Chapter 2

Literature Survey

The purpose of this chapter is to provide an overview of the research fields relevant to this thesis. Three fields will be surveyed; they are: (1) 3D surface texture synthesis, (2) 2D texture synthesis and (3) surface representation methods for relighting. These research fields will be reviewed in more detail later in the thesis when required by the context.

As introduced in chapter 1, *3D surface texture synthesis* techniques can synthesise new texture images under different viewing and lighting conditions. The input sample data for 3D surface texture synthesis can be a set of intensity images or representations of the sample texture. The synthesised results can be relit using illumination directions and viewing angles that differ from those used in original sample images. Few publications are available so far in this research area.

In contrast to 3D surface texture synthesis, the terminology *2D texture synthesis* is exclusively used in this thesis to refer to synthesising a large image from a small intensity image of the sample texture. Thus, this term is equivalent to *texture synthesis*, which is commonly used in computer vision and graphics communities. In this thesis, we also use *texture synthesis* to refer to *2D texture synthesis* since the former appeared in most relevant literature. There are many publications in this research area.

We use the terminology *surface representation methods for relighting* to refer to the techniques that can extract surface representations from a set of images and relight (render) these representations using illumination conditions that differ from those of the original. We also use the term *surface relighting representations,* *surface representations for relighting,* or *surface representation maps* to refer to the extracted representations.

2.1. Three-dimensional surface texture synthesis

Since the main objective of this thesis is to develop inexpensive and reliable approaches for the synthesis of real-world 3D surface textures, we first present a detailed review in this area. There are only five publications that can be classified into this area. They are [Zalesny2000], [Zalesny2001], [Liu2001], [Tong2002] and [Leung2001].

Zalesny and Van Gool's work

Zalesny and Van Gool in [Zalesny2001] present a multi-view texture model which can synthesise new texture images under different viewpoints. These synthesised images can catch the effect of foreshortening due to changing viewpoints. They propose a compact model that captures the first and second order statistics of different pixel pairs, which are named *cliques* [Zalesny2000]. For each clique type, the histogram of pixel value difference is calculated. The sample texture is first modelled for a single viewpoint, typically a fronto-parallel one. The result image is initialised by an independent noise with pixel values uniformly distributed in the range of sample image. Then, different clique types are collected to form a neighbourhood structure. In order to synthesis a texture image with a novel viewpoint, the neighbourhood structure is deformed by contracting and stretching according to the angle between the two views. Clique types in the deformed neighbourhood structure are used to extract new statistical parameters—difference histograms—from the sample image with the desired viewpoint. Finally, these statistical parameters are combined with the deformed neighbourhood structure to generate the result image. During synthesis process, statistics of each clique type in the neighbourhood structure are forced to keep consistent between the result image and the sample image.

Their work did produce a compact multiview texture model that can capture viewpoint dependencies in the appearance of textures. They do not however, consider varying illumination which is the focus of this thesis.

Leung and Malik's work

The earliest publication that considers varying illumination in 3D surface texture synthesis probably is [Leung2001], in which Leung and Malik use 3D textons to represent the visual appearance of real-world surface textures. They first apply a set of linear Gaussian derivative filters on 20 images of a sample 3D surface texture with different viewing/lighting conditions (from CUReT database [Dana1999a]). Then they generate 3D textons that associate with appearance vectors containing the outputs of the filters. Each pixel in any sample image can be labelled with a 3D texton that associates with an appearance vector in a 960 dimensional space. The 3D textons can be used to reconstruct novel images under varying lighting/viewing conditions. Although they did mention that 3D textons can be used in the synthesis of 3D surface textures by modifying the 2D texture synthesis algorithm proposed in [Efros1999], the computation is very expensive because synthesis has to be performed in the 960 dimensional space. Furthermore, the algorithm in [Efros1999] uses Sum of Square Differences (SSD) as the similarity measurement, which produces large errors when matching is performed in a high-dimensional space. Few synthesis results are shown in their paper.

