

Noise in 3D Laser Range Scanner Data

Xianfang Sun*
Cardiff University, UK
Beihang University, China

Paul L. Rosin†
Cardiff University, UK

Ralph R. Martin‡
Cardiff University, UK

Frank C. Langbein§
Cardiff University, UK

Abstract

This paper discusses noise in range data measured by a Konica Minolta Vivid 910 scanner. Previous papers considering denoising 3D mesh data have often used artificial data comprising Gaussian noise, which is independently distributed at each mesh point. Measurements of an accurately machined, almost planar test surface indicate that real scanner data does not have such properties. An initial characterisation of real scanner noise for this test surface shows that the errors are not quite Gaussian, and more importantly, exhibit significant short range correlation. This analysis yields a simple model for generating noise with similar characteristics. We also examine the effect of two typical mesh denoising algorithms on the real noise present in the test data. The results show that new denoising algorithms are required to effectively remove real scanner noise.

Keywords: 3D laser scanner, correlation analysis, Fourier analysis, noise modeling.

Index Terms: G.1.2 [Numerical Analysis]: Approximation—Approximation of surfaces and contours; I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations

1 Introduction

Noise is ubiquitous in measured data. Surface mesh models built using measurement data obtained using 3D range scanners necessarily contains some type of noise. To remove the noise in surface mesh models, many mesh denoising algorithms have been developed, for example [1, 3, 4, 5, 6, 8, 9, 13, 14, 17, 19, 20, 21, 22, 23, 24, 25, 26, 27] and other references therein. In evaluating the effectiveness of denoising algorithms both visual and numerical comparisons are used [20]. For meshes corresponding to real scanned data, however, we can in most cases only perform visual comparisons, because ground truth data required for evaluation are almost always unavailable. However, to provide a more objective evaluation and more thorough testing, numerical comparisons are required. Having synthetic models of real scanner noise would be useful for evaluating denoising algorithms.

Any synthetic model used for evaluating denoising algorithms should take an exact model surface and add noise, which is to be removed by the algorithms. The noise should have the same characteristics as noise in real measurement data. Experimentally, measurement noise is often assumed to be Gaussian in wide range of disciplines. Thus, various mesh denoising algorithms have been developed based on the Gaussian noise assumption [1, 6], while others used synthetic models with Gaussian white noise (that is, independent Gaussian noise per mesh vertex) in evaluation of their algorithms [3, 13, 20, 26, 27]. However, real 3D laser scanner noise is, in practice, not quite Gaussian according to our observations, and

is correlated at adjacent mesh points. Shen [18] has also previously observed that real measurement noise on *edges* does not exactly follow the pattern prescribed by a random number generator, and has given a new method of generating synthetic edge noise.

One approach to understanding system noise is to combine noise models for the components of the measurement system. For example, an optical sensor is used, and intensity noise models for optical sensors have been discussed in the literature [10]. However, depth values are not computed directly from intensity values, but instead triangulation is used after locating pixels of *maximum* intensity. Other sources of noise might include optical components such as the lens, the mirror (and its drivers), and laser, as well as electronic noise. Furthermore, numerical errors may also arise in the proprietary software used to calculate positions. Ultimately, especially as we do not have access to all the components of a commercial system, and even if we did, the noise coming from some components may not be well characterised, we suggest that such a componentwise modeling approach is not practical. Instead, therefore, we simply consider the overall noise present in the output of the scanner.

In this paper, we show that the measurement noise in a *flat area* does not follow the common assumption of Gaussian white noise, and we analyse real measurement noise from a 3D range scanner. We also give a method to synthesise noise with similar characteristics. We then discuss the different results obtained by applying certain denoising algorithms to remove real scanner noise, and synthetic Gaussian white noise. The aim of this paper is to bring the attention of the mesh processing community to the fact that **real scanner noise is not independent Gaussian noise per mesh vertex**—any algorithm evaluation and development should be based on more realistic assumptions about 3D scanner noise.

The remainder of this paper is organised as follows. Section 2 discusses our approach to acquisition of the test data we used to characterize noise. We explain the choice of test specimens and various data preprocessing procedures, and discuss our assumptions about the shape of the underlying test surface. In Section 3 we analyse the noise in the test data based on a quasi-statistical approach to describe the distribution and correlation of the scanner noise. We further employ the discrete Fourier transform to characterise the scanner noise. Section 4 describes artificial noise synthesis based on this analysis using the inverse discrete Fourier transform. Section 5 presents the results of using two typical mesh denoising algorithms to demonstrate the difficulties of removing real scanner noise, compared to removing synthetic Gaussian white noise. Finally, Section 6 concludes this paper.

2 Data Acquisition and Preprocessing

2.1 Data Acquisition

We describe in this section the methods used to acquire noisy data using a commercial scanner, and a standard test object. The approach was to use a highly accurate test object, so that any noise in the measurements would be attributable to the scanner, and not to the test piece itself.

