Capture and Synthesis of 3D Surface Texture

International Journal of Computer Vision (VISI), 62(1-2), 2005, pp177-194.

J. Dong and M. J. Chantler

Abstract

We present and compare five approaches for capturing, synthesising and relighting real 3D surface textures. Unlike 2D texture synthesis techniques they allow the captured textures to be relit using illumination conditions that differ from those of the original. We adapted a texture quilting method due to Efros and combined this 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 in isolation. We developed a qualitative test for the assessment of the complete synthesis systems. Ten observers were asked to rank the images obtained from the five methods using five real textures. Statistical tests were applied to the rankings. The six-base-image eigen method produced the best quantitative relighting results and in particular was better able to cope with specular surfaces. However, in the qualitative tests there were no significant performance differences detected between it and the other two top performers. Our conclusion is therefore that the cheaper gradient and three-base-image eigen methods should be used in preference, especially where the surfaces are Lambertian or near Lambertian.

Capture and Synthesis of 3D Surface Texture

Abstract - We present and compare five approaches for capturing, synthesising and relighting real 3D surface textures. Unlike 2D texture synthesis techniques they allow the captured textures to be relit using illumination conditions that differ from those of the original. We adapted a texture quilting method due to Efros and combined this 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 in isolation. We developed a qualitative test for the assessment of the complete synthesis systems. Ten observers were asked to rank the images obtained from the five methods using five real textures. Statistical tests were applied to the rankings.

The six-base-image eigen method produced the best quantitative relighting results and in particular was better able to cope with specular surfaces. However, in the qualitative tests there were no significant performance differences detected between it and the other two top performers. Our conclusion is therefore that the cheaper gradient and three-base-image eigen methods should be used in preference, especially where the surfaces are Lambertian or near Lambertian.

1 Introduction

The majority of texture synthesis research is concerned with learning and generating 2D (photographic) images of texture [Heeger95, Zhu00, DeBonet97, Wei00, Efros99, Bar-Joseph99, Simoncelli98, Xu01, Efros01 & Ashikhmin01]. If the subjects are 3D surface textures (such as brick, woven or knitted textiles, embossed wallpapers, etc.) then 2D techniques cannot provide the information required for rendering under other than the original illumination. This presents limitations for realistic rendering of textures in augmented and virtual reality applications. Figure 1 illustrates the dramatic effect that varying illumination direction can have on the appearance of a surface texture.

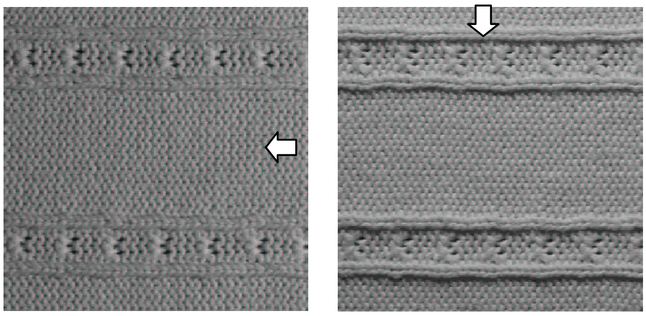

Figure 1. Two images of a 3D surface texture imaged under differing illumination (arrows indicate direction of illumination).

1.1 Related Work

Few publications are available on 3D surface texture synthesis *per se*. Zalesny and Van Gool, in [Zalesny00 & Zalesny01] present a multi-view texture model which can synthesise new viewpoints. They do not however, consider varying illumination which is the focus of this paper.

Shum and his colleagues[Liu01] used the CUReT database [Dana99] to develop a method for the generation of bi-directional texture functions (BTFs). They applied a shape-from-shading algorithm to recover surface height and albedo maps of samples assuming Lambertian reflectance. These are used to synthesise larger heightmap and image *templates.* The final images are synthesised by matching and copying blocks between template images and the reference sample images.

In [Leung01] Leung and Malik proposed the use of 3D textons to synthesise new images under arbitrary viewpoints and illuminations. In later work Shum et. al. also exploited the idea of 'textons' and coupled this with a modified 2D texture synthesis algorithm[Tong02]. As 3D texons use multi-element vectors the computation is expensive. They therefore transform 3D textons to surface textons, which are dot products of base vectors in 3D texton space, to provide an efficient implementation.

1.2 Our Approach

There are many ways in which image relighting may be performed [Basri01, Blinn78, Cook82, Epstein95, Malzbender01, Nayar91, Nishino01, Ramamoorthi02, Rushmeier97, Shashua92, Woodham81 and Zhang98]. In essence relighting takes multiple images of a subject, obtained under different illumination conditions, and compresses these data into some lower order $\mathbf{R}^m$ space. Synthesis of 3D surface texture can therefore simply be viewed as extending an existing $\mathbf{R}^1$ algorithm to $\mathbf{R}^m$ and coupling it to an appropriate relighting scheme.

In this way we develop five 'novel' and inexpensive 3D synthesis methods each of which produces texture representations that can be rendered in real-time in a modern desktop personal computer.

We compare these 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.

1.3 Organisation of Paper

This paper is organised as follows. Section 2 presents the overall 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 – it is derived from an efficient 2D algorithm. Section 5 describes each approach in detail and their corresponding assessments are presented in section 6.

2 Basic Framework

The synthesis of 3D-surface descriptions naturally deals with more information than its 2D counterpart. The later requires consistent texture patterns to be generated in a single image that has no perceptual difference from the original. Synthesis is therefore performed in an $\mathbf{R}^1$ (monochrome) or $\mathbf{R}^3$ (colour) space. In contrast, 3D texture synthesis requires generation of an underlying geometric pattern in addition and may involve the production of varying reflectance functions.

A single sample image does not normally provide enough information regarding surface topology, albedo, and reflectance. Relighting techniques therefore employ multiple images which are used to derive either implicit or explicit representations of these surface properties. Furthermore, we would like the sample data to be captured using off-the-shelf digital cameras, and the results to be capable of being rendered in real-time on current desktop machines. Representations should therefore be of low dimension and preferably capable of per-pixel-rendering using linear combinations.

Our framework for 3D surface texture synthesis (Figure 2) therefore comprises of three parts:

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 synthesise a description of a larger area of surface texture.
3. Rendering (or relighting) of the surface representation according to the specified set of lighting conditions.

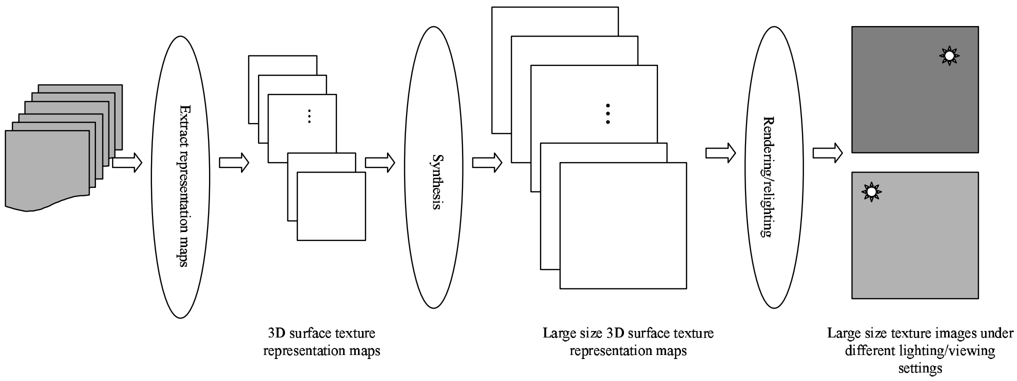

Figure 2. The basic framework for 3D surface texture synthesis.

3 Surface Representations for Relighting

The first stage of our framework abstracts a surface representation of the sample given multiple images of that sample. The choice of these representations has significant impact both on the computational requirements and the quality of the final result.

3.1 Available Representations

The subject of 3D surface texture representation and relighting has attracted many researchers. Bi-directional reflectance distribution functions (BRDF) and surface geometry provide what are perhaps the most sophisticated of these representations. Dana and Nayer [Dana99] present a method for measuring BTFs (Bi-directional Texture Functions) and BRDFs of a texture. Their database (CUReT) was built using an extensive set of real-world textures imaged under a wide variety of illumination and viewing conditions. It has been used for texture synthesis and modelling despite the fact that the majority of the images require registration before use in relighting and shape-from-shading algorithms.

Leung and Malik [Leung01] used the K-means algorithm to obtain a vocabulary of 3D textons from 20 CUReT textures. For each material, 20 images obtained under different viewing/lighting conditions were selected and processed to obtain sets of "3D textons" (combinations of simple image texture elements). This vocabulary was used to represent the sample textures.

