Chapter 3

Framework

3.1. Introduction

The goal of this thesis is to develop inexpensive approaches for synthesis and relighting of 3D surface textures. In chapter 2, we presented a survey and showed that few publications are available regarding 3D surface texture synthesis. However, many surface representation techniques and 2D texture synthesis methods have been published in recent years. The aim of this chapter is therefore to propose an overall framework for the synthesis and relighting of 3D surface textures. The framework will be capable of combining 2D texture synthesis methods with surface representation techniques in a methodical manner. Based on this framework, we will define the data environment that we need for all experiments in the thesis.

The overall framework comprises three parts:

1. extraction of a 3D surface representation from multiple images of the texture sample;
2. use of the representation to synthesise a description of a larger area of the surface texture; and
3. rendering (or relighting) of the synthesised surface representation according to a specified set of lighting conditions.

For assessment purposes, we employ a set of images selected from the PhoTex database [McGunnigle2001]. The database contains many images per texture that have been captured under varied illumination directions. These images form the basic data environment for this thesis.

The chapter is organised as follows. Section 3.2 introduces the framework. Section 3.3 describes the data environment that we use for assessment. Finally we summarise the content of the whole chapter in section 3.4.

3.2. A framework for the synthesis of 3D surface textures

In this section, we introduce the overall framework for the synthesis of 3D surface textures.

3.2.1. Framework

The synthesis of 3D surface textures naturally deals with more information than its 2D counterpart. The latter 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 surface texture synthesis either implicitly or explicitly requires generation of surface geometric information and reflectance properties. A single sample image, which is used as input in 2D texture synthesis, does not normally provide enough information regarding surface topology and reflectance. Thus we have to employ multiple images (or their representations) of the sample 3D surface texture, which contain enough information regarding geometric and reflectance properties as the input data. Furthermore, we would like the sample image data set to be captured in an inexpensive way, i.e. using off-the-shelf digital cameras, and the synthesis results to be capable of being rendered in real-time on current desktop machines.

Our main goal is to develop inexpensive approaches for the synthesis and relighting of 3D surface texture. However, the original image set is normally of large dimension. It is impractical to directly synthesise the original image set in a large dimensional space, because the time consumed by texture synthesis algorithms increases linearly with the dimensionality of input sample vectors [Xin2002][Efros1999][Turk2001]. On the other hand, surface representation techniques can be used to convert the original image set into relighting representations of low dimension. For example, Malzbender *et. al.* used a method to convert 50 images into six Polynomial Texture Maps [Malzbender2001]. A PCA-based method suggests 5 ± 2 eigenimages are sufficient to represent arbitrary lighting for the Lambertian and specular lobe although complex reflectance will increase the dimensionality [Epstein1995]. Thus, 3D surface texture synthesis can use the representations of the sample texture which are of lower dimension as input.

Therefore, our framework for the synthesis of 3D surface textures comprises 3 stages:

Stage 1: Extraction of the 3D surface representation

The first stage is to extract a suitable relighting representation of the sample 3D surface texture. Surface representation (i.e. image-based relighting) methods can be applied to a set of images captured under different illumination directions. The representation should be of a low dimension and preferably capable of per-pixel-rendering using a linear combination, which requires less computation and is compatible with modern graphics systems. Many techniques may be used for this stage, as reviewed in chapter 2. A simple example is that we may generate two surface gradient maps and an albedo map by applying traditional photometric stereo techniques [Woodham1981] and use them to represent a Lambertian surface. Thus, each pixel location will be represented by a multi-dimension vector.

Stage 2: Synthesis of the 3D surface representation

In the second stage, we need to select and modify a 2D texture synthesis algorithm so that it can use multi-dimensional vectors as input and perform the synthesis in $\mathbf{R}^{\mathbf{n}}$ space, where $\boldsymbol{n}$ is the dimensionality of the surface representation. In the case of a Lambertian surface, we can use two surface gradient maps $p(x, y)$, $q(x, y)$ and an albedo map $al(x, y)$ to represent the sample 3D surface texture; the synthesis will take the three dimensional vector $V = (p(x, y), q(x, y), al(x, y))$ as input, where (x, y) is the pixel co-ordinate.

During the synthesis process, it is important to preserve the correlation between representation maps. Suppose we are synthesising surface gradient and albedo maps and the pixel values in the result representations originate from those in the samples. At a pixel location (x'_0, y'_0) in the result map, the surface gradient values $p(x'_0, y'_0)$, $q(x'_0, y'_0)$ and albedo value $al(x'_0, y'_0)$ must derive from the sample surface gradients $p(x_0, y_0)$, $q(x_0, y_0)$ and albedo $al(x_0, y_0)$ at the same location (x_0, y_0) on all input sample maps.

