Capture and Synthesis of 3D Surface Texture

Junyu Dong and Mike Chantler[1]

Abstract - **This paper presents and compares six novel approaches for capturing, synthesising and relighting real 3D surface textures. Unlike 2D texture synthesis these techniques allow the captured textures to be relit using illumination conditions, and viewing angles, that differ from those of original. Our approaches each comprise two stages: synthesis and relighting. Synthesis can be applied either before or after relighting. The relighting stage is implemented in three different ways: using image-based, gradient-based, and height-based approaches. Thus there are a total of six different ways in which we may combine these functions. We present a representative set of results selected from our experiments with 30 textures. The best images are obtained when image-based or gradient-based relighting is used after synthesis.**

I. Introduction

Research into texture synthesis is normally concerned with learning and generation of 2D images of texture [1,2, 3, 4, 5, 6, 7, 8, 14, 17]. 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 and viewpoint conditions. This presents obvious limitations for high-fidelity rendering of textures in augmented and virtual reality applications. Fig. 1 illustrates the dramatic effect that varying illumination direction can have on images of a 3D surface texture.

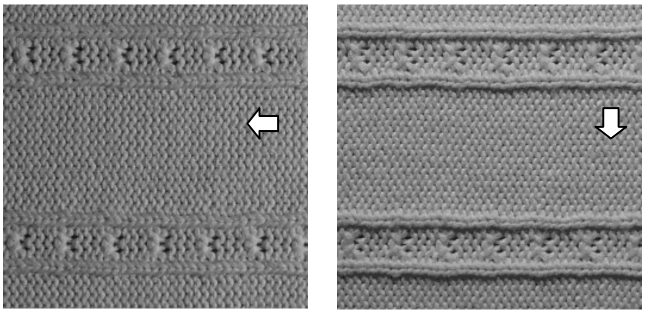

Fig. 1. Two images of a 3D surface texture imaged under differing illumination

Our objective is to develop inexpensive and reliable techniques for the capture, synthesis, and relighting of 3D surface textures for use in standard computer graphics applications. With the exception of [18, 15, 9] we believe that no other work has been published on this subject. Zalesny and Van Gool in [18] presented a multi-view texture model which can synthesise viewpoint independent new images. Leung and Malik [15] proposed the use of "3D textons" to synthesise new images with arbitrary viewpoints and illuminations. The 3D textons are derived from CUReT database images [10]. This database contains many images of texture samples taken using different view points and illuminations. In [9], Shum and his colleagues also exploited the CUReT database and present a method to generate bidirectional texture functions (BTFs) from a set of images. In contrast, we use three images per texture and a simpler (and also more limiting) bump-map representation. Our motivation for using three images is that it provides the necessary data set for classic photometric stereo [11] and that the bump-map representations [12,16] can be imported into standard graphics packages, which can be used to generate images from different viewpoints and under different illumination conditions.

Given three images of a sample texture there are two major tasks to perform: texture synthesis, and relighting. In principle these may be performed in either order. We use three different surface representations:

1. the three image photometric set;
2. the surface normal and albedo maps; and
3. the height and albedo maps.

The gradient and height maps are estimated using photometric stereo and frequency domain integration respectively. We present six major ways in which we may combine these algorithms together.

The major novelty of this paper therefore, is in proposing and evaluating these six different approaches to the synthesis and relighting of surface texture.

The remainder of this paper introduces the synthesis and relighting methods that we use, describes the six approaches that we have evaluated, and presents results and conclusions based on experiments with 30 real surface textures.

II. Capture, Synthesis and Relighting of Surface Textures

In order to generate realistic images of 3D surface textures from small samples we need to:

1. capture images of the sample texture that will enable us to extract a suitable representation of the 3D surface,
2. use the surface representation to synthesise a description of a larger surface of the required dimensions, and
3. render (or relight) the surface representation in accordance with a specified set of lighting and viewing conditions.