Nayar and Dana [Dana99] 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 [Dana99, Epstein95, Nishino01 and Zhang98]. It is widely used by many researchers. Based on experiments, Epstein et. al. in [Epstein95] suggested that five base images (plus or minus two) can be effectively used to represent arbitrary lighting for many different objects. They concluded that this approach can accurately model Lambertian surfaces with specular lobes, while specular spikes, small shadows and occluders can be treated as residuals. Naturally both the specularity and the complexity of surface geometry increases the number of base images required.

For Lambertian surfaces, images obtained for the purposes of 3-image photometric stereo [Blinn78 and Woodham81] can be used to implicitly represent surface normal and albedo maps. For example, Shashua[Shashua92] proposes that a linear combination of three images can be used to generate images of the surface illuminated from a new directions; and by analysing the spherical harmonics of lighting, Basri and Jacobs[Basri01] showed that a 9D linear subspace can be used to accurately approximate the set of images obtainable from a Lambertian object.

Malzbender et. al. [Malzbender01] introduced Polynomial Texture Maps(PTMs), which use a biquadratic polynomial to model luminance. The coefficients of the biquadratic polynomial are stored per pixel, and used to reconstruct the surface colour under varying lighting conditions. PTMs can capture variations due to surface self-shadowing and interreflection and have given impressive looking results.

3.2 Selection of Approaches

As previously discussed one of our aims is to develop techniques that can be provide real-time rendering when implemented in a desk top machine. 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. BTFs are too expensive in terms of the dimensionality required - we have therefore selected five other methods listed below for further investigation:

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° [Shashua92].

Gradient: The second method uses surface gradient and albedo maps derived using photometric stereo [Woodham81 and Rushmeier97].

PTM: This approach uses Polynomial Texture Maps (PTM), due to Malzbender et. al. [Malzbender01].

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

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

These approaches are described in detail in Section 5.

4 Synthesis methods

The second stage of our framework synthesises representations of larger surfaces given suitable representations of the sample textures. As discussed in the introduction this form of texture synthesis can be viewed as an extension of existing 2D algorithms. That is we simply extend synthesis in $\mathbf{R}^1$ or $\mathbf{R}^3$ to synthesis in $\mathbf{R}^m$.

Many 2D approaches are based on statistical models or assumptions [Heeger95, Zhu00, DeBonet97, Wei00, Efros99, Bar-Joseph99, Simoncelli98, Xu01, Efros01 & Ashikhmin01]. However, Efros and Leung [Efros99] presented a simple method which exploits the Markov random field assumption. It selects output pixels by matching neighbourhoods between the sample and synthesised images to produce realistic results for many textures. Wei and Levoy [Wei00] used multi-resolution images to accelerate this synthesis process. More recently, two *quilting* algorithms [Xu01 & Efros01] have been shown to be capable of producing remarkable results while requiring little computation. We have therefore selected Efros and Freeman's 2D image quilting method [Efros01] as the basis for our synthesis approach.

Efros and Freeman's method synthesizes a new image by 'stitching' together small patches from the sample image. The new image is generated in raster order. First, a block is randomly selected from the sample image and pasted into the new image beginning at the first row and the first column. Then another block is selected as a candidate neighbour. This is placed next to the first block so that they overlap one another. The overlapping area between the two blocks is used to test the goodness of fit of the candidate using an L2 norm. This is repeated for different candidates and the one with the minimum SSD is selected as the final neighbour. A minimum error boundary cut is calculated in the overlapping area so that the boundary looks smooth. Both vertical and horizontal overlapping areas are used for selecting best-matched blocks inside the new image. This whole process is repeated until an output image of the required size has been generated.

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 [Ashikhmin01]. 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

This section describes each of the five approaches: *3I, Gradient, PTM, Eigen3* and *Eigen6,* in detail. First however, we describe the image capture stage which is common to all of these approaches.

5.1 Image Capture

We use a fixed 12-bit monochrome CCD camera to capture a 36 images of each sample. A light-source located approximately at 1m from the sample is used at three slant angles and twelve tilt angles. The tilt angles are separated by intervals of 30°. All images used in this paper are available from our online image database. (www.cee.hw.ac.uk/texturelab/database) With the exception of the *3I* (three image) method all of the approaches use the full set of 36 images.

5.2 The *3I* method

Under the assumption of Lambertian reflectance, Shashua[Shashua92] proposes that a linear combination of three base images can be used to generate new images under different illuminant directions. According to the Lambertian reflectance law:

$$I_{(\tau,\sigma)}(x,y) = \lambda\rho \frac{-p\cos\tau\sin\sigma - q\sin\tau\sin\sigma + \cos\sigma}{\sqrt{p^2+q^2+1}} \qquad (1)$$

where:

- I is the intensity of an image pixel at position *x, y*
- λ is the incident intensity to the surface
- ρ is the albedo value of the Lambertian reflection
- τ is the tilt angle of illumination
- σ is the slant angle of illumination
- p and q are the partial derivatives of the surface height function in the *x* and *y* directions respectively

From (1) we can express $I_{(\tau,\sigma)}$ an image captured under an illuminant direction of (τ,σ) as a linear sum of three images captured using non-colinear illuminant vectors. Thus:

$$\begin{aligned} I_{(\tau,\sigma)}(x,y) = & \left(\frac{\cos\tau\sin\sigma}{2\sin 45^\circ} - \frac{\sin\tau\sin\sigma}{2\sin 45^\circ} + \frac{\cos\sigma}{2\cos 45^\circ}\right)\cdot I_{(0,45)}(x,y) \\ & + \frac{\sin\tau\sin\sigma}{\sin 45^\circ}\cdot I_{(90,45)}(x,y) \\ & + \left(\frac{\cos\sigma}{2\cos 45^\circ} - \frac{\cos\tau\sin\sigma}{2\sin 45^\circ} - \frac{\sin\tau\sin\sigma}{2\sin 45^\circ}\right)\cdot I_{(180,45)}(x,y) \end{aligned} \qquad (2)$$

We use 3 images captured with illumination separated by 90° in tilt. That is the illumination is provided at a common slant (45° in our case) and at tilt angles of 0°, 90° and 180°. These images are used directly in equation 2 for rendering, no intermediate images are used.

The complete synthesis procedure and sample images are shown in Figure 3. The three photometric images (*a, b & c*) of the sample are captured. These are used to synthesise (in $\mathbf{R}^3$ space) three larger images (*a', b' & c'*) which form the base images for relighting using equation 2.

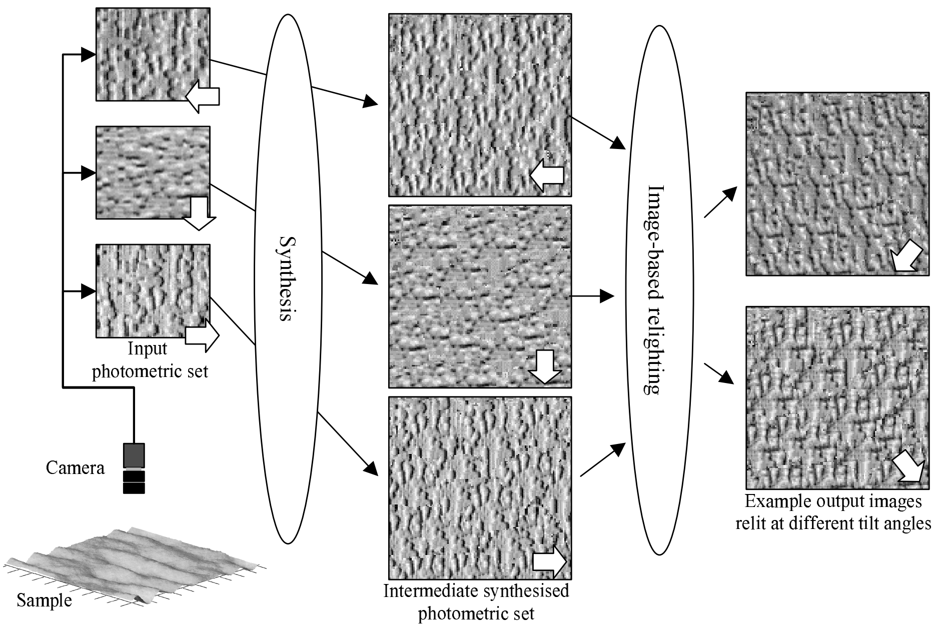

Figure 3. Synthesis directly using 3 photometric sample images (the 3I method)

5.3 The *Gradient* Method

Photometric stereo commonly uses three images to estimate the gradient and albedo maps of a Lambertian surface. Additional images lead to an over-constrained system, which may be solved using least squares techniques to provide potentially more accurate solutions. Given 36 images we obtain 36 equations that are expressed in matrix product form in (3).

$$\begin{bmatrix} -\cos\tau_1\sin\sigma_1 & -\sin\tau_1\sin\sigma_1 & \cos\sigma_1 \\ -\cos\tau_2\sin\sigma_2 & -\sin\tau_2\sin\sigma_2 & \cos\sigma_2 \\ \vdots & \vdots & \vdots \\ -\cos\tau_{36}\sin\sigma_{36} & -\sin\tau_{36}\sin\sigma_{36} & \cos\sigma_{36} \end{bmatrix} \begin{bmatrix} al(x,y)p(x,y) \\ al(x,y)q(x,y) \\ al(x,y) \end{bmatrix} = \begin{bmatrix} I_{(\tau_1,\sigma_1)}(x,y) \\ I_{(\tau_2,\sigma_2)}(x,y) \\ \vdots \\ I_{(\tau_{36},\sigma_{36})}(x,y) \end{bmatrix} \quad (3)$$

where:

al(x,y) is the product of albedo $\rho(x,y)$ and the constant incident intensity λ

Thus we can use surface gradient and albedo maps to represent 3D surface texture. By synthesizing and relighting surface gradient and albedo maps, we can generate new images under arbitrary illumination. We call this the *Gradient* method. The complete synthesis process is similar to the method shown in Figure 3.