The scanner we used in our test was a Konica Minolta Vivid 910, which has a specified point accuracy of $50\mu\text{m}$. Boehnen and Flynn [2] reviewed several different 3D scanners, and found that that Konica Minolta Vivid 910 was the most accurate scanner in a

*e-mail: Xianfang.Sun@cs.cardiff.ac.uk

†e-mail: Paul.Rosin@cs.cardiff.ac.uk

‡e-mail: Ralph.Martin@cs.cardiff.ac.uk

§e-mail: F.C.Langbein@cs.cardiff.ac.uk

face scanning scenario. Here we take it as a representative commercial scanner using laser triangulation to capture range data. To acquire noisy data from this scanner, some specimen must be measured. Clearly, the higher the precision of the specimen, the more certain we can be that the noise is due to the scanning process. We first tried to use a standard gauge block as the specimen. While it was very flat, it was too polished and almost mirror-like, with the result that the scanner was unable to acquire satisfactory measurements. It was thus clear that a *slightly* rough surface would be needed to obtain satisfactory measurements. (The standard alternative, coating the test part with a white powder, would lead to an unacceptably uneven surface due to the nature of the powder spraying process).

We chose to use the N1 specimen from a Microsurf 315 test surface produced by Rubert & Co Ltd as our test surface. Such test plates are notionally produced for the purpose of testing surface roughness measuring machines. This specimen is a $22.5 \times 15\text{mm}^2$ rectangular flat metal plate with mean roughness $R_a = 0.025\mu\text{m}$ and mean roughness depth $R_z = 0.29\mu\text{m}$. Here R_a and R_z are parameters of surface roughness as defined by ISO Standard 4287/1:1984 [11]. R_a is the arithmetic average of the absolute values of the roughness profile ordinates, and R_z is the arithmetic mean value of the single roughness depths, taken over consecutive sampling lengths [16]. The sampling length used to compute R_a and R_z was given as 0.8mm according to Rubert & Co Ltd. Since both R_a and R_z are significantly smaller than the point accuracy specified for the Konica Minolta Vivid 910 scanner, in the rest of the paper, we may simply consider any sub-area on the test surface with diameter smaller than the sampling length as planar, and ignore noise due to the surface roughness. However, over sub-areas larger than 0.8mm in extent, we must be careful not to assume that the surface is exactly planar—on the scale of the whole test piece, it may exhibit non-negligible curvature.

Having chosen the test specimen, we collected measurement data as follows. During the scanning, the test plate was fixed at about 585mm from the focus of the scanner, with its normal approximately aligned with the scanner's optical axis (unless stated otherwise). (Although we did not use an instrument to align the optical axis precisely with the normal to the test surface, the actual angle between the optical axis and the normal to the test plate could be computed accurately from the test data itself). Unless otherwise stated, we aligned the long side of the N1 specimen with the x direction (this is relevant as the N1 specimen is produced by grinding, so its surface finish is anisotropic, albeit on a very fine scale). For each complete scan we obtained an array of 3D point coordinates $\{x_i, y_i, z_i\}$ in the scanning coordinate system, where x and z axes are along the scan line and the optical axis, respectively, and the x, y, z axes form a right-hand coordinate system. Here, $\{x_i, y_i\}$ are coordinates on a grid with an interval of approximate 0.1734mm between successive points in both x - and y -directions, and z_i is the distance in the z -direction from a surface point i to the focus of the scanner.

To analyse the properties of the noise produced by the scanner, we captured data from the test plate in several different ways. Firstly, in order to test the repeatability, we fixed the test plate in the same place and repeatedly scanned it several times. Test results showed that the noise properties were repeatable in the statistical sense that the noise mean, standard deviation, and correlation varied little between scans, and visually, the scans had similar raised and sunken areas as shown in Fig. 1. This repeatability justifies the use of a single set of test data in our analysis except where more data are required.

Secondly, to analyse the effect of surface orientation on noise, we rotated the test plate through varying angles about the x -axis and scanned the plate to obtained different data. Test results showed that varying orientation had little effect on the results—only minor changes were found in standard deviation and correlation, and ob-

vious raised and sunken areas still existed. We thus only present our analysis on the data measured when the optical axis was orthogonal to the test plate.

Thirdly, to assess the effect of the scan line direction on the noise, we took two scans with the test plate aligned so that first its long side ran approximately in the x -direction, then in the y -direction. This test allowed us to verify that the anisotropy of the test surface does not affect the noise.

2.2 Data Preprocessing

After capturing the raw measurement data we preprocessed the data to separate the noise from the background surface. We performed the following preprocessing steps:

S1: Trim off data coming from near the edges of the test piece, as the noise in such regions has quite different characteristics from that at interior points, and needs to be analysed specifically [18]. Our trimmed data consist of 9375 gridded points with 125 columns in x -direction and 75 rows in y -direction, and the intervals between successive points in x - and y -directions are 0.1735mm and 0.1733mm, respectively.

S2: Fit a smooth surface around each measured surface point i in 3D space using the measurement data in its neighbourhood N_i , $\{(x_j, y_j, z_j) : j \in N_i\}$. We represent the fitted surface as

$$z = f(x, y, \mathbf{p}), \quad (1)$$

where \mathbf{p} is a parameter vector of the surface. We will discuss shortly the choice of surface model used and neighbourhood size.

