Chapter 6

Synthesis and Relighting

6.1. Introduction

In chapter 4, we introduced five methods for representing and relighting surface textures. In chapter 5, we selected a 2D texture synthesis algorithm. In this chapter we present five approaches that combine the surface representation methods with the 2D texture synthesis algorithm to synthesise images of 3D surface textures under arbitrary lighting directions. We will compare these synthesis approaches according to the quality of their output results. The criterion for the comparison is the resemblance, as perceived by human vision, between output results and input samples. The work described in this chapter corresponds to the final stage in our overall framework, as highlighted in Figure 6.1.1.

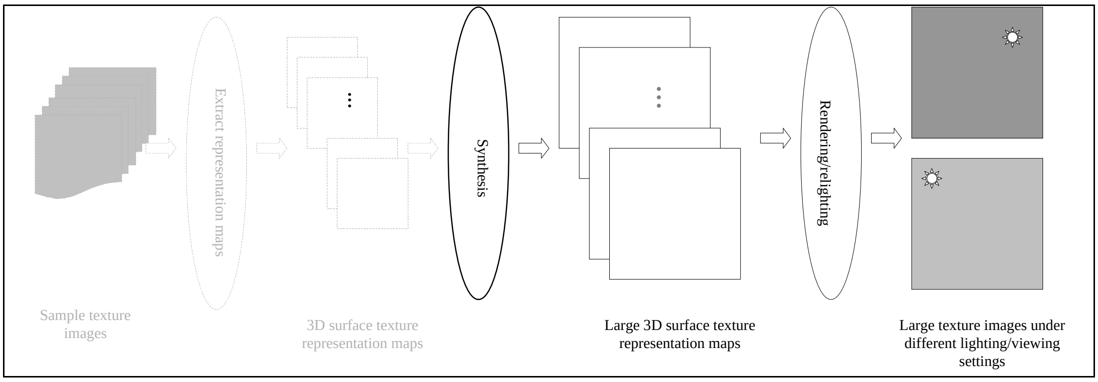

Figure 6.1.1 The final stage of the overall framework

We modify the 2D texture synthesis algorithm selected in chapter 5 so that it can take sample surface representations as input and perform synthesis in multi-dimensional space. We propose five approaches for the synthesis of 3D surface textures that correspond to the five surface representation and religing methods introduced in chapter 4:

The *3I* synthesis approach: This approach uses three images of the sample texture as input and it relights the synthesised images using the *3I* relighting method. Synthesis is performed in $\mathbf{R}^3$ space.

The *Gradient* synthesis approach: The second approach uses surface gradient and albedo maps as input and it relights the synthesised surface gradient and albedo maps using the *Gradient* relighting method. Synthesis is also performed in $\mathbf{R}^3$ space.

The *PTM* synthesis approach: This approach uses Polynomial Texture Maps (PTM) as input and it relights synthesised PTMs using the *PTM* relighting method. Synthesis is performed in $\mathbf{R}^6$ space.

The *Eigen3* synthesis approach: The fourth approach uses the first three eigen base images as input and it relights synthesised base images using the *Eigen3* relighting method.

The *Eigen6* synthesis approach: This is identical to the previous approach except that it uses the first six base images as input. Thus, synthesis is performed in $\mathbf{R}^6$ space.

For our experiments we use the same 23 textures as those used in chapter 4.

We are also interested in the performances of these five approaches concerning the quality of their synthesised results. In chapter 4, we performed a quantitative assessment of the five surface representation and relighting methods. However, we can not perform a similar quantitative comparison here because ground truth data is not available. We therefore qualitatively assess the five synthesis approaches. We perform psychophysical experiments to rank these five approaches based on human perception. Based on the rank data, we use Fredman's nonparametric two-way Analysis of Variance followed by a multi-comparison method to test their significance. The conclusion is that the *Gradient* and *Eigen3* approaches outperform any of the other approaches if both the synthesised results and computational cost are considered.

The chapter is organised as follows. Section 6.2 introduces the five synthesis approaches. Section 6.3 describes the psychophysical experiments for the qualitative comparison of the five approaches. Finally in section 6.4 we draw conclusions from the results of this chapter.

6.2. The five synthesis approaches

This section introduces five synthesis approaches: *3I*, *Gradient*, *PTM*, *Eigen3* and *Eigen6*. They employ the same basic algorithm—the modified Efros and Freeman's 2D texture synthesis algorithm. However, they use different inputs, which comprise different multi-dimensional vectors that represent a sample surface texture under arbitrary illumination directions. During the synthesis process, each pixel location on the sample surface is represented by multi-dimensional vectors that are extracted using the surface representation methods introduced in chapter 4. The synthesis algorithm uses the multi-dimensional vectors as input to synthesise new surface representation maps. They are finally relit using the relevant relighting methods to obtain new images under different illumination directions.