Phase 1 above must obviously be performed first, however, phases 2 and 3 can be reversed. We will therefore first present each of these phases in isolation. The next section describes 6 possible ways of combining these phases in detail.

A. Image capture

We capture three images of each sample taken under different illumination conditions. This allows us to estimate the surface gradient and albedo fields using photometric stereo. It also allows us to employ a simple linear interpolation scheme for image-based relighting

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

that is described later. The images are captured using a light-source located approximately at 1m from the sample. The illuminant vectors' slant angle (angle with the camera axis) is kept at 50° while the tilt angle (the angle that the illuminant vector makes when it is projected onto the plane on which the texture sample sits) is switched through 0°, 90° and 180°. The camera's position and orientation are not changed.

B. Synthesis of 3D Surface Textures

The synthesis of 3D-surface descriptions naturally deals with more information than its 2D counterpart. The former requires that both albedo and surface normal information be represented, whereas 2D-texture synthesis only requires intensity data to be encoded. In principle any 2D-synthesis method may be extended to deal with 3D-surface descriptions. Thus for the monochrome case, the 2D scalar field is replaced by either a two or a three element vector field, i.e. we are extending from an $\mathbf{R}^1$ vector space to $\mathbf{R}^2$ and $\mathbf{R}^3$ spaces. An $\mathbf{R}^2$ vector space is required to encode height and albedo data, whereas $\mathbf{R}^3$ is required to store the other two surface representations. Monochrome techniques can be generalised to colour approaches.

Our Approach. Our synthesis approach is based on Efros's image quilting method [14], which can produce remarkably good results at a smaller computational cost. The method synthesises a new image by stitching together small patches of existing sample image. The new image is generated in raster order. First, a randomly selected block from the sample image is pasted into the new image beginning at the first row and the first column. Then, for its neighbor block in raster order, an overlapping area between these two blocks is introduced to measure similarity. The neighbor block is selected by finding the minimum distance between the overlapping areas of sample and existing blocks. Finally, 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.

There are two small modifications in our approach. First, instead of locating the best-mached block using search and a distance measure, we select the corresponding neighbor of last selection. This is based on a simple idea that when a best-matched block is selected from the sample image, its neighbouring block will be its best-matching neighbor in the result image, providing that this sample neighbor block exists. This modification is similar to the method described in [17]. It has produced good results when tested on our 30 sample textures and it also reduces the computation required.

The other modification is that we perform the synthesis in $\mathbf{R}^1$, $\mathbf{R}^2$, or $\mathbf{R}^3$ space. $\mathbf{R}^1$ space synthesis is the conventional monochrome image-based synthesis - we use this to synthesise images from relit samples. $\mathbf{R}^2$ space synthesis is used to synthesise combined albedo/height data, while $\mathbf{R}^3$ space synthesis is used to synthesise triple image photometric sets, or surface normal/albedo data.

We use Euclidean distance as the basis for our similarity measures and propagate 1, 2, or 3 element vector values appropriate to the representation space.

C. Texture Relighting Methods

Given three images of a sample captured under photometric lighting conditions there are at least three ways in which the texture may be relit. These three methods are discussed below.

Image-based Relighting. Our image-based relighting method uses a simple linear combination of images that assumes a 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. In our case we have used a slant angle of 50° and tilt angles of 0°, 90° and 180°. Thus:

$$\begin{aligned} I_{(\tau,\sigma)}(x,y) = {} & \left(\frac{\cos\tau\sin\sigma}{2\sin 50^\circ} - \frac{\sin\tau\sin\sigma}{2\sin 50^\circ} + \frac{\cos\sigma}{2\cos 50^\circ}\right)\cdot I_{(0,50)}(x,y) \\ & + \frac{\sin\tau\sin\sigma}{\sin 50^\circ}\cdot I_{(90,50)}(x,y) \\ & + \left(\frac{\cos\sigma}{2\cos 50^\circ} - \frac{\cos\tau\sin\sigma}{2\sin 50^\circ} - \frac{\sin\tau\sin\sigma}{2\sin 50^\circ}\right)\cdot I_{(180,50)}(x,y) \end{aligned}$$

Gradient-based Relighting. We use photometric stereo [11] which exploits (1) to estimate the surface gradients *p*(*x*,*y*) and *q*(*x*,*y*) from the three images. This provides a 'bump-map' representation [12,16]. In our experiments we have used simple Lambertian rendering of these gradient fields to relight the surface under a particular set of illumination conditions.

