Photometric Stereo and Painterly Rendering

G. McGunnigle M.J. Chantler

G. McGunnigle and M.J. Chantler are with the Department of Computing and Electrical Engineering, Heriot Watt University, Riccarton, Edinburgh EH14 4AS, United Kingdom. E-mail: gmg@cee.hw.ac.uk

Abstract

We present a technique that enhances an existing algorithm for painterly rendering. The technique models the topography of the painted surface and incorporates it into the rendering. Photometric stereo (PS) is used to estimate the shape of the painted surface. The statistics of real samples painted with different brush strokes are measured. The analysis was used to improve a non-photorealistic rendering algorithm. We found PS to be an effective method of surface recovery. Different stroke styles were found to have distinctive point statistics. We conclude that PS is an effective way to recover the topography of oil paintings; that the point statistics distinguish the types of brush stroke in a concise and meaningful way; and that a surface model based on the power spectrum can be used to improve the performance of a painterly rendering algorithm.

I. Introduction

Painterly rendering is the processing of a real or synthetic image to give it the appearance of a painting, Figure 1. It is an area both of research, e.g.[1],[2] and of commercial exploitation e.g.[3]. Although a painting's texture is an important component of the work we are unaware of any schemes that incorporate the topography of the painting. In this paper we introduce a statistical description of the surface into a modified form of Cabral's algorithm [4]. We believe the results are more realistic and pleasing to the eye than the original rendering.

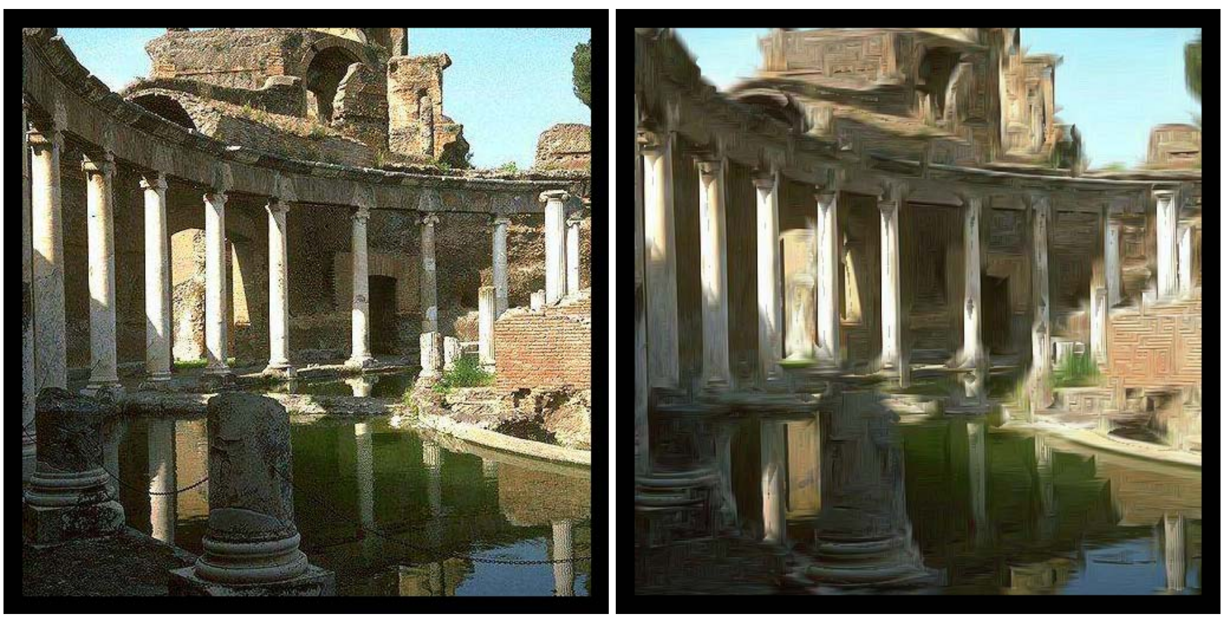

Fig. 1. Original image and Cabral's rendering.