6.2.1. The general algorithm for the synthesis of surface texture representations

The general algorithm for the synthesis of surface texture representations is an extension of the 2D synthesis algorithm that we selected in chapter 5. The algorithm synthesises a result representation by 'stitching' together small blocks from a sample representation. It uses a Sum of Absolute Differences (SAD) as the metric for selecting best-matched blocks in the sample. For 2D texture synthesis, the calculation of SAD only uses pixel intensity values. In the case of 3D surface texture synthesis, each pixel location on the sample surface is expressed as a multi-dimensional vector. The general algorithm therefore uses multi-dimensional vectors. The SAD that we use for multi-dimensional surface representations is:

$$\text{SAD}=\sum_{i=1}^{n}\sum_{\substack{(x,y)\in\Omega_j\\(x',y')\in\Omega_j}}\left|m_i(x,y)-m_i'(x',y')\right| \tag{6.2.1}$$

where:

(x, y) represents a sample pixel location

(x', y') represents a result pixel location

$m_i(x, y)$ is a pixel value at (x, y) in the i^{th} sample representation map

$m_i'(x', y')$ is a pixel value at (x', y') in the i^{th} result representation map

Ω_j is an overlapping area covered by block j

n is the dimensionality or the total number of sample representation maps.

The best-matched blocks are found by minimising the SAD between the overlapping windows of the sample and result representation maps.

The sample surface and output representations are stored as multiple images. The number of images is equal to the dimension of the representations. Thus synthesis in $\mathbf{R^3}$ space involves three input images and three output images, as shown in Figure 6.2.1.

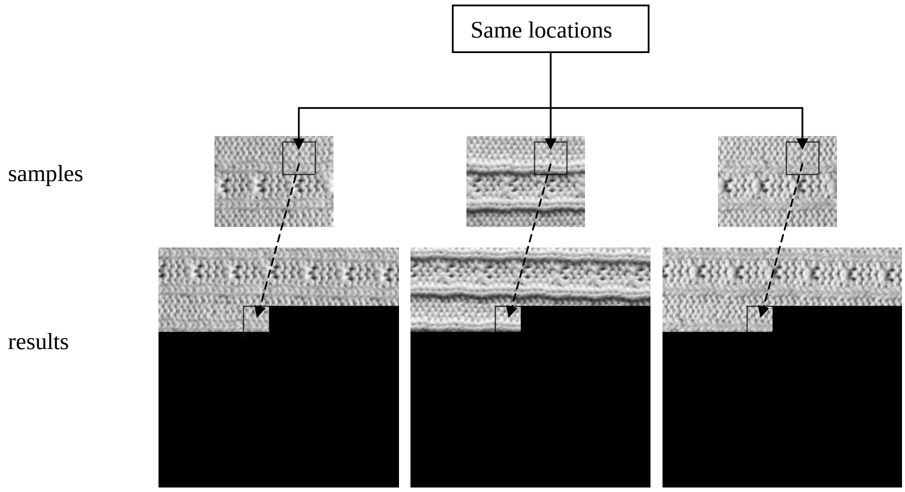

Figure 6.2.1 Each group of best-matched blocks in synthesised results comes from the same location in samples

The synthesised representation maps are then relit using corresponding relighting methods to produce the final results.

Matching errors

It should be noted that *matching errors* exist during the selection of best-matched blocks by calculating the minimum SAD in $\mathbf{R^n}$ space. Suppose we are observing two synthesis processes. The first process synthesises only one representation map in $\mathbf{R^1}$ space using pixel values as input; the second synthesises all representation maps simultaneously in $\mathbf{R^n}$ space. All other parameters are identical. At the same locations of two output representation maps, the best-matched block obtained in $\mathbf{R^1}$ space might be different from its counterpart in the group of best-matched blocks that are produced simultaneously in $\mathbf{R^n}$ space (using *n*-dimensional vectors as input). In the other words, the group of best-matched blocks produced in $\mathbf{R^n}$ space does not guarantee each individual in the group is the same as the best-matched block produced in $\mathbf{R^1}$ space. Figure 6.2.2 illustrates this process. Each large image (output) in Figure 6.2.2 (a) is synthesised independently in $\mathbf{R^1}$ space. For the framed blocks in output images, their best-matched blocks in the samples have different locations. These locations also differ from those in the sample images of (b), in which synthesis is performed in $\mathbf{R^3}$ space. In (b), all framed blocks in output images lie in the same location.

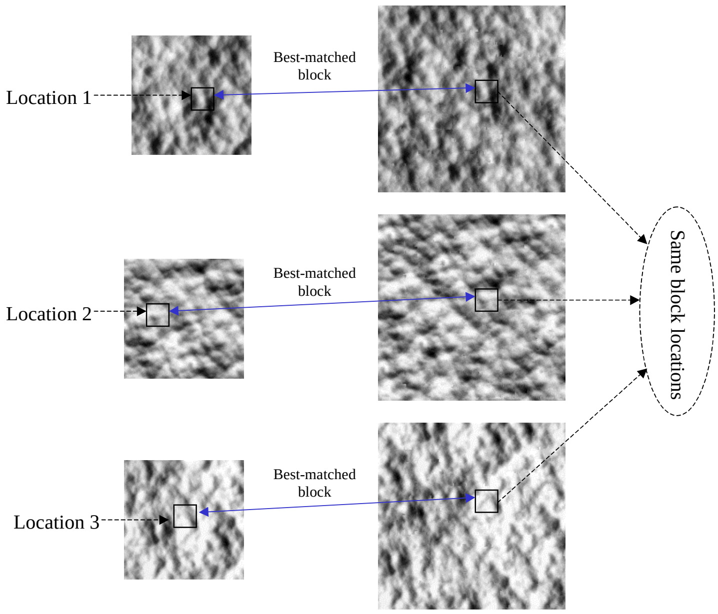

(a)