S3: Estimate the measurement noise e_i corresponding to each surface point i by

$$e_i = z_i - f(x_i, y_i, \mathbf{p}). \quad (2)$$

In the second preprocessing step above, we need to choose a surface model for fitting, and also the size of the neighbourhood used. For this we examined the following three assumptions about the shape of the underlying test surface:

A1: A plane gives an adequate *global* fit to the data;

A2: A quadratic surface gives an adequate *global* fit to the data; or

A3: No assumption is made about the *global* shape. Instead a procedural model is used based on iterative fitting of planes to data in local regions taking into account the ISO noise characterization parameters.

If we make Assumption A1, we fit a plane to the whole set of measurement data. Function f in Eqn. (1) is defined as

$$f(x, y, \mathbf{p}) = a + bx + cy. \quad (3)$$

where the parameter vector is $\mathbf{p} = [a, b, c]$.

Fig. 1 shows the measured noise using Eqn. (2) in accordance with Assumption A1. To help visualise the noise surface, we have enlarged the noise values e_i 50 times, and used Phong shading to produce Fig. 1(a); *actual* noise heights (in mm) are given by the associated colour bar, that is, the shading effects have been exaggerated while the colouring is unaltered. Similar Figures discussed later have also been processed in a similar way. Fig. 1(b) simply shows the signs of the noise values as binary values: black (negative) and white. From Fig. 1, it can be seen that the noisy surface has obvious raised and sunken areas, and is obviously curved—its top-right and bottom-left corners are visibly higher than other parts of the surface.




Figure 1: Extracted noise given Assumption A1, shown with Phong shading: (a) measured noise (colour scale in mm), and (b) the sign of the noise: positive, white; negative, black.




Figure 2: Extracted noise given Assumption A1 (specimen rotated 90° relative to Fig. 1 before scanning), else as per Figure 1.

Firstly, we note that bumps of the magnitude present in the noise surface would easily be visible to anyone examining the test plate visually, if they were actually present. The test plate does not show an appearance anything like that in Figure 1, but instead, a very fine pattern of x -direction scratches at a much smaller size (due to the grinding process used to make the test plate). It is clear then that the short distance scale structure visible in Figure 1 is due to measurement errors, and not due to surface bumps in the test plate itself. This is backed up by the observation that the measured bumps are much greater than the manufacturer's claimed roughness of the test plate.

A second issue is whether the apparent curvature of the noise surface is due to actual curvature of the test surface (again, it is only guaranteed to be *locally* flat), or whether it is due to some systematic error on a long scale produced by the scanner.

To try to resolve the second issue, we rotated the specimen by 90° around the z -axis, and obtained another data set. Fig. 2 shows the extracted noise in this case, again using Assumption A1. It can be seen that while the detailed noise pattern is somewhat differ-




Figure 3: Extracted noise given Assumption A2, as per Figure 1.




Figure 4: Extracted noise given Assumption A3, as per Figure 1.

ent, the overall curvature in the noise surface has also rotated with it—its top-left and bottom-right corners are now visibly higher than other parts of the surface. This leads us to believe that the curvature of the noise surface comes from the specimen itself, and is not due to systematic structuring of noise produced by the scanner—in the latter case, the noise distribution would not have rotated. As a result, we can thus rule out Assumption A1 as being unsatisfactory, and do not use it for further analysis.

We also note that while in Fig. 1 the noise shows a diagonal structuring tendency, in Fig. 2 this is less clear, and even in some places tends perhaps more to a horizontal structuring. This is probably due to interaction between the laser striping system, the anisotropy of the test surface, and the relative angles of laser and camera systems. Apart from this question of directionality, however, the general structure of the noise appears similar in both orien-

tations. Thus, future experiments were restricted to data measured when the long side of the specimen was aligned with the a -axis.

Next, we consider Assumption A2. In this case, we fit a quadratic surface to the whole set of measurement data. In this case the function f is defined as

$$f(x, y, \mathbf{p}) = a + bx + cy + dx^2 + exy + fy^2. \quad (4)$$

where the parameter vector is $\mathbf{p} = [a, b, c, d, e, f]$.

Fig. 3 shows the estimated noise using Assumption A2. It can be seen that the noise surface again has obvious large scale raised and lowered areas with a similar distribution as in Fig. 1, except that the overall curvature in Fig. 1 is no longer present. Using the same reasoning as before, we can again say that the magnitude of the noise is due to the scanner.

We still need to answer the question of whether a quadratic surface can sufficiently represent the underlying surface of the specimen. To answer this question, we have also fitted higher order surfaces, e.g. cubic surfaces, to the measurement data, and extracted the noise using Eqn. (2). Visually, the distribution pattern of the extracted noise in each case is quite similar to that shown in Fig. 3. Furthermore, differences in the residual errors when fitting quadratic, and cubic or other higher order surfaces, are less than the specified roughness heights R_a and R_z of the test specimen. Thus, we can safely say that a quadratic surface sufficiently represents the underlying surface of the specimen, and the estimated noise surface in this case is very close to the real noise due to the scanner.