Liu et. al. *'s work*

Liu *et. al.* in [Liu2001] also exploit the CUReT database to develop a method for generating Bidirectional Texture Functions (BTFs). They firstly select and register four sample images from the CUReT image database, and then apply a shape-from-shading algorithm to recover the sample surface height and albedo maps by assuming the Lambertian reflectance. These are used to synthesise a larger height map and image *templates* by applying the 2D texture synthesis algorithm proposed in [Efros1999]. In order to produce the final image with a novel viewing/lighting condition, a reference image with the same viewing/lighting condition is selected from the BTF database and transformed into a grey scale image with the histogram equalised to that of the template image. Finally, the result image is synthesised by matching and copying blocks between the sample reference image and the template image.

Several limitations exist in Liu *et. al.* 's method [Liu2001]. Firstly, the method requires the registration of images because images in CUReT database are not registered. This is never a trivial task and can not guarantee every texture in the database can be successfully registered. Secondly, they assume the Lambertian reflectance on the surface texture in order to perform shape-from-shading. Consequently, some real-world textures with non-Lambertian reflectance can not be used as input due to this assumption. Furthermore, applying shape-from-shading assumes integratibility on the surface, which does not always hold for real-world surfaces [Tong2002]. Finally, a sample reference image has to be used to provide pixel values for the output synthesised BTFs with the desired viewing and lighting conditions. This requires additional computation and memory space to store the sample reference image. Nevertheless, this paper shows realistic rendering results for Lambertian surfaces and is the most relevant to our work described in this thesis. We show the flow chart of this work in Figure 2.1.1.

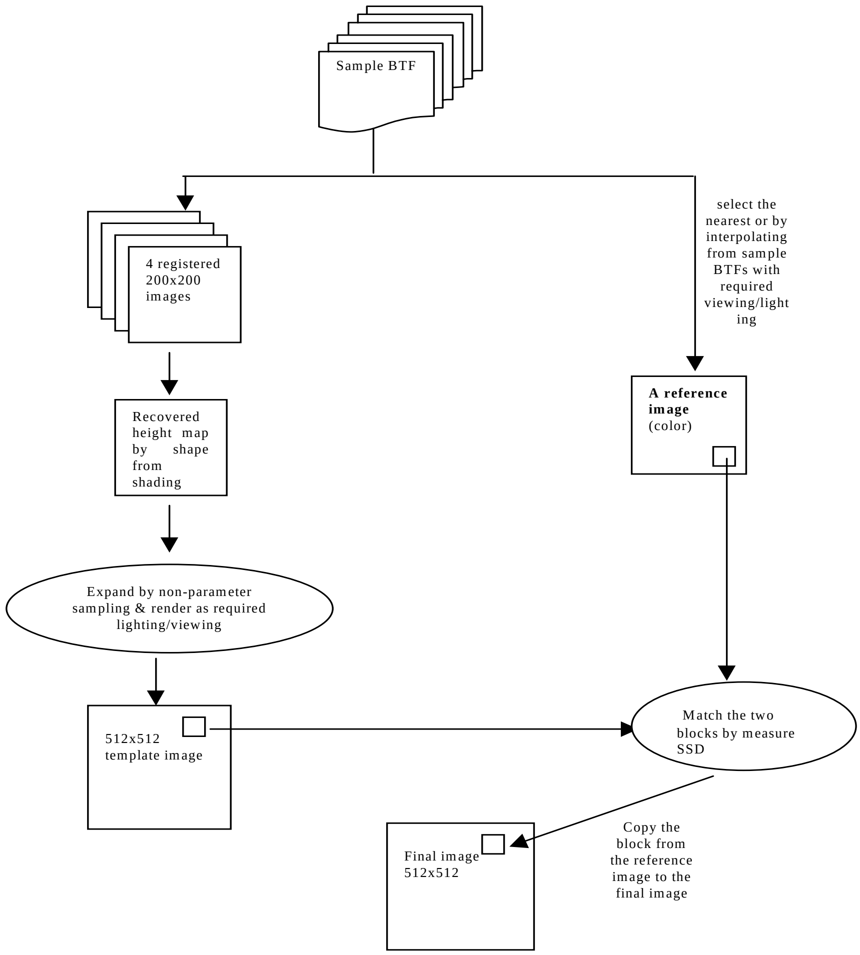

Figure 2.1.1 The flowchart of the method introduced in [Liu2001]

