Chapter 4

The Imaging Process

4.1 Introduction

In the previous chapters we have considered the modelling of rough surfaces and their interaction with light to form a textured image. However, strictly, machine vision algorithms do not operate on images. An algorithm is applied to a data set. The aim of this chapter is to describe the link from the image, whose formation was described in the previous chapters to the data set, on which the classifier discussed in the next chapter will operate.

The motivation for the development of this model is two-fold. First, a recurring theme in this thesis is that texture analysis cannot be properly investigated in isolation, but must be observed in the context of the physical process of which it is part. The need for an analytical description of the imaging process follows from this premise. Since CCD cameras are the most common imaging devices in machine vision, it is logical that we concentrate on these devices. There is however a second, more practical, motivation: in the latter part of this thesis several algorithms will be proposed. In order that their performance may be assessed under realistic conditions, *yet in a controlled manner*, we must develop a 'simulation engine' which characterises noise and distortion in a realistic way. This engine will be based on the work reported in this chapter.

4.1.1 Noise and Texture

We note that noise has largely been ignored in the texture analysis field. Those papers which do consider the effect of noise, such as [Liu90], [Szirani96] and [Porter97], do not cite experimental evidence to support their models. Liu and Jernigen state that their noise is random, additive and, the author presumes, white. They use the model to select a subset of measures from a feature set defined in terms of the power spectrum, and then

evaluate their performance for various S/N ratios. Szirani investigates the use of neural networks for the tasks of deblurring and texture segmentation. In the case of texture segmentation he investigates the effect of additive noise and the number of quantisation bits upon the classification accuracy. Porter and Canagarajah use additive white noise, with a signal to noise ratio in the range 0-25dB, to compare the robustness of Markov models, Gabor filters and wavelet approaches, though again with no physical justification of either the form or the magnitude of the noise model. Outwith texture analysis, Terrillon, [Terrillon96] adopts a more analytical approach, relating specific noise mechanisms, i.e. speckle and film grain noise to rotation-invariant classification. We note that unlike many of the effects observed in this chapter, these effects are either isotropic, or rotate with the image.

4.1.2 Organisation of this Chapter

The aim of this chapter is the development of a spectrally defined model of the imaging process. We noted in the previous section that where texture analysis does account for noise, it is usually considered to be white and additive, without any theoretical or experimental justification. In this thesis, we have taken a physically-based approach to the effects considered in the belief that this approach is more rigorous and leads to a more accurate and more flexible understanding of how surfaces can be classified. In this context, we do not believe that it is appropriate to treat noise as being additive and white without obtaining supporting evidence.

This support may come from two direct sources, the literature, or experiment. Unfortunately the literature is fragmented and deals mostly with mechanisms, rather than the spectral effects which are of relevance to this thesis. Similarly, the experimental approach is also limited, since many of the effects are difficult to resolve from the correct image. Where the effect is not easily separable from the correct image, we may use a third approach—simulation. This allows us to take a mechanism from the literature and observe the nature and magnitude of its effect for a textured image of the type considered in this thesis.

The process of developing an imaging model consists of four stages (*Figure 4.1.1)*. In the first, the literature is surveyed in order to identify noise mechanisms and, where possible, their spectral effects. At the end of this section, the noise mechanisms will be categorised as belonging to four functional classes. These classes are organised on the

basis of the experimental approach which will be required to characterise and measure them in sections 4.4 and 4.5. In the experimental section, we set about measuring those categories of imaging artefact which can be directly observed. Those which cannot be directly measured are further divided into categories we are unable to analyse, and those which can be investigated using simulation. In the former case we are limited to a qualitative description of the effect, while in the latter, a model of the effect *specific to an exemplar texture* will be developed. The descriptions obtained will be integrated with those from the simulation section to form a complete noise model.

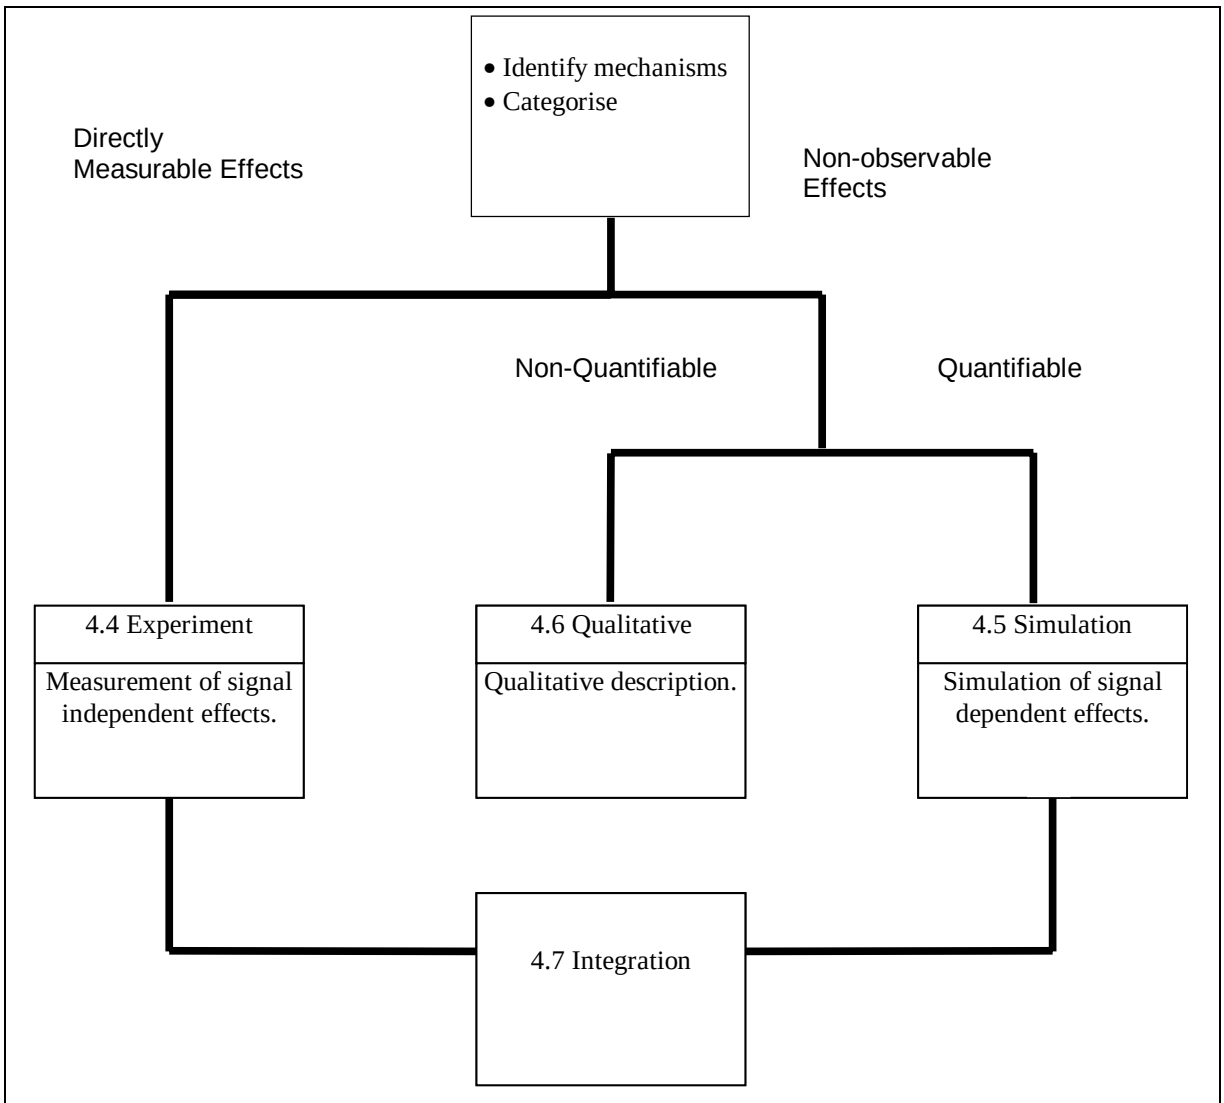

Figure 4.1.1 The organisation of the noise investigation.

Since this chapter is the first point at which we can directly observe the data set, in the final section we compare both the first order statistics and the measured spectra with those predicted by our models. This represents our first opportunity to use our surface, rendering and imaging simulations in order to verify our models.

4.2 Literature

The aim of this section is firstly to identify noise mechanisms from the literature and then to categorise them into classes according to the experimental approach required to quantify them.

4.2.1 Overview

Though several authors do give descriptions of the entire process, e.g. [Healey94] and [Lenz90], most of the literature relevant to the imaging process is usually limited to a single effect. This specialisation is problematic in our application; we are interested in the overall effect, rather than the individual components. This section surveys the literature to obtain a description of the mechanisms which form components in the transition from image to data set. The organisation of this section is therefore arranged by mechanism, rather than perceived effect, and follows the order in which the signal encounters these mechanisms. We use three categories; system optics, the CCD array, and the framestore (*Figure 4.2.1)*

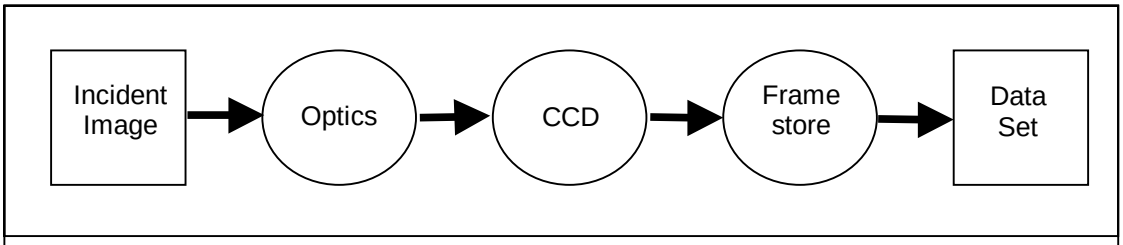

Figure 4.2.1 Section approach.

While this organisation is imposed by the nature of the literature, it is of limited use in our application. This section therefore concludes by categorising noise mechanisms by the experimental approach which will be used to quantify them.

4.2.2 System Optics

We retain the assumption of orthographic projection, but now consider sub-optimal focusing. Drawing on the work carried out in the field of shape from focus, we adopt the term Point Spread Function (PSF) which represents the image *i(x,y)* resulting

from a single point source imaged by the camera. This is equivalent to the two dimensional impulse function of systems theory. Although usually associated with optics, the PSF may be used to characterise the transfer function of the entire imaging system.

Subbarao and Lu use paraxial geometric optics and the assumption of spatial invariancy to develop an cylindrical PSF [Subbarao94]:

$$b(x,y) = \frac{4}{\pi d_m^2} \qquad \text{where } x^2 + y^2 \le \frac{d_m}{4}$$

and (4.2.2a)

$$b(x,y) = 0 \qquad \text{where } x^2 + y^2 > \frac{d_m}{4}$$

where d_m is the diameter of the blur circle in millimetres. As the diameter of the blur circle decreases, the image becomes sharper; as the diameter approaches zero, the impulse response of the optical system approximates the ideal impulse function. Pentland argues that the PSF is best modelled as a two dimensional Gaussian function $b(r,\sigma_b)$ where σ_b is a spatial constant proportional to the diameter of the blur circle [Pentland87], (in fact Subbarao also provides the option for a Gaussian PSF in his simulation engine). This model is also used by Nayar and Nakagawa [Nayar94] who state it both in the spatial and frequency domains:

$$b(x,y) = \frac{1}{2\pi\sigma_b^2}\exp\left[-\left(\frac{x^2+y^2}{2\sigma_b^2}\right)\right] \tag{4.2.2b}$$

$$B(u,v) = \exp\left[-\left(\frac{u^2+v^2}{2}\right)\sigma_b^2\right] \tag{4.2.2c}$$

where $B(u,v)$ is the Fourier transform of $b(x,y)$.