Height-based Relighting. Height-maps are probably the most popular representations for bump-maps. We use a global integration in Fourier space [13], which minimises the problems, associated with integrating errors. This is an efficient, non-recursive approach, but suffers problems when the gradient data has been affected by shadows. These height-data can be used for synthesis or relit directly. Both require derivative estimation and application of (1) for relighting.

III. Six Combinations of Texture Synthesis and Relighting

Having briefly described the synthesis and relighting algorithms that we use, we now describe the ways in which we have combined them to provide six different approaches to 3D-surface texture synthesis.

Approach 1: Synthesis Followed by Image-based Relighting (Fig. 2). In this approach the three photometric images of the sample are first used to synthesise three larger photometric images of the target (using the method described in section 2.2). Image-based relighting then uses this larger image-set to produce the output image under user specified illumination conditions.

Approach 2: Synthesis Followed by Gradient-based Relighting. Theoretically this approach produces results identical to approach 1, only the data representation differs. The synthesis is performed using the images of the sample, but photometric stereo is used to provide surface gradient and albedo estimates which are then rendered using gradient-based relighting.

Table 1. Summary of the 6 approaches

	1st phase	2nd phase	3rd phase
1	$\mathbf{R}^3$ *synthesis* (produces large photometric images)	*Image-based relighting* (produces final image)	n/a
2	$\mathbf{R}^3$ *synthesis* (produces large photometric images)	*Photometric stereo* (produces large gradient and albedo maps)	*Gradient-based relighting*
3	$\mathbf{R}^3$ *synthesis* (produces large photometric images)	*Photometric stereo + integration* (produces large height and albedo maps)	*Height-based relighting*
4	*Image-based relighting* (produces small relit image)	$\mathbf{R}^1$ *synthesis* (produces final image)	n/a
5	*Photometric stereo* (produces small *gradient and albedo* maps)	$\mathbf{R}^3$ *synthesis* (produces large gradient and albedo maps)	*Gradient-based relighting*
6	*Photometric stereo + integration* (produces small height and albedo maps)	$\mathbf{R}^2$ *synthesis* (produces large height and albedo maps)	*Height-based relighting*

Approach 3: Synthesis Followed by Height-based Relighting. Again this approach is very similar to the previous approach, except that the gradient data is integrated using global integration to provide a height-map. Relighting naturally requires derivative estimation which we also perform using a global algorithm.

Approach 4: Image-based Relighting of the Sample Followed by Texture Synthesis. In this approach we use the photometric image-set of the sample to relight the sample under user specified illumination conditions. This single image is then used as the basis for the synthesis stage of the larger output image.

Approach 5: Gradient-based Synthesis Followed by Gradient-based Relighting. This approach first uses photometric stereo to provide estimates of the gradient and albedo fields for the sample. Synthesis algorithm is applied to these data, and the resulting gradient and albedo maps are used with gradient-based relighting to provide the final image.

Approach 6: Height-based Synthesis Followed by Height-based Relighting. The estimated gradient field of the sample obtained using photometric stereo is integrated to provide an estimated height-map. This is combined with the albedo data in the synthesis algorithm to provide larger height and albedo maps, which are themselves used in height-based relighting.

Summary of approaches. For convenience the six approaches described above are summarised in **Table 1**.