We have also used assumption A3 as an alternative method of verifying the above assertion: A3 makes no assumptions about the global shape of the underlying test surface. Under Assumption A3, we only need to consider the parameters of roughness. Since the parameters of roughness R_a and R_z of the specimen are $0.025\mu\text{m}$ and $0.29\mu\text{m}$, respectively, which are both significantly smaller than the scanner accuracy of $50\mu\text{m}$, we can simply consider the surface as a plane over areas of diameter smaller than the sampling length used to define R_a and R_z , and ignore the influence of the surface roughness.

This sampling length is 0.8mm . So for each measured point on the surface, we chose a circular neighbourhood of diameter 0.8mm , and fitted a local plane to the measured data within this neighbourhood. The point was then projected onto the plane, and the projected point was taken as the estimated position of the real point on the underlying test surface. Because the surface fitted to the noisy measured points has a roughness exceeding the range designated by R_a and R_z of the specimen, unfortunately it is not a good approximation to the real surface. Thus, we iteratively used the approximate points to fit local planes and provide new projected approximate points until the roughness of the surface formed by the new approximate points reached the range designated by R_a and R_z (it required 323 iterations to do so). The final approximate points can be considered to be good approximations to the real surface points, and the differences between the noisy points and their corresponding approximate points can then be considered to be our estimate of noise. Note that because of the iteration process, the surface may have shifted—in fact, the mean of the differences is $-2.3\mu\text{m}$, so we subtract the mean from each difference, to get the final estimated noise with zero mean.

Fig. 4 shows the extracted noise surface using the above method. The noise pattern can be seen to be very similar to that estimated using Assumption A2 shown in Fig. 3, although the actual values differ slightly. The independent assumptions made for approaches A2 and A3 give us confidence that the similar results obtained are useful estimates of the actual scanner noise surface, and that this noise really does contain the structures shown in Fig. 3.

We have computed the standard deviation of the differences between noise values estimated using A2 and A3. It is $4.9\mu\text{m}$, which is only 10% of the specification accuracy of the scanner. While




Figure 5: Differences between Fig. 3 and Fig. 4.




Figure 6: Histogram of the estimated noise magnitude, the red line shows the estimated Gaussian distribution.

these differences are not negligible, the difference surface is slowly varying over a long distance with low amplitude (see Fig. 5). We also computed the standard deviations of the noise based on A2 and A3, which are $16.2\mu\text{m}$ and $16.6\mu\text{m}$, respectively. They are very close and both are less than the specified accuracy of the scanner.

For purposes of analysing correlations in noise in x - y , both assumptions A2 and A3 provide reasonable models of the real scanner noise; however, assumption A2 leads to a much cheaper computation, so in the rest of the paper, our analysis is based on Assumption A2 unless otherwise stated.

We now make perhaps the most important point in the whole paper, which is clear from a visual inspection of Figures 3 and 4: **noise at each measurement site is not independent**. There is a clear correlation between noise at adjacent measurement sites, and indeed on longer length scales, too. This is contrary to a common assumption adopted by many papers on noise removal methods for mesh data: the model most often assumed for scanner noise is independent Gaussian noise at each measurement point [1, 13, 27, 3, 6, 26].

3 Noise Analysis

In this section, we now analyse the properties of the estimated noise. We carry out both a statistical analysis, and a Fourier analysis, of the noise. Because the extracted noise is only an approximation to the real scanner noise, we call our analysis a *quasi-statistical* analysis of the real scanner noise. We use the noise estimate based on Assumption A2 throughout this section.

3.1 Quasi-Statistical Analysis

3.1.1 Is the Noise Gaussian?

We start our quasi-statistical analysis with a discussion of the distribution of the magnitude of the estimated noise. Fig. 6 shows a histogram of the estimated noise taken at all points of the test surfaces, together with a zero-mean Gaussian distribution with the same standard deviation as the estimated distribution. This histogram has 80 equally spaced bins. The noise histogram appears to agree fairly well with the Gaussian distribution, but to decide this issue more definitely, we perform a statistical test.

The standard deviation of the estimated noise is $\sigma = 16.2\mu\text{m}$. We need to verify the null hypothesis:

$$H_0 : f(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x^2}{2\sigma^2}}$$

with the alternative hypothesis

$$H_1 : f(x) \neq \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x^2}{2\sigma^2}}$$

where $f(x)$ is the probability density function of the estimated noise. For this we use Pearson's chi-square test [12]. Consider the statistic

$$K = \frac{1}{N} \sum_{i=1}^m \frac{n_i^2}{p_i} - N, \quad (5)$$

where N is the number of observations (the number of measurement points, in our case here, $N = 9375$), m is the number of bins in the histogram (see later), n_i is the number of observations in the i^{th} bin, and p_i is the probability content of the i^{th} bin. Under the hypothesis H_0 , the distribution of K is generally accepted as being close enough to $\chi^2(m-1)$ if the expected numbers of events per bin, Np_i , is greater than five in each case [12]. If the probability $P(\chi^2(m-1) > K) < \alpha$, the null hypothesis is said to be rejected in the significance level α . Commonly used α are 5%, 1%, and 0.1%. The smaller the α , the stronger the evidence. In the rest of the paper, we use $\alpha = 5\%$ as the significance level.