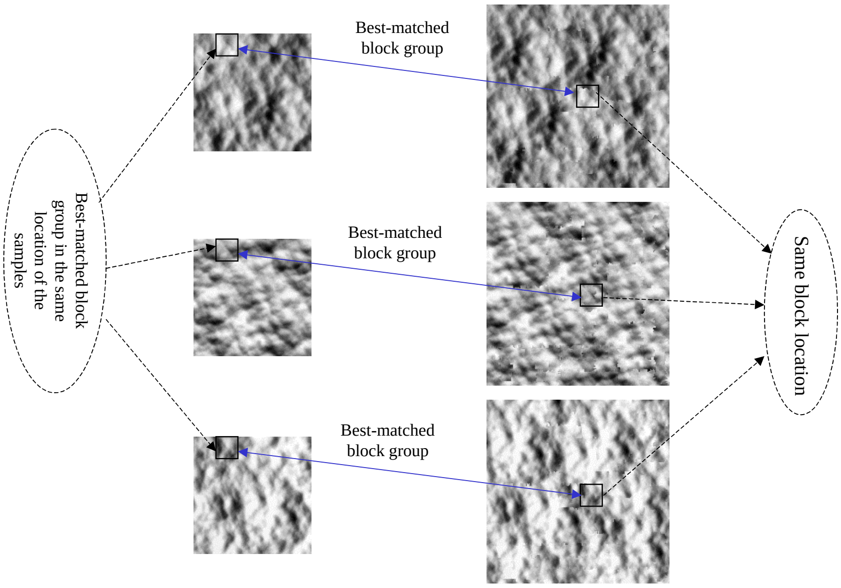

(b)

*Figure 6.2.2 The group of best-matched blocks produced in **R**^{3} space does not guarantee each individual in the group is the same as the best-matched block produced in **R**^{1} space. (a) Each large image (output) is synthesised separately in **R**^{1} space; all framed blocks in the output images lie in the same location but their best-matched blocks have different locations in the samples. (b) Synthesis in **R**^{3} space. All framed blocks lie in the same location in output images and are identical to those in (a), but their best-matched block group has the same location in the samples. This location differs from each of those in (a).*

The reason for producing matching errors is that the minimum SAD, which decides the best-matched blocks, is normally greater than zero when synthesising real-world surface texture representations. Thus, the following mathematical statement is obvious:

$$\underset{\Omega_j}{\text{Min}}\ \{\sum_{\substack{(x,y)\in\Omega_j \\ (x',y')\in\Omega_j}}\sum_{i=1}^{n}\left|m_i(x,y)-m_i'(x',y')\right|\}\geq\sum_{i=1}^{n}\underset{\Omega_j}{\text{Min}}\{\sum_{\substack{(x,y)\in\Omega_j \\ (x',y')\in\Omega_j}}\left|m_i(x,y)-m_i'(x',y')\right|\}\quad(6.2.2)$$

The left side of equation (6.2.2) represents the minimum SAD calculated using *n*-dimensional vectors, while the right side is the sum of the minimum SAD calculated in **R**^{1} space. The *matching error* can be seen as the difference between the two sides of equation (6.2.2). The higher the dimensionality of input vectors is, the larger the *matching errors* might be. *Matching errors* will introduce discontinuities in the result representation maps.

6.2.2. The *3I* synthesis approach

The *3I* synthesis approach first synthesises three output images from three sample *photometric images*, which are captured under linearly independent illumination directions. The synthesis is therefore performed in $\mathbf{R}^3$ space. The three synthesised *photometric images* are then relit to generate new images under arbitrary illumination directions using a linear combination—the *3I* relighting method, as introduced in chapter 4. Figure 6.2.3 shows the process in $\mathbf{R}^3$ space.

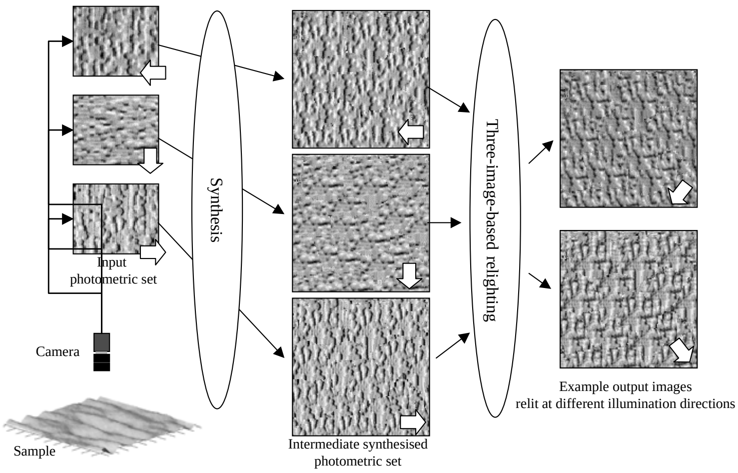

Figure 6.2.3 The 3I *synthesis approach*

6.2.3. The *Gradient* synthesis approach

The *Gradient* synthesis approach synthesises output surface gradient and albedo maps from sample maps. These are generated using the *Gradient* representation method. Synthesis is also performed in $\mathbf{R}^3$ space. Since pixel values in surface gradient maps are normally smaller than those in the albedo map, all pixel values are transformed into same scale during synthesis process. This gives the surface gradient and albedo maps the same weight when calculating Sum of Absolute Difference (SAD). However, the synthesised surface gradient and albedo maps still use pixel values from the corresponding original sample maps. They are relit using the Lambertian model to generate final images under arbitrary illumination directions. Figure 6.2.4 shows the whole synthesis process.

