Comparison of Five 3D Surface Texture Synthesis Methods

Junyu Dong and Mike Chantler[1]

Abstract - We present and compare five approaches for synthesizing and relighting real 3D surface textures. We adapted Efros's texture quilting method and combined it with five different relighting representations, comprising: a set of three photometric images; surface gradient and albedo maps; polynomial texture maps; and two eigen based representations using 3 and 6 base images. We used twelve real textures to perform quantitative tests on the relighting methods. We develop a systematic qualitative test for the assessment of the complete synthesis systems. Our conclusion is that the cheaper gradient and three-base-image eigen methods should be used in preference to the other methods, especially where the surfaces are Lambertian or near Lambertian.

1 Introduction

There is a growing interest in the synthesis and relighting of 3D surface textures. Zalesny and Van Gool, in [23 & 24] present a multi-view texture model which can synthesise new viewpoints. Shum and his colleagues [14] used the CUReT database [4] to develop a method for the generation of bi-directional texture functions (BTFs). In [13] Leung and Malik proposed the use of 3D textons to synthesize new images under arbitrary viewpoints and illuminations. In later work, Tong *et. al.* also exploited the idea of 'textons' and coupled this with a modified 2D texture synthesis algorithm[19]. In [26], Dong and Chantler proposed six inexpensive methods for the synthesis of 3D surface texture based on the Lambertian assumption. Essentially, it extended $\mathbf{R}^1$ synthesis techniques to deal with three $\mathbf{R}^m$ space relighting representations: height-based, gradient-based and image-based.

In this paper we extend the comparison to include one polynomial and two Eigen-based relighting methods. In addition we develop a systematic two-part evaluation process. First, we quantitatively assess the relighting methods. Second we assess the complete approaches using psychophysical experiments coupled with statistical tests.

This paper is organised as follows. Section 2 briefly describes the framework that we use for the synthesis of 3D surface texture images. Section 3 surveys and selects the five relighting representations that we use. Section 4 introduces the basic synthesis technique that we employ. Section 5 describes each approach in detail and their corresponding assessments are presented in section 6.

2 Basic Framework

Our framework for 3D surface texture synthesis consists of three stages:

1. Extraction of a suitable representation of the 3D surface texture sample from a set of input images.
2. Use of the representation of the sample to synthesize a description of a larger area of surface texture.
3. Rendering (or relighting) of the surface representation according to a specified set of lighting conditions.

3 Surface Representations for Relighting

For the first stage of our system we need a simple and cheap method of obtaining a parsimonious representation of the sample 3D surfaces. This is a popular field of research that has attracted many researchers. Dana and Nayer [4] present a method for measuring BTFs (Bi-directional Texture Functions) and BRDFs of a texture. Leung and Malik [13] used the K-means algorithm to obtain a vocabulary of 3D textons from 20 CUReT textures. Nayar and Dana [5] proposed three BTF derived models for 3D surface texture: a histogram model, a correlation model and a principal component analysis (PCA) model. Using Principal Components to represent and relight 3D surface texture has an advantage that it makes no assumptions about texture surface reflectance [5 & 10]. For Lambertian surfaces, images obtained for the purposes of 3-image photometric stereo [2 & 21] can be used to implicitly represent surface normal and albedo maps. Shashua[18] proposes that a linear combination of three images can be used to generate images of the surface illuminated from a new directions. Malzbender *et. al.* [15] introduced Polynomial Texture Maps(PTMs), to capture variations due to surface self-shadowing and interreflection and have given impressive looking results.

Our aim is to develop techniques that can provide real-time rendering when implemented in a consumer level PC or laptop. This limits us to relighting approaches that (a) use low dimensional representations, and (b) use simple or common graphics calculations such as weighted sums of base images. We have therefore selected the following five methods, from the above, for further study:

3I: This method uses three images of the sample texture taken at an illumination slant angle of 45° and tilt angles of 0°, 90° and 180° [18].

Gradient: The 2nd method uses surface gradient and albedo maps derived using photometric stereo [17].

PTM: This approach uses Polynomial Texture Maps (PTM), due to Malzbender *et. al.* [15].

Eigen3: The fourth method uses the first three PCA base images [10].

Eigen6: This is identical to the previous method except that it uses the first six base images [10].