We generated histograms using $m_0 = 3, \dots, 100$ equal-width bins, except that we combined neighbouring bins having small expected numbers of events into larger bins such that each was expected to have more than five events.

We then used Eqn. (5) to compute K and computed the probability of $\chi^2(m-1) > K$, where m was the number of bins after merging small bins. The result shows that except for the cases $m_0 = 4, 5, 7, 8, 10$ (corresponding to $m = 4, 5, 5, 6, 7$, respectively), the probabilities are all less than 0.05, and furthermore, except for $m = 4$ ($m_0 = 4$) and $m = 6$ ($m_0 = 8$), there always exists another division of m bins which has a probability of less than 0.05. This means that except for two special choices of m , we should reject H_0 at the given level of significance $\alpha = 5\%$. In simple words, **we should reject the hypothesis that the estimated noise distribution is Gaussian**. Nevertheless, Fig. 6 makes it clear that the distribution of the estimated noise is quite Gaussian-like.

We also note that, as well as the obvious interpretation that the noise itself is not Gaussian, other explanations are available. Errors in the estimated noise due to the estimation procedure, or the presence of small bumps on the surface, may also explain the small deviation between the estimated noise and a Gaussian model.

3.1.2 Noise Correlation

We now consider noise autocorrelation in the x - and y -directions. For convenience of description, we assume that the data points are indexed from left to right (x), and from top to bottom (y). We use $\{i, j\}$, $i = 1, \dots, N_x$, $j = 1, \dots, N_y$, to denote the point in the i^{th} column and the j^{th} row from top, and $e(i, j)$ to denote the estimated noise at this point. The actual values were $N_x = 125$ and $N_y = 75$ for our measured data.

We also performed a statistical hypothesis test to evaluate the noise autocorrelation. If the noise distribution were Gaussian, we could simply use linear correlation coefficient and perform t -test to assess the autocorrelation. However, because the extracted noise is not Gaussian as we showed in the last section, it is more appropriate to use a nonparametric test. We have used the popular Spearman rank-order correlation coefficients, which does not require the knowledge of the probability distribution of the noise [15].




Figure 7: Spearman autocorrelation showing probability $P(t(M-2) > |T|)$ in x -direction (top) and y -direction (bottom).

Let $k = 1, \dots, N_x - 1$ be the autocorrelation length to be considered in the x -direction, and $R_{x1k}(i, j)$ and $R_{x2k}(i, j)$ be the rank values of $e(i, j)$ among all the other $e(m, n)$ in the square areas ($m = 1, \dots, N_x - k$, $n = 1, \dots, N_y$) and ($m = k + 1, \dots, N_x$, $n = 1, \dots, N_y$), respectively. Furthermore, let \bar{R}_{xk} be the mean of the rank value in either square area, given by $((N_x - k)N_y + 1)/2$. Define $\tilde{R}_{x1k}(i, j) = R_{x1k}(i, j) - \bar{R}_{xk}$ and $\tilde{R}_{x2k}(i, j) = R_{x2k}(i, j) - \bar{R}_{xk}$. Then the rank-order autocorrelation coefficient in the x -direction is computed by

$$\rho_{Sx}(k) = \frac{\sum_{j=1}^{N_y} \sum_{i=1}^{N_x-k} \tilde{R}_{x1k}(i, j) \tilde{R}_{x2k}(i+k, j)}{\sqrt{\sum_{j=1}^{N_y} \sum_{i=1}^{N_x-k} \tilde{R}_{x1k}^2(i, j)} \sqrt{\sum_{j=1}^{N_y} \sum_{i=1}^{N_x-k} \tilde{R}_{x2k}^2(i+k, j)}}. \quad (6)$$

Similarly in the y -direction, the rank-order correlation coefficient is computed by

$$\rho_{Sy}(k) = \frac{\sum_{i=1}^{N_x} \sum_{j=1}^{N_y-k} \tilde{R}_{y1k}(i, j) \tilde{R}_{y2k}(i, j+k)}{\sqrt{\sum_{i=1}^{N_x} \sum_{j=1}^{N_y-k} \tilde{R}_{y1k}^2(i, j)} \sqrt{\sum_{i=1}^{N_x} \sum_{j=1}^{N_y-k} \tilde{R}_{y2k}^2(i, j+k)}}. \quad (7)$$

The notations in the above formulae have corresponding meanings to those in Eqn. (6).

Having obtained the rank-order correlation coefficients, we can perform the t -test to evaluation the autocorrelation of the estimated noise. Our null hypothesis is that the noise at each measurement location is uncorrelated, corresponding to

$$H_0 : \rho_{Sx}(k) = 0 \text{ for } k = 1, \dots, N_x - 1,$$

and similarly for y . The alternative hypothesis is that

$$H_1 : \rho_{Sx}(k) \neq 0,$$

and similarly for y .