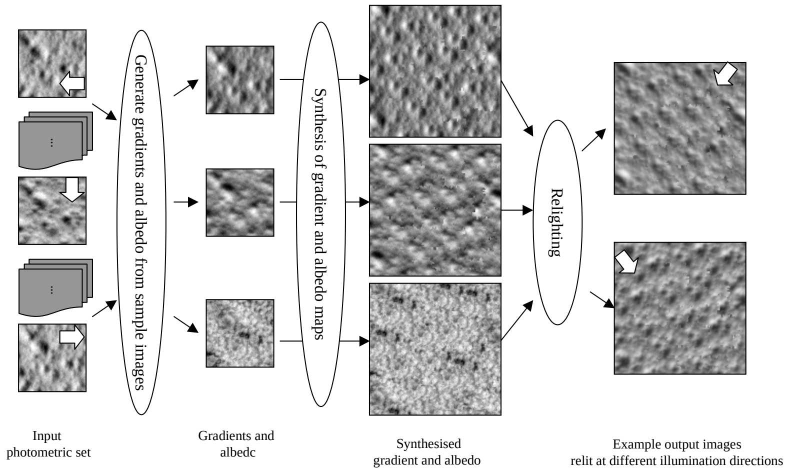

Figure 6.2.4 The Gradient *synthesis approach*

6.2.4. The *PTM* synthesis approach

This *PTM* synthesis approach performs synthesis in $\mathbf{R}^6$ space. The six-dimensional sample Polynomial Texture Maps are also transformed into same scale so that they have the same weight when calculating SAD. The synthesised PTMs are relit using the *PTM* relighting method [Malzbender2001] to produce final images under different illumination directions.

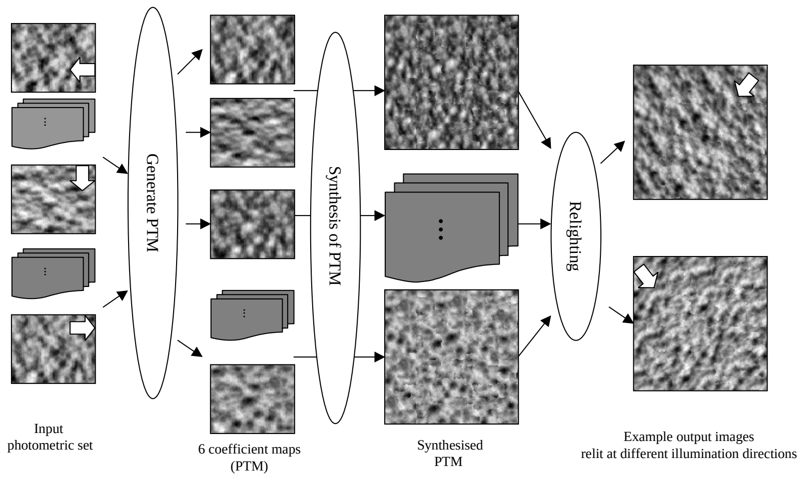

Figure 6.2.5 The PTM *synthesis approach*

6.2.5. The *Eigen3* and *Eigen6* synthesis approaches

The *Eigen3* or *Eigen6* approach uses the first 3 or 6 eigen base images as input to synthesise output eigen base images. The sample eigen base images are generated using the *Eigen3* or *Eigen6* surface representation method. They are also transformed into the same scale during synthesis process so that they have equal weight in calculating SAD between samples and results. The synthesised base images are relit using a bilinear interpolation—the eigen-based relighting methods described in chapter 4 to generate new images under varied illumination directions.

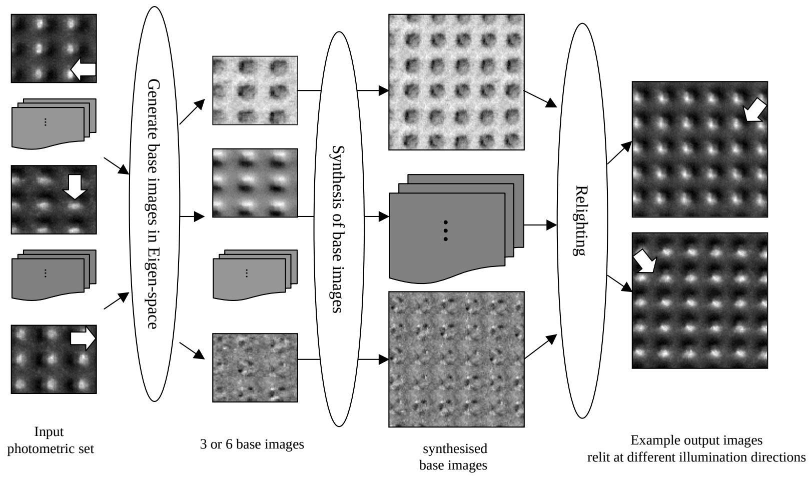

Figure 6.2.6 The Eigen3 *and* Eigen6 *approaches*

6.2.6. Summary

We have presented five approaches for the synthesis and relighting of 3D surface textures. They use surface representation maps extracted from a set of sample images as input to synthesise new surface representations. The synthesised representations are then relit using the corresponding relighting methods to generate final result images under arbitrary illumination directions. We summarise the five approaches in Table 6.2.1. Synthesis results of 23 textures with illumination angles of (τ =60°, σ =60°) and (τ =120°, σ =60°) are shown in Appendix B.