5.4 The *PTM* Method

Malzbender proposed the use of a quadratic function (4) as the base representation for relighting surfaces[Malzbender01].

$$L(x,y;l_x,l_y) = a_0(x,y)l_x^2 + a_1(x,y)l_y^2 + a_2(x,y)l_xl_y + a_3(x,y)l_x + a_4(x,y)l_y + a_5(x,y) \quad (4)$$

where:

(l_x, l_y) is the projection of the normalized light vector into the local texture coordinate system (x, y)

L is the resultant surface luminance,

($a_0(x,y) - a_5(x,y)$) are coefficients stored as spatial maps and named Polynomial Texture Maps(PTMs).

For each sample texture, 6 coefficient maps (PTM) are generated by using SVD to solve the over-determined system (i.e. the 36 equations in the form of 4) for every location (x, y). This method can produce realistic results for those textures with cast 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 illuminations are obtained by using the new lighting vector (l_x, l_y) to relight of the output PTMs according to the equation (4). The whole process is shown in Figure 4.

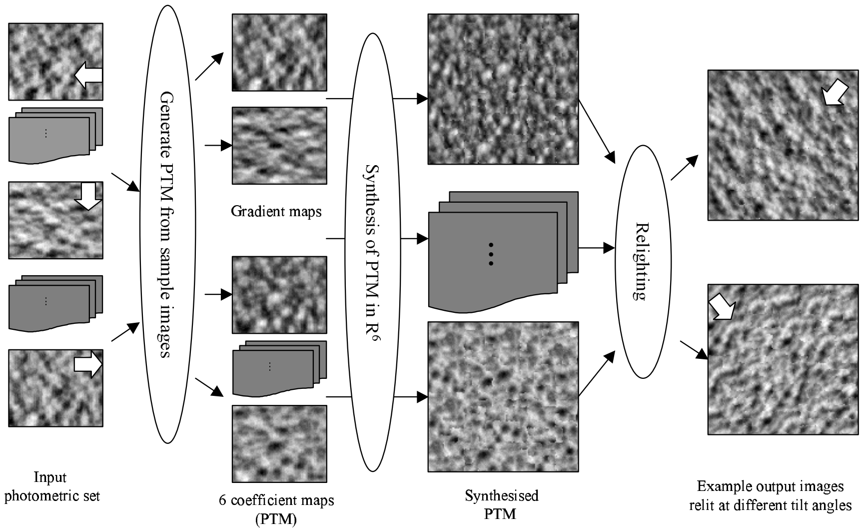

Figure 4. Synthesis ousing Polynomial Texture Maps (the PTM method).

5.5 The Eigen3 & Eigen6 Methods

In these approaches we use 3 or 6 base images in eigen-space to represent and synthesise 3D surface textures. We perform Singular-Value Decomposition (SVD) on the 36D sample image space.

Let ***M*** be a ***mxn*** real matrix, where ***m*** is the number of images and ***n*** is the number of pixels per image; by SVD we can write

$$\mathrm{M} = U\Sigma V^T \tag{5}$$

where **U** is a ***mxm*** orthonormal matrix, Σ is ***mxm*** diagonal matrix, and ***V*** is a ***mxn*** orthonormal matrix. Let $\Sigma = diag(\sigma_1, \sigma_2, ..., \sigma_m)$, where σ_i is the singular value of matrix M and $\sigma_i \geq \sigma_{i+1}$. By tolerating a small information loss, we may use $\Sigma' = diag(\sigma_1, ..., \sigma_k, 0, ..., 0)$ to approximate $\Sigma = diag(\sigma_1, \sigma_2, ..., \sigma_m)$ providing the singular values decrease quickly. Then each row of matrix $\Sigma' V^T$ represents a base image. 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. Figure 5 shows the complete procedure.

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

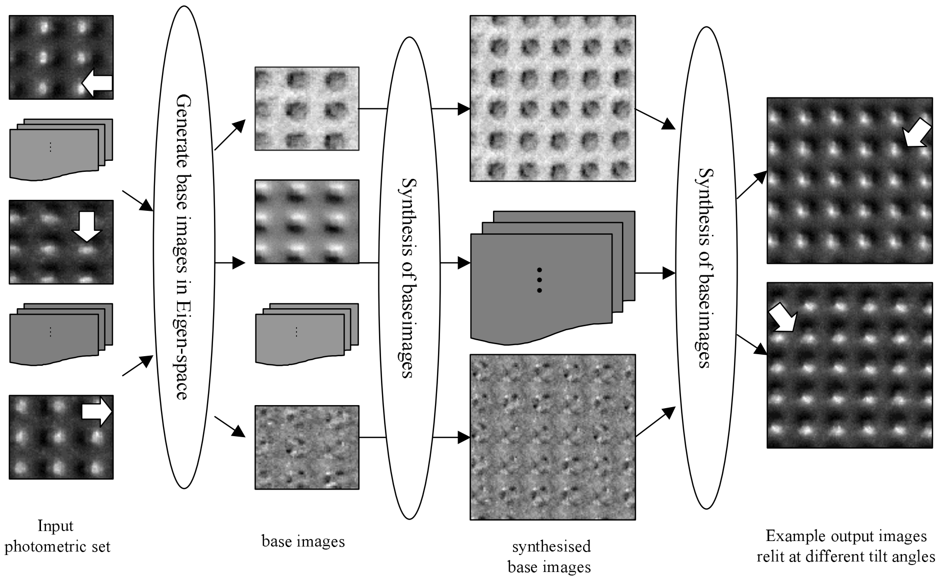

Figure 5. Synthesis using base images in eigen-space. (The example texture used in the figure is "aci", which contains specular reflectance.)

5.6 Summary of Approaches

This section provides a tabular summary (Table 1) of the five approaches that we have developed and will report on in the next section.

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	Generate sample Polynomial Texture Maps	$\mathbf{R}^6$ synthesis (produces large Polynomial Texture Maps)	PTM- based Relighting
Eigen3	Generate 3 base images of sample in eigen-space	$\mathbf{R}^3$ synthesis (produces large eigen base images)	Eigen-based relighting
Eigen6	Generate 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

If we remove the synthesis stage from our framework 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. These texture samples are shown in Figure 6.

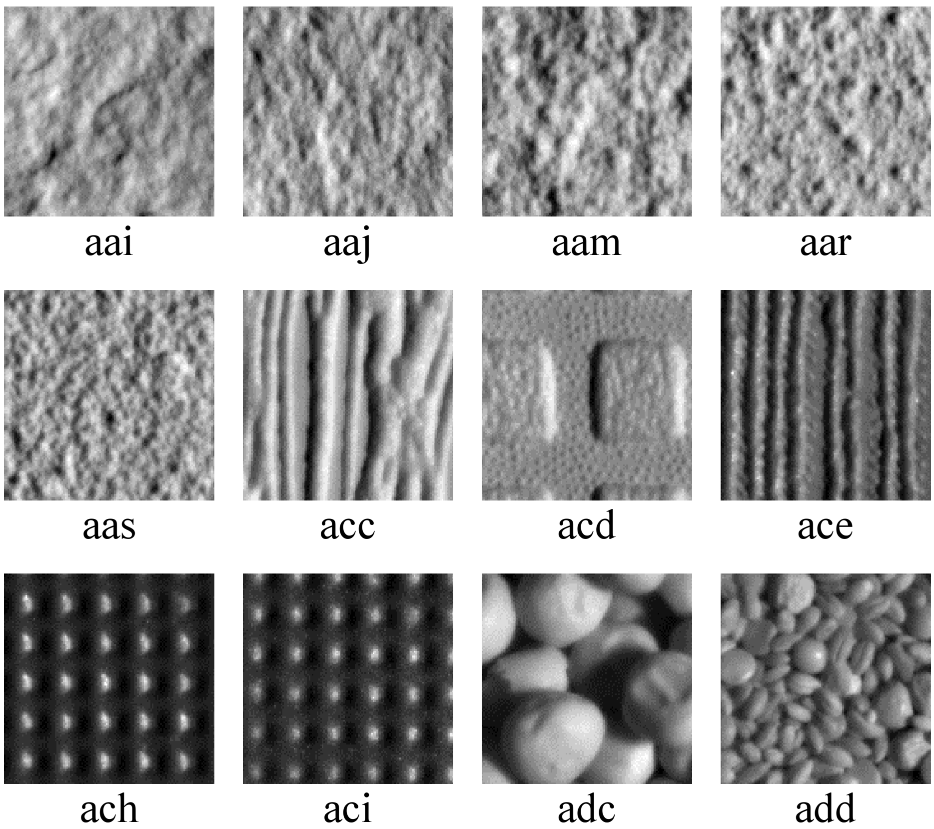

Figure 6. Sample textures.