The topography of a painting is an important component of its appearance and an artist may deliberately manipulate the texture of the paint in his own style. The use of raking light to accentuate the appearance of topography is an established technique in art conservation and authentication. However, its physical significance is not generally considered and its interpretation is subjective. This paper applies a technique from computer vision—Photometric Stereo (PS) to estimate the shape of the surface, and random process theory to model it.

We found that, using a surface model recovered with PS, we were able to accurately model the appearance of the surface—indicating that our model of the surface and its properties do fit the data. The point statistics of the different samples fit our intuitive descriptions of the surfaces. We found that the power spectrum does not significantly improve our ability to distinguish brush strokes, but is useful for modelling the surface texture. Finally, we found that by combining the surface model obtained from our analysis with a painterly rendering technique from the literature with we can improve the appearance and realism of the rendering.

II. Surface Recovery

First we recover the surface shape. Photometric stereo (PS) infers the slope of a surface patch by measuring how the intensity of the patch changes as the direction of the light source is varied. This assumes the relationship between slope, lighting and intensity is known. Algorithms differ in the complexity and type of reflectance model, however, most are based on Lambert's law. The object of most research is to make PS robust to deviations from Lambert's law, e.g. glossy highlights and shadowing.

One aim of this paper is to test if PS is suitable for oil paintings. Lambert's law describes matt reflectance, yet oil paintings are glossy. Furthermore, a region with very rough texture may have significant shadowing. Both of these effects will degrade the slope estimate, however both are affected by the zenith of the light source. Illuminating from a shallow angle, i.e. raking light, will highlight subtle surface texture and will cause most of the glossy reflection to miss the camera, however it will also increase the amount of shadowing. The optimum zenith for a particular surface will be a tradeoff between these effects. In this section we will test how well PS can cope with these effects.

We use a modified form of the photometric algorithm proposed in [5]. The azimuth of the light source is varied in 30° steps to give 12 images. The intensity of each facet is expressed as a function of azimuth. The Fourier series of each facet's function is calculated. Ideal Lambertian reflectance is confined to the mean and fundamental components of the series. By suppressing the higher harmonics we can attenuate the non-Lambertian components of reflection. That is we can reduce the effect of glossy highlights on the photometric estimation.

Since we do not know the surface shape we cannot directly assess the accuracy of estimation. However we can use our estimate of the surface to predict its appearance under given lighting conditions. By comparing our predictions with the actual images we can assess the accuracy of the surface estimate and the reflectance model, Table I. These results show that the algorithm is able to accurately predict the appearance of the surface under various lighting conditions. That is, the combination of the estimated surface and the rendering model is a good model of the observed situation.

Tilt Angle	0°	30 °	60°	90°	120°	150°
S/N (dB)	22.37	19.89	21.00	21.90	20.78	20.69
Tilt Angle	180 °	210 °	240 °	270 °	300 °	330 °
S/N (dB)	22.36	21.19	20.52	21.09	21.07	20.61

TABLE I

Accuracy of prediction

III. Surface Description

The recovered surface must be quantified. Photometric stereo estimates the surface derivatives, or slopes, of small patches of the surface. These may be integrated to estimate the surface height function, however, in our experience it is often more effective to work in terms of slope. Slope is a vector quantity consisting of two components: the rates of change of height with respect to the x and y-axes. Both components are stochastic—that is they exhibit both random and deterministic behaviour. In this section we will analyse the statistical behaviour of these variables for five surface samples, Figure III.