Table 6.2.1 Summary of the 5 approaches

Approach	**1st phase**	**2nd phase**	**3rd phase**
3I	No processing required in this phase as the three images are used directly	$\mathbf{R}^3$ synthesis (produces 3 output photometric images)	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 output gradient and albedo maps)	Gradient-based relighting
PTM	Generates sample Polynomial Texture Maps	$\mathbf{R}^6$ synthesis (produces output Polynomial Texture Maps)	PTM- based Relighting
Eigen3	Generates 3 base images of sample in eigen-space	$\mathbf{R}^3$ synthesis (produces output eigen base images)	Eigen-based relighting
Eigen6	Generates 6 base images of sample in eigen-space	$\mathbf{R}^6$ synthesis (produces output eigen-base images)	Eigen-based relighting

6.3. Qualitative assessment of the five approaches

Section 6.2 described five approaches for the synthesis and relighting of 3D surface textures. This section evaluates the performances of these methods concerning the quality of their synthesis results. In chapter 4, we have quantitatively assessed the surface representation and relighting methods. The conclusion is that the *3I* representation method produces the worst performance and the *Eigen6* method produces the best. The $\mathbf{R}^6$ *PTM* representations perform better than $\mathbf{R}^3$ *Gradient* representations, although it can not be considered superior to the computationally cheaper *Eigen3* representations in $\mathbf{R}^3$ space. We are interested in whether the qualitative performance[1] of the five synthesis approaches is consistent with the quantitative assessment results of relighting methods.

Despite the significant quantity of research on texture synthesis approaches little has been published concerning their assessment. The majority of researchers therefore simply display their results alongside those of their competitors and leave

[1] Note that unlike the assessment of surface representation and relighting methods, we can not perform a quantitative comparison because no ground-truth data is available.

the comparison to readers [DeBonet1997, Wei2000, Efros1999, Xu2001, Efros2001 and shikhmin2001]. 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 [Copeland2001] but very little has been published on the systematic qualitative assessment of texture synthesis results *per se*. In this section, we introduce a simple qualitative approach which uses nonparametric statistical tests and psychophysical experiments.

6.3.1. Design of the psychophysical experiments

Since we are interested in comparing the performances of the five synthesis approaches concerning the quality of synthesis results, we use rank (ordinal) data as the scale of statistic measurement. An ordinal scale of measurement represents an ordered series of relationships or rank order. In our case, we wish to know which methods outperform others or which one can achieve the best, second, or third performance. Unlike precise measurement, rank data is suitable for qualitative measurement. Furthermore, the advantage of using rank data is that it can be simply obtained from observation.

We asked a set of ten human observers to rank different synthesis approaches by comparing output images with input samples. The main concern is the resemblance between the samples and results under multiple illumination directions. In order to avoid distraction from other effects during comparison, we simply place the sample images alongside results with same illumination conditions. Although we have performed the synthesis on 23 sample textures and we can generate images with arbitrary illumination directions, we only select a representative subset from the results for the psychophysical experiments so that observers are relieved from exhaustive comparison. The subset comprises eleven textures (near 50% of all textures) with two illumination directions. These textures include surfaces that exhibit near Lambertian reflectance, Lambertian reflectance with shadows and interreflections, and specular reflectance. These textures also include surfaces with stochastic and structured patterns.

For each texture, we used each of the five approaches to synthesise two output images under illumination angles of (τ =60°, σ =60°) and (τ =120°, σ =60°). These images are shown in Table 6.3.1 and labelled as “aaj”, “aas”, “ace”, “adc”, “add”, “aar”, “acd”, “aai”, “ach”, “aci” and “abj”. Observers were asked to compare real sample images with synthesised images and rank the results for each of the eleven textures from the best to the worst. The illumination directions are indicated by block arrows in the figure. No other instructions were given concerning as to what qualities to look for when comparing methods. Thus we collect 110 sets of rankings (10 observers x 11 textures).

Table 6.3.1 **Synthesis and relighting results from the five methods for 11 textures. The small images in each cell are the samples; the large images are synthesis results. Arrows indicate illumination directions (τ =60° and τ =120°).**