An example of the effect of the blurring function is shown in *Figure 4.2.2*.

Figure 4.2.2 Focused and defocused (σ_b=0.04) images.

We briefly note one other effect of the optical system: vignetting is an artefact of the compound lens system which causes a roll-off in sensitivity with distance from the centre of the image. The effect is most significant when using wide angle lenses or when using a lens at full aperture. Kolb et al. [Kolb95] give an example where intensity at the edges of the sensor plane is one third that of the centre. However, like Subbarao we explicitly ignore this effect, and assume it to be corrected for in the set up of the system.

4.2.3 CCD Array

We now consider the CCD array itself, a clear description of the noise mechanisms and observed effects is difficult since each mechanism may contribute to one or more observed effect. Each effect in turn may be due to a combination of noise sources which are not necessarily confined to the array. Since researchers generally investigate only part of the process, or to varying degrees of depth, taken as a whole, the literature often appears to be inconsistent. The aim of this section is to present a coherent model of effects occurring on the array in a form that is relevant to texture analysis.

We begin with dark current noise: as with any silicon device at ambient temperature, thermally generated electrons are produced within the body of the sensor. If these electrons wander to the sensing surface they may be captured by a photoelectron site, giving rise to noise. This effect is stochastic: the generation of an electron is a random event. However, due to the inhomogeneous nature of the lattice, electrons are more likely to be generated at some locations than others. While the subsequent movements of the electrons will also be random, some sites will be more prone to electron

capture than others [Talmi80]. In this sense dark current noise may be argued to contribute to both fixed pattern and spatially independent noise (as shot noise). It should also be noted that this effect is highly dependent on temperature; each 6.7° rise in temperature will cause a doubling in the dark current, so changes in temperature that imperceptible to the human observer may cause significant changes in both the fixed pattern and spatially independent noise.

A fundamental limitation of all imaging devices is the quantum nature of light: the number of photons which arrive at a site in a given period is a Poisson process and may be described as shot noise [Healey94]. In most practical applications the large number of incident photons will cause the distribution to tend towards a Gaussian distribution with variance equal to its mean. The larger the number of incident electrons, the less significant the standard deviation will be relative to the mean.

Two other effects occurring at this stage are blooming and non-uniform sensitivity. Blooming occurs when some sites are intensely illuminated causing a large number of electrons to be generated at a site, some of which then overflow to adjacent sites; we will neglect this effect. Non-uniform sensitivity is due to variations in the quantum efficiency of sites caused by fabrication effects [Healey94], we shall consider this to be fixed pattern noise.

4.2.4 Framestore

The previous sections described the process of measuring image intensity over a discrete grid, this section is concerned with the transfer to numerical representation. The process may be described in two steps: first a serial readout of the rows, with low pass filtering of the pulse train and the addition of the synchronisation signal. In the second stage, the signal is stripped of synchronisation information, resampled and quantised before being stored in memory. We describe three problems associated with this process identified in the literature:

1. serial filtering effects,
2. errors in synchronisation, and
3. quantisation errors.

Serial Filtering

Having collected the image, it is necessary to convert it to a continuous, interlaced, serial signal. In their paper Boie and Cox develop a camera model in electronic terms [Boie92].

The circuit model consists of both serial and parallel components, the parallel noise mechanisms i.e. leakage currents and Johnson noise are modelled by the i_n term, the serial component models noise associated with amplification and is also modelled as a current source, $e_n 2\pi\omega C_t$. These components are combined in (4.2.4a) to model the noise PSD due to the camera:

$$2G_n(\omega) = \overline{i_n^2} + \overline{e_n^2}(2\pi\omega)^2 C_t^2 \qquad (4.2.4a)$$

where

$G_n(\omega)$ is the total noise current power density,

i_n is the total noise current,

e_n is the equivalent input noise voltage,

C_t is the total camera capacitance,

The model is high pass in character, and predicts that noise power will increase with the square of frequency. However, since the camera is being readout in a serial manner this characteristic only applies to the rows, the noise at different points on the same column will be uncorrelated and the noise signal along columns is therefore predicted to be white. In fact Boie found both spectra to be essentially white and goes onto argue that the bandlimiting in the preamplifier acts as a prewhitening filter.

Synchronisation

The next stage we consider covers the removal of synchronisation information. The CCD array is read out in rows, each of which corresponds to a low pass filtered signal bounded by a blanking interval. If the rows are not perfectly synchronised, sampling points will be out of alignment with those of adjacent rows. Wang and Chen report this error is within the bounds ±0.4 pixels [Wang96].

Quantisation

The final stage considered is the quantisation of the signal. Most video analogue to digital converters (ADC) have a resolution of 8 bits, giving a random error uniformly distributed between $\pm\frac{1}{512}$ of the full scale deflection. The quantisation noise may be treated as a white noise process with variance $\frac{0.5^2}{12}$ [Healey94] . In practice quantisation noise is more significant than this figure would suggest since the full dynamic range of the ADC is rarely utilised and the presence of both random and fixed pattern noise reduces

the proportion of the ADC range which is available to the signal.

4.2.5 Conclusions

In this section, we have discussed the noise mechanisms described in the literature. In order to incorporate these mechanisms into a spectral model of the observed effects, we must:

- characterise the spectral shape, and
- measure its magnitude.

The approach we take will depend on the nature of the mechanism. We define four functional categories of artefact; transfer function, fixed pattern noise, temporal noise—which we define to be variations in the image of the same scene over consecutive frames, and signal dependent noise. A given mechanism may contribute to one or more of these categories.

The predominant effect of the optical system considered in the literature is the blur function, we adopt Pentland's model and treat it as a Gaussian transfer function.

Dark current noise was considered to be stochastic, contributing to both temporal and fixed pattern noise. Shot noise was described as being temporal, having a white spectrum and conforming to a Poission distribution, which tends to the Gaussian case as irradiance increases.

We have been unable to find any work on the spectral effects of line jitter, and hence we are forced to resort to simulation in order to estimate the effects line jitter will introduce into our model.

Serial filtering will affect the transfer function, however, the effect may not be noticeable due to the dominance of optical blur. The filtering effect may manifest itself in the filtering of noise generated subsequent to the optical system, most obviously in temporal noise.

Quantisation noise has on the one hand been described as being signal dependent [Lim, p.592] and on the other, as being white [Healey94]. While these descriptions are not necessarily incompatible, we believe it is necessary to verify that the white noise assumption is safe for the type of images used in this thesis. We again use simulation to assess the character and degree of the effect.

	Transfer function	Fixed Pattern Noise	Temporal Noise	Signal Dependent Noise
Blur	√			
Dark noise		√	√	
Shot noise			√	
Filtering	√		√	
Jitter			√	√
Quantisation				√

Table 4.2.1 Noise categories.

If we assume a linear transfer function, the system transfer function is easily measured using calibration images. Temporal noise is also relatively easy to measure, using disparities in multiple images of the same scene. Fixed pattern noise can be observed by preventing light from entering the lens and observing the resulting image. However, if the camera does not have its autogain disabled, it is not possible to measure the effect relative to a textured image. Unfortunately, we found that the autogain was required for stable operation of the camera/framestore, and we are unable to measure the effect, though we do characterise the effect in section 4.4. Signal dependent noise is also difficult to measure and model. We adopt the following approach. The effect is simulated with a synthetic exemplar texture, the difference between the correct image and the noisy image is treated as being additive noise, (we shall describe it as pseudo-additive). In this way we treat imaging artefacts as being either a transfer function or as (pseudo) additive.

4.3 The Case for Sub-sampling

4.3.1 Introduction

The easiest form of noise to identify is that disclosed by temporal variation of the image. Using multiple images of the same scene we observe the spectral characteristics of temporal noise relative to that of the static part of the image. The spectral characteristics of the S/N ratio lead us to implement sub-sampling in the test images. In the next section we estimate the overall transfer function of the system and the remaining temporal noise, subsequent to pre-processing.

4.3.2 Justification of Sub-sampling

We average 10 frames to produce a 'clean' image and produce residual error images from the original images, we then average the PSD of these images, *Figure 4.3.1*. We note that the temporal noise is clearly coloured. This immediately calls into question the white noise assumption made in most papers which deal with the effects of noise on texture.

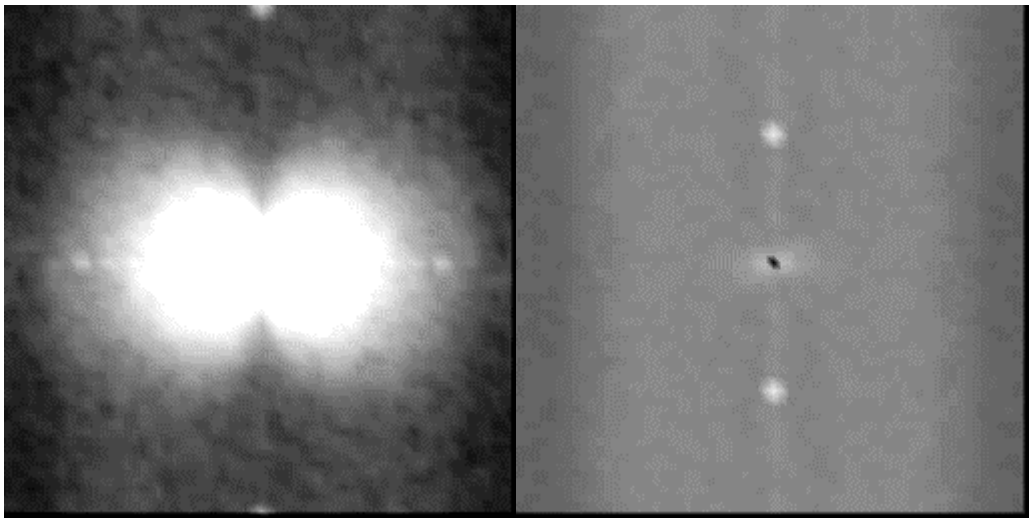

(a.) (b.)

Figure 4.3.1 Spectra of Time averaged image (a) and averaged residual spectra (b)

We define two nominal, possibly overlapping, classes of noise: spectral artefacts and temporal variation. Taking the PSD of a captured rock image, *Figure **4.3.4***, we identify three suspected artefacts of the imaging process in *Figure **4.3.3***: a wide vertical-running noise plateau (1), and 2 groups of spikes (2) and (3). We confirm that these are artefacts of the imaging process by rotating both the light source and the surface by 90°, *Figure **4.3.2***. While the characteristic illumination effects rotate, the artefacts (1)

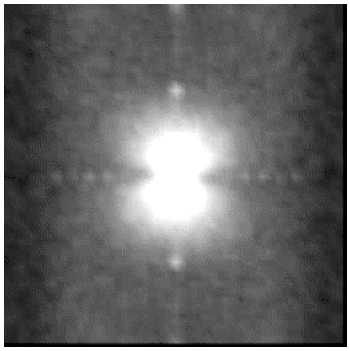

Figure 4.3.2 Rotated Rock PSD

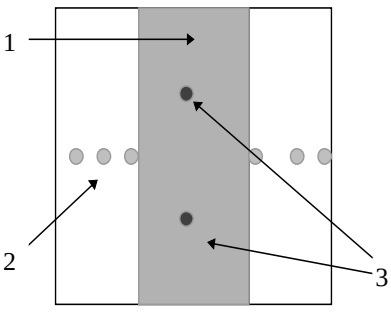

Figure 4.3.3 Key Figure

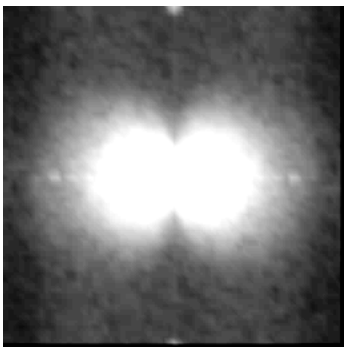

4.3.4 Rock Image Spectrum

and (2) remain static, showing that these effects are associated with the imaging process.