[1] Junyu Dong and Mike Chantler are with the Texture Lab, Heriot-Watt University, Edinburgh, Scotland(M.J.Chantler@hw.ac.uk)

4 Synthesis methods

The 2nd stage of our framework synthesizes representations of larger surface areas given a representation of a sample texture. This achieved by extending a 2D algorithm. That is we simply extend synthesis in $\mathbf{R}^1$ or $\mathbf{R}^3$ to synthesis in $\mathbf{R}^m$.

We have selected Efros and Freeman's 2D image quilting method [8] as the basis for our synthesis approach.

We have made three small modifications to this quilting algorithm. First, instead of locating the best-matching block using search, we more often select the corresponding neighbour of last selection. This simplification dramatically increases the speed of the algorithm without apparently affecting the output. It can be seen as a simple extension of the algorithm in [1]. Second, we perform synthesis in $\mathbf{R}^m$ space, where m is the dimensionality of the surface representation we are using. Third, we use an error metric based on a sum of absolute differences rather than more expensive L2 norm.

5 The Five Methods

The 3rd stage of our system takes the synthesized representation of the larger area and relights it under specified illumination conditions. This section briefly describes each of the five candidate methods identified in Section 3. They are the *3I, Gradient, PTM, Eigen3* and *Eigen6* methods. (With the exception of the *3I* method, all of the approaches use 36 input images.)

The *3I* method

Under the assumption of Lambertian reflectance, Shashua[18] proposes that a linear combination of three base images can be used to generate new images under different illuminant directions. Thus, the three sample images can be used for synthesizing and relighting of 3D surface texture. This method was first introduced in [26].

The *Gradient* Method

Photometric stereo commonly uses three images to estimate the gradient and albedo maps of a Lambertian surface [21]. Additional images lead to an over-constrained system, which may be solved using least squares techniques to provide potentially more accurate solutions. Thus we use surface gradient and albedo maps produced by 36 images using SVD to represent 3D surface textures. By synthesizing and relighting surface gradient and albedo maps, we can generate new images under arbitrary illumination. We call this the *Gradient* method.

The *PTM* Method

Malzbender proposed the use of a quadratic function (4) as the base representation for relighting surfaces [15]. For each sample texture, 6 coefficient maps (PTM) are generated by using SVD to solve the over-determined system for every location (x, y). This method can produce realistic results for those textures with self-shadows and inter-reflections.

We use all 36 images to generate Polynomial Texture Maps (PTMs) of the sample. These sample PTMs are used to synthesise a new set of output PTMs (in $\mathbf{R}^6$ space). New images under arbitrary illumination are obtained by using the new lighting vector (l_x, l_y) to relight the output PTMs.

The *Eigen3 & Eigen6* Methods

In these approaches we use 3 or 6 base images in eigen-space to represent and synthesize 3D surface textures [10]. We perform Singular-Value Decomposition (SVD) on the 36D sample image space. The first 3 (or 6) base images of the sample are used in $\mathbf{R}^3$ (or $\mathbf{R}^6$) space to synthesize larger base images. The relighting process then simply consists of generating linear combinations of these new base images.

The advantage of using an eigen-space approach is that we may synthesize textures with arbitrary reflectance functions, although specular spikes will require large numbers of base images.

Table 1. **Summary of the 5 approaches**

Approach	1st phase	2nd phase	3rd phase
3I	No processing required in this phase as the three (a, b, c) images are used directly	$\mathbf{R}^3$ synthesis (produces 3 large photometric images a', b', c')	Image-based relighting (produces final image)
Gradient	Produces sample gradient(p,q) and albedo maps (al) using all sample images	$\mathbf{R}^3$ synthesis (produces large gradient and albedo maps)	Gradient-based relighting
PTM	Generates sample Polynomial Texture Maps	$\mathbf{R}^6$ synthesis (produces large Polynomial Texture Maps)	PTM- based Relighting
Eigen3	Generates 3 base images of sample in eigen-space	$\mathbf{R}^3$ synthesis (produces large eigen base images)	Eigen-based relighting
Eigen6	Generates 6 base images of sample in eigen-space	$\mathbf{R}^6$ synthesis (produces large eigen-base images)	Eigen-based relighting

6 Assessment of results

We compare the five approaches in two stages. First, we quantitatively assess the relighting methods. Second, we assess the complete synthesis approaches using psychophysical experiments coupled with statistical tests.