3I	***Gradient***	***PTM***	***Eigen3***	***Eigen6***
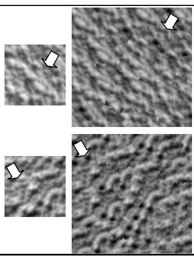	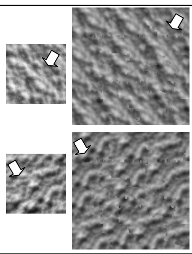	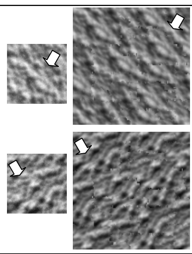	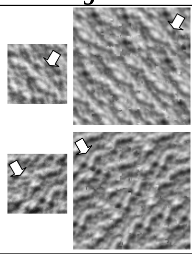	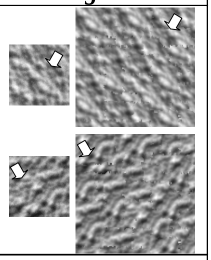
Texture "aaj"				
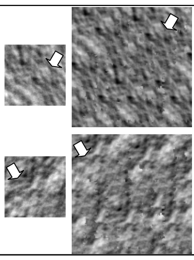	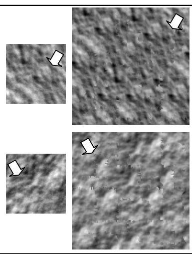	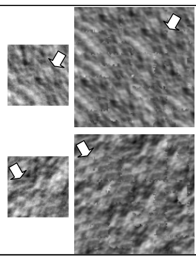	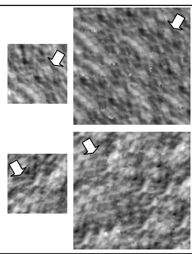	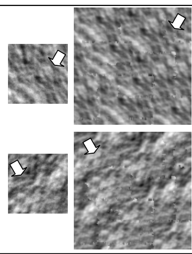
Texture "aas"				
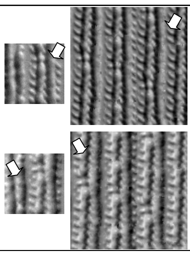	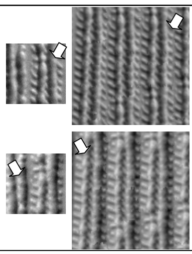	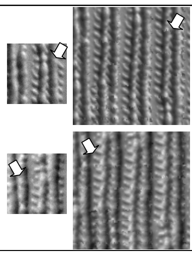	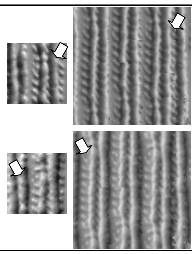	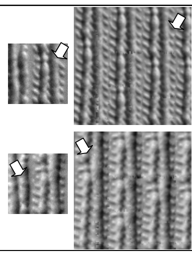
Texture "ace"				

3I	*Gradient*	*PTM*	*Eigen3*	*Eigen6*

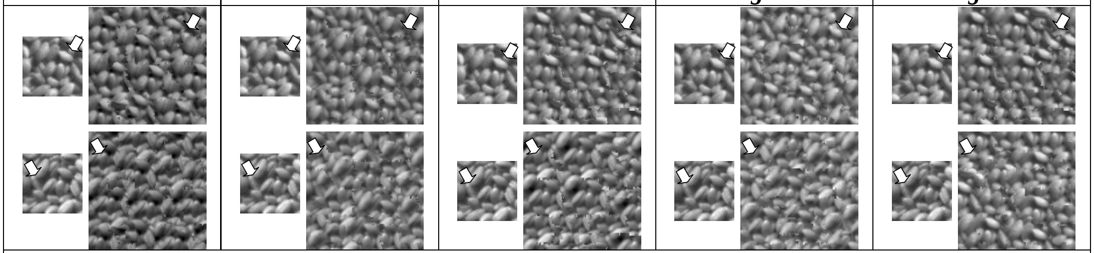

Texture “adc”

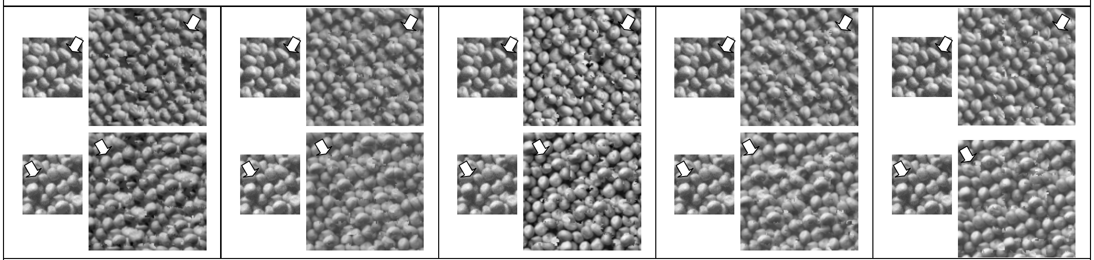

Texture “add”

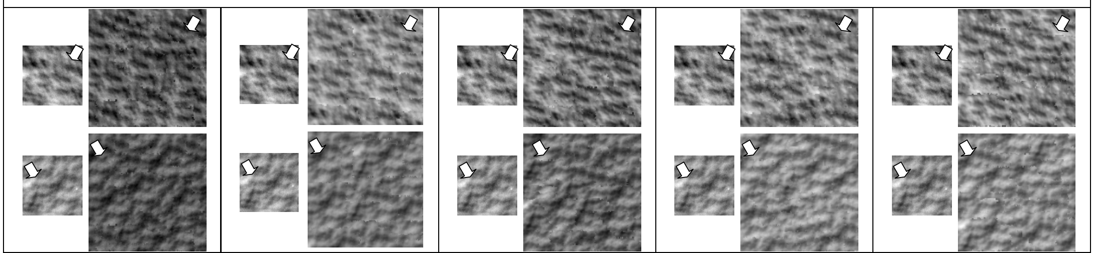

Texture “aai”