We observed from *Figure 4.3.1* that the temporal noise spectra is not white. While the spectra is uniform in the vertical direction, there is a pronounced "plateau" running in the horizontal direction. In *Figure 4.3.5* we plot the profile of the 1D spectrum in the horizontal averaged over the spectrum columns. There is a clear trend which is remarkably uniform for all the textures.

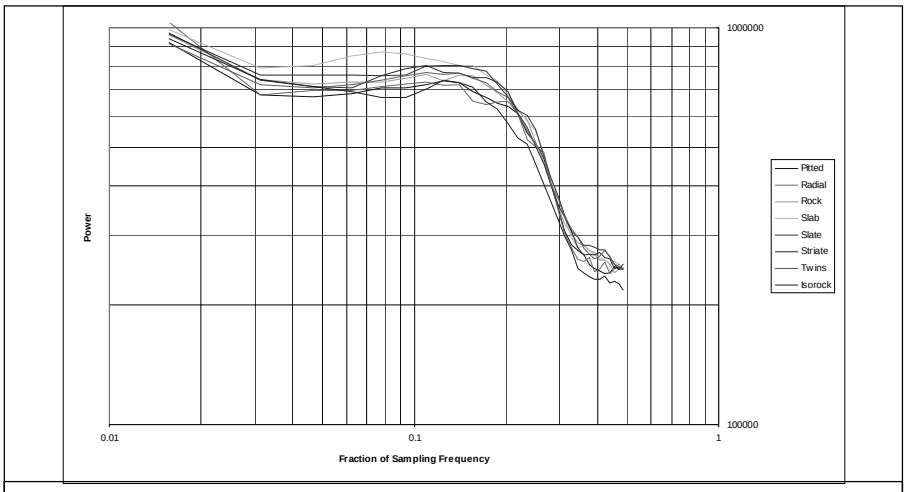

Figure 4.3.5 Horizontal Profiles of Residual Power Spectra

We now consider noise in the context of imaged test surfaces. We measure the signal to temporal noise ratios of several textures along both the vertical and horizontal axes of the frequency plane. The textures are illuminated from the vertical direction, and the signal is at its maximum in this direction. According to our models, the signal along the horizontal axis is composed of the non-linear components and noise. The frequency dependent signal to temporal noise ratio is defined below:

$$E_i(\omega) = 10\log_{10}\left(\frac{\sqrt{I_i(\omega)^2 - E_r(\omega)^2}}{E_r(\omega)}\right)$$

where

$I_i(\omega)$ = the power of the average measured image at frequency ω.

$E_r(\omega)$ = the power of the residue image.

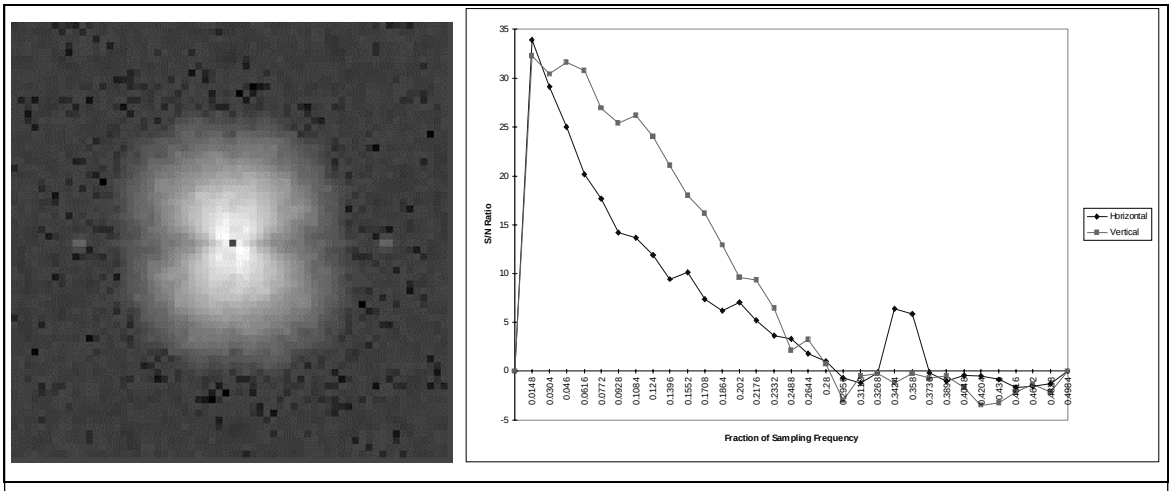

Figure 4.3.6 Signal to Temporal Noise Spectrum for Rock Texture

The signal to temporal noise ratio, *Figure 4.3.6*, shows a clear frequency dependency as we would expect from the fractal nature of the image. High frequencies are characterised by low signal to noise ratios of 0dB or less. This leads us to believe that little information, useful for classification, will be found at these frequencies. This immediately commends low pass filtering and decimation for two reasons.

1. Low pass filtering will remove the majority of noise spikes and the associated attenuation of high frequency, vertical running components will mean that the temporal noise process can be effectively described as being white.
2. The suppression of high frequencies will allow decimation of the images with the associated reduction in processing time for the classification algorithm.

This procedure was therefore adopted, and all the images were low pass filtered with a 7th order Butterworth filter implemented in a 25X25 mask prior to decimation. A linear image rescaling was also carried out at this stage.

4.4 Experimental Investigation

4.4.1 Introduction

This section is concerned with the experimental measurement of image artefacts. We estimate the overall transfer function of the system, including pre-processing, using a calibration image, and fitting an experimental model to the observed data. As with the previous section we use multiple images of the same scene to identify temporal variation before quantifying the effect in terms of its first and second order statistics.

4.4.2 Transfer Function

Having adopted the down-sampling procedure, we now investigate the overall effect of the blur function. The blur function is assumed to be isotropic and Gaussian in accordance with Pentland [Pentland87]. We compare a calibration image with a synthetic image and a blurred synthetic image *Figure 4.4.1*, the spectral results are shown in *Figure 4.4.2*.

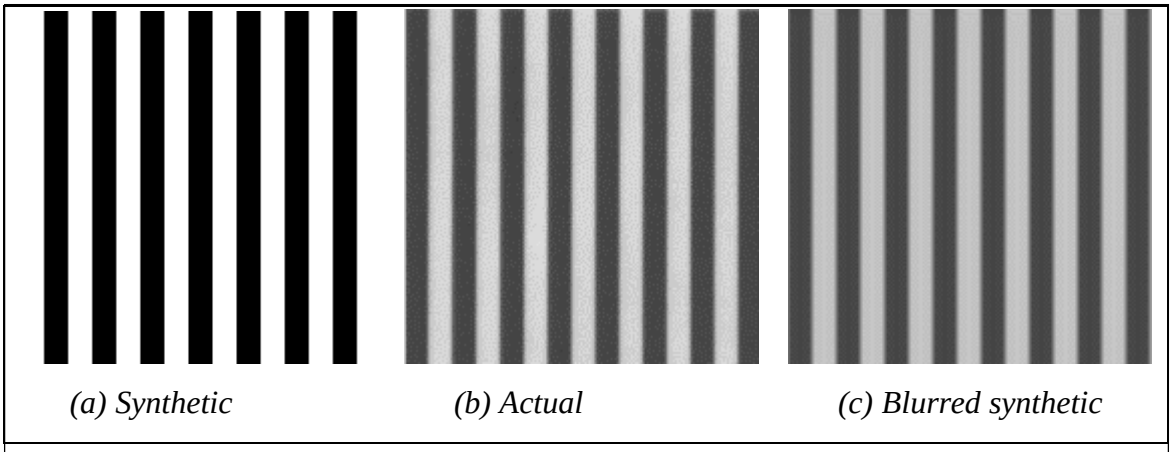

(a) Synthetic *(b) Actual* *(c) Blurred synthetic*

Figure 4.4.1 Synthetic, actual and blurred calibration images.

The blur function with σ_b=0.02 accurately predicts the magnitude of the peaks in the calibration spectra, with the exception of the very low frequencies, i.e. below 0.025 f_s. We therefore adopt this as our optics model.

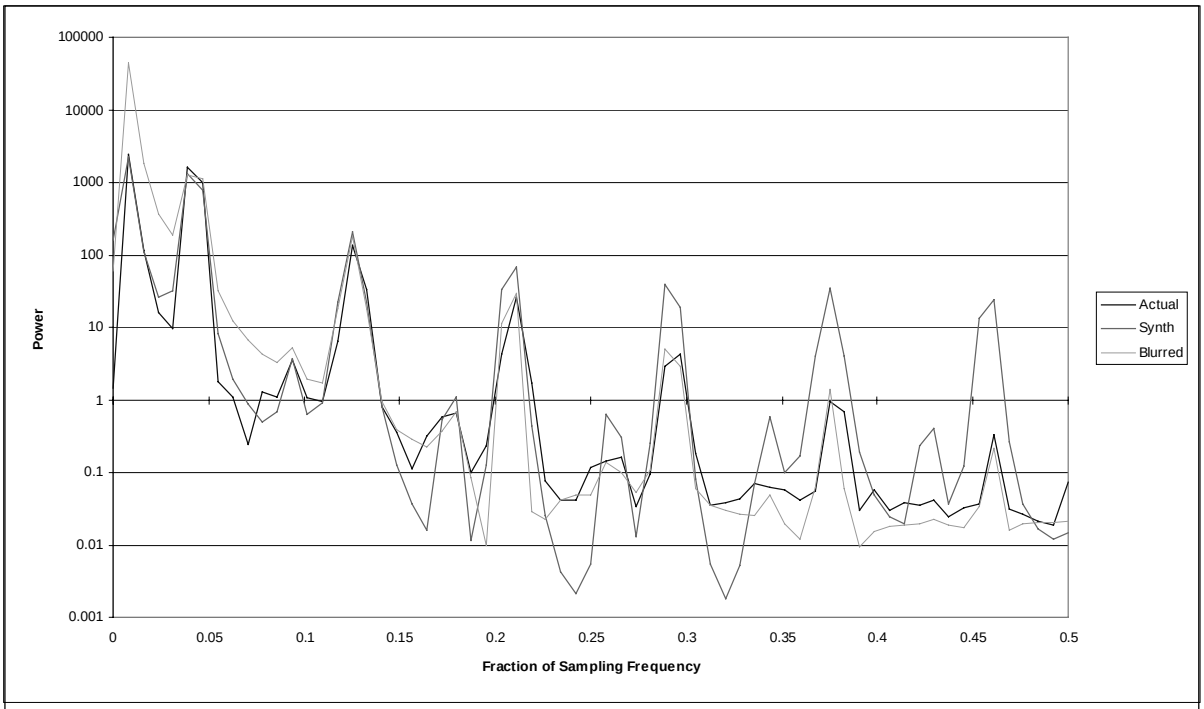

Figure 4.4.2 Comparison of Actual, Synthetic and Blurred Synthetic Calibration Spectra

4.4.3 First Order Statistics of Temporal Noise

Now, let us consider temporal noise. We treat any variation over consecutive frames of the image of a given scene as being noise. We call this temporal noise and model it as being additive. We begin by estimating the first order statistical properties of this noise. Averaging a series of temporal images will reduce temporal noise. The additive noise will be of the same order of magnitude as the quantisation noise and a disparity histogram calculated from a mean image will have a low standard deviation, making it difficult to form a good estimate of its properties. Instead we mitigate this problem by subtracting two images of the same scene; this is equivalent to the subtraction of two random noise processes.