IV. Results

We used a database of 30 textures for testing the six approaches described above[2]. Each photometric set of a texture comprises three images taken at illumination tilt angles of 0°, 90° and 180° as shown in Fig. 2a. This figure also shows the intermediate photometric image set generated in approach 1. It clearly shows that the illumination effects have been preserved by the texture synthesis. The output images were obtained using the linear relighting algorithm described in section 2.2. These again show effect of different tilt angles.

A. Comparison of Approaches

Fig. 3.a shows output images for five different textures for each of the six methods. For comparison purposes they have all been synthesised using an illuminant tilt of 45°. From these results we make the following observations:

- Approaches 1, 2 and 3 produce the very similar results. This is not surprising given that the only difference between them is that the second and the third approach contain additional functions: non-linear transform (photometric stereo and integration) followed by its inverse (gradient-based rendering and height-based rendering). However, they do allow the use of a bump-map surface representation.
- Heavily shadowed textures might cause problems in Approaches 2 and 3. Furthermore, since the integration uses a global approach, we find that if the differentiation process used is a local operation - as one might expect to be used in a rendering algorithm, the rendering results for some textures are not as good as those produced by using a global differentiation.
- In general it seems to be better to apply the synthesis stage before any other processing. That is it is better to 'grow' the photometric image samples into larger images straightaway (as in approaches 1-3) rather than apply synthesis to an intermediate representation as is the case for approaches 5-6.
- In approach 4, synthesis is performed in an $\mathbf{R}^1$ intensity space. Each time in order to get a new image under a different illumination, the synthesis process has to be performed. Because the synthesis algorithm introduces random patches from the sample, those new result images might look like taken from different patches of the sample texture, instead of images from the same patch under different illuminations.
- In approach 5, synthesis is performed in gradient/albedo space - here the higher frequencies, and hence noise, are enhanced.
- In approach 6, synthesis is performed in height/albedo space which is likely to propagate integration errors.

Therefore the overall conclusion is that using synthesis before relighting, and avoiding integration (approaches 1 & 2) provides the best results. Fig. 4. shows result images with different slant and tilt illuminant angles by using approach 1.

V. Conclusions

We have proposed six different approaches for the synthesis and relighting of 3D surface texture. They are straightforward to implement and unlike [15, 9] only require a photometric set of three images. However, they are limited to Lambertian surfaces. We have evaluated these approaches using samples of 30 real surface textures.

Our preliminary conclusions are:

[1] Some of the textures are available at http://www.cee.hw.ac.uk/texturelab/database/dbase/. The remainder will be made available shortly.

1. that synthesis using photometric data is to be preferred over synthesis from intermediate representations as the propagation of errors introduced by integration etc. is reduced;
2. that it makes no difference if a gradient/albedo or three-image photometric representation is used in for relighting - except that bump-maps are widely accepted, and
3. that the best results were obtained using approaches 1 and 2, that is synthesis followed by either image or gradient-based relighting. However, the bump map representation allows the surface texture to be applied to 3D models (Fig.3b) and to be viewed from different directions.

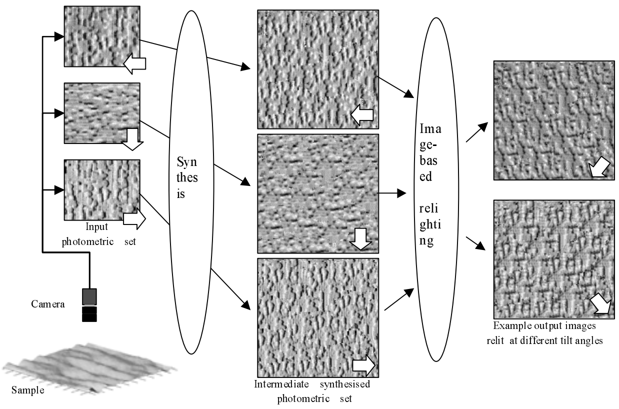

(a)