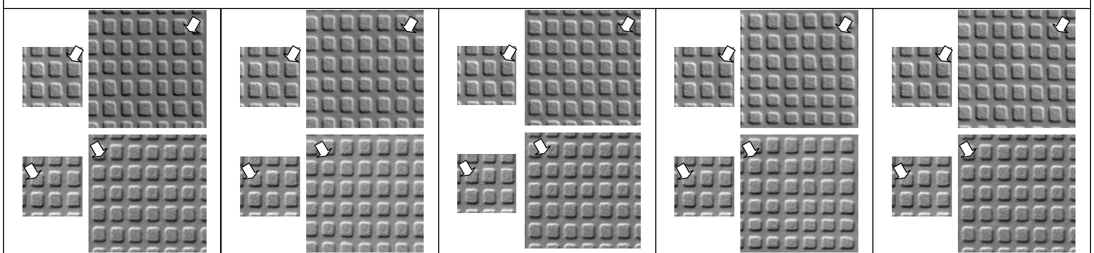

Texture “acd”

3I	*Gradient*	*PTM*	*Eigen3*	*Eigen6*
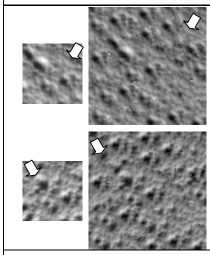	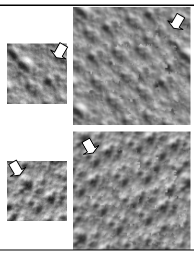	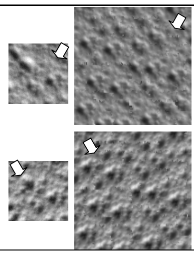	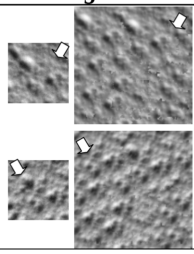	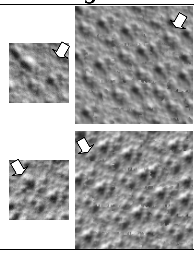
Texture “aar”				
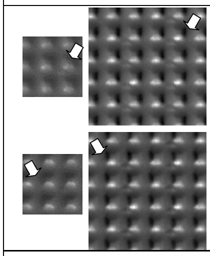	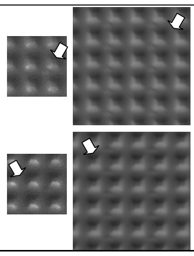	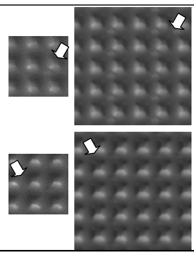	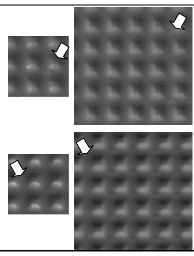	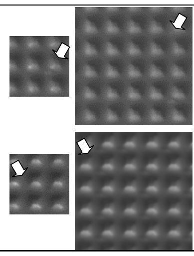
Texture “ach”				
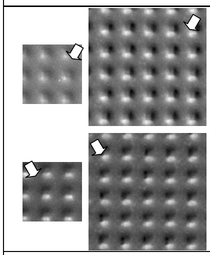	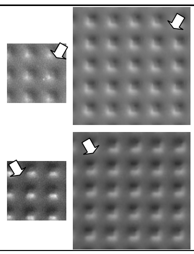	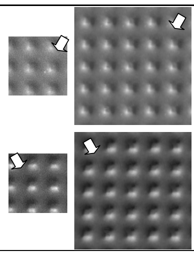	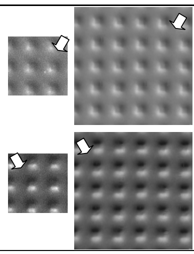	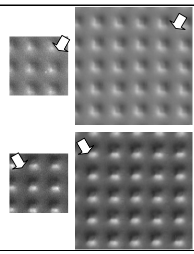
Texture “aci”				
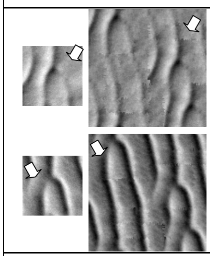	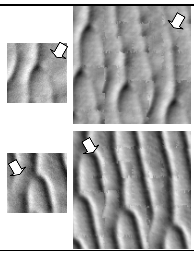	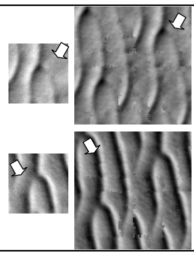	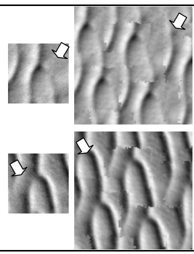	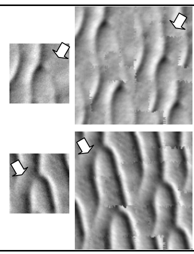
Texture “abj”				

6.3.2. The test of significant difference—Friedman's nonparametric two-way Analysis of Variance

We firstly would like to know whether there are significant differences between the performances of these approaches according to the rankings. Since observers performed their rankings independently, we use Friedman's nonparametric two-way Analysis of Variance (ANOVA) to test for significance.

Friedman's nonparametric two-way Analysis of Variance (ANOVA) is designed to determine if we may conclude from sample evidence that there are differences between treatment effects (which in our case are the five approaches). We therefore construct a matrix which contains one column for each approach. Each column contains 110 rank data (10 observers x 11 textures). Friedman's test compares the means of these columns (see [Daniel1990] for more details). The null hypothesis H_0 is that there are no significant differences between the five methods, while the alternative hypothesis H_1 is that at least one is different. The test statistic is defined as:

$$\chi_r^2 = \frac{12}{bk(k+1)} \sum_{j=1}^{k} \left[R_j - \frac{b(k+1)}{2} \right]^2 \qquad (6.3.1)$$

where:

b is total number of rank data for each method (110)

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

R_j is the sum of rank data for each method.