We relight these textures using same illumination conditions that were used to capture the input images. The results are compared using a *normalised root mean* *square difference* (ε) metric.

$$\varepsilon = \frac{1}{36}\sum_{k=1}^{36}\frac{e_k}{\sigma_k} \tag{6}$$

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 Figure 7in (a) and (b) .

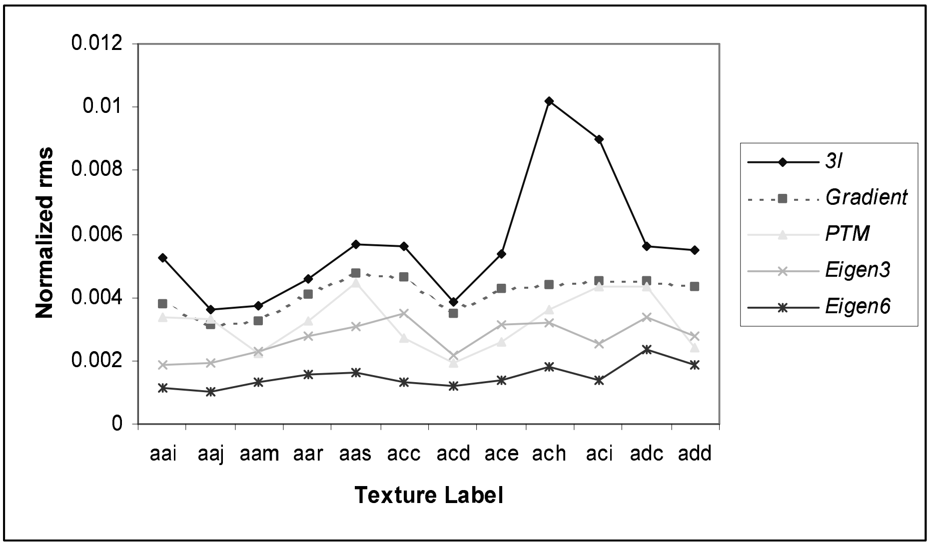

(a) Relighting error vs texture for the five approaches

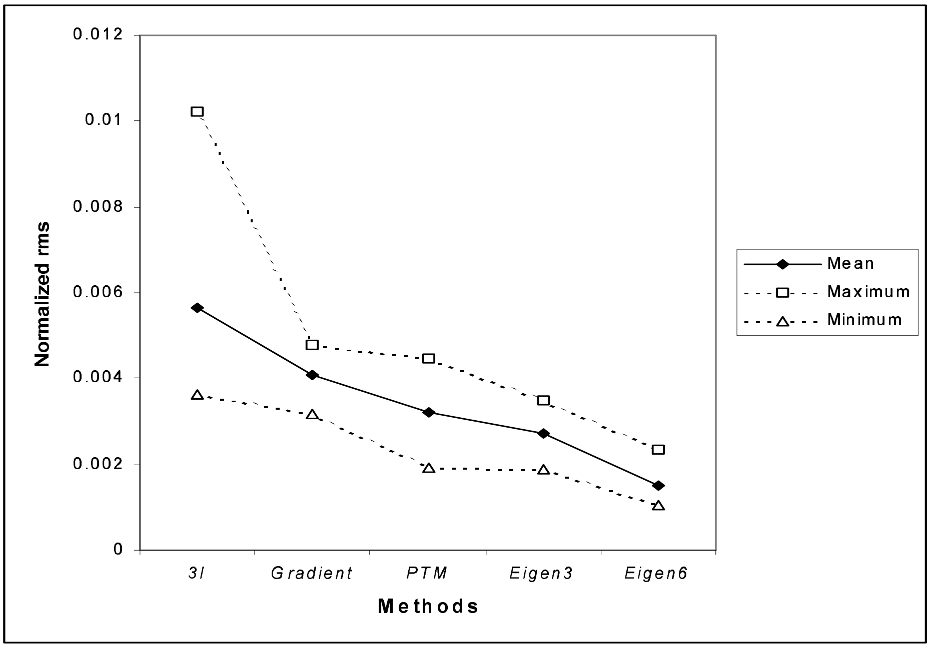

(b) Mean, Maximum, and minimum relighting errors calculated across textures

Figure 7. Quantitative relighting results for the five methods

Example output images and their absolute difference images are shown in Figure 8. These three sets of results are representation of the images we obtained for Lambertian, Lambertian with shadows, and specular surfaces.

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 that can only produce accurate results when the textures have pure Lambertian surfaces with no shadowing.

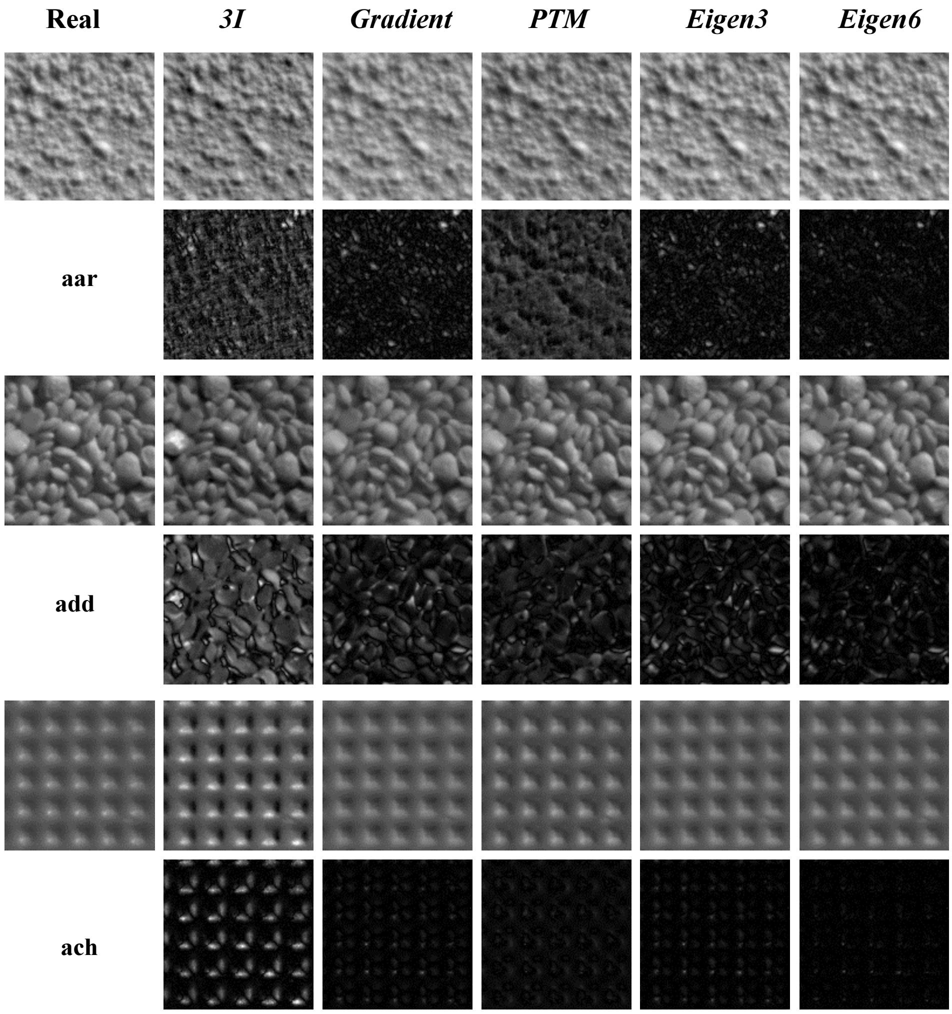

Figure 8. Example output from the five methods for the three textures "aar", "add" and "ach". (For each texture, the first image of the first row is the real image, and the second to the fifth images are synthesis result; the second row contains difference images between real and synthesised images.) ***look at this again??***

Of the remaining methods, two (*Eigen6* & *PTM*) use more expensive $\mathbf{R}^6$ representations while *Gradient* & *Eigen3* use $\mathbf{R}^3$. We would therefore expect the first pair of techniques to outperform the latter, and on aggregate the *PTM* method does indeed provide 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.