Texture	Mean Disparity	Standard Deviation of Disparity	Texture	Mean Disparity	Standard Deviation of Disparity
Isorock	0.45	1.78	Slab	0.03	1.69
Pitted	0.08	1.58	Slate	0.40	1.68
Radial	-0.02	1.59	Striate	-0.02	1.67
Rock	-0.35	1.64	Twins	-0.07	1.62

Table 4.4.1 Table of image disparities.

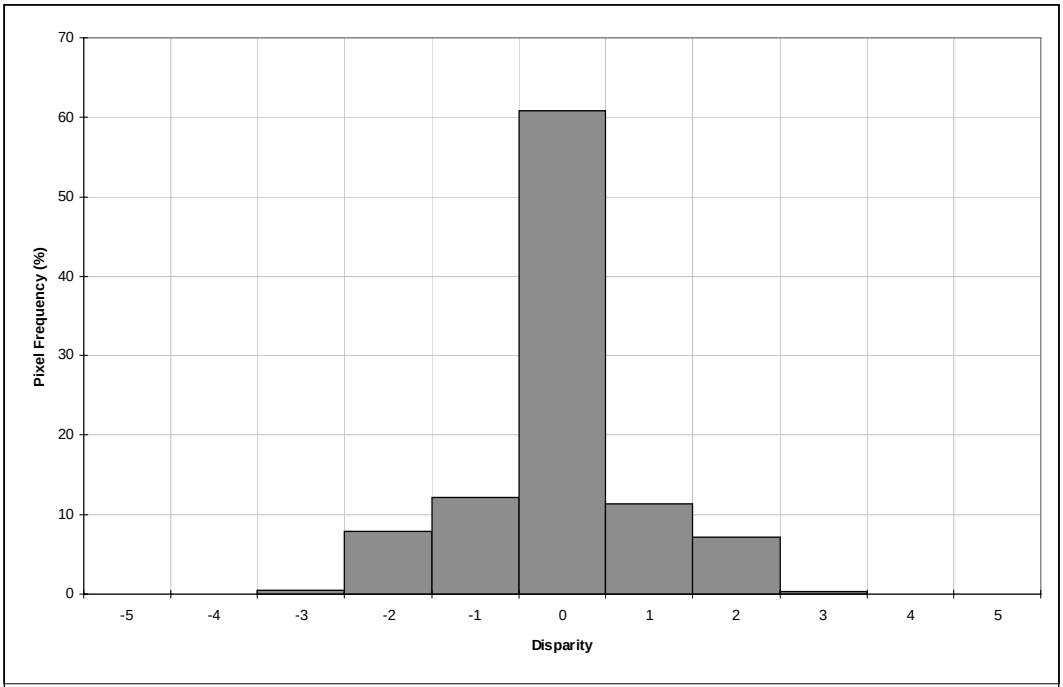

Figure 4.4.3 Disparity histogram for rock texture.

If we assume that the temporal noise distributions, which underlie the disparity distribution, are identical and Gaussian, we may infer their standard deviations from that of the disparity histogram.

$$\sigma_1 = \sigma_2 = \sqrt{\frac{\sigma^2_{disparity}}{2}}$$

This results in an estimated standard deviation (σ_t) of 1.17 averaged over all the textures. We may also usefully state the temporal noise level in terms of the signal to noise residue.

Texture	**S/R (dB)**	**Texture**	**S/R (dB)**
Rock	25.90	Slate	15.90
Striate	24.08	Pitted	19.63
Isorock	22.67	Twins	20.76
Slab	26.23	Radial	17.20

Table 4.4.2 Signal to temporal noise levels for the test textures.

4.4.4 Spectral Characteristics of Temporal Noise

Having observed the statistics of the disparity images, and inferred those of the temporal noise process, we now consider the spectral characteristics of the temporal noise. We do so by observing the spectra of the same disparity images which will be the scaled image of the noise spectra.

In section 4.3 we advocated the use of sub-sampling for the images used in this thesis, we now make a polar plot of the temporal noise spectra of the subsampled *Rock* image (*Figure 4.4.4*), for comparison we also show the polar plot for the original Rock image prior to sub-sampling. We can see that the noise power in the subsampled image is far more evenly distributed throughout the angular range than that of the original image. With the exception of the peak at θ=90°, the plot shows the noise power distribution to be essentially isotropic in the case of the sub-sampled image.

If we now plot power against radial frequency for the *Rock* image (*Figure 4.4.5*), we note a gentle fall in power with increasing frequency. A least squares fit over the entire frequency range gives an estimated β=0.70, if the beta parameter is measured over the midband range f_s=0.14 to 0.38, a much lower estimate of β=0.39 is made.

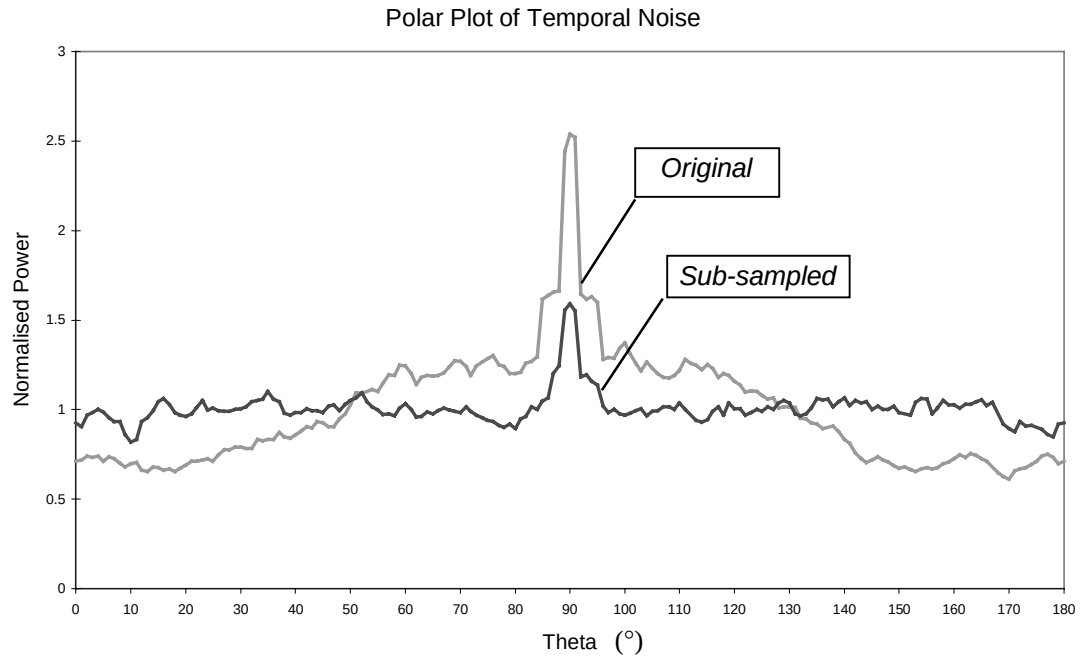

Figure 4.4.4 Polar plot of temporal noise spectra of original and sub-sampled images.

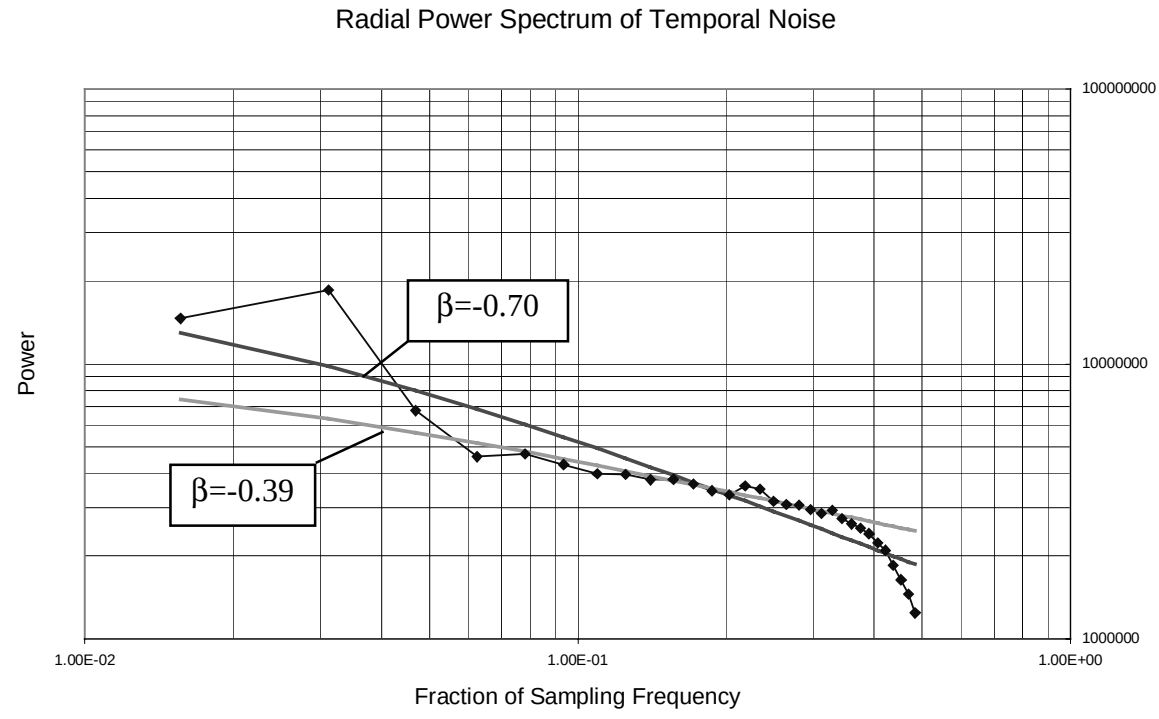

Figure 4.4.5 Radial power spectra of temporal noise associated with (sub-sampled)Rock image.

The result of the polar and angular plots is the description of the noise process as being of isotropic and mildly low pass in character. Given the relatively low estimates of the beta parameter, and the convenience of the assumption, we will assume in this thesis that the temporal noise is white.

Comparison of the radial noise spectra for different textures (*Figure 4.4.6)* show them to be similar. Some differences are apparent at low frequencies, though we note that these correspond to only a few spectral components. Overall, we believe the similarity of the spectra vindicate our assumption that temporal noise can be treated as being signal independent, at least within the class of textures considered in this thesis.

We therefore conclude that the spectra of temporal noise associated with the textures of this thesis may be modelled as being realisations of the same noise process, independent of the texture characteristics, and, furthermore, that the process may be well approximated by a white noise process.