The test results indicated that there is at least one method which performs significantly differently from the others at a confidence level of 100%.

6.3.3. The multiple comparison

Since there is significant difference between the performances of these approaches, we are interested in which approaches perform better than others. We therefore use 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 [Hochberg1987]. The multiple comparison compares each pair of approaches and outputs the confidence interval for the difference at certain confidence level.

We use the Statistic Toolbox in Matlab to perform the multiple comparison. The result is shown in Figure 6.3.1. 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.

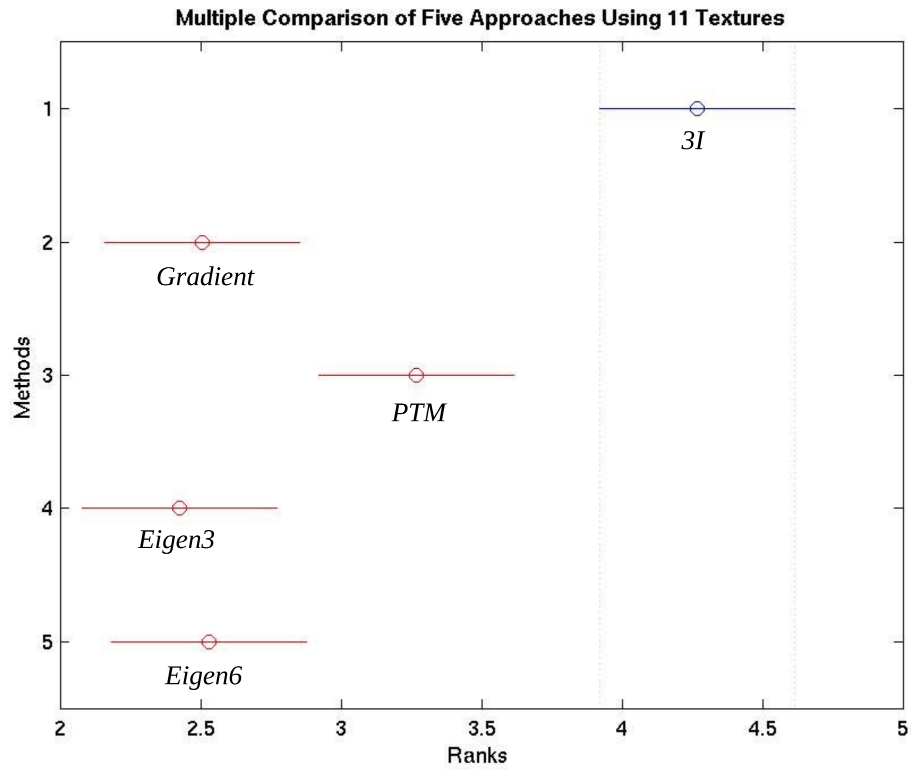

Figure 6.3.1 Multiple comparison test of the five approaches. Small circles and lines represent the group means and their intervals. The horizontal axis indicates rank values. Two means are significantly different if their intervals are disjoint.

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* approaches. However, each of these methods does outperform both *3I* and *PTM*, while the *PTM* method outperforms the *3I*.

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 is 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. Consequently, when the samples and results with same illumination directions are being compared, the effect due to discontinuities might counteract the good performance produced in relighting. Therefore, the overall performance of *Eigen6* is lowered to the same level as *Eigen3* and *Gradient* in the qualitative assessment. Correspondingly, although *PTM* performed better than *Gradient* in the relighting assessment, it failed to outperform *Gradient* in the qualitative comparison of synthesis results.

If we take computation complexity into account, we find that synthesis in $\mathbf{R}^6$ space is of course the most expensive. It exactly doubles the computation time compared with $\mathbf{R}^3$ synthesis. Thus we conclude that the *Gradient* and *Eigen3* approaches on average offer as good a performance as of any of the other methods and incur low computational cost. However, if image-acquisition requirements have to be kept low then the *3I* synthesis approach, which uses only three photometric images, provides relighting at the cost of lower quality output.

6.4. Conclusion

In this chapter, we proposed five approaches for the synthesis and relighting of 3D surface texture. The five approaches—*3I, Gradient, PTM, Eigen3* and *Eigen6* use the corresponding surface representations of a sample texture as input to a modified version of Efros and Freeman's image quilting method. The synthesised surface representations are relit to produce new images under arbitrary illumination directions. For the *3I, Gradient,* and *Eigen3* approaches, synthesis is performed in $\mathbf{R}^3$ space, while the *PTM* and *Eigen6* approaches perform synthesis in $\mathbf{R}^6$ space.

We qualitatively compared the five approaches by employing psychophysical experiments. We asked ten observers to rank different synthesis approaches by comparing output images with input sample images. The ranked data were first tested using Friedman's nonparametric two-way Analysis of Variance. The test suggests that there is at least one significant difference between the performances of these five approaches. A multiple comparison was then applied to determine which approaches outperform others. The conclusion is that, at the confidence level 99%, the *Gradient*, *Eigen3* and *Eigen6* approaches perform better than *3I* and *PTM*. If computation complexity is taken into account, the *Gradient* and *Eigen3* approaches are preferable.