6.1 Quantitative Assessment of Relighting Methods

We can quantitatively assess relighting methods by directly comparing relit images with their corresponding real (input) images. We use 12 textures with reflectance properties ranging from diffuse to strongly specular. Some include shadows and interreflections.

We evaluate the ability of these five methods in predicting new images with illumination conditions that differ from those used for the extraction of surface representations. We employ a *leave-one-out* method, which leaves one image out of the 36 images that we have captured for each texture and tests it as an unknown. We produce 36 relit images in total, which are compared with 36 original images to calculate the normalised *rms* error.

The results are compared using a *normalised root mean square difference* (ε) metric.

$$\varepsilon = \frac{1}{36}\sum_{k=1}^{36}\frac{e_k}{\sigma_k}$$

where:

$$e_k = \frac{1}{NM}\sqrt{\sum \left(r(x,y) - i(x,y)\right)^2}$$

σ_k is the standard deviation of image k,

NM is the size of the images in pixels,

$i(x,y)$ is the x,y^{th} pixel of an input image,

$r(x,y)$ is the x,y^{th} pixel of a relit image,

The results of performing these comparisons are shown in Figure 1.

From these results it can be seen that the *3I* method produces the worst performance. This is not surprising given that it uses three input images whereas the other four methods use 36. The reason is that 3 images can only produce accurate results when the textures have pure Lambertian surfaces with no shadowing. Of the remaining methods, two (*Eigen6* & *PTM*) use more expensive $\mathbf{R}^6$ representations while *Gradient* & *Eigen3* use $\mathbf{R}^3$. On aggregate the *Eigen6* method provides the best figure. However, the performance of the *PTM* approach can not really be separated from that of its cheaper *Eigen3* competitor. It must be cautioned however, that these numerical results may not necessarily agree with qualitative judgements.

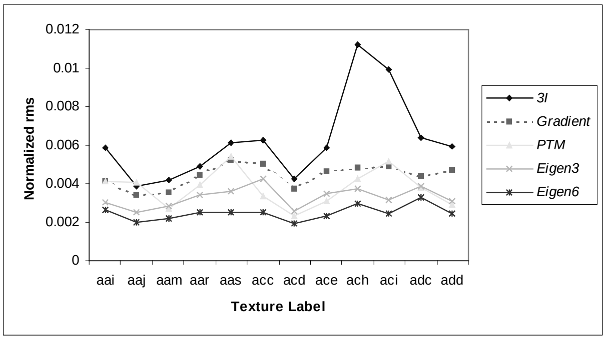

Figure 1. Relighting results for the five methods

6.2 Qualitative Assessment of 5 Approaches

Despite the significant quantity of research on texture synthesis approaches little has been published concerning their assessment. Direct numerical comparison is difficult, as the output textures have no conventional ground-truth. The majority of researchers therefore simply display their results alongside those of their competitors and leave the comparison to readers [7, 20, 8, 22, 9 & 1]. Few provide any experimental support. Copeland *et. al.* did use a psychophysical experiment with ten observers to assess the ability of a numerical error metric to model the perceptual differences between texture patterns [3] but very little has been published on the systematic qualitative assessment of texture synthesis results *per se*. We have therefore developed a simple qualitative approach which uses nonparametric statistical tests and psychophysical experiments.

Five representative textures of different reflectance functions and topology were selected. They included surfaces that exhibited pure Lambertian reflectance, Lambertian reflectance with shadows, and interreflections. For each texture, we used each of the five methods to synthesize two output images under illumination of (tilt angle τ =60°, slant angle σ =60°) and (tilt angle τ =120°, slant angle σ =60°). These images are shown at the end of this paper (aaj, aas, ace, adc, add).

Ten human observers were asked independently to rank the results for each of the five textures from the best to the worst. No other instructions were given concerning as to what qualities to look for when comparing methods. We used Friedman's nonparametric two-way Analysis of Variance (ANOVA) and a multi-comparison method to test their significance.

In our experiments we firstly wished to decide whether there was any significant difference between the performance of the methods. We therefore constructed a matrix which contains one column for each method. Each column contains 50 rank data (10 observers x 5 textures). Friedman's test compares the means of these columns (see [6] for more details). The null hypothesis H_0 is that all five methods make no significant difference for synthesis of 3D surface texture, while the alternative hypothesis H_1 is that at least one is different. The test statistic is defined as:

$$\chi_r^2 = \frac{12}{bk(k+1)}\sum_{j=1}^{k}\left[R_j - \frac{b(k+1)}{2}\right]^2$$

where:

b is total number of rank data for each method (50)

k is the number of methods to be compared (5), and

R_j is the sum of rank data for each method.