Stage 3: relighting

This is the final stage of the overall framework. We relight the synthesised representations using a specified set of lighting conditions to produce desired texture images. Relighting can be seen as an inverse process to the extraction of the representation maps. Recall the previous example. Since we extract surface gradient and albedo maps using the Lambertion reflectance model, the relighting will use the same model as well.

Figure 3.2.1 illustrates the overall framework.

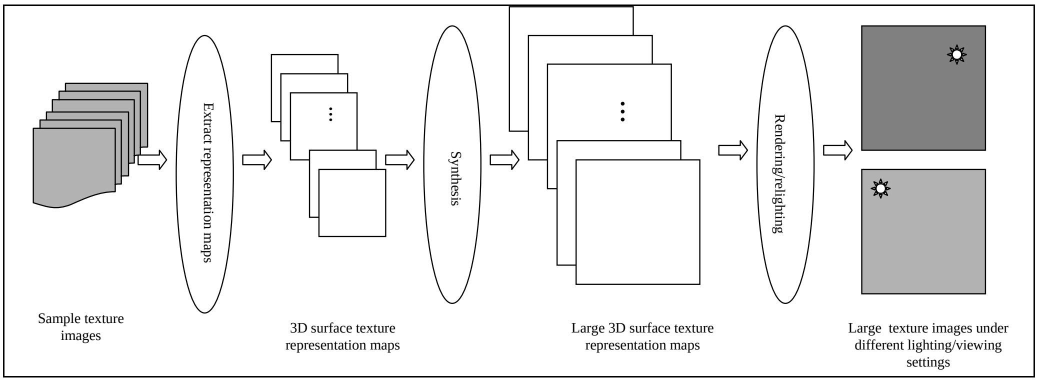

Figure 3.2.1 The overall framework for the synthesis of 3D surface textures

3.2.2. An example: the approach of [Liu2001]

Previous work on 3D surface texture synthesis can be represented within the proposed framework. We present an example by using the method introduced in [Liu2001].

Shum and his colleagues [Liu2001] exploited the CURet database to develop a method for the synthesis of Bi-directional Texture Functions. In the first stage, they applied a shape-from-shading algorithm to recover sample surface height and albedo maps under the assumption of Lambertian reflectance. Next, the 2D texture synthesis algorithm proposed in [Efros1999] is used to produce a large height map. In the final (relighting) stage, the synthesised height map is rendered to generate image *templates.* A reference image with desired lighting/viewing conditions is selected from the sample BTFs. The result image is obtained by matching and copying blocks between template images and the reference sample images. Figure 3.2.2 shows how this method relates to our framework.

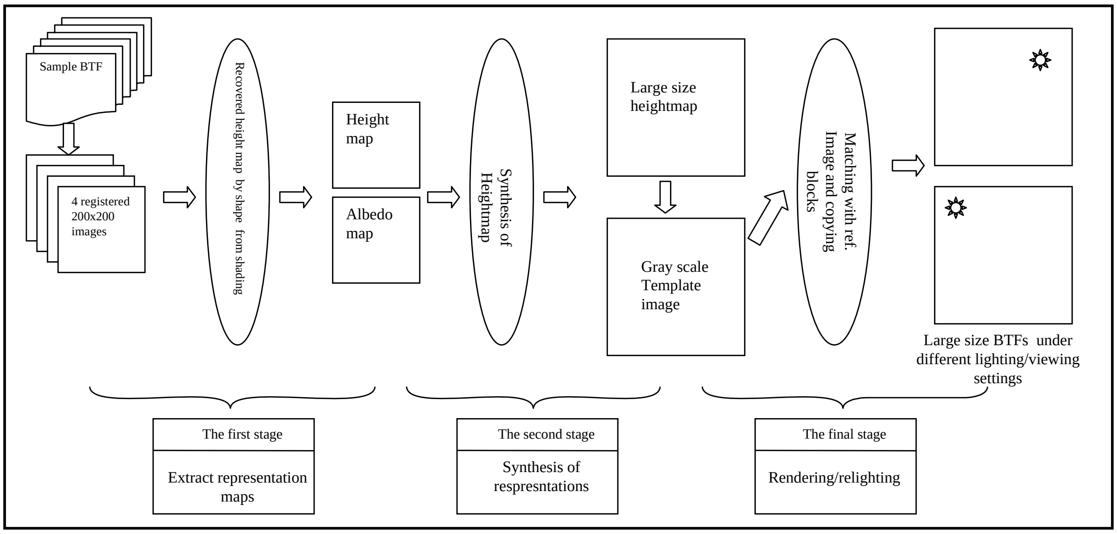

Figure 3.2.2 The method introduced in [Liu2001] can be represent within our framework.