Given a correlation coefficient ρ , (here ρ could be $\rho_{Sx}(k)$ or $\rho_{Sy}(k)$), the statistic

$$T = \sqrt{M-2} \frac{\rho}{\sqrt{1-\rho^2}}, \quad (8)$$

is close to $t(M-2)$ (t -distribution with $M-2$ degrees of freedom) under hypothesis H_0 . (Here M is the number of sample points used in computing ρ .) Thus, if the probability $P(t(M-2) > |T|) < \alpha$, the null hypothesis is said to be rejected in the significance level α .

Fig. 7 shows how the probability $P(t(M-2) > |T|)$ varies with k with the green line at $P = 5\%$. From the figure it can be seen that $P(t(M-2) > |T|) < 5\%$ for all correlation lengths less than 11 in the x -direction and 47 in the y -direction, respectively; there are also




Figure 8: Power spectrum along x -direction (top) and y -direction (bottom); the green vertical lines are at frequency $1/0.8\text{mm}$.

other greater correlation lengths for which $P(t(M-2) > |T|) < 5\%$. Thus it is highly suggestive that the noise is autocorrelated in both x - and y -directions, at various length scales.

We have also performed the t -test based on linear autocorrelation coefficients, which are greatly dependent on the assumption that the noise is Gaussian. Its result is quite close to the above even though the Gaussian assumption is not exactly correct as we showed in the last section. In summary, we can thus say that **the estimated noise is really autocorrelated**.

3.2 Fourier Analysis

The above section provides a quasi-statistical analysis of the estimated noise. Because this analysis is based on Assumption A2, and hence an estimate of the noise, we can only state that the above analysis provides a qualitative indication. In this section, we perform an alternative Fourier analysis of the measurements which is not based on any assumption of the specimen shape.

We first performed a 1D Fourier analysis on the measurement data (not the extracted noise). Fig. 8 shows a log-log scale plot of the power spectrum along the x - and y -direction respectively. From the figure it can be seen that the measurement data are correlated since the power spectra are not constant in any band. However, the correlation here could be from either the measurement noise or the surface signal itself of the specimen. Considering that the specimen surface is quite smooth and its fluctuations are quite small, the signal will mainly exist at low frequencies, and the high-frequency components will mainly consist of measurement noise. The vertical green line shown in Fig. 8 is at the frequency $f = 1/0.8\text{mm}$. Note that 0.8mm is the sampling length given for calculating the surface roughness by Rubert & Co Ltd. As we discussed before, the surface roughness below the sampling length is negligible, and thus we can be sure that the spectral power for frequencies $f > 1/0.8\text{mm}$ (to the right of the green line) is essentially due to scanner noise. Since the power spectra are also not constant for frequencies $f > 1/0.8\text{mm}$, we are sure that the measurement noise is not white noise.

To consider the noise as a 2D process, we further performed a 2D Fourier analysis on the measurement data. Fig. 9 shows its 2D discrete Fourier transform. It can be seen that, as was shown by the 1D Fourier transforms, the magnitude of the 2D discrete Fourier transform also decreases with the frequency, and it varies little in the high frequency range. It can also be seen that the phase of the 2D discrete Fourier transform varies randomly with little regularity. The 2D Fourier transform shown here also tells us that the measurement data are correlated, and furthermore, the measurement noise is not white noise.




Figure 9: 2D discrete Fourier transform (zero-frequency at the centre). Top: logarithmic magnitude of the spectrum. Bottom: phase of the spectrum, in degrees.

4 Noise Synthesis

From the quasi-statistical analysis in Section 3.1, we see that the noise is not white and its distribution is not Gaussian, thus we cannot generate synthetic noise using straightforward pseudo-random noise sequences. However, using the Fourier analysis in Section 3.2, we may generate synthetic noise via use of inverse Fourier transforms as we now discuss.

When performing Fourier analysis in Section 3.2, we used the measured data, and thus the result of the Fourier transform is not purely due to noise—it is a mixture of both the noise and the signal (variations in shape of the specimen surface from planar). Because the signal exists mainly at low frequencies while the noise mostly has high frequencies, we may simply set the low frequency components of the Fourier transform to zero, and use the remaining components to model the noise—although again this is not an exact description of the real noise. Usually, a sharp cut-off filter as used here can introduce ringing artifacts. However in practice, these were negligible and so were not considered further.

The top of Fig. 10 shows the inverse Fourier transform with the low frequency components (the lowest 5 components in both x - and y -directions) set to zero. Clearly, using the inverse Fourier transform in this way preserves the general structure of the original noise shown in Figs. 1, 3, and 4, while removing the obvious curvature of the specimen surface shown in Fig. 1.

