Online Research @ Cardiff

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/5279/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Sun, Xianfang ORCID: <https://orcid.org/0000-0002-6114-0766>, Rosin, Paul L. ORCID: <https://orcid.org/0000-0002-4965-3884>, Martin, Ralph Robert and Langbein, Frank Curd ORCID: <https://orcid.org/0000-0002-3379-0323> 2009. Noise analysis and synthesis for 3D laser depth scanners. Graphical Models 71 (2) , pp. 34-48. 10.1016/j.gmod.2008.12.002 file

Publishers page: <http://dx.doi.org/10.1016/j.gmod.2008.12.002>
<<http://dx.doi.org/10.1016/j.gmod.2008.12.002>>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies.

See

<http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.

Noise Analysis and Synthesis for 3D Laser Depth Scanners

Xianfang Sun ^{a,b,*}, Paul L. Rosin ^a, Ralph R. Martin ^a,
Frank C. Langbein ^a

^a*School of Computer Science, Cardiff University, UK*

^b*School of Automation Science and Electrical Engineering, Beihang University, China*

Abstract

This paper analyses the noise present in range data measured by a Konica Minolta Vivid 910 scanner, in order to better characterise real scanner noise. Methods for denoising 3D mesh data have often assumed the noise to be Gaussian, and independently distributed at each mesh point. We show via measurements of an accurately machined almost planar test surface that real scanner data does not have such properties: the errors are not quite Gaussian, and more importantly, exhibit significant short range correlation. We use this to give a simple model for generating noise with similar characteristics. We also consider how noise varies with such factors as laser intensity, orientation of the surface, and distance from the scanner. Finally, we evaluate the performance of three typical mesh denoising algorithms using real and synthetic test data, and suggest that new denoising algorithms are required for effective removal of real noise.

Key words: 3D laser scanner, scanner noise analysis, noise modeling, 3D surface denoising

1 Introduction

Surface mesh models built using data obtained from 3D laser depth scanners necessarily contain some noise. To remove such noise from surface mesh mod-

* Corresponding author. Tel.: +44 (0) 29 2087 9355; fax: +44 (0) 29 2087 4598.

Email addresses: Xianfang.Sun@cs.cardiff.ac.uk (Xianfang Sun),
Paul.Rosin@cs.cardiff.ac.uk (Paul L. Rosin),
Ralph.Martin@cs.cardiff.ac.uk (Ralph R. Martin),
F.C.Langbein@cs.cardiff.ac.uk (Frank C. Langbein).

els, many mesh denoising algorithms have been developed—see [1–19] and other references therein. In evaluating the effectiveness of such denoising algorithms, both visual and numerical comparisons are used [12]. For meshes corresponding to real scanned data, however, in most cases only visual comparisons are performed, because exact ground truth data are almost always unavailable. However, more objective evaluation requires numerical comparisons, and this is generally done using synthetic data—known meshes to which noise has been added, following a simple model. Better characterisation of real scanner noise would allow more realistic artificial test data to be produced for such purposes.

Many authors have assumed, either when formulating mesh denoising algorithms [1,5], or testing them [2,8,12,18,19], that the measurement noise is Gaussian white noise (independently distributed per measurement point). However, real 3D laser scanner noise is, as we will show later, not quite Gaussian, and more significantly, is strongly correlated at adjacent mesh points. Shen [20] has previously observed that real measurement noise on *edges* does not exactly follow a pattern expected based on a random number generator, and gave a method for generating synthetic edge noise. More recently, Van Gestel et al [21] evaluated the performance of laser line scanners mounted on coordinate-measuring machines (CMMs). They mainly analysed the influence of various factors on random and systematic errors of the scanner; they tested a different scanner arrangement to the one we use, which is fixed in position rather than being mounted on a CMM. In this paper, we focus on characterising scanner noise from the point of view of denoising.

One approach to characterising 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 [22]. 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 potentially include optical components such as the lens, the mirror (and its actuators), and laser, as well as electronic noise. Furthermore, numerical errors may arise in the proprietary software used to calculate positions. Ultimately, we do not have access to all the components of a commercial system, and even if we did, the noise due to some components may not be well understood. Thus, a component-wise noise modelling approach is not practical.

Instead, we analyse overall noise in the measurements provided by a 3D range scanner, taken over a *flat area*. The main 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**. Algorithm evaluation and development should be based on more realistic assumptions about 3D scanner noise. We also discuss the different results obtained when

applying various denoising algorithms to real scanner data and synthetic data with added Gaussian noise. We furthermore give a method to synthesise noise with similar characteristics to that observed.