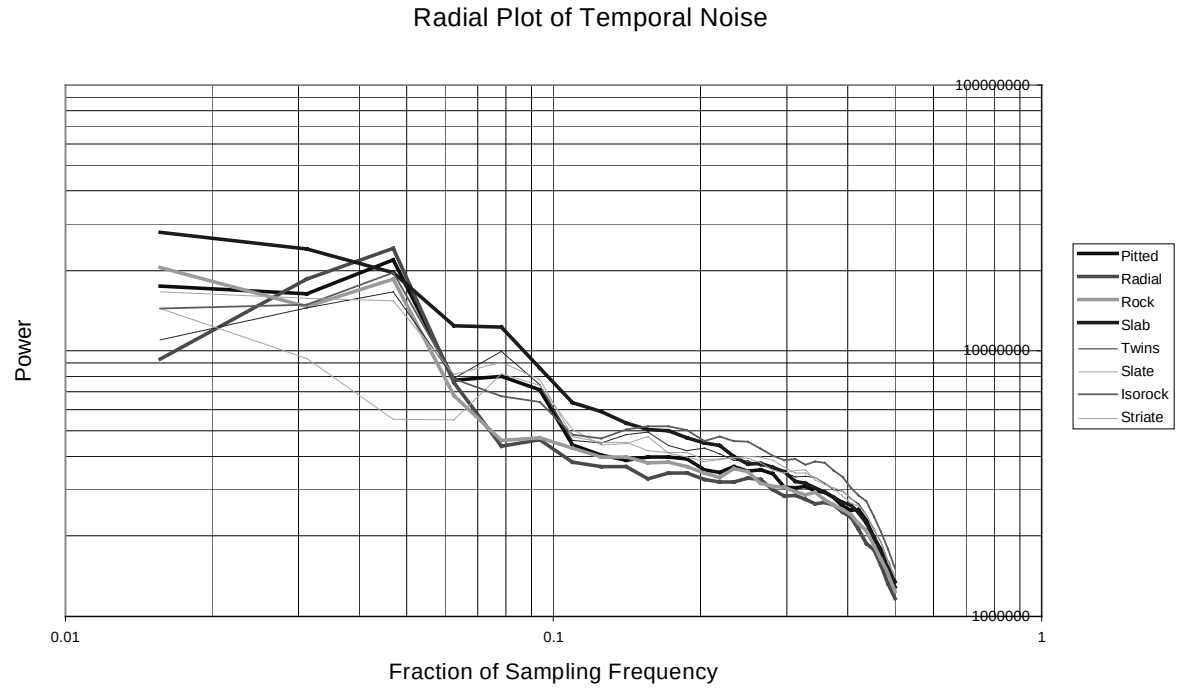

Figure 4.4.6 Radial temporal noise spectra.

4.4.5 Summary

In this section we have found the transfer function to be well approximated by a Gaussian with parameter σ_b = 0.02. It was also concluded that the temporal noise associated with the subsampled image is reasonably approximated as white noise of

variance 1.17.

4.5 Use of Simulation to Investigate imaging phenomena

In the literature section, several noise mechanisms were described which cannot be easily resolved from the uncorrupted image. In the cases considered here, it was not possible to either observe the mechanism or measure the effect directly. We therefore use simulation to estimate the magnitude and nature of the spectral effects caused by the noise sources. It is, however, important to note the limitations of simulation as an investigative tool. We cannot prove causal links between a mechanism and an observed effect. Instead, we are limited to stating that the results predicted by the model are compatible with the observed results.

In order to make the experiment as relevant to texture analysis as possible, we will use a synthetic test surface with realistic parameters. A fractal surface will be illuminated with the empirical reflectance map. Noise simulations will then be carried out on this image and the measured and simulated spectra compared.

4.5.1 Quantization Noise

While most cameras/framestores are 8-bit, in practice, due to lighting variations and the field of view not consisting entirely of texture, it is likely only a fraction of this range will be utilised to convey useful texture information. To illustrate this point we show the standard deviations of the pixel intensities for the test surfaces illuminated from $\tau = 0°$ (*Table 4.5.1)*. In the third column, the number of bits required to quantise the range from -2σ to $+2\sigma$ is shown.

Texture	σ	**Bits**	**Texture**	σ	**Bits**
Isorock	9.48	6	Slate	6.19	5
Radial	6.61	5	Striated	12.20	6
Rock	11.37	6	Twins	8.45	6
Slab	13.10	6	Pitted	6.68	5

Table 4.5.1 Effective quantisations of data sets.

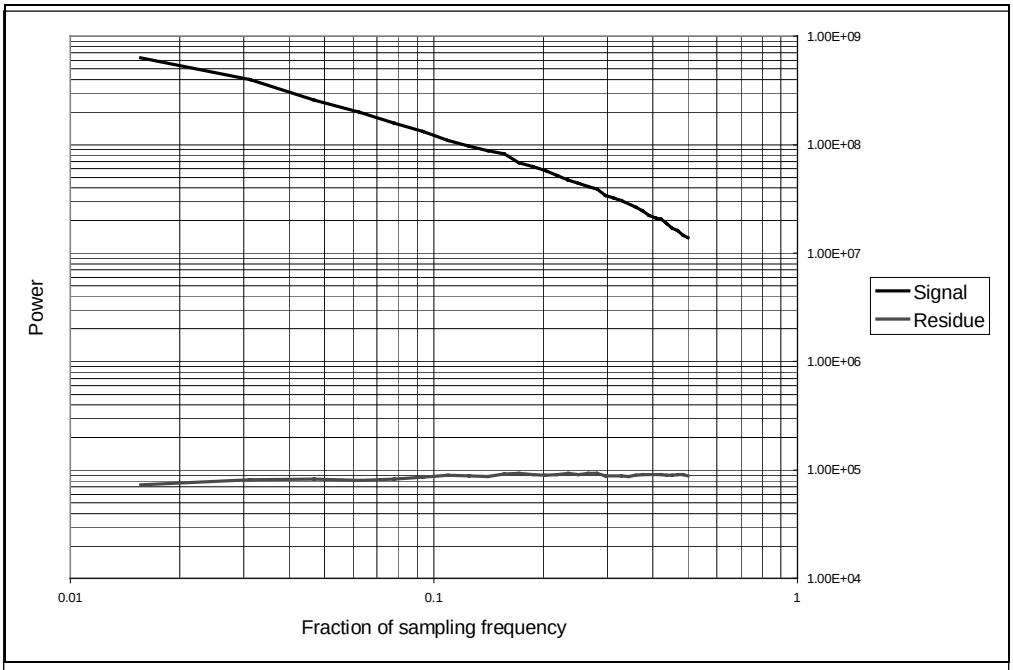

Figure 4.5.1 Power spectra of original signal and residue signal.

If we assume that the active region ranges from -2σ to $+2\sigma$, all of our images are effectively quantised in 5 or 6 bits. This was clearly an area of concern and the effect was simulated by comparing the original rendering with that of the image quantised to 6 bits. The spectrum of the residual image was found to be small and effectively white, *Figure 4.5.1*. This is consistent with Szirani who found that at least four bits were required for reliable classification.

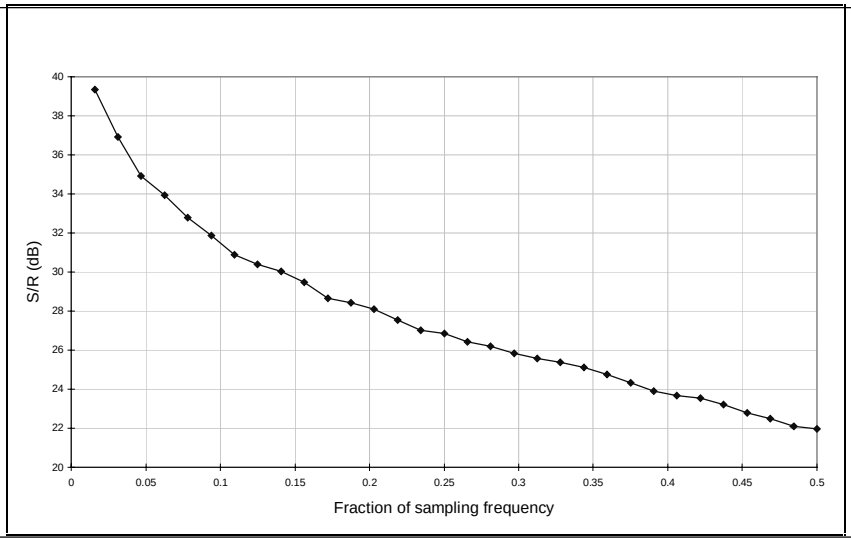

Figure 4.5.2 Signal to residue ratio vs. frequency

We therefore conclude that it is safe to assume that the residual spectrum due to quantisation is white. Furthermore in our simulations the S/R ratio varies from

approximately 40dB at low frequencies to 22dB at high frequencies for a 6 bit quantisation.

4.5.2 The Effect of Jitter

We again use our experimentally based synthetic surface to generate a blurred image which is then supersampled (X10) in a raster scan with a random jitter drawn from a uniform distribution of width -0.2 to 0.2 pixels being added at the beginning of each line. The serial signal is then subsampled to its original rate and the power spectrum examined for artefacts. The resulting residual spectrum is shown in *Figure 4.5.3.*

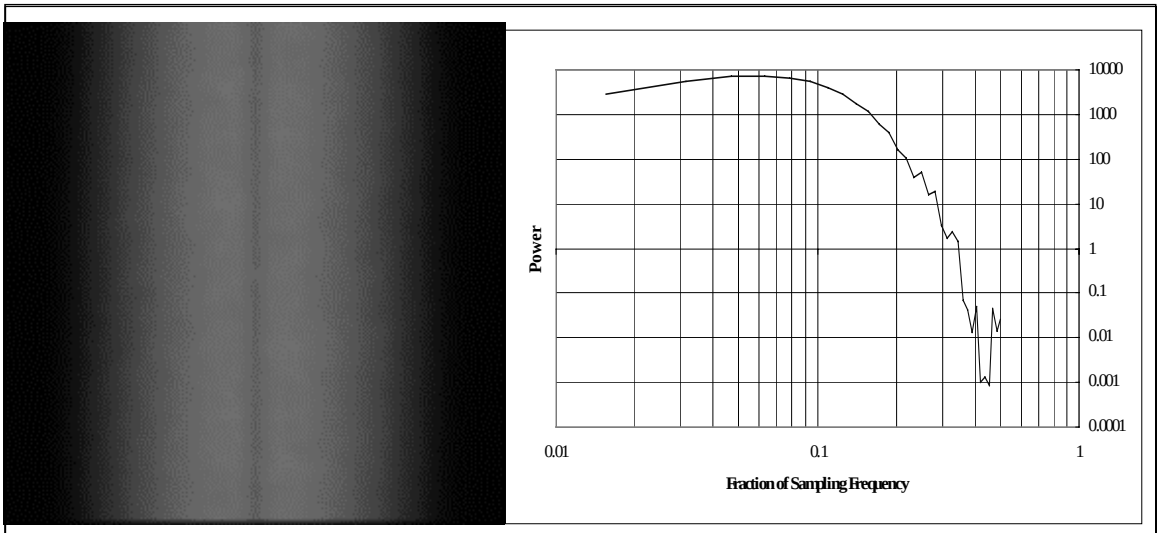

(a) Linear frequency scale, log power scale. (b) Log frequency;log power scale.

Figure 4.5.3 Jitter induced 2D residual spectrum (a) and mean row spectrum (b).

The simulated effect is similar to the plateau effect observed for temporal noise prior to subsampling.