Having obtained the Fourier transform of the measurement data, we can use a simple model to fit the magnitudes with respect to frequency. We tried several models, and finally chose one model which has minimum fitting residual among these models. The model we use is

$$Z(i, j)^{\frac{1}{4}} = \begin{cases} a_0 + b_0i + c_0j + d_0i^2 + e_0ij + f_0j^2 & i \neq 0, j \neq 0, \\ a_1 + b_1i + d_1i^2 & j = 0, \\ a_2 + b_2j + f_2j^2 & i = 0, \end{cases} \quad (9)$$

where $Z(i, j)$ is the magnitude at frequency $f_x = i/L_x$, $f_y = j/L_y$, and L_x and L_y are the specimen side lengths in the x - and y -directions respectively. The parameters obtained by finding a best fit are $a_0 = 1.5600$, $b_0 = -0.0185$, $c_0 = -0.0176$, $d_0 = 0.0001$, $e_0 = 0.0003$, $f_0 = 0.0000$, $a_1 = 1.9134$, $b_1 = -0.0417$, $d_1 = 0.0005$, $a_2 = 1.8352$, $b_2 = -0.0530$, $f_2 = 0.0008$, and the squared sum of the residual errors is 2.21% of the squared sum of the magnitudes. Fig. 10, middle, shows the result of computing an inverse




Figure 10: Generating noise by inverse Fourier transform with low frequency components set to zero. Top: original magnitude and phase except for low frequency components; Middle: synthetic noise with fitted magnitude and original phase except for low frequency components; Bottom: synthetic noise with fitted magnitude and random phase except for low frequency components.

Fourier transform using the modelled magnitude together with the original phase data. This generated noise can be seen to have a structure quite close to the original structure, verifying the utility of this model.

Because the phase spectrum is rather unstructured as shown in Fig. 9, we cannot easily fit a model to the phase spectrum. Instead, we may use random phases for generating synthetic noise. Fig. 10, bottom, shows noise synthesised by inverse Fourier transform of the modelled magnitudes and random phases. The synthetic noise in this case can be seen to have a somewhat different structure from the original one, which is to be expected as there is *some* structure in the phase of the noise, as shown in Fig. 9. Nevertheless, the structure of this synthetic noise is generally fairly similar to the original noise at least with respect to the magnitudes, sizes, shapes and density of the bumps. In practice, synthetic noise is not necessarily an exact copy of any original noise, and we consider the noise generated using inverse discrete Fourier transforms of fitted magnitudes and random phase as reasonable. It is certainly closer to real scanner noise than white Gaussian noise.

5 Denoising experiments

Section 3.1 has shown that real scanner noise is neither strictly Gaussian, nor independent at each location or even white in nature, contrary to the most common assumptions adopted by many papers

considering denoising or other processing of scanner data. In this section, we consider how various algorithms perform on real noise as opposed to synthetic independent Gaussian noise often used to test such algorithms in the papers where the authors have described them and assessed their performance. Here, we consider two typical denoising algorithms: the original Laplacian algorithm [24] and a simple feature-preserving denoising algorithm [20]. We compare the denoising results achieved by these two algorithm, firstly on real scanner data, and secondly on data generated by adding Gaussian white noise with the same standard deviation as that of scanned data on the best quadratic surface fitted to the scanning data.

Fig. 11 shows the denoising results after 1, 5, 10, 20, and 50 iterations, as commonly used in the literature. It is quite clear that fewer iterations of these denoising algorithms are needed to remove Gaussian white noise than to remove real scanner noise. Furthermore, the denoising algorithms are less successful at completely removing the bumpy structures present in real scanner noise.

Because a feature-preserving denoising algorithm like that from [20] can preserve features better than a non-feature-preserving algorithm, such as the original Laplacian algorithm [24], the former has more difficulty than the latter in remove structural features in real scanner noise.

From the denoising experiments, we can see that neither non-feature-preserving nor feature-preserving algorithms can effectively deal with structured nature of real scanner noise. We thus conclude that **many previous papers claiming good smoothing results based on experiments with synthetic noise are over-optimistic in their assessment of the ability of the reported methods to remove real scanner noise.**

There would seem to be plenty of scope to carefully design new algorithms which take into account the real, rather than assumed, nature of scanner noise.

6 Conclusions

In this paper, we have investigated the noise characteristics in measurement data obtained by a Konica Minolta Vivid 910 scanner. Visual demonstrations and a quasi-statistical analysis of the estimated noise has shown that real scanner noise is neither Gaussian, nor independently distributed. Fourier analysis has further demonstrated that such noise is not white noise. Based on the Fourier analysis, we have proposed a noise synthesis method using the inverse discrete Fourier transform on a fitted model-generated power spectrum and randomly generated phase spectrum. Visually, we have shown such synthetic noise has broadly the same structural features as real scanner noise.

We have also discussed the effectiveness of two typical mesh denoising algorithms on removing noise from real measurement data and synthetic data with Gaussian white noise. Our experiments have shown that it is more difficult to remove noise from real measurement data than from synthetic data with Gaussian white noise. New denoising algorithms are required which take into account the real nature of scanner noise; we intend to investigate such methods. Robust estimation algorithms such as robust moving least-squares fitting [7] may be of use here.

Acknowledgements

We are grateful to Dr. Shenglan Liu for helpful discussions and preparation of the measurement data. This work was supported by EPSRC Grant EP/C007972 and NSFC Grant 60674030.

References

- [1] M. Alexa. Wiener filtering of meshes. In *Proc. Shape Modeling International*, pages 51–57, Washington, DC, USA, 2002. IEEE Computer Society.




Figure 11: Denoising results using original Laplacian algorithm (LA) and a feature-preserving algorithm (FPA) on real scanner data (RSD) and data with Gaussian white noise (GWD). Rows from top to bottom show the original data and the denoising results after 1, 5, 10, 20, and 50 iterations.