This paper is an extended version of a paper presented at the SMI'08 conference [23]. Here, we additionally analyse the variation of noise with respect to ambient lighting conditions, scanner laser intensity, surface distance from the scanner, and orientation of the surface. We also present a linear correlation analysis to add further evidence of autocorrelation in the noise. Denoising experiments using synthetic noise generated by our model are also presented, and a further denoising algorithm is tested, making the experimental results more representative.

The remainder of this paper is organised as follows. Section 2 discusses acquisition of data for our experiments, considering choice of test specimens, data preprocessing procedures, and the shape of the underlying test surface. Section 3 analyses noise in the test data using a quasi-statistical approach; we also employ the discrete Fourier transform to characterise the scanner noise. Section 4 describes synthesis of artificial noise based on this analysis, using the inverse discrete Fourier transform. Section 5 demonstrates the difficulties of removing real scanner noise, as opposed to synthetic Gaussian white noise, for three typical mesh denoising algorithms. Section 6 concludes the paper.

2 Data Acquisition and Preprocessing

2.1 Data Acquisition

We now describe the scanner used, the test piece, and our methods of acquiring data.

The scanner used in our tests was a Konica Minolta Vivid 910 with ‘TELE’ lens (focal distance = 25mm), which has a specified accuracy of $100\mu\text{m}$ in the z (depth) direction in ‘FINE’ mode. Boehnen and Flynn [24] evaluated several different 3D scanners, and found that Konica Minolta Vivid 910 was the most accurate for use in a face scanning scenario. Here we take it as a representative commercial scanner using optical triangulation of a laser stripe to capture range data.

Clearly, the higher the precision of the test specimen, the more certain we can be that any measured noise is due to the scanning process and not surface roughness of the specimen. We first considered use of a standard gauge block as the specimen. While very flat, it was too polished and almost mirror-like,

and the scanner was unable to acquire satisfactory measurements. Instead, a *slightly* rough surface was 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 as our main test surface the N1 specimen from a Microsurf 315 set of test surfaces produced by Rubert & Co Ltd. Such test plates are produced for the purpose of testing surface roughness measuring machines. This specimen is a $22.5\text{mm} \times 15\text{mm}$ rectangular flat metal plate with mean roughness $R_a = 0.025\mu\text{m}$ and mean roughness depth $R_z = 0.29\mu\text{m}$ (see later). To give confidence that the results obtained were not specific to this particular test specimen, we also tested the N8 specimen from the same Microsurf 315 set, which is much rougher than N1 with $R_a = 3.2\mu\text{m}$ and $R_z = 15.6\mu\text{m}$. Analysis results were similar to those from the N1 specimen, so we do not report details, but simply show its noise surface in Fig. 3. This visually demonstrates that the noise from this specimen has similar characteristics, although the correlation length appears shorter. While it might have been useful to also consider curved test pieces, we did not do so because of the difficulty in obtaining one machined to a known roughness. However, we measured the flat plate at different orientations, which gives some insight into what might be expected when scanning a curved surface.

R_a and R_z are parameters of surface roughness as defined by ISO Standard 4287/1:1984 [25]. 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 [26]. The sampling length for these test pieces was 0.8mm . Note that for the N1 specimen, both R_a and R_z are significantly smaller than the depth accuracy of $100\mu\text{m}$ claimed for the Konica Minolta Vivid 910 scanner. Henceforth, we 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.

We collected measurement data as follows. During 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), and the centre of the test plate approximately on the optical axis. (Although we did not use an instrument to align the optical axis precisely, 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—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 aligned with scan lines and the optical axis, respectively, and the x, y, z axes form a right-hand coordinate system. The $\{x_i, y_i\}$ coordinates lie on a grid, and z_i measures distance of a surface point i to the focal plane of 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 standard deviation of the noise varied little between scans (see Section 3.1.1), and visually, the scans had similar raised and sunken areas as are apparent in Fig. 1. We also tried scanning under different ambient lighting conditions, and again the results were similar. This repeatability justifies the use of a single set of test data in our analysis, except as specified later.

Secondly, to analyse the effect of surface orientation on noise, we rotated the test plate through varying angles about the x -axis before scanning it. The results showed that varying orientation had little effect—only minor changes were found in standard deviation of noise and correlation, and obvious raised and sunken areas still existed. We thus only present our main analysis for data measured with the test plate approximately orthogonal to the optical axis.