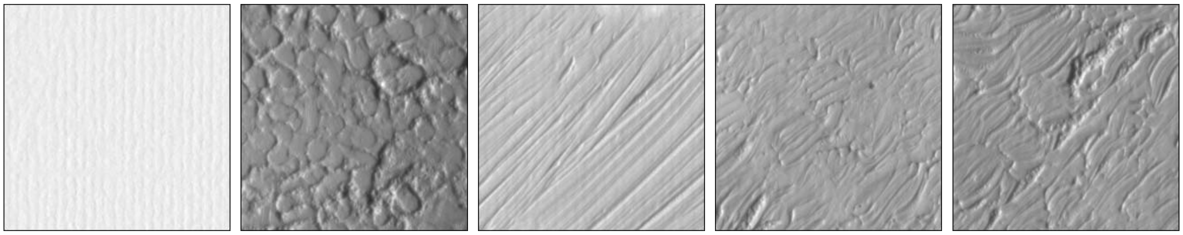

Fig. 2. Samples: paper, iso, longstroke, shortstroke1 and shortstroke2

A. Point Statistics

The point statistics of a vector can be described using the joint distribution of their components. This can be described analytically using the covariance matrix, or graphically using a scatter plot. The scatter plot of slopes reflects the properties of the surface: an isotropic surface will have a circular cluster; a directional surface will have a less symmetrical distribution; a rough surface will have a wider distribution of points than a smooth surface. The scatter plots of the test surfaces are shown in Figure 3. The relatively smooth *paper* surface has the tightest cluster, the rough *Iso* texture has the widest dispersion. This distribution is the closest to being circular, i.e. this is the most isotropic surface. The *long stroke* surface, the most directional, has the most elliptical distribution. The *Short stroke* surfaces are intermediate.

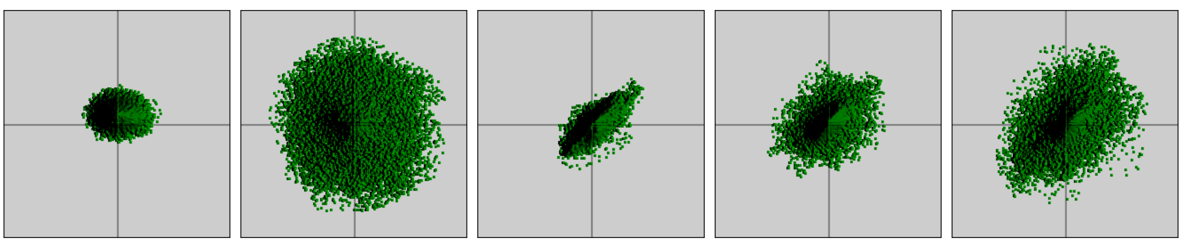

Fig. 3. Surface derivative scatter plots of the test samples.

The orientation of the distribution will follow the direction of the brush stroke. It is useful to have statistics that describe *how* directional brushstrokes are, but which are independent of the orientation of the brush stroke. We use Principal Components Analysis to decorrelate the measurements: in effect rotating the axes until they are aligned with the major axes of the distribution. The transformed distribution has a diagonal covariance matrix of the form shown in (3), where λ_1 and λ_2 are the dimensions of the major and minor axes.

$$C = \begin{bmatrix} \overline{p^2} & \overline{pq} \\ \overline{pq} & \overline{q^2} \end{bmatrix} \tag{1}$$

$$C' = \begin{bmatrix} {\lambda_1}^2 & 0 \\ 0 & {\lambda_2}^2 \end{bmatrix} \tag{2}$$

$$M = \sqrt{\lambda_1^2 + \lambda_2^2} \tag{3}$$

$$D = \frac{\lambda_1}{\sqrt{\lambda_1^2 + \lambda_2^2}} \tag{4}$$

These can be used to define a pair of measures that describe the roughness (3) and the degree of directionality of the surface (4). Values of the measures are shown for the samples in Table II, they agree with our intuitive ideas about the roughness and directionality of the surfaces. These measures are useful for quantifying the subjective impressions of the observer.

Surface	Magnitude	Directionality
Paper	0.021	0.70
Iso	0.22	0.90
Long stroke	0.092	0.43
Short stroke	0.098	0.59
Short stroke2	0.149	0.59

TABLE II

ROUGHNESS AND DIRECTIONALITY PARAMETER VALUES FOR TEST SAMPLES.