3.2.3. Discussion

It should be noted that we could use many alternative techniques in different stages of the framework. In the first stage, many methods are available to directly obtain relighting representations of the sample surface (without applying image-based techniques) [Rushmiere1998]. For example, 3D surface geometry can be acquired by using a 3D scanner; BRDF data of the sample texture might be measured although it is particularly difficult. Once we have the sample geometry and BRDF data, we can synthesise new geometry and corresponding BRDF data for a large surface area. Then the illumination techniques may be applied to produce the final images.

3.3. The image data environment for the thesis

According to the proposed framework, we firstly need to extract a suitable representation of the sample texture from a set of pre-recorded images. This section will therefore introduce the image data sets that we will use throughout the thesis.

In this thesis, we use a set of images selected from the PhoTex database (http://www.cee.hw.ac.uk/texturelab/database/photex/). The database contains many rough surface texture samples. For each texture sample, many images were captured using a fixed camera (a Vosskühler CCD 1300LN) while the rough surface was illuminated from various directions. Figure 3.3.1 shows the experimental set-up. The origin of the co-ordinate is at the centre of the sample surface. The camera is perpendicular to the sample surface, which is globally flat and lies in the *x*-*y* plane. The camera's line of sight (axis) overlaps with the *z* axis. Thus, the illuminant direction is defined by a slant angle and a tilt angle. Slant is the angle between the *z* axis and the illuminant vector; tilt is the angle between the *x* axis and the vector produced by projecting the illuminant vector onto the *x*-*y* plane. More details about the PhoTex database can be found on its website.

All images in the database are monochromatic with 1280x1024 resolution. Each pixel is stored in 12 bits. We call the images in the database *photometric* images, where the term *photometric* is in the context of photometric stereo: i.e. inferring information about a static scene by altering the lighting conditions [Woodham1981]. With image data sets selected from this database, surface representations can be extracted using various methods.

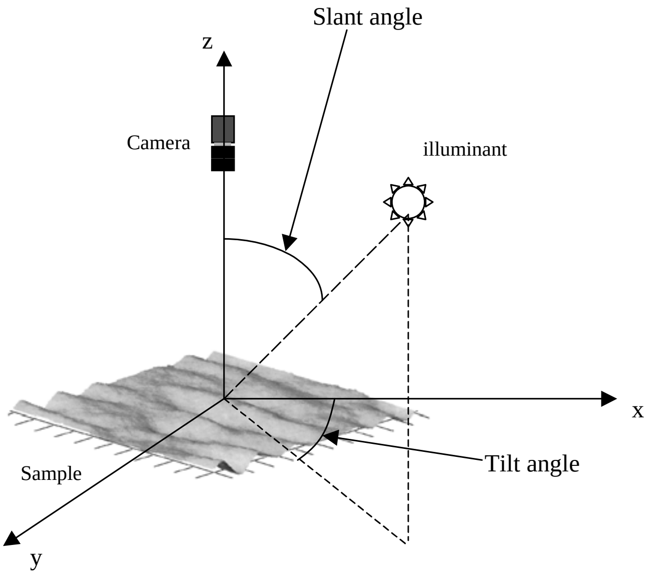

Figure 3.3.1 The imaging set-up and definitions of the slant and tilt angle.

We select image data sets according to two criteria. One is that the sample should comprise suitable granularity. Obviously, surface appearance can only be described by *texture* at certain scales [Dana1999a]. Large granularity in the sample tends to be perceived as individual elements. Suppose we are interested in the texture of rough rock surface. We would like the captured surface to cover a large area with small granularity on the rock. It makes no sense to focus on a tiny surface patch that has large granularity in the image. Thus, if original images in the database contain large granularity, it may be necessary to downsample. The other criterion is that selected texture types should cover a wide range. This is especially important if we want to evaluate our methods on real textures. The selected textures should include:

1. rough and smooth surfaces,
2. glossy and matte surfaces,
3. non-shadowing and shadowing surfaces, and
4. near-regular and stochastic patterns.

In Appendix A, we show sample images of all the selected textures.

3.4. Conclusion

In this chapter, we proposed a practical framework for the synthesis of 3D surface textures and introduced the image data environment for the thesis. The framework can combine surface representation techniques and 2D texture synthesis methods in a methodical manner to synthesise 3D surface textures. It comprises three stages: the first stage converts a set of pre-recorded images into a surface representation of a lower dimension; the second stage employs a 2D texture synthesis algorithm and extends it to $\mathbf{R}^{n}$ space; the final stage renders the synthesised representations according to a set of lighting conditions. Many surface representation techniques and 2D texture synthesis methods can be used in the framework.

Based on the framework, we introduced the image data environment used for all experiments in this thesis. We exploit the PhoTex database and select image data sets according to two criteria: the granularity and the need for a range of texture types.