6.2 Qualitative Assessment of the Five Synthesis 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 [DeBonet97, Wei00, Efros99, Xu01, Efros01 & Ashikhmin01]. 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 [Copeland01] but there is 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.

We asked a set of ten human observers to rank different synthesis methods by comparing output images with input images. These results were tested to see if the rankings were statistically significant.

Five textures representative of different reflectance functions and topology were selected. They included surfaces that exhibited pure Lambertian reflectance, Lambertian reflectance with shadows, and interreflections. These images are shown at the end of this paper (aaj, aas, ace, adc, add). For each texture, we used each of the five methods to synthesise two output images under illuminations of (τ =60°, σ =120°) and (τ =60°, σ =120°) Observers were asked 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. Thus we collected 50 sets of rankings (10 observers x 5 textures).

Since each observer performed their ranking independently we used Friedman's nonparametric two-way Analysis of Variance (ANOVA) and a multi-comparison method to test their significance.

Friedman's nonparametric two-way Analysis of Variance (ANOVA) is designed to determine if we may conclude from sample evidence that there is difference among treatment effects. 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 [Daniel90] 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 [Hochberg87]. The result is shown in Figure 9. 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.

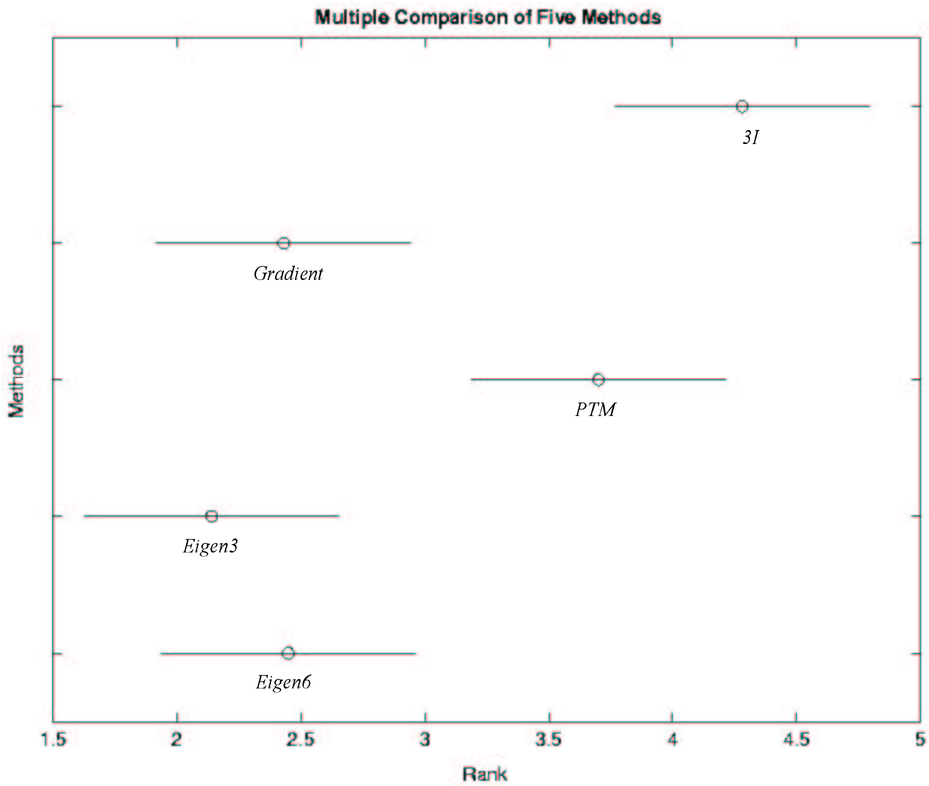

Figure 9. Multiple comparison test of the five methods. Small circles and lines represent the group means and their intervals. The horizontal axis indicates rank values. Two means are significantly different if their intervals are disjoint.

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.

If we take computation complexity into account, we find that synthesis in $\mathbf{R}^6$ space is of course the most expensive - in our approaches, it exactly doubles the computation time compared with $\mathbf{R}^3$ synthesis.

The conclusion therefore from this limited test is that *Gradient* and *Eigen3* would appear on average offer as good a performance as of any of the other methods and to provide it for low computational cost. In theory the more expensive *Eigen6* should offer better performance for non-Lambertian surfaces as was indicated by the quantitative relighting tests. 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.

7 Conclusions

We have presented and compared five approaches for capturing, synthesising and relighting real 3D surface textures. We adapted a texture quilting method due to Efros and combined this with five different relighting representations:

3I – three photometric images,

Gradient – gradient and albedo maps

PTM – Polynomial texture Maps

Eigen3 – Three base eigen images

Eigen6 – Six base eigin images

We performed quantitative tests on the relighting methods, and we developed a qualitative test for the assessment of the complete synthesis systems.

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 successfully.

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.

Table 2. Synthesis and relighting results from the five methods for 12 textures. The first five textures("aaj","aas","ace","adc","add") are used in the qualitative tests. The left most images are from the sample, the remainder are synthesis results. Arrows indicate illumination direction.