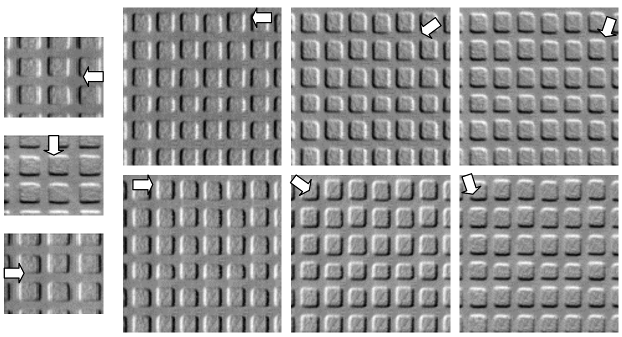

(b)

Fig. 2. (a) shows the input photometric samples, the synthesised photometric image set, and finally two output images relit at different tilt angles by using approach 1; (b) shows 3 samples and 5 final output images with 5 different illuminant tilt angles. (The illuminant tilt angles are shown by the white block arrows)

References

1. D.J. Heeger, J.R. Bergen. “Pyramid-based texture analysis/synthesis” , *Proceedings. International Conference on Image Processing (Cat. No.95CB35819). IEEE Comput. Soc. Press. Part vol.3, 1995, pp.648-51 vol.3*.
2. S.C. Zhu, X.W. Liu, Y.N. Wu. “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, June 2000, pp.554-69.*
3. J.S. De Benet. “Multiresolution sampling procedure for analysis and synthesis of texture images”, *Computer Graphics Proceedings, SIGGRAPH 97. ACM. 1997, pp.361-8.*
4. L. Wei, M. Levoy. “Fast texture synthesis using tree-structured vector quantization”, *Computer Graphics Proceedings. Annual Conference Series 2000. SIGGRAPH 2000. Conference Proceedings. ACM. 2000, pp.479-88.*
5. A.A. Efros, T.K. Leung. “Texture synthesis by non-parametric sampling”, *Proceedings of the Seventh IEEE International Conference on Computer Vision. IEEE Comput. Soc. Part vol.2, 1999, pp.1033-8 vol.2*.
6. Z. Bar-Joseph, “Statistical Learning of Multi-Dimensional Textures”, *MSc thesis, Institute of Computer Science, The Hebrew University of Jurusalem, June, 1999*.
7. J. Portilla and Eero P. Simoncelli, “Texture Modeling and Synthesis using Joint Statistics of Complex Wavelet Coefficients”, *Proceedings of the IEEE Workshop on Statistical and Computational Theories of Vision, Fort Collins, CO, June,1999.*
8. Y. Xu, S.C. Zhu, B. Guo, H.Y. Shum. “Asymptotically Admissible Texture Synthesis”, *Second International Workshop of Statistical and Computational Theories of Vision, Vancouver, Canada, July 13, 2001.*
9. X. Liu, Y. Yu and H.Y. Shum, “Synthesizing Bidirectional Texture Functions for Real-World Surface”, *SIGGRAPH2001*.
10. K. J. Dana, B. Van Ginneken, S. K. Nayar, J. J. Koenderink, “Reflectance and Texture of Real-World surfaces”, *ACM Transactions on Graphics, Vol. 18, No. 1, January 1999, Pages 1-34*.
11. R. Woodham, “Analysing images of curved surfaces”, *Artificial Intelligence, 17:117-140, 1981.*
12. J.F. Blinn. “Computer Display of Curved Surfaces”, *Ph.D. Thesis, Uversity of Utah, 1978.*
13. R. T. Frankot and R. Chellappa, “A Method for Enforcing Integrability in Shape from Shading Algorithms”, *IEEE Transactions on Pattern Analysis and Machine Intelligence, VOL. 10, No. 4, July 1988, Page(s):439-451.*
14. A. A. Efros and W. T. Freeman, “Image Quilting for Texture Synthesis and Transfer”, *Proceedings of SIGGRAPH ’01, Los Angeles, California, August, 2001.*
15. T. Leung and J. Malik, “Representing and Recognizing the Visual Appearance of Materials Using Three-dimensional Textons,” *International journal of Computer Vision, Vol 43, Number 1, pp 7-27, June 2001.*
16. H. Rushmeier, G. Taubin and A. Gueziec, “Applying Shape from Lighting Variation to Bump Map Capture," *Proceedings of the Eighth Eurographics Rendering Workshop, Eds. J. Dorsey and P. Slussallek, Saint-Etienne, France, June 16-18, 1997, pp. 35-44.*
17. M. Ashikhmin, “Synthesizing Natural Textures”, *The proceedings of 2001 ACM Symposium on Interactive 3D Graphics, Research Triangle Park, NorthCarolina March 19-21, pp. 217-226.*
18. A. Zalesny and L. Van Gool, “Multiview Texture Models”, *IEEE Computer Soc. Conf. on Computer Vision and Pattern Recognition (CVPR’01), December 2001.*