B. Surface Spectra

The point statistics are clearly distinctive for different test surfaces, however they are limited in their ability to *model* a surface. The surface spectra quantifies the relationship between *pairs* of pixels and, for Gaussian surfaces, it gives a complete description. Although our samples are not Gaussian, they may still have characteristic spectra.

The surface spectrum is two dimensional: it is convenient to resolve this into radial and polar spectra. The radial spectra of selected samples are shown in Figure 4 (left). Differences between the curves are mostly in terms of amplitude rather than form: that is they could be described more concisely using point statistics. In fact this is true of all the samples, except paper which has a well defined peak at approximately 600 m^{-1}. The polar spectra of the samples are more distinctive, Figure 4 (right), though since they only have a single peak, these too are adequately defined by point statistics.

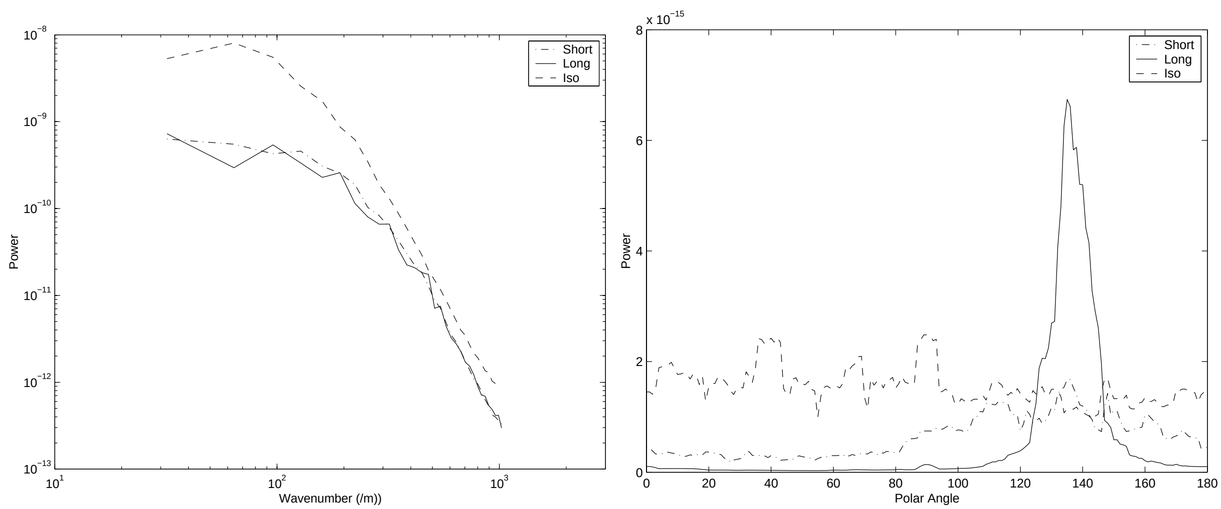

Fig. 4. Radial (left) and polar (right) power spectra of 3 samples.

IV. Painterly Rendering

Painterly rendering is a non-photorealistic computer graphics technique to form images that look as if they had been painted. The algorithm may interact with a human operator to create an image, e.g. [9], or it may transform an existing image, e.g. [4]. The algorithm may be applied to still [8] or moving [6] images. In this paper we will consider only the transformation of existing, still images. Although there is a lot of research in this area, we are unaware of anyone incorporating the surface texture of the painting in the rendering. The aim of this paper is to see whether a painterly rendering can be improved by incorporating surface texture.

There has been a great deal of work on the rendering of existing images. Researchers address the issues of brushstroke shape, placement and direction. Statistical techniques [7], gradient information from the original image [8], heuristics and even the use of arbitrary images [1] have been used as cues. Hertzmann has also pioneered the use of multiresolution techniques in this context [2] as cues. In this paper we will concentrate on one technique for painterly rendering and extend it to include surface texture.

