.

SANDER, P. V., GORTLER, S. J., SNYDER, J., AND HOPPE, H. 2002. Signal-specialized parametrization. In *Proceedings of the 13th Eurographics workshop on Rendering*, 87–98.

SURAZHSKY, V., AND GOTSMAN, C. 2003. Explicit surface remeshing. In *Proceedings of the Eurographics/ACM SIGGRAPH symposium on Geometry processing*, 20–30.

TEWARI, G., SNYDER, J., SANDER, P., GORTLER, S., AND HOPPE, H. 2004. Signal-specialized parameterization for piecewise linear reconstruction. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, 55–64.

TURK, G. 1992. Re-tiling polygonal surfaces. In *ACM SIGGRAPH Conference Proceedings*, 55–64.

VOLEVICH, V., MYSZKOWSKI, K., KHODULEV, A., AND KOPYLOV, E. A. 2000. Using the visual differences predictor to improve performance of progressive global illumination computation. *ACM Trans. Graph. 19*, 2, 122–161.

WALTER, B., PATTANAIK, S. N., AND GREENBERG, D. P. 2002. Using perceptual texture masking for efficient image synthesis. *Comput. Graph. Forum 21*, 3, 393–400.

WANG, L., GU, X., MUELLER, K., AND YAU, S.-T. 2005. Uniform texture synthesis and texture mapping using global parameterization. *The Visual Computer 21*, 8-10, 801–810.

WILLIAMS, N., LUEBKE, D., COHEN, J. D., KELLEY, M., AND SCHUBERT, B. 2003. Perceptually guided simplification of lit, textured meshes. In *Proceedings of the 2003 symposium on Interactive 3D graphics*, 113–121.

WINDSHEIMER, J. E., AND MEYER, G. W. 2004. Implementation of a visual difference metric using commodity graphics hardware. In *Human Vision and Electronic Imaging IX. Proceedings of the SPIE, Volume 5292*, 150–161.

YUAN, X., NGUYEN, M. X., ZHANG, N., AND CHEN, B. 2005. Stippling and silhouettes rendering in geometry-image space. In *Proceedings of Eurographics Symposium on Rendering*, 193–200.

(e)

(f)

Figure 7: Images (e) and (f) are rendered comparison of curvature driven remeshing and perceptually driven remeshing given a vertex budget. Our perceptually driven remeshing algorithm allocates less vertices to the top part of the model covered by the leaf texture. Notice that more details can be seen in the eyes, nose, and mouth area in image (f).

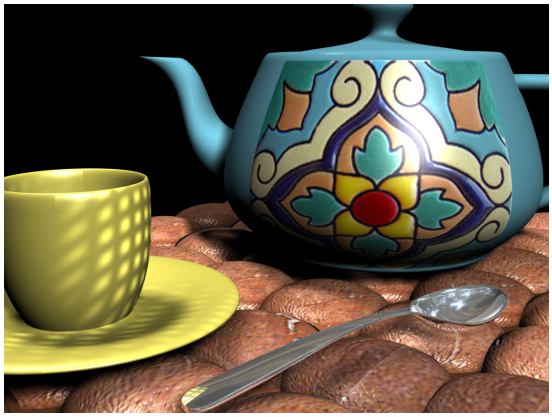

(a)

Figure 10: Image (a) shows a complex scene where multiple surface signals exist, including diffuse color textures, normal maps, surface reflection, and shadow pattern. Our perceptually driven remeshing algorithm can take all these surface signals into account.