Tong et. al.*'s work*

Later work by Tong *et. al.* can synthesise BTFs on arbitrary surfaces by using surface textons [Tong2002]. Surface textons are defined by linear combinations of appearance vectors associated with 3D textons [Leung2001]. Tong *et. al.* suggest in [Tong2002] that the method proposed by [Liu2001] is not suitable for the synthesis of BTFs on arbitrary surfaces, because it is time consuming to reconstruct/render the appearance from the recovered sample geometry for all lighting and viewing settings. In addition, they suggest that it is impractical to directly synthesise 3D textons and reconstruct BTFs [Leung2001] on the surface of a 3D model because of the huge memory space required for storing appearance vectors. Thus, they pre-calculate the dot product for each pair of appearance vectors and store the results in a matrix. This matrix is then used for searching the best-matched pixel in sample BTFs for each vertex while the appearance vectors are discarded. Nevertheless, they still apply a fast searching algorithm for acceleration. The typical time consumed by their algorithm is 45 minutes for generating 3D textons and 21 minutes for synthesising a 96×96 image with 250k vertices on a 700Mhz Pentium III.

To summarise:

We have reviewed five available publications related to 3D surface texture synthesis. Zalesny and Van Gool present a multi-view texture model which can synthesise new texture images under different viewpoints with a fixed illumination direction [Zalesny2000 & Zalesny2001]. Leung and Malik propose the use of 3D textons to synthesise new images under arbitrary viewpoints and illuminations with expensive computation [Leung2001]. Liu *et. al.* apply a shape-from-shading technique to recover the surface heightmap under the Lambertian assumption and then use it for the synthesis of BTFs [Liu2001]. In later work, Tong *et. al.* introduce a method to synthesise BTFs on arbitrary surfaces by using 3D textons [Tong2002]. However, these techniques are computationally complex.

In contrast to previous work, our main objective in this thesis is to develop inexpensive approaches for the synthesis of 3D surface textures under varying illumination directions. We wish the synthesised texture representations to be capable of being loaded into graphics hardware and rendered in real-time on a modern desktop personal computer.

2.2. Two-dimensional texture synthesis

Although very few publications are available in the research field of 3D surface texture synthesis, many 2D texture synthesis techniques have been published during the past two decades. This section presents a brief survey of these 2D synthesis techniques. We will further review the relevant publications in more detail in chapter 5.

In [Xu2001], Xu *et. al.* present a short review on recent 2D texture synthesis approaches based on the underlying stochastic mechanisms employed by the sampling algorithms. Following their work, we also divide available publications on 2D texture synthesis into two groups according to sampling strategies. The first group employs global sampling strategies, which decide result pixel values by matching global statistics between the sample and result images in feature space. The second group uses local sampling strategies, which decide result pixel values by matching local statistics. Many different techniques have been used by the two sampling strategies. These techniques produce significantly different synthesis results and synthesis speeds. In later chapters, we will show that the taxonomy of 2D texture synthesis literature is related to the development of inexpensive approaches for 3D surface texture synthesis.

2.2.1. Texture synthesis methods based on global sampling strategies

A *global sampling strategy* means a texture synthesis algorithm generates result pixel values by matching global statistics between the sample and result images in feature space. The feature space is normally the multi-dimensional space spanned by feature images, which are produced by imposing a set of filters on the sample image; it may also be the 1D real space in which the pixel intensities of the sample image lies. This sampling strategy is called *ensemble sampling* in [Xu2001].