4.5.3 Serial Filtering

The image is read from the CCD array in discrete, serial form. This pulse train is then low pass filtered to yield a continuous waveform. A one-dimensional filter applied to this signal will produce an attenuation with frequency of the signal in the rows, while leaving the columns unaffected. We hypothesise that a white noise process filtered in this way will produce a spectrum with a characteristic vertically running plateau.

We test this hypothesis by simulation, using a Butterworth filter (arbitrarily designed to be 10th order) with breakpoint at 0.25 f_s, we filter a two dimensional white

noise process in serial form. The resulting spectrum *(Figure 4.5.4)* is of similar form to that observed for the jitter simulation.

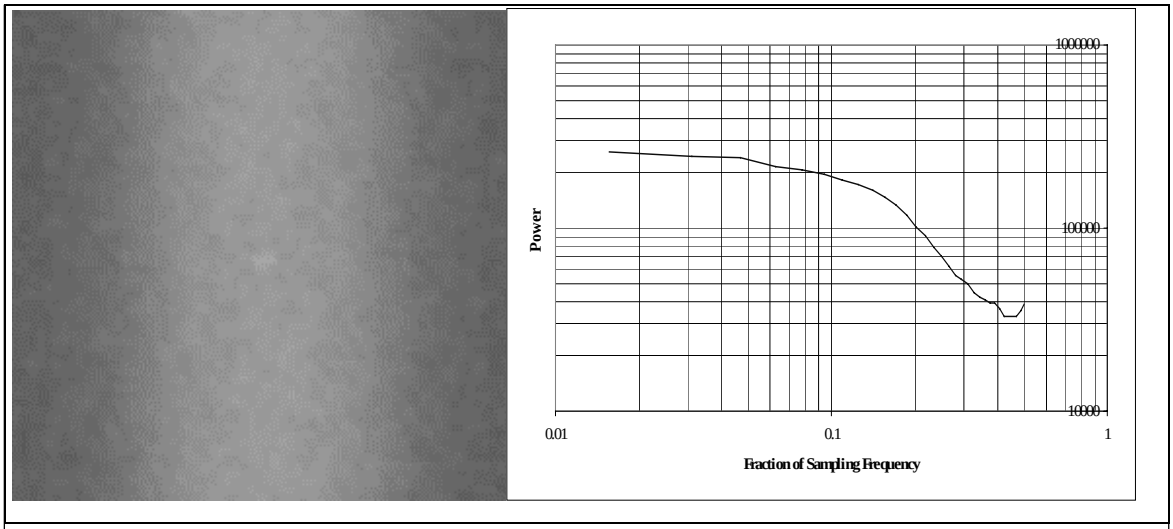

Figure 4.5.4 The spectral effects of an anti-aliasing filter.

Like the jitter simulation this produces a residue spectrum of the “plateau” form. The fact that two mechanisms may form rival hypotheses for a particular observed phenomenon serves to underline the inability of simulation to establish causal relationships. In any case, the results of the simulation indicate that the effect can be reduced to a white noise source if subsampling is employed.

4.5.4 Summary

In this section we found that the members of the data set were effectively quantised in 5 to 6 bits. The residual signal associated with quantisation was found to be white, with the result that quantisation noise is most significant at high frequencies. This section also disclosed that serial filtering of a white noise source and line jitter both give rise to the 'plateau' effect, and are possible sources of this plateau. However, as with the real data set after subsampling, these are effectively white noise sources.

4.6 Non-Quantifiable Noise Mechanisms

The aim of this chapter is to develop and parameterise a noise model. There are some noise mechanisms which are not readily measured either by direct experimentation, nor by simulation. Since we cannot measure these mechanisms in a useful way, we are unable to include them in a noise model. We do, however, briefly discuss two effects for completeness.

4.6.1 Camera Non-linearity

We conclude this section by briefly noting the effect of camera non-linearity. While many cameras may be set to give a linear response, the large range of intensity values, and the limited number of bits, mean that it is often more convenient to use a camera with γ=0.5[1]. *Figure 4.6.1* shows the spectra of a "Chessboard" calibration image. While the measured and synthetic spectra share the same spectral peaks, particularly at low frequencies, both the measured spectra have additional peaks due to non-linearity-induced harmonics.

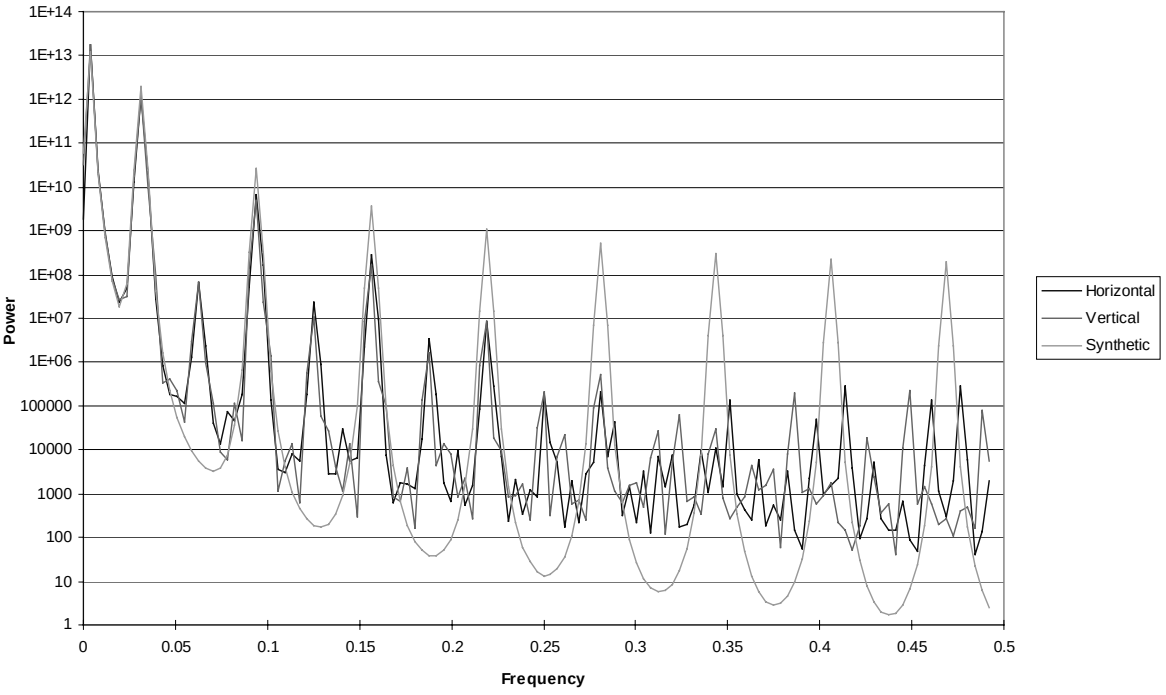

Figure 4.6.1 The generation of harmonics due to camera non-linearity.

In the previous chapter we measured the reflectance function in terms of camera response, effectively lumping together camera and reflectance non-linearity. We note that unlike reflectance, the effects of camera non-linearity are dependent on the incident intensity. As a consequence of this, the amplitude of the harmonics *relative to the fundamental* will change if the incident intensity is varied. A change in the intensity of illumination may therefore have serious effects on the classification of textures which cannot be corrected for using standard normalisation techniques.

[1] The camera amplifier can be modelled using the relationship $i_{out} = k.i_{in}^{\gamma}$,, where *k* is a constant. In linear operation γ=1.0, more commonly however γ=0.5 to allow quantisation of a large dynamic range in

We also note that a non-linear camera means that even the additive noise measured earlier will be signal dependent in effect. We justify our use of an additive model by measuring the effect of noise on the data set itself, we are in effect forming a linear approximation around the actual operating point. The uniformity of the temporal noise spectrum measured in *Figure 4.3.5* supports the validity of this assumption.

4.6.2 Fixed Pattern Noise

Associated with several of the noise sources considered in section 4.2 is the phenomenon of fixed pattern noise, that is an imaging effect which is present regardless of the image content. Given a camera without autogain and with a linear gain function, it is possible to measure fixed pattern noise by preventing light from entering the lens. However, we found both the autogain and the non-linear gamma function to be essential when imaging textures. Consequently, while it is possible to observe the effects of fixed pattern noise, it is not possible to predict its effect when imaging textures.

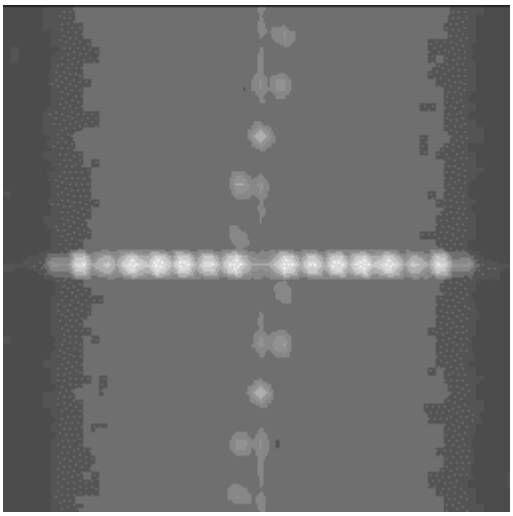

Figure 4.6.2 Fixed Pattern Noise Spectrum

We are therefore unable to quantify or model the effect of fixed pattern noise and while we do illustrate the spectral properties in *Figure 4.6.2* we do not include it in our noise model. Comparison with *Figure* ***4.3.4*** shows that several of the peaks observed in the FPN spectrum are also apparent in the imaged texture spectrum. We note however, that these are only obvious at high frequencies, and the effect will be reduced by subsampling.

---

a limited number of bits.

4.6.3 Summary

This section has briefly discussed those noise effects which we are unable to integrate into our model. Ideally it would be possible to use a linear camera without autogain. However, in practice the required dynamic range make this difficult to use with only eight quantisation bits. Due to non-linearity it is important that a constant level of incident intensity be maintained.

Fixed pattern noise is not in itself a problem to classification, though its non-linear interaction with the correct image may generate spurious spectral components.

4.7 An Imaging Model

Up until this point we have concentrated on describing and quantifying the noise and distortion associated with the imaging process. We now draw together the strands of this investigation to form a model of the imaging process. This model will serve as both an analytical tool and as the basis of simulations of the imaging process.

4.7.1 Integrating noise components

Two of the simulated effects, serial filtering and line jitter, are temporal in nature and their effect will be apparent in the temporal noise observed in section 4.5. The effect of quantisation, conversely, is static and is not accounted for in the temporal measurements. We now combine the temporal noise with quantisation noise to form an estimate of the error between the measured and actual values.

The pdf of a random process which is the sum of two random variables may be obtained by convolving the distributions of the two variables [Peebles, p.102]. We form expression (4.4.2) from the convolution of the uniform distribution with the postulated Gaussian distribution. Assuming a zero mean Gaussian process:

$$p_x(e) = \frac{1}{\sqrt{2\pi\sigma_t^{\,2}}} \int_{t=-0.5}^{0.5} e^{\left(\frac{(e-t)^2}{2\sigma_t^{\,2}}\right)} dt \qquad (4.4.2)$$

The parameter σ_t, controls the shape of the distribution, *Figure 4.7.1*. For temporal noise with a relatively large variance (σ_t =1), the resulting distribution resembles a Gaussian curve, smaller variances (σ_t =0.1) give rise to distributions closer to the uniform case.