The test result indicated that there is at least one method which performs significantly differently from the others at a confidence level of (1.0 - 2.3e-14) x 100% (effectively 100%).

We therefore used a multiple comparison test of means that is designed to provide an upper bound on the probability that any comparison will be incorrectly found to be significant [12]. The result is shown in Figure 2. Each group mean is represented by a small circle within an interval. Two means are significantly different if the associated intervals are disjoint, and are not significantly different if their intervals overlap.

Based on the results of this test in which the confidence levels of the intervals are 99% ($\alpha = 0.01$) we make the following observation. There are no significant differences between the performances of the *Gradient*, *Eigen3*, and *Eigen6* methods. However, each of these methods does outperform both *3I* and *PTM*. Although *Eigen6* produced the best quantitative relighting results, its qualitative performance in the synthesis experiments was not significantly better than its two nearest competitors: *Gradient* and *Eigen3*. This maybe because synthesis is performed in $\mathbf{R}^6$ space which is more prone to matching errors. These errors often introduce discontinuities, which are particularly noticeable to human observers. Finally, if low computational and image-acquisition requirements have to be kept low then the *3I* method, that uses only three photometric images, provides relighting at the cost of lower quality output.

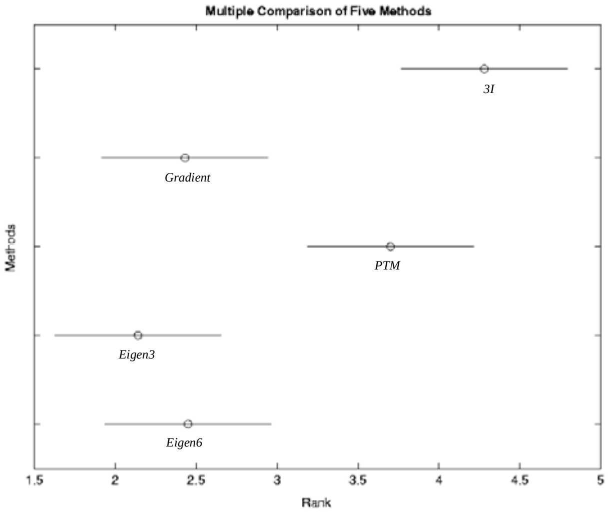

Figure 2. Multiple comparison test of the five methods. Small circles and lines represent the group means and their intervals. The horizontal axis indicates rank values. Two means are significantly different if their intervals are disjoint.

7 Conclusions

In this paper, by adding *PTM* and *Eigen* based methods, we have extended the range of relighting techniques that we have employed from 2 to 5. In addition we developed a systematic qualitative test using a set of 12 real surface textures.

All the methods used thirty-six images except for *3I* that only uses three. This reduced data usage was reflected in the performance of this method, which is only capable of rendering unshadowed Lambertian surfaces. The six-base-image eigen method produced the best quantitative relighting results and in particular it was shown to be better at relighting specular surfaces. However, in the qualitative tests, no significant performance differences were detected between it and the other two top performers: *Eigen3* and *Gradient*. However, the computational complexity of *Eigen6* is approximately twice that of these two $\mathbf{R}^3$ based competitors.