Two-dimensional texture synthesis is highly related to modelling a sample texture in terms of texture perception, which was pioneered by Julesz's conjecture. Julesz suggested that the Nth-order joint empirical densities of image pixels, e.g. the co-occurrence statistics for intensities, can statistically characterise a sample texture

[Julesz1962]. This has promoted a great deal of research in texture synthesis that employs global sampling strategies. These texture synthesis methods synthesise an output image according to statistical models. The models are derived from the sample image and employ a set of statistics. The output image is generated using the same statistics as those of the sample.

The majority of texture synthesis approaches employing global sampling strategies combine the use of statistical models with a bank of filters and multiresolution image representations. The multiresolution representations can capture long-range and nonlinear spatial interactions and therefore reduce the computational complexity. The sample image is first transformed into a multiresolution representation, and then the result image is synthesised by matching statistics across all resolutions. Heeger and Bergen use the steering pyramid and the Laplacian pyramid for texture synthesis by matching histograms between the sample and result pyramids [Heeger1995]. Their method fails to synthesise textures with distinguishable features, e.g. highly structured textures. De Bonet uses the Laplacian pyramid and analyses the input texture by computing the joint occurrence across multiple resolutions in the feature space [De Bonet1997]; the output texture is generated by sampling successive spatial frequency bands from the input texture, conditioned on the similar joint occurrence of features at all lower spatial frequencies. Van Nevel develops a texture synthesis method that relies on matching the first and second order statistics of wavelet subbands [Van Nevel 1998]. Based on joint statistics of complex wavelet coefficients in the multiresolution framework, Portilar and Simoncelli propose a parametric texture model that can synthesise a wide range of artificial and natural textures [Portilla2000]. In [Copeland2001], Copeland *et. al.* use the gray-level co-occurrence (GLC) model coupled with multiresolution data structure for texture synthesis. They also employed ten human observers to test the correlation between the synthesis results and their texture similarity metric by performing psychophysical experiments. In [Campisi2002], the Circular Harmonic Functions are used to develop a mutiresolution approach for texture synthesis. It essentially extends previous work in [Jacovitti1998] by using multiresolution decomposition.

There are also several texture synthesis methods that employ statistical models derived from filtered images without explicitly using multiresolution image representations. Eom proposes a 2D moving average (MA) model for texture synthesis and analysis [Eom1998], and the result image is generated in frequency domain by using estimated parameters of the 2D MA model. Jacovitti *et. al.* use hard-limited Gaussian process to develop a twin stage texture synthesis-by-analysis [Jacovitti1998]. Zhu *et. al.* present a definition of texture as the *Julesz ensemble*, which is the set of all images sharing identical statistics, and texture synthesis is achieved by sampling the ensemble using a Markov chain Monte Carlo algorithm [Zhu2000]. Histograms of feature images are employed in their approach.

To summarise:

For 2D texture synthesis, a *global sampling strategy* decides result pixel values by matching global statistics between the sample and result images in feature space. Among 2D texture synthesis approaches employing global sampling strategies, the majority apply multiresolution decomposition techniques and impose filters in multiresolution domain to generate the statistical descriptions of the sample image. The synthesis is then performed by matching statistics across multiple resolutions in feature space [Heeger1995, De Bonet1997, Van Nevel1998, Portilla2000, Copeland2001 & Campisi2002]. Only few methods directly apply a bank of filters on the sample image without explicitly using multiresolution decomposition; the result image of these methods is synthesised by matching statistics in feature space [Eom1998, Jacovitti1998 & Zhu2000]. Table 2.2.1 shows the summary of typical texture synthesis methods employing global sampling strategies.