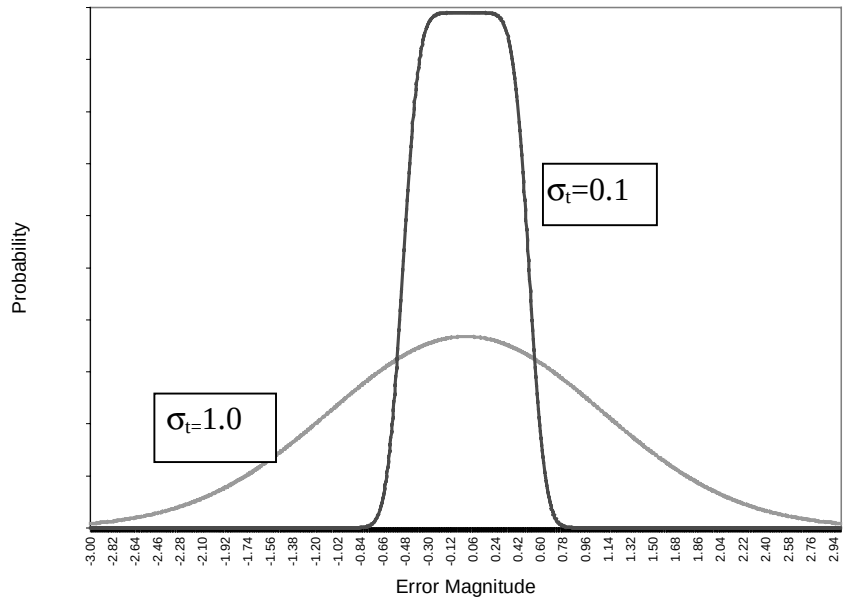

Figure 4.7.1 The effect of the sigma parameter on the PDF of a linear combination of Gaussian and uniform processes.

If we combine the estimated temporal noise distribution with the quantisation distribution (which exists in the interval ±1.2 due to rescaling due to pre-processing) we can form the distribution shown in *Figure 4.7.2.* We find that the resulting distribution can be modelled as being Gaussian of standard deviation (σ_n) 1.36.

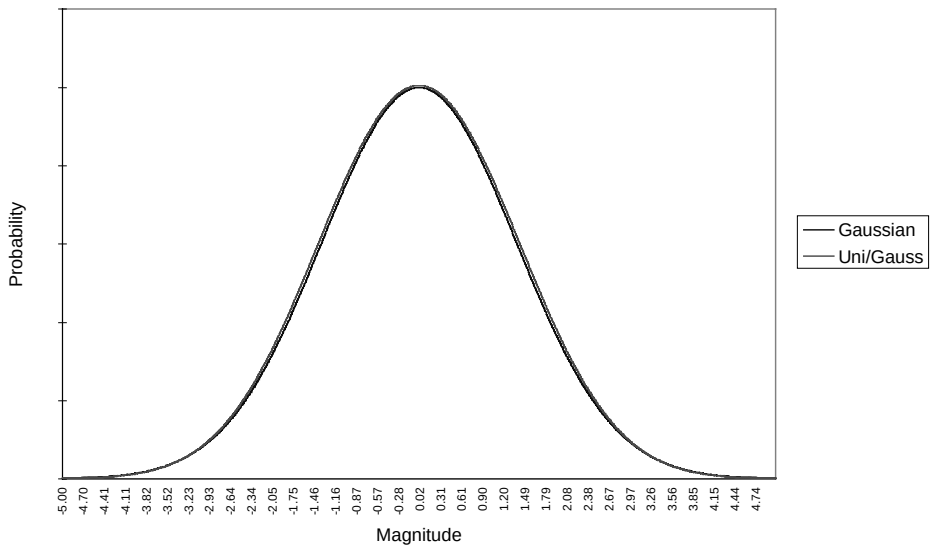

Figure 4.7.2 Estimated distribution of combined temporal and quantisation noise.

Since both the subsampled temporal and the quantisation residual spectra are considered to be white, the linear combination of the two will also be white.

4.7.2 Development of a Noise Engine

We now use the work carried out in this chapter to develop a noise engine which will be used later in this report. This engine forms a computational component in our system model.

The decision to use downsampled images allows several simplifications to be made to our model, such as the elimination of noise spikes at $0.25f_s$, the reduction of the relative effect of the blur function and the suppression of those directional noise mechanisms that only become apparent at high frequencies.

The model itself, *Figure 4.7.3*, consists of only two components, a blur function and an additive white noise process, both of which are considered to be global in character. The estimated blur function was found to fit well with the Gaussian model presented by Pentland. While localised distortions were experienced in the experimental work, these occurred in regions of the image which were not utilised. Their absence in our working data and the relative difficulty of their description lead us to omit them from our model.

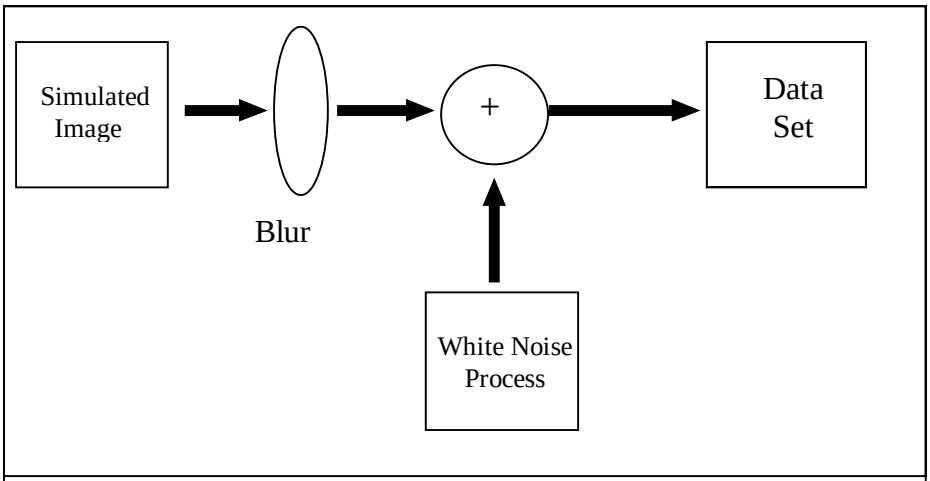

Figure 4.7.3 Noise model for sub-sampled images.

The second component of our noise model, additive white noise, lumps together several noise mechanisms. Dark noise, shot noise and quantisation noise are considered as being white in accordance with the literature, while jitter noise and filtered noise are also considered to be white as a consequence of subsampling.

Since we only consider texture which is caused by topography with no albedo variation, we will assume that camera non-linearity is lumped together with the reflectance function and modelled implicitly with the reflectance map.

Noise Mechanism	Description	Notes	Modelled
Dark noise	Temporal	Subsampled temporal noise found to be white.	Gaussian White
Shot Noise	Temporal	Subsampled temporal noise found to be white.	Gaussian White
Jitter	Temporal	Approximately white when subsampled	Gaussian White
1,2 Noise	Narrow band noise at 0.25 fs	Suppressed by lpf prior to sub-sampling.	Ignored
Quantisation		Simulation confirms accepted white noise assumption for broadband textures. While quantisation error is itself drawn from a uniform distribution, when lumped together with temporal noise of the magnitude observed here, the resulting distribution is Gaussian in character.	White noise.

Table 4.7.1 Summary of noise mechanisms and their effects.

The engine will have default parameters as follows: a Gaussian blur function with σ_b=0.02 and a noise function with standard deviation 1.36, quantisation will be carried out to 6 bits. The effect on a synthetic fractal image is shown in *Figure 4.7.4*.

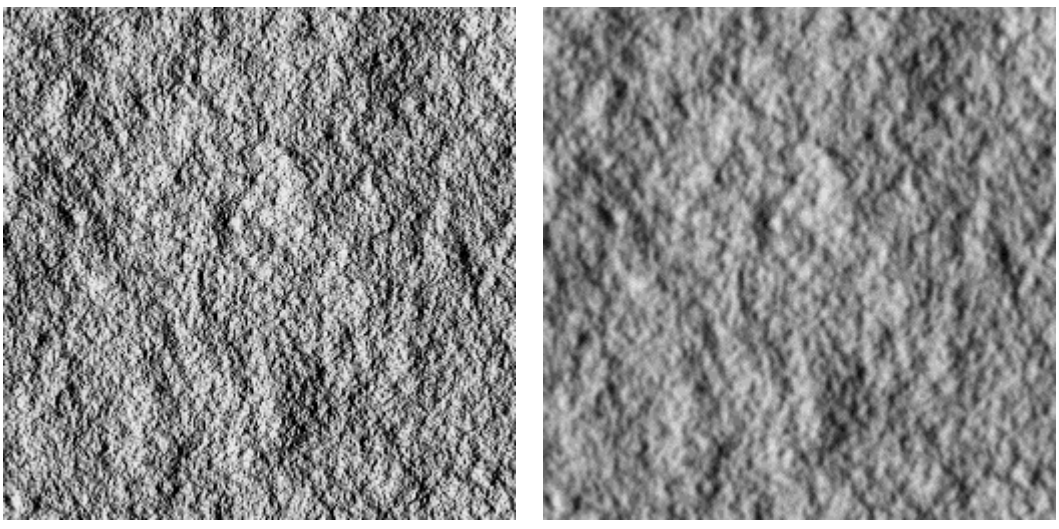

Figure 4.7.4 Original and processed images.

4.8 Assessment of the Physically Based Model of Texture Images

The work carried out in this thesis has so far consisted of modelling the formation of textured images from rough surfaces. We have now reached a stage in the imaging

process at which we are able to compare the predictions made by our models with actual measurements. Comparison of measurement and prediction will give an indication of the accuracy of our models.

4.8.1 First Order Statistics

The theoretical work carried out in the previous chapters has been performed with the (non-essential) assumption that the data had, at all stages, a Gaussian distribution. This was justified on the basis of the height distribution of our surface models, and the linearity of the subsequent stages of the process. We have now reached the point where we experimentally evaluate this assumption in a direct way.

The standard test for testing whether a distribution conforms to a model is the Chi square test (used in the phase section). This will give the probability that the observed measurements are drawn from a particular distribution. If the sampled distribution really is drawn from the model distribution we would expect the measured pdf to converge to the models as the number of samples increase. Accordingly, the probability of a fit for *a given pdf* will decrease as the number of samples increases. Since in this application we use a very large number of samples the measured pdf must be almost exactly Gaussian.

In fact, the measured distributions are not Gaussian, but may be described as near-Gaussian. The large sample size means that the Chi square test will return a vanishingly small probability for distributions which are near-Gaussian. In this case we must rephrase the question, instead of asking *is this distribution Gaussian*? we must ask *how close is the observed distribution to the Gaussian form?* We therefore seek a figure of merit which operates on the pdf. and gives a smoother transition as the actual pdf. moves away from the model.

The requirement for a smooth transition is shared by criteria used in parameter estimation. With this in mind we define a modified form of the L_1-norm [Siva, p.67]. We define the figure in equation 4.9.

$$L_1 = \sum_{i=0}^{N} \left| p_o(i) - p_g(i) \right|$$

where p_o is the observed probability density function, and p_g is the pdf of a Gaussian distribution with the same mean and standard deviation.

In order to relate the figure to a meaningful form, we show the pdf of the textures with the lowest and highest measures, (*Figure 4.8.1*) before tabulating the criterion in *Table 4.8.1*.

In all cases the Gaussian model is parameterised by mean and standard deviation measured in the usual way, with no parameter optimisation.