Cabral proposes a technique for the visualisation of dense vector fields, however he also applies the technique to painterly rendering. The image is differentiated and a needle map derived from the vector field is convolved with the original image. We have implemented a simplified version, where the needles are aligned with image axes. The original image, and the painterly rendering are shown in Figure 1. The length of the needle controls the directionality of the operator, Figure IV.

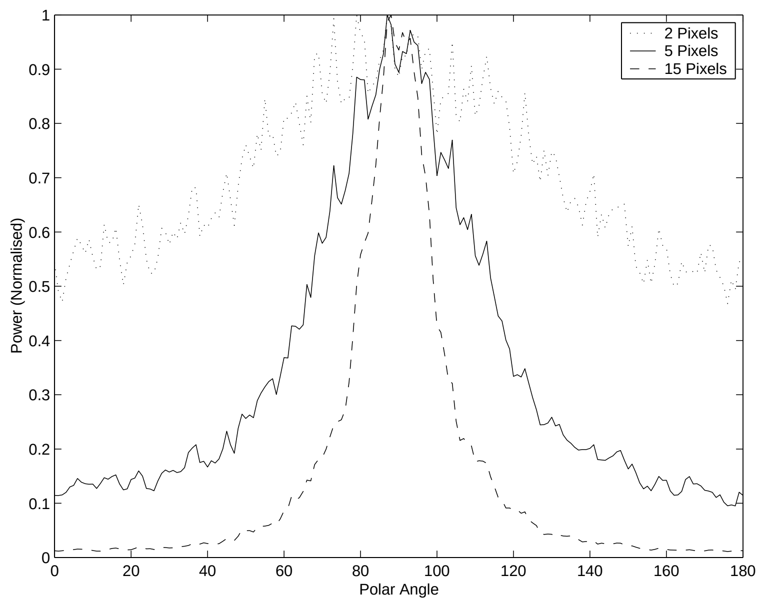

Fig. 5. Polar distribution of power with the LIC operator of various lengths.

The novel approach of this paper is to treat the output as a coloured three dimensional surface, rather than a coloured plane. We must generate a landscape that is not only of the same style as the transformed image, but is also synchronous with the brushstrokes. That is we need to generate topography with the correct spectrum, and with the same directionality and phase as the transformed image. Our approach is simple: the vector field is calculated from and convolved with the original image using the simplified Cabral method; we use the radial spectrum measured earlier to generate an isotropic random field;

finally, we convolve the vector field with the random field—in effect adaptively filtering the random field to confer the correct directionalities and phases. The transformed image is treated as a surface colouring and is rendered along with the filtered height field using Lambert's law, Figure 6 (right). We believe the resulting image is both more realistic and more aesthetically pleasing. By altering the length of the needle, and the radial spectrum, we can simulate different types of brushstroke, Figure 7.

Because the result of the rendering is a three dimensional surface, physical phenomena can now be taken into account. Both the lighting conditions and the reflectance function influence the appearance of the rendered image and we can control these to alter the image. Increasing the zenith of the illuminant "flattens" the image; decreasing the zenith accentuates the surface texture. We have used Lambert's Law—this affects the intensity of a facet but not its hue. If specularity is modelled we can alter the colour characteristics of the image: as we increase the amount of specular reflection the image becomes whiter. By using a more physical model of the painting we can introduce lighting effects that allow physically meaningful controls over the image's appearance.

V. Conclusions

We found photometric stereo to be effective at recovering the surface texture from oil paintings. The point statistics of the surface are distinctive for different types of brushstroke. They are an objective yet intuitive way of distinguishing between brushstrokes. None of the painted samples were found to have particularly distinctive power spectra, however, the power spectrum is useful for modelling the radial frequency characteristics of the painted surface. We used this information, along with phase and directional information from a real image to generate a painterly rendering of the input image. We believe the topographical component makes the rendering more realistic and attractive.