Table 2.2.1Characteristics of typical global sampling methods

Reference	Global statistics	Number of pyramid levels	Iterat-ions	Complexity/ time-consumed/ speed
[Heeger1995]	Marginal histograms	4	5	Faster than [Portilla2000]
[DeBonet1997]	Joint occurrence of features	Depends on the sample size	1	Slower than [Heeger1995]
[Eom1998]	Moving average model parameters, elongation and orientation parameters	1	1	unspecified
[Nevel1998]	Mean, histograms and the correlation matrix	3	1	2 minutes for 400 largest entries in the correlation matrix using a Sun UltraSparc
[Portilla2000]	Marginal Statistics, coefficient correlation, magnitude correlation and cross-scale phase statistics	3	50	20 minutes for a 256x256 image using 500Mhz Pentium workstation
[Zhu2000]	Marginal histograms of filtered responses	Unspecified	20 to 100	Slower than [Portilla2000] according to [Xu2001]
[Copeland2001]	Co-occurrence matrix	3	5 (spin-flip algorithm)	2.5 minutes using a Silicon Graphics Indy with a IP22 processor
[Campisi2002]	First and second order statistics	3 to 7	2 or 3	Computational complexity depends on the number of filters and iterations. Time-consumed is not specified.

2.2.2. Texture synthesis methods based on local sampling strategies

A *local sampling strategy* means the texture synthesis algorithm generates result pixel values solely by using local information in the sample and result images. A typical example is to compute local conditional distributions using certain neighbourhoods and synthesise pixels in the result image in raster order. The majority of texture synthesis methods with local sampling strategies make certain statistical assumptions. We further divide these synthesis approaches into two sub-classes. One sub-class explicitly uses parametric statistical models for the synthesis. The other uses non-parametric methods.

Representative texture synthesis approaches using local sampling strategies and parametric models include [Cross1983, Popat1993, Bader1995, Zhu1998, Zhang1998b and Kokaram2002]. These methods first estimate the parameters of the assumed statistical models for the input sample image, and then synthesise the result image using the statistical models. Cross and Jain use Markov random field models to represent the sample image [Cross1983]. Popat and Picard present a method that first performs clustering analysis on the sample data and then calculates the probability mass function using Gaussian parameters for texture synthesis [Popat1993]. Bader *et. al.* propose the use of scalable data parallel algorithms for the 2D texture synthesis using Gibbs random fields [Bader1995]. Zhu *et. al.* develop a Markov random field model based on feature images, which are produced by a bank of filters with large image lattice; the result image is synthesised by using a Gibbs sampler [Zhu1998]. Zhang *et. al.* exploit the wavelet autoregressive model and radial basis functions in a multiresolution domain for texture synthesis [Zhang1998b]. Kokaram estimates the parameters of 2D autoregressive models and uses the models to synthesise missing gaps in images [Kokaram2002].

Non-parameteric texture synthesis approaches have the advantage that the estimation of parameters in statistical models is not necessary. Thus, the computational complexity is normally lower compared with their parametric counterparts. In particular, the method proposed by Efros and Leung is widely used in texture synthesis research [Efros1999][1]. It assumes a Markov random field model and calculates the conditional distribution of a pixel given all its neighbours by querying the sample image and finding all similar neighbourhoods. It further inspired the work in [Wei2000], which improved the performance of the original algorithm by employing a multiresolution image representation and an accelerating algorithm. The methods in these two publications can produce excellent results while simplifying the whole synthesis process. Based on these two algorithms, several texture synthesis approaches have been developed and applied in different areas [Hertzmann2001, Efros2001, Parada2001, Ashikhmin2001, Harrison2001, Tonietto2002, Zelinka2002, Cohen2003, Zhang2003 and Nealen2003].