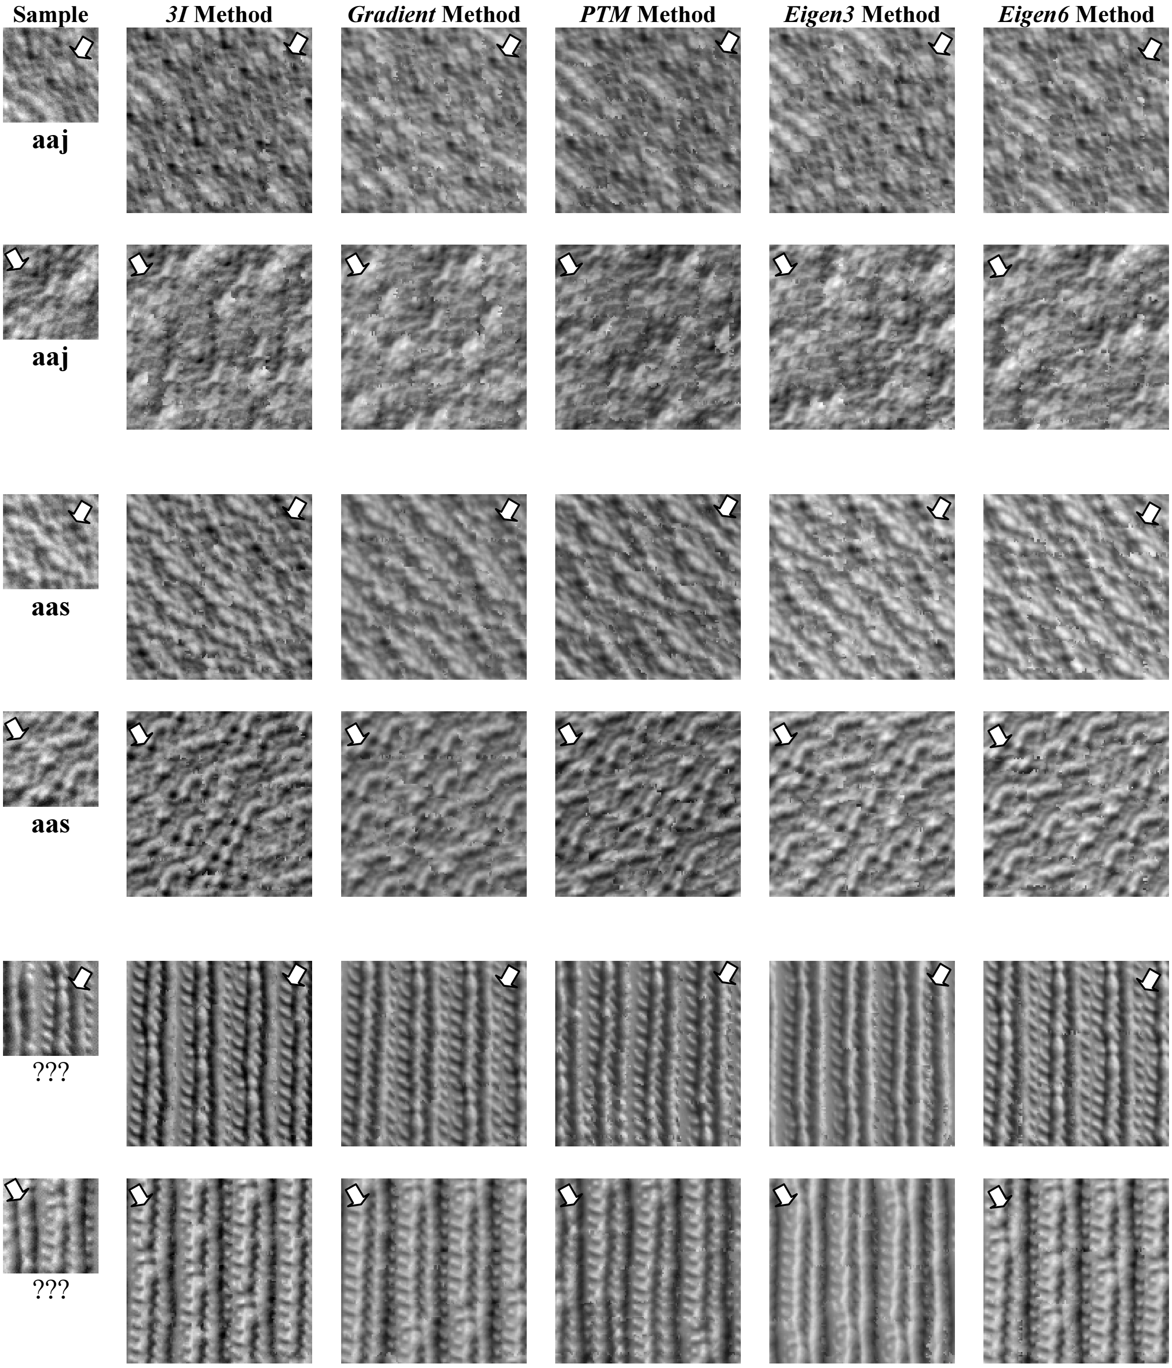

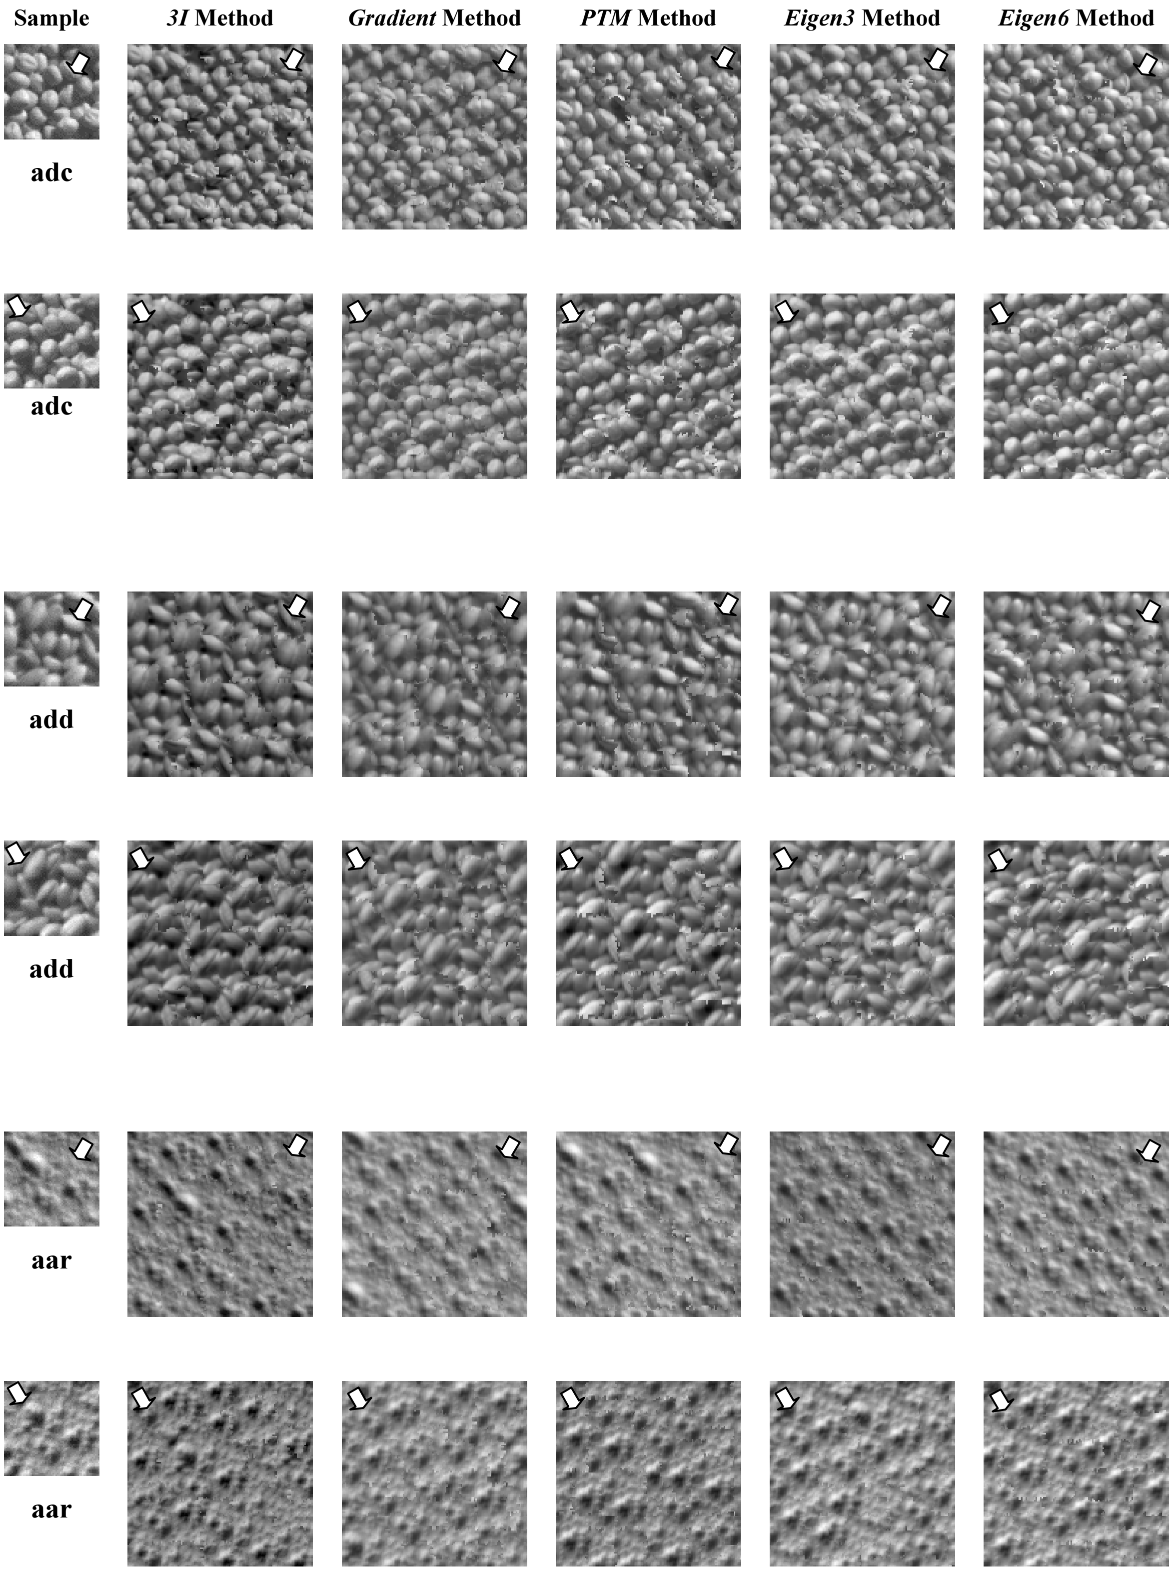
Sample
3I Method
Gradient Method
PTM Method
Eigen3 Method
Eigen6 Method
adc
adc
add
add
aar
aar