- [2] C. Boehnen and P. Flynn. Accuracy of 3d scanning technologies in a face scanning scenario. In *3DIM '05: Proceedings of the Fifth International Conference on 3-D Digital Imaging and Modeling*, pages 310–317, Washington, DC, USA, 2005. IEEE Computer Society.
- [3] C.-Y. Chen and H.-Y. Cheng. A sharpness dependent filter for mesh smoothing. *Computer Aided Geometric Design*, 22(5):376–391, 2005.
- [4] U. Clarenz, U. Diewald, and M. Rumpf. Anisotropic geometric diffusion in surface processing. In *Proc. Conf. Visualization*, pages 397–405. IEEE Computer Society, 2000.
- [5] M. Desbrun, M. Meyer, P. Schröder, and A. H. Barr. Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proc. 26th Conf. Computer Graphics and Interactive Techniques*, pages 317–324, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co.
- [6] J. R. Diebel, S. Thrun, and M. Brünig. A Bayesian method for probable surface reconstruction and decimation. *ACM Trans. Graphics*, 25(1):39–59, 2006.
- [7] S. Fleishman, D. Cohen-Or, and C. T. Silva. Robust moving least-squares fitting with sharp features. *ACM Trans. Graph.*, 24(3):544–552, 2005.
- [8] S. Fleishman, I. Drori, and D. Cohen-Or. Bilateral mesh denoising. *ACM Trans. Graphics*, 22(3):950–953, 2003.
- [9] K. Hildebrandt and K. Polthier. Anisotropic filtering of non-linear surface features. *Computer Graphics Forum*, 23(3):391–400, 2004.
- [10] Y. Hwang, J.-S. Kim, and I.-S. Kweon. Sensor noise modeling using the skellam distribution: Application to the color edge detection. In *CVPR*, pages 1–8, 2007.
- [11] ISO. *Surface Roughness - Terminology - Part 1: Surface and its Parameters*, International Standard ISO 4287/1, 1984.
- [12] F. James. *Statistical Methods in Experimental Physics*. World Scientific Publishing Ltd, New Jersey, 2nd edition, 2006.
- [13] T. R. Jones, F. Durand, and M. Desbrun. Non-iterative, feature-

- preserving mesh smoothing. *ACM Trans. Graphics*, 22(3):943–949, 2003.
- [14] Y. Ohtake, A. Belyaev, and H.-P. Seidel. Mesh smoothing by adaptive and anisotropic Gaussian filter applied to mesh normals. In *Vision, Modeling, and Visualization*, pages 203–210, 2002.
- [15] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, Cambridge, U.K., 2nd edition, 1992.
- [16] Rubert & Co Ltd. *Roughness Parameters*. <http://www.rubert.co.uk/Ra.htm>.
- [17] O. Schall, A. Belyaev, and H.-P. Seidel. Feature-preserving non-local denoising of static and time-varying range data. In *Proc. ACM Symp. Solid and Physical Modeling*, pages 217–222, New York, NY, USA, 2007. ACM Press.
- [18] J. Shen. Synthesization of edge noises for touch probe and laser sensor measurement. *Computer-Aided Design & Applications*, 4(1–4):247–256, 2007.
- [19] J. Shen, B. Maxim, and K. Akingbehin. Accurate correction of surface noises of polygonal meshes. *Int. J. Numerical Methods in Engineering*, 64(12):1678–1698, 2005.
- [20] X. Sun, P. L. Rosin, R. R. Martin, and F. C. Langbein. Fast and effective feature-preserving mesh denoising. *IEEE Trans. Visualization and Computer Graphics*, 13(5):925–938, 2007.
- [21] X. Sun, P. L. Rosin, R. R. Martin, and F. C. Langbein. Random walks for mesh denoising. In *Proc. ACM Symp. Solid and Physical Modeling*, pages 11–22, New York, NY, USA, 2007. ACM Press.
- [22] G. Taubin. A signal processing approach to fair surface design. In *Proc. SIGGRAPH'95*, pages 351–358, 1995.
- [23] G. Taubin. Linear anisotropic mesh filtering. IBM Research Report RC22213(W0110-051), IBM T.J. Watson Research Center, October 2001.
- [24] J. Vollmer, R. Mencl, and H. Müller. Improved Laplacian smoothing of noisy surface meshes. *Computer Graphics Forum*, 18(3):131–138, 1999.
- [25] H. Yagou, Y. Ohtake, and A. G. Belyaev. Mesh smoothing via mean and median filtering applied to face normals. In *Proc. Geometric Modeling and Processing*, pages 124–131, 2002.
- [26] S. Yoshizawa, A. Belyaev, and H.-P. Seidel. Smoothing by example: Mesh denoising by averaging with similarity-based weights. In *Proc. Shape Modeling and Applications*, page 9, Washington, DC, USA, 2006. IEEE Computer Society.
- [27] Y. Yu, K. Zhou, D. Xu, X. Shi, H. Bao, B. Guo, and H. Shum. Mesh editing with Poisson-based gradient field manipulation. *ACM Trans. Graphics*, 23(3):644–651, 2004.