Other typical non-parametric approaches unrelated to the two algorithms proposed in [Efros1999 and Wei2000] include [Paget1998, Ashlock1999, Bar-Joseph2001, Xu2001, Liang2001 and Gousseau2002]. In [Paget1998], Paget and Longstaff propose a non-causal, non-parametric and multiscale Markov random field model for 2D texture synthesis; they employ the Parzen-window density to estimate the frequency of occurrence. In [Ashlock1999], generic algorithms are used to track the basic texture elements and produce a non-parametric partially ordered Markov random field model for texture synthesis. In [Bar-Joseph2001], Bar-Joseph *et. al.* construct a tree representation of the input signal in multiresolution domain and generate a new tree representation by learning and sampling the conditional probabilities of the paths in the original. Their method can synthesise static and time-varying textures. In [Xu2001], a patch-pasting algorithm is introduced for the fast texture synthesis. Later work in [Liang2001] extends it by sampling patches according to a non-parametric estimation of the local conditional MRF density function; the performance is also improved. More recently, Gousseau presents a texture synthesis method by sampling from level sets [Gousseau2002].

To summarise:

Texture synthesis approaches based on local sampling strategies have attracted much attention in recent years. Several parametric methods have been proposed to firstly model the sample image and then synthesise the result using the

[1] Note: in [Efros2001], it has been pointed that a nearly identical algorithm was proposed in [Garber1981] but discarded due to its then computational intractability.

parameters [Popat1993, Bader1995, Zhu1998, Zhang1998 and Kokaram2002]. However, many researchers employ non-parametric methods that are capable of producing promising results with less computation [Efros1999, Wei2000, Hertzmann2001, Efros2001, Parada2001, Ashikhmin2001, Tonietto2002, Bar-Joseph2001, Xu2001, Liang2001, Gousseau2002, Zelinka2002, Cohen2003 and Nealen2003]. In particular, the algorithms in [Efros1999 and Wei2000] have promoted further work in different research directions.

2.2.3. Summary

In section 2.2.1 and 2.2.2, we reviewed recent publications on 2D texture synthesis. These publications can be divided into two classes depending on whether global or local sampling strategies are used. Most texture synthesis approaches with global sampling strategies synthesise a result image by matching global statistics in feature space and multiresolution domain. Among texture synthesis methods with local sampling strategies, both parametric models and non-parametric models can be used. Recent publications suggest that some non-parametric texture synthesis methods can produce good synthesis results with less computation.

2.3. Surface representation methods for relighting

As introduced in chapter 1, varying the illumination directions can produce significant effects on images of a 3D surface texture. These images can exhibit remarkable differences, which present challenges in both computer vision and computer graphics. It is therefore important to extract surface representations of the sample texture under arbitrary illumination directions. Once the representations are available, they can be relit to generate new images with arbitrary lighting conditions. This section briefly reviews relevant publications in this research area, which involves reflectance distribution modelling, model-based and image-based relighting (rendering) techniques.

2.3.1. Extracting surface relighting representations using reflectance models

The most accurate surface relighting representations can be described by Bidirectional Reflectance Distribution Functions (BRDF) [Nicodemus1977]. With full BRDF data, images of the sample surface or object under arbitrary illumination can be produced. However, full BRDF data are difficult to obtain because the measurement of BRDF is very expensive and time-consuming. Various local-based reflectance models have been used in computer vision and computer graphics to describe how lights are reflected from a surface and reach to the observer. Commonly used models include the Lambertian model, the Torrance-Sparrow model [Torrance1967], the Phong model [Phong1975], the Cook-Torrance model [Cook1982], the Nayar model [Nayar1991] and other models [He1991 and Oren1994]. Obviously, extracting surface representations using reflectance models is equivalent to estimating the models' parameters, which normally represent surface geometric and material properties. However, these models can only be seen as approximations of the ground-truth, as the physics of light reflection involves extremely complicated nonlinear processes.