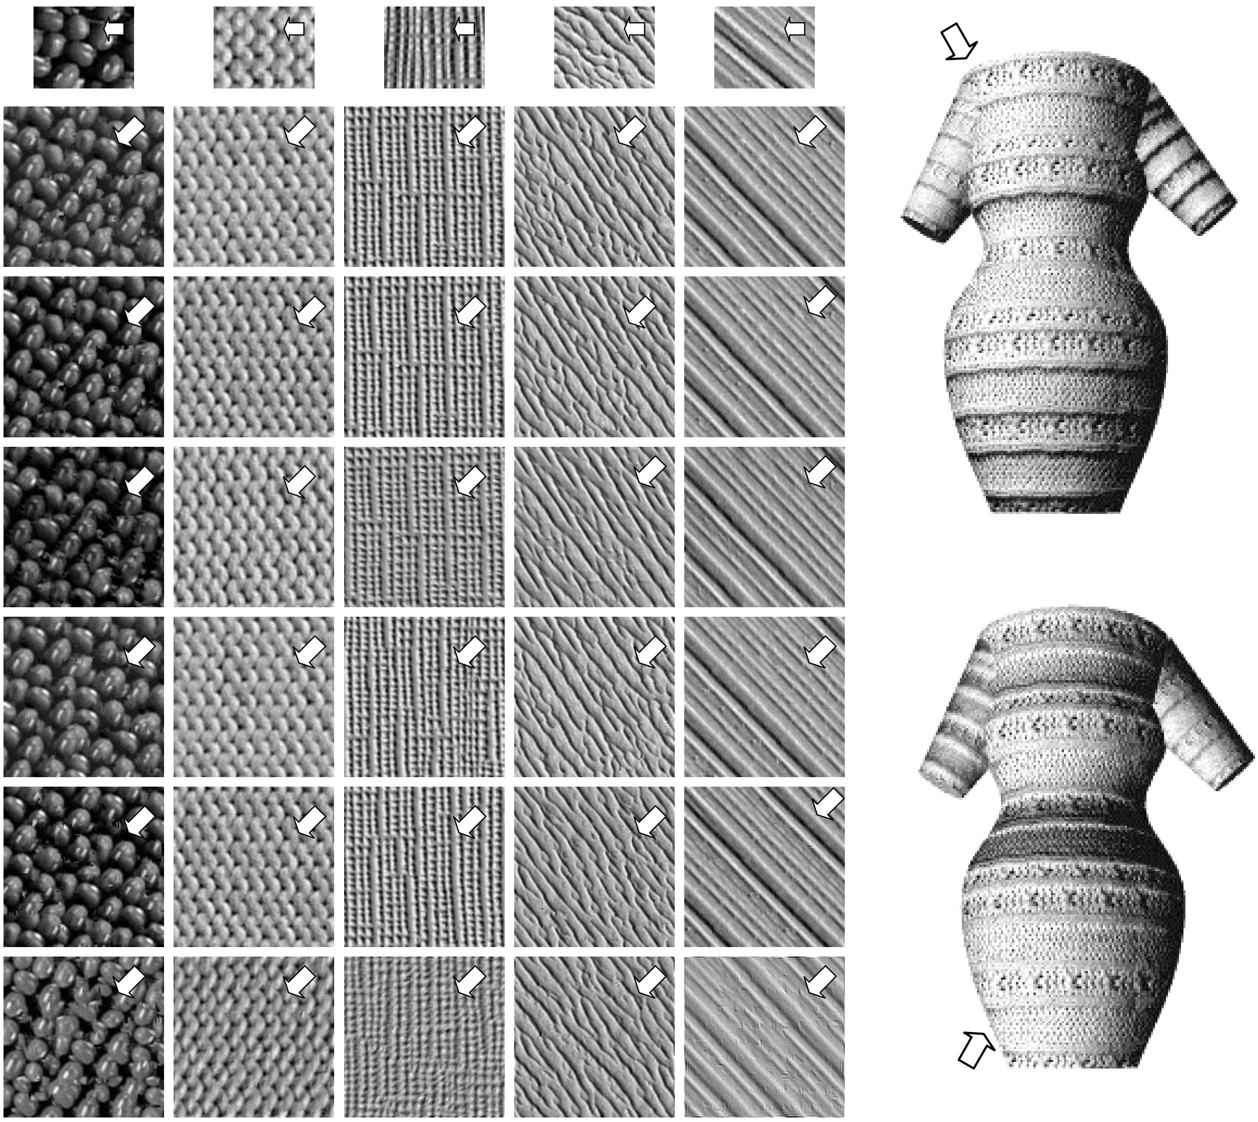

(a) (b)

Fig. 3.

(a) Comparison of output images of the 6 approaches for 5 example textures. The first row shows the 0° tilt image taken from the input sample photometric set for each texture. The next 6 rows show the corresponding outputs synthesised by using approaches 1 to 6 for a tilt angle of 45° (as indicated by the block arrows).

(b) Results obtained from overlaying height and albedo maps, synthesised using approach 3, onto a 3D model. The result was rendered using Micrografx Simply 3D. Block arrows show the different illumination directions.