References

1. Ashikhmin, M. 2001. Synthesizing Natural Textures. In the proceedings of 2001 ACM Symposium on Interactive 3D Graphics, Research Triangle Park, NorthCarolina March 19-21, pp. 217-226.
2. Blinn, J.F. 1978. Computer Display of Curved Surfaces. Ph.D. Thesis, Uversity of Utah.
3. Copeland, A. C., Ravichandran, G. and Trivedi, M. M. 2001 Texture synthesis using gray-level co-occurrence models, Opt. Eng. 40, pp. 2655.
4. Dana, K. J., Van Ginneken, B., Nayar, S. K. and Koenderink, J. J. 1999. Reflectance and Texture of Real-World Surfaces. ACM Transactions on Graphics, Vol. 18, No. 1, Pages 1-34.
5. Dana, K.J. and Nayar, S.K. 1999. 3D Textured Surface Modeling. Nayar WIAGMOR Workshop CVPR '99
6. Daniel, W.W. 1990. Applied nonparametric statistics, 2nd ed. PWS-LENT Publishing Company. Boston.
7. De Benet, J.S. 1997. Multiresolution sampling procedure for analysis and synthesis of texture images, SIGGRAPH 97. pp.361.
8. Efros, A. A. and Freeman, W. T. 2001. Image Quilting for Texture Synthesis and Transfer. 28th annual conf. on Computer graphics and interactive techniques, pp 341 - 346 Los Angeles, California
9. Efros, A.A. and Leung, T.K. 1999. Texture synthesis by non-parametric sampling. ICCV99. V2, pp.1033-8.
10. Epstein, R.; Hallinan, P.W.; Yuille, A.L. 1995. 5/spl plusmn/2 eigenimages suffice: an empirical investigation of low-dimensional lighting models. Workshop on Physics-Based Modeling in Computer Vision. p108.
11. Heeger, D.J. and Bergen, J.R. 1995. Pyramid-based texture analysis/synthesis. Int. Conference on Image Processing (Cat. No.95CB35819). IEEE Comput. Soc. Press. Part vol.3, pp.648-51
12. Hochberg, Y. and Tamhane, A. C. Multiple Comparison Procedures, 1987, Wiley.
13. Leung, T. and Malik, J. 2001. Representing and Recognizing the Visual Appearance of Materials Using Three-dimensional Textons. International journal of Computer Vision, Vol 43, No. 1, pp 7-27.
14. Liu, X., Yu, Y. and Shum, H.Y. 2001. Synthesizing Bidirectional Texture Functions for Real-World Surface. Conf on Computer graphics and interactive techniques**,** pp97–106. LA, California.
15. Malzbender, T., Gelb D. and Wolters, H. 2001. Polynomial Texture Maps, Rroc. of the 28th annual conference on Computer graphics and interactive techniques, pp: 519 – 528. Los Angeles, California.
16. Nayar, S. K., Ikeuchi, K. and Kanade, T. Surface Reflection: Physical and Geometrical Perspectives. PAMI 1991, pp. 611-634.
17. Rushmeier, H., Taubin, G. and Gueziec, A. 1997. Applying Shape from Lighting Variation to Bump Map Capture. In Proceedings of the Eighth Eurographics Rendering Workshop, Eds. J. Dorsey and P. Slussallek, Saint-Etienne, France, pp. 35-44.
18. Shashua, A. 1992. Geometry and Photometry in 3D visual Recognition. *PhD thesis, MIT.*
19. Tong, X.; Zhang, J.; Liu L.; Wang X.; Guo B. and Shum, H.-Y. 2002. Synthesis of bidirectional texture functions on arbitrary surfaces. ACM Transactions on Graphics (TOG). SIGGRAPH02, Volume 21 Issue 3. Pages: 665 – 672.
20. Wei, L. and Levoy, 2000. M. Fast texture synthesis using tree-structured vector quantization. In Computer Graphics Proceedings. Annual Conference Series 2000. SIGGRAPH 2000. Conference Proceedings. ACM. pp.479-88.
21. Woodham, R. 1981. Analysing images of curved surfaces. Artificial Intelligence, 17:117-140.
22. Xu, Y., Zhu, S.C., Guo, B. and Shum, H.Y. 2001. Asymptotically Admissible Texture Synthesis, In Proc. 2nd International Workshop of Statistical and Computational Theories of Vision, Vancouver.
23. Zalesny, A. and Van Gool, L. 2000. A compact model for viewpoint dependent texture synthesis. SMILE 2000. Workshop on 3D Structure from Images, Lecture Notes in Computer Science 2018, M Pollefeys et. al. (Eds.), pp.123-143.
24. Zalesny, A. and Van Gool, L. 2001. Multiview Texture Models. IEEE Conf. on Computer Vision and Pattern Recognition, p615.
25. Dong, J. and Chantler, M. 2002. Capture and synthesis of 3D surface texture. Proceeding of the 2nd International Workshop on Texture Analysis & Synthesis. 1 June, Copenhagen, Denmark, pp.41-45.

Table 2. *Synthesis and relighting results from the five methods for 12 textures. The left most images are the samples, the remainder are synthesis results. Arrows indicate illumination directions (τ =60° and τ =120°).*