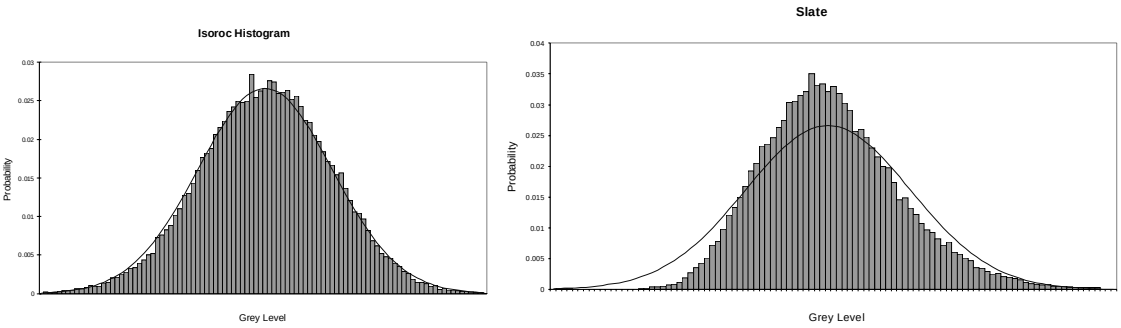

Figure 4.8.1 Histogram distributions of grey levels for most and least Gaussian distributions.

Texture	**L_1 Statistic**	**Texture**	**L_1 Statistic**
Isoroc	0.039	Pitted	0.095
Rock	0.077	Radial	0.173
Slab	0.053	Slate	0.173
Striate	0.061	Twins	0.099

Table 4.8.1 Degree of 'normality' in grey level distribution.

The distributions range from the very Gaussian-like *Isoroc* to the positively skewed, *Slate* and *Radial* textures. However, even the *Slate* and *Radial* textures are broadly Gaussian and we shall retain the assumption of normality in the theoretical models used in this work.

4.8.2 Second Order Statistics

We have described the data set as being the output of a linear system with input conforming to one of the surface models described in Chapter 2.

$$I(\omega,\theta) = S(\omega,\theta).\omega.k.\cos^2(\theta-\tau).B(\omega,\theta) + W$$

We plot the power spectra of the images of the test surfaces illuminated at τ 0° in *Figure 4.8.2.*

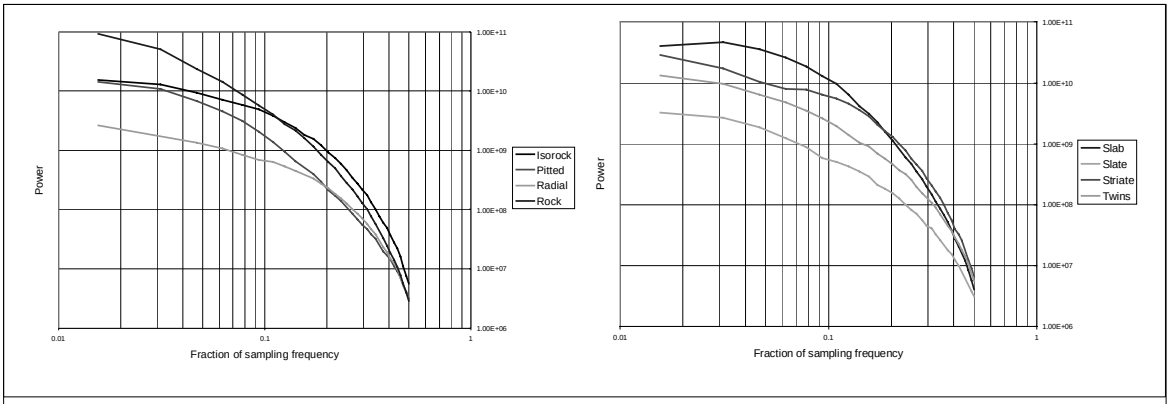

Figure 4.8.2 Power spectra of test textures illuminated at τ=0°.

The first point we note from *Figure 4.8.2* is that the power spectra are no longer fractal, i.e. they do not have a linear roll-off when plotted on a log-log graph. We believe this to be primarily due to the Gaussian transfer function, associated with blur acting as a low pass filter. In *Figure 4.8.3* we plot the image power spectra against that of the fractal surface rendered and 'imaged' using our models. The spectra are not identical, however, they are in reasonable agreement and this result does give some verification of the models used in this thesis.

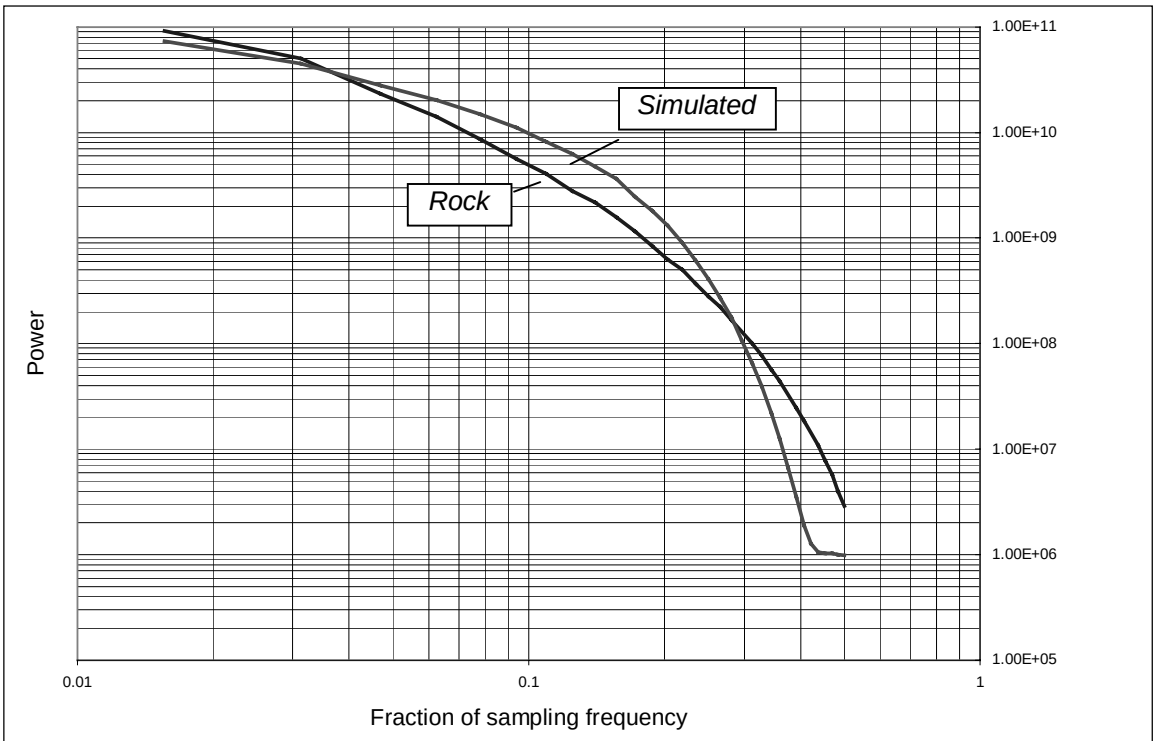

Figure 4.8.3 Comparison of Rock spectra with the blurred fractal spectra.

The result discussed above while analytical, does not give a very intuitive idea of the accuracy of our models. In *Figure 4.8.4* we show the original *Rock* image as well as a second semi-synthetic image. The second image was obtained by imposing the predicted power spectrum on the *Rock* image while retaining the original's phase spectrum.

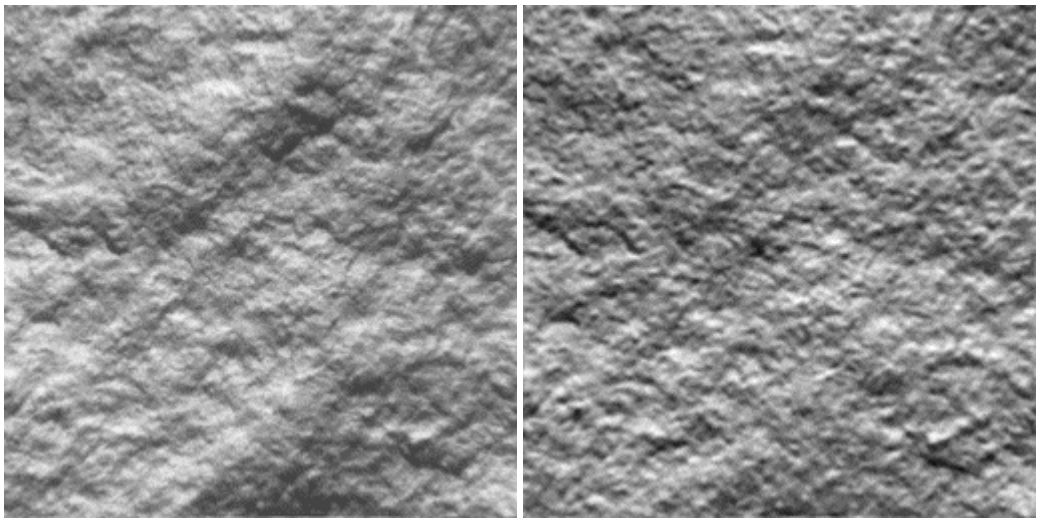

Figure 4.8.4 Comparison of original rock surface with modified PSD reconstruction.

4.9 Conclusions

Summary

In this chapter the case for down-sampling the test image has been demonstrated on grounds of noise reduction and the advantages of subsampling to classification, rotation invariant classification and image modelling have been highlighted.

The chapter has consulted both the literature and experimental data to model the transition of the signal from incident image to data set. The form and parameter values of the model have been estimated for the rock surfaces used in the work described in this thesis.

Blurring was found to be the dominant artefact of imaging, and a Gaussian model of the imaging system transfer function was adopted. We have shown that the common assumption of a white noise model is not justified, in our case at least, without down-sampling.

We have established signal to noise ratios for the textures considered in this thesis, giving a measure of the quality of these images. Furthermore, we have measured the S/N ratio as a function of sampling frequency, allowing the effects on spectrally-definable

classification feature measures to be considered.

We have also compared the statistics of the imaged textures with those predicted by our models. We found that the first order statistics while not *exactly* Gaussian could *reasonably* be described as such. We also found that the spectra were not fractal, though the measured spectra were reasonably approximated with our surface, rendering and imaging models.

Implications for Texture Analysis

The fractal nature of the signal and the observed effects of blur and white noise discourage the use of high frequencies as consistent sources of discriminatory information. While the blur function reduces the feature mean and variance, particularly at high frequencies, the additive white noise increases the texture's variance in feature space. The magnitude of the degradation in classification will depend on the proximity of the textures in feature space as well as the level of the noise itself.

While the absolute power of noise seems to be similar in all images, the power relative to the texture varies depending on which texture is used. Consequently, we cannot apply Wiener filtering to the image prior to classification. Use of the related technique of spectral subtraction is possible, albeit with the introduction of serious distortions. In line with the philosophy of this thesis, processing beyond sub-sampling will be avoided.

The directionality of noise and imaging effects for the original image makes the unprocessed image unsuitable for rotation invariant algorithms. Fortunately, these directional effects can largely be suppressed by subsampling the image and we recommend the operation prior to the use of a rotaion-invariant algorithm.

The findings of this chapter are relevant to the remainder of this thesis. The results affect the design of the classifier in chapter 5 and its modelling in chapter 6. The most significant contribution of this chapter is, however, to chapter 8, where an experimental algorithm is evaluated using the simulation engine developed in this chapter.