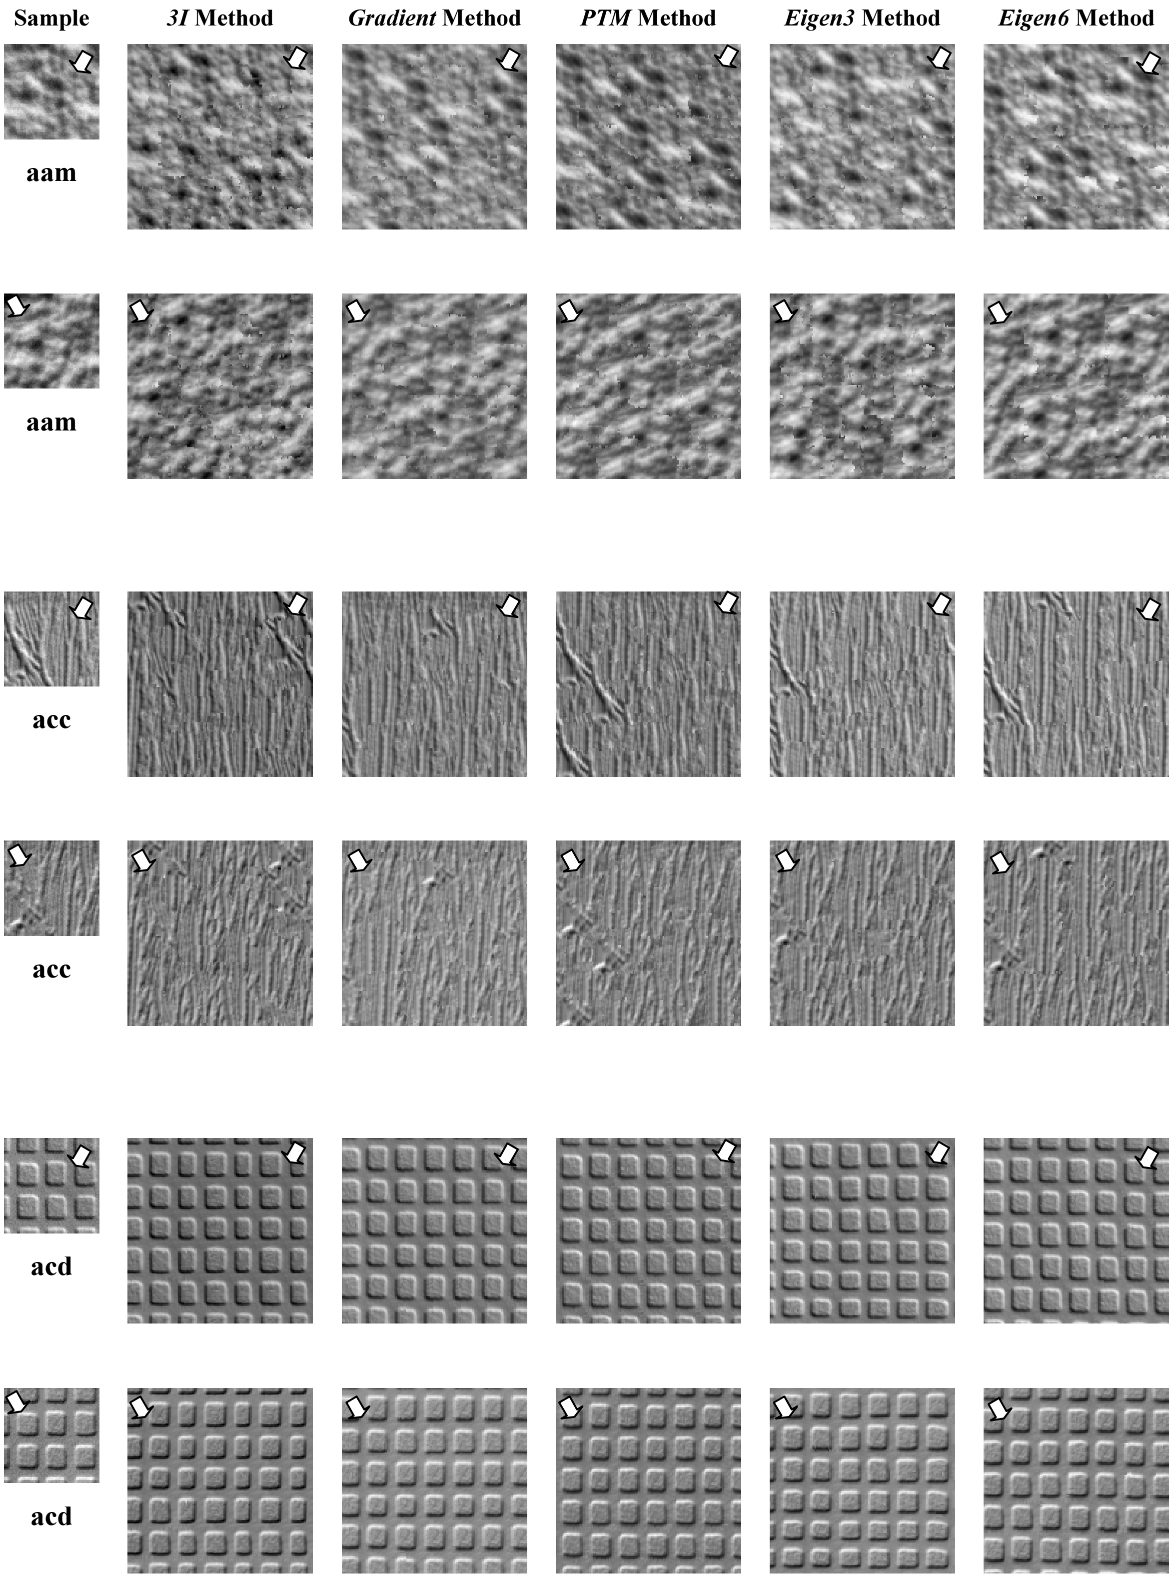
Sample
3I Method
Gradient Method
PTM Method
Eigen3 Method
Eigen6 Method
aam
aam
acc
acc
acd
acd

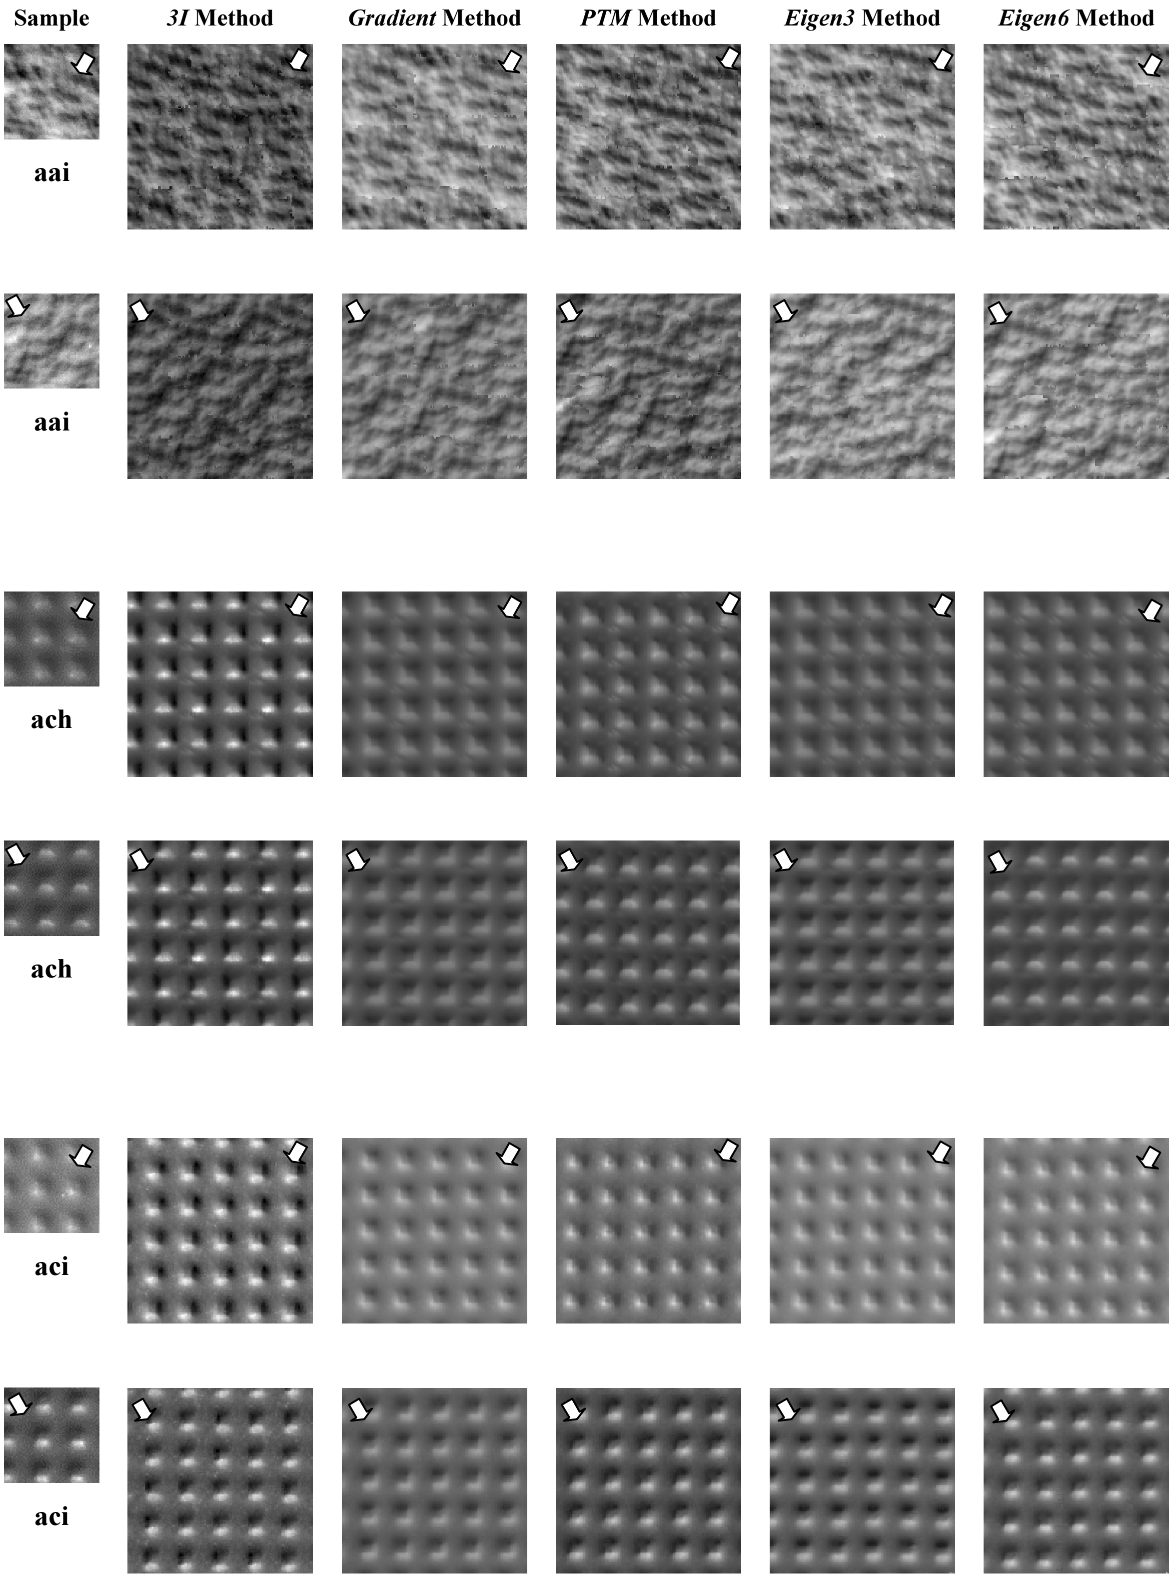
Sample
3I Method
Gradient Method
PTM Method
Eigen3 Method
Eigen6 Method
aai
aai
ach
ach
aci
aci