Methods for estimating the parameters of reflectance models has been extensively investigated in recent years. Photometric stereo is one of the major techniques used to obtain surface geometric and material properties [Woodham1981, Horn1989, Nayar1990, Kay1995, Rushmier1997, Saito1996 and Lin1999]. This approach requires a fixed camera, several lighting conditions and a static object. Traditional photometric stereo methods assume the Lambertian reflectance function and use three images to obtain surface gradient maps and an albedo map [Woodham1981 and Horn1989]. If the sample surface exhibits both diffuse and specular components, more complex reflectance models are required. Consequently, more images are needed in order to estimate the parameters [Nayar1990, Kay1995, Rushmier1997, Saito1996 and Lin1999]. By firstly separating diffuse and specular components, both diffuse and specular parameters can be estimated using photometric stereo techniques. The combined use of range and intensity data is another popular technique that can be used to extract surface relighting representations from existed reflectance models [Ikeuchi1991, Lu1995, Sato1997, Ramamoorthi2001 and Nishino2001]. For example, Sato *et. al.* use multiple range images to recover surface shape and then estimate reflectance parameters of the Torrance-Sparrow model [Sato1997]. Polarisation techniques can also be use to separate reflection components so that surface representations can be estimated [Nayar1996].

The surface geometric representations estimated from reflectance models are usually expressed as surface normals or surface gradient maps. Extracting surface normals from an intensity image is also the aim of shape-from-shading [Horn1989]. Integration techniques can be further used to obtain the depth information or the height map from surface normals [Klette1996]. Both local and global integration approaches have been proposed in the past [Coleman1982 and Frankot1988]. Global approaches are more robust to noise than local approaches [Gullón2002].

2.3.2. Extracting surface relighting representations using other techniques

There are also a great number of other techniques that can be used to obtain surface relighting representations without directly employing reflectance models. The surface relighting representations derived from these techniques are not, in general, geometrical and material properties of the surface.

Image-based relighting (rendering) techniques can generate realistic images from pre-recorded images without using complex rendering processes as in geometry-based computer graphics [Kang1997, McMillan1999, Koudelka2001, Lin2002, Matusik2002 and Wong2002]. In [Kang1997], Kang presents a survey on early image-based rendering techniques. In [Matusik2002], Matusik *et. al.* introduce a system that can acquire and display high quality graphical models of objects using opacity hulls; both effects produced by changing view and illumination conditions are considered. In [Wong2002], Wong *et. al.* define the plenoptic illumination function that can relight images while supporting view interpolations. However, many image-based rendering techniques can only synthesise new images under different viewpoints, while the illumination remains fixed [Chen1995, Levoy1996 and Gortler1996].

The representation of varied BRDF on a surface requires numerous sample images. Researchers have developed several methods to approximate this model by projecting these images into general base functions so that the representation is more compact for practical applications [Lalonde1997, Lafortune1997 and McAllister2002]. Lalonde and Fournier use wavelet coefficients to represent large anisotropic BRDF data sets [Lalonde1997]. The Lafortune representation consists of a diffuse component and several specular lobes which are generalised Phong lobes [Lafortune1997]. McAllister *et. al.* employ the Lafortune representation for rendering the Spatial BRDFs using register combiners in an Navidia Geforce 4 graphics card [McAllister2002].

Eigen-based methods are broadly used to extract surface relighting representation [Epstein1995, Zhang1998a, Georghiades1999 and Nishino2001]. These methods apply principal component analysis (PCA) or singular value decomposition (SVD) on a set of pre-recorded images and extract base images as the surface relighting representations. New images under arbitrary illumination directions can be generated by linearly combining these base images. Obviously, eigen-based approaches also belong to the class of image-based techniques. In addition, they can be used in pattern recognition and image impression [Nishino2001, Turk1991 and Belhumeur1997].

In the literature regarding surface representation methods, many other mathematical models are also exploited to express the sample images as linear or nonlinear combinations of a set of base functions, such as Fourier Series [Huang1984 and McGunnigle2001], spherical harmonics [Basri2001 and Ramamoorthi2001] and steering functions [Ashikhmin2002]. These base functions normally form base images and can be used to synthesise new images under arbitrary illumination conditions.