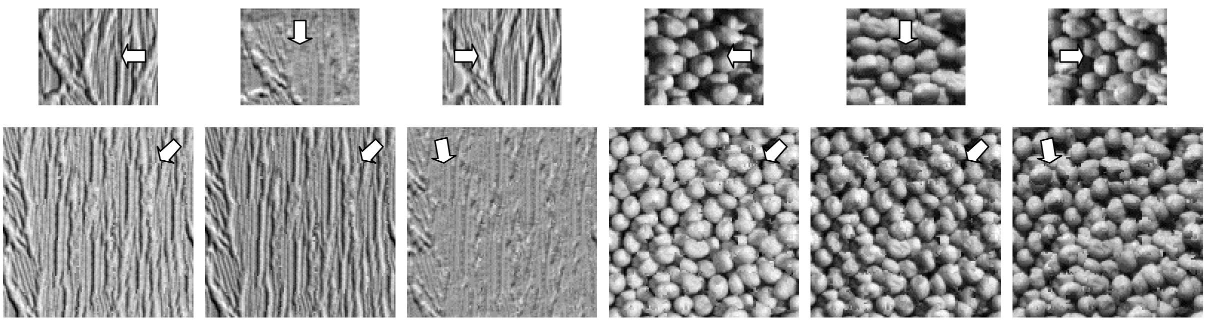

Fig.4. Results of different illuminant tilt and slant angles. The first row shows 6 sample images of two textures. Each 3 samples for one texture are taken from 0°, 90° and 180° respectively. The second row, from left to right for each texture, shows the synthesised images with illumination tilt angle 45° and slant angle 20°, illumination tilt angle 45° and slant angle 45°, illumination tilt angle 100° and slant angle 45° (tilt angles are indicated by the block arrows).