References

[Ashikhmin01] 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.

[Bar-Joseph99] Bar-Joseph, Z. 1999. Statistical Learning of Multi-Dimensional Textures. MSc thesis, Institute of Computer Science, The Hebrew University of Jurusalem.

[Basri01] Basri, R. and Jacobs, D. 2001. Lambertian reflectance and linear subspaces. In Proceedings. Eighth IEEE International Conference on Computer Vision. ICCV 2001. Volume: 2 , 2001. Page(s): 383 –390.

[Blinn78] Blinn, J.F. 1978. Computer Display of Curved Surfaces. Ph.D. Thesis, Uversity of Utah.

[Cook82] Cook, R. L. and Torrance, K. E. 1982. A Reflectance Model for Computer Graphics. ACM Transactions on Graphics (TOG). Volume 1, Issue 1. Pages: 7 – 24.

[Copeland01] Copeland, A. C., Ravichandran, G. and Trivedi, M. M. 2001 Texture synthesis using gray-level co-occurrence models: algorithms, experimental analysis, and psychophysical support, Opt. Eng. 40, pp. 2655.

[Dana99] 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.

[Dana99] Dana, K.J. and Nayar, S.K. 1999. 3D Textured Surface Modeling. Nayar WIAGMOR Workshop CVPR '99

[Daniel90] Daniel, W.W. 1990. Applied nonparametric statistics, 2nd ed. PWS-LENT Publishing Company. Boston.

[De Benet97] De Benet, J.S. 1997. Multiresolution sampling procedure for analysis and synthesis of texture images. In Computer Graphics Proceedings, SIGGRAPH 97. ACM. pp.361.

[Efros01] Efros, A. A. and Freeman, W. T. 2001. Image Quilting for Texture Synthesis and Transfer. In Proceedings of the 28th annual conference on Computer graphics and interactive techniques , Pages: 341 - 346 Los Angeles, California

[Efros99] Efros, A.A. and Leung, T.K. 1999. Texture synthesis by non-parametric sampling. In Proceedings of the Seventh IEEE International Conference on Computer Vision. IEEE Comput. Soc. Part vol.2, pp.1033-8.

[Epstein95] Epstein, R.; Hallinan, P.W.; Yuille, A.L. 1995. 5/spl plusmn/2 eigenimages suffice: an empirical investigation of low-dimensional lighting models. In Proceedings of the Workshop on Physics-Based Modeling in Computer Vision. Page(s): 108.

[Frankot98] Frankot, R. T. and Chellappa, R. 1998. A Method for Enforcing Integrability in Shape from Shading Algorithms. IEEE Transactions on Pattern Analysis and Machine Intelligence, VOL. 10, No. 4, Page(s):439-451.

[Heeger95] Heeger, D.J. and Bergen, J.R. 1995. Pyramid-based texture analysis/synthesis. In Proceedings International Conference on Image Processing (Cat. No.95CB35819). IEEE Comput. Soc. Press. Part vol.3, pp.648-51

[Hochberg87] Hochberg, Y. and Tamhane, A. C. Multiple Comparison Procedures, 1987, Wiley.

[Leung01] 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, Number 1, pp 7-27.

[Liu01] Liu, X., Yu, Y. and Shum, H.Y. 2001. Synthesizing Bidirectional Texture Functions for Real-World Surface. In Proceedings of the 28th annual conference on Computer graphics and interactive techniques , Pages: 97 - 106 Los Angeles, California.

[Malzbender01] Malzbender, T., Gelb D. and Wolters, H. 2001. Polynomial Texture Maps, In Proceedings of the 28th annual conference on Computer graphics and interactive techniques, Pages: 519 –

528.
[Nayar91] Nayar, S. K., Ikeuchi, K. and Kanade, T. 1991. Surface Reflection: Physical and Geometrical Perspectives. IEEE Transactions on Pattern Analysis and Machine Intelligence. Vol. 13, No. 7, pp. 611-634.

[Nishino01] Nishino, K.; Sato, Y. and Ikeuchi, K. 2001. Eigen-Texture method: appearance compression and synthesis based on a 3D model. Pattern Analysis and Machine Intelligence, IEEE Transactions on , Volume: 23 Issue: 11, Page(s): 1257 –1265

[Ramamoorthi02] Ramamoorthi, R. 2002. Analytic PCA construction for theoretical analysis of lighting variability in images of a Lambertian object . Pattern Analysis and Machine Intelligence, IEEE Transactions on , Volume: 24 Issue: 10, Page(s): 1322 –1333

[Rushmeier97] 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.

[Shashua92] Shashua, A. 1992. Geometry and Photometry in 3D visual Recognition. *PhD thesis, MIT.*

[Simoncelli98] Simoncelli, E.P and Portilla, J. 1998 Texture characterization via joint statistics of wavelet coefficient magnitudes In Proceedings. 1998 International Conference on , Volume: 1 , Page(s): 62 -66 vol.1

[Tong02] 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). Proc. of the 29th annual conference on Computer graphics and interactive techniques, Volume 21 Issue 3. Pages: 665 – 672.

[Wei00] 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.

[Woodham81] Woodham, R. 1981. Analysing images of curved surfaces. Artificial Intelligence, 17:117-140.

[Xu01] Xu, Y., Zhu, S.C., Guo, B. and Shum, H.Y. 2001. Asymptotically Admissible Texture Synthesis, In Proc. Second International Workshop of Statistical and Computational Theories of Vision, Vancouver, Canada.

[Zalesny00] 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.

[Zalesny01] Zalesny, A. and Van Gool, L. 2001. Multiview Texture Models. In Proc. IEEE Computer Soc. Conf. on Computer Vision and Pattern Recognition, Volume: 1, Page(s): 615 –1180.

[Zhang98] Zhang, Z. 1998. Modeling Geometric Structure and Illumination Variation of a Scene from Real Images", In Proc. International Conference on Computer Vision(ICCV'98), Bombay,India. Page(s): 1041 –1046.

[Zhu00] Zhu, S.C., Liu, X.W. and Wu, Y.N. June 2000. Exploring texture ensembles by efficient Markov chain Monte Carlo-Toward a "trichromacy" theory of texture. IEEE Transactions on Pattern Analysis & Machine Intelligence, vol.22, no.6, pp.554-69.