2.3.3. Extracting 3D surface texture representations for relighting

Rough surface textures can be seen as a finer scale geometric description with regular or random components. In theory, methods surveyed in section 2.3.1 and 2.3.2 can all be used to extract relighting representations of 3D surface textures.

Nevertheless, researchers have proposed special methods to represent 3D surface textures under arbitrary illumination directions.

Representing the appearance of 3D surface textures only received attention in recent years [Koenderink1996, Stavridi1997, Dana1999a, Dana1999b, Leung2001, Malzbender2001 and Ashikhmin2002]. In [Dana1999a], Dana *et. al.* define Bidirectional Texture Function (BTF) that can represent 3D surface textures under varied illumination and viewing directions; they construct the CUReT database that contains many images from over 60 samples. Dana and Nayar further investigate three BTF models, including the histogram model, the correlation model and PCA models [Dana1999b]. Leung and Malik exploit the CUReT database and employ a bank of 48 filters coupled with clustering analysis to derive 3D textons, which can be used to represent and recognise the visual appearance of 3D surface textures [Leung2001]. Malzbender *et. al.* propose a quadratic lighting model that uses Polynomial Texture Maps(PTM) to reconstruct the surface colour under varying lighting conditions [Malzbender2001]. Ashikhmin uses a set of steering basis functions for relighting bumpy surfaces [Ashikhmin2001].

2.3.4. Summary

We have presented a brief review of methods that can be used to extract surface relighting representations from a set of pre-recorded images. As the most compact representations, surface geometric and material properties can be obtained by estimating the parameters of various locally-based reflectance models [Woodham1981, Horn1989, Nayar1990, Kay1995, Rushmier1997, Saito1996, Lin1999, Ikeuchi1991, Lu1995, Sato1997, Ramamoorthi2001 and Nishino2001]. They can then be relit using corresponding reflectance models to generate new images under different illumination conditions. Image-based relighting/rendering are also commonly used techniques that can convert the pre-recorded images into relighting representations [Kang1997, McMillan1999, Koudelka2001, Lin2002, Matusik2002 and Wong2002]. Other methods employ mathematical models to express a set of sample images using linear or nonlinear combinations of basis functions, such as eigen-based methods [Epstein1995, Zhang1998a, Georghiades1999, and Nishino2001], Fourier serious [Huang1984] and spherical harmonics [Basri2001].

In recent years, special interest has been given to the research into representing the appearance of 3D surface textures. Several methods have been proposed and shown great promise in computer vision and computer graphics [Koenderink1996, Stavridi1997, Dana1999a, Dana1999b, Leung2001, Malzbender2001 and Ashikhmin2002].

2.4. Conclusion

This chapter has briefly reviewed the related research fields to this thesis. These comprise the literature on:

(1) 3D surface texture synthesis approaches,
(2) 2D texture synthesis approaches, and
(3) surface representation methods for relighting.

Based on this survey, we conclude that very few publications are available regarding 3D surface texture synthesis, while there are a great number of methods in the fields of 2D texture synthesis and the extraction of surface representations for relighting.

Among the 3D surface texture synthesis approaches, Zalesny and Van Gool's work can only synthesise new images with varied viewpoints, while the illumination direction is fixed [Zalesny2000 and Zalesny2001]. Liu *et. al.* use a 2D texture synthesis algorithm based on [Efros1999] and Lambertian surface representations for the synthesis of BTFs [Dana1999a]; In [Tong2002] and [Leung2001], a 2D texture synthesis algorithm based on [Efros1999] and the 3D texton representations are combined for the synthesis of BTFs. However, these methods require expensive computation.

In contrast, our main objective in this thesis is to develop inexpensive approaches for the synthesis and relighting of 3D surface textures. In next chapter, we will introduce a basic framework that can combine 2D synthesis approaches with surface representation methods in a methodical manner to synthesise new texture images under arbitrary illumination directions.