References

[1] P. Haeberli *Paint by Numbers: Abstract Image Representations*, Computer Graphics, SIGGRAPH Annual Conference Proceedings 1990, pp.207-214

[2] A. Hertzmann *Painterly Rendering with Curved Brush Strokes of Multiple Sizes*, Computer Graphics (SIGGRAPH 98 Proceedings), pp. 453-460

[3] Adobe Systems, *Photoshop 4.0*

[4] B. Cabral and L.C. Leedom *Imaging Vector Fields Using Line Integral Convolution*, SIGGRAPH 93, Anaheim, CA, ACM SIGGRAPH..

[5] G. McGunnigle and M.J. Chantler *Segmentation of Rough Surfaces Using Reflectance*, British Machine Vision Conference 2001, Vol. 1 pp.323-332

[6] B.J. Meier *Painterly Rendering for Animation*, SIGGRAPH : ACM Special Interest Group on Computer Graphics and Interactive Techniques ACM Press New York, NY, USA pp. 477 - 484, 1996

[7] S. Treavett and M. Chen *Statistical Techniques for the Automated Synthesis of Non-Photorealistic Images.*, Proc. $15^{th} Eurographics U.K. Conf. March$1997.

[8] M. Kass and A. Witkin *Analyzing Oriented Patterns.* , Computer Vision, Graphics and Image Processing, 37, 1987.

[9] B. Baxter, V. Scheib, M. Lin and D. Manocha *DAB: Interactive Haptic Painting with 3D Virtual Brushes* Proc. ACM SIGGRAPH, 2001

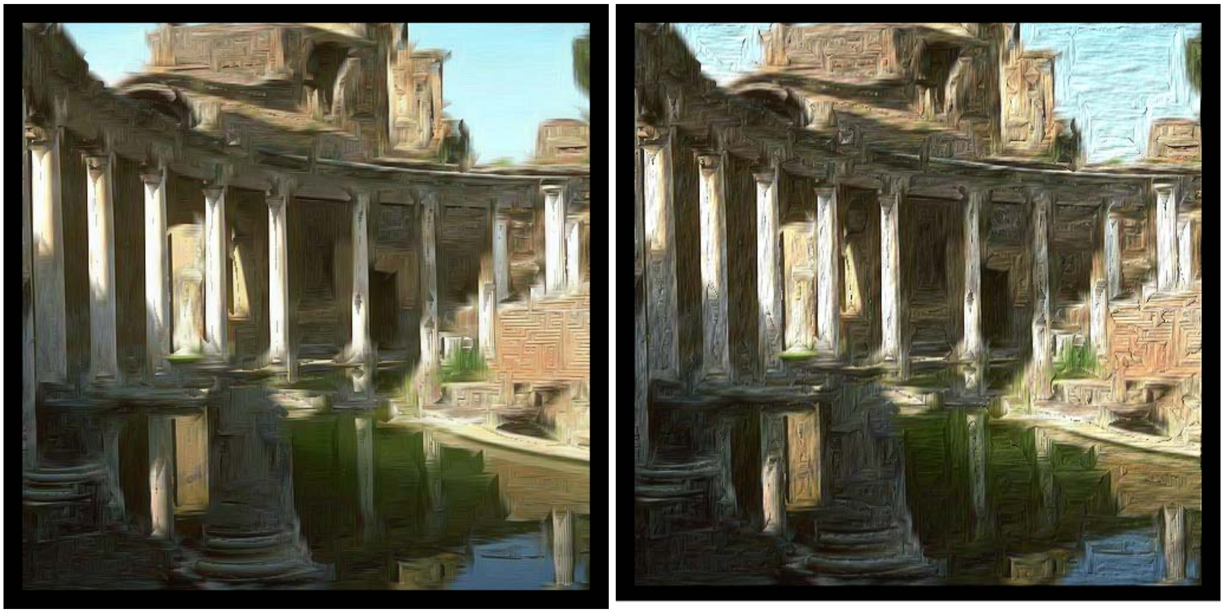

Fig. 6. Cabral's simplified rendering (left) and the effect of surface texture..

Fig. 7. Long Stroke (left) and Short stroke rendering (right).