Thirdly, to assess the effect of the scan line direction on the noise, given that our sample was anisotropic, 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 allowed us to verify that the anisotropy of the test surface did not have a significant effect on the noise.

2.2 Data Preprocessing

After capturing the raw measurement data we preprocessed the data to separate the noise from the background surface, using the following steps:

- S1:** Trim off measured data 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 [20]. Our primary trimmed data consists in each case of 9375 gridded points with 125 columns in the x -direction and 75 rows in the y -direction, with intervals between successive points in the x - and y -directions of 0.1735mm and 0.1733mm, respectively.
- S2:** Fit a smooth surface around each measured surface point i in 3D using the measurement data in its neighbourhood N_i , $\{(x_j, y_j, z_j) : j \in N_i\}$. The

fitted surface is represented as

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

where $\mathbf{p}$ is a parameter vector defining the surface. Choice of surface model and neighbourhood size are discussed below.

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

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

The second preprocessing step requires choice of a surface model and size of surface fitting neighbourhood. We considered three possible 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 the data in local regions, taking into account how the ISO standard characterizes noise.

Under Assumption A1, we fit a plane to the whole set of measurement data. Function f in Eq. (1) is then becomes

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

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

Fig. 1 shows the estimated measured noise using Eq. (2) together with Assumption A1. To help visualise the noise surface, we have magnified 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: shading effects have been exaggerated while the colouring is unaltered. A similar technique has also been used later for other Figures. Fig. 1(b) simply shows the *sign* of the noise, black being negative. Fig. 1 shows that the noise surface based on this assumption has obvious raised and sunken areas, and is clearly curved—its top-right and bottom-left corners are higher than other parts of the surface.

Firstly, we note that bumps of the magnitude indicated by the noise surface would easily be seen by anyone visually examining the test plate, 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). Thus, the local structure visible in Figure 1 must be due to measurement errors, and not surface bumps in the test plate. This deduction is backed up by noting that the measured bumps are much greater than the manufacturer's claimed roughness of the test plate.

Fig. 1. Extracted noise for N1 sample 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.

Secondly, the apparent curvature of the noise surface may be due to actual curvature of the test surface (which is only guaranteed to be *locally* flat), or due to systematic error over a long scale produced by the scanner.

To resolve this 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 under Assumption A1. It can be seen that while the detailed noise pattern is somewhat different, the overall curvature in the noise surface has also rotated with it—its top-left and bottom-right corners are now higher. Thus, the curvature of the noise surface comes from the specimen itself, and is not due to systematic errors produced by the scanner—in the latter case, the noise distribution would not have rotated. We also conclude that Assumption A1 is unsatisfactory, and a better surface model is needed.

Fig. 2. Extracted noise for N1 sample given Assumption A1 (specimen rotated 90° relative to Fig. 1 before scanning), else as per Figure 1.

A further observation is 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 and the anisotropy of the test surface. Apart from this question of directionality, however, the general *structure* of the noise appears similar in both orientations. Thus, subsequent experiments were restricted to data measured with the long side of the specimen aligned with the x -axis.

Fig. 3 shows the noise for the N8 test piece under Assumption A1. The N8 specimen is clearly curved in a different way to the N1 specimen. Again, there are obvious bumps on the noise surface, with a similar general pattern as for the N1 noise surface, although the characteristic size and spacing of bumps is smaller than for the N1 sample. We thus concentrated solely on the N1 specimen in our further analysis.

Next, we considered Assumption A2 as a potentially better model for the shape of the N1 test piece, and fitted a quadratic surface to the whole set of measurement data. In this case the function f takes the form

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

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

(a)

(b)

Fig. 3. Extracted noise for N8 sample given Assumption A1, as per Figure 1.

Fig. 4 shows the estimated noise using Assumption A2. Again, the noise surface again has obvious large scale raised and lowered areas with a similar distribution to that seen in Fig. 1, except that the overall curvature in Fig. 1 is no longer present. Using the same reasoning as before, we again claim that the magnitude of the noise is due to the scanner.

We still needed to decide whether a quadratic surface could sufficiently represent the underlying surface of the N1 specimen. To answer this question, we also fitted higher order surfaces, e.g. cubic surfaces, to the measurement data, and extracted the noise using Eq. (2). Visually, the distribution pattern of the extracted noise in each case is similar to that shown in Fig. 4. Furthermore, the 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

Fig. 4. Extracted noise for N1 sample given Assumption A2, as per Figure 1.

noise surface with this assumption is close to the real noise due to the scanner.

We 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. As the roughness of the specimen (measured by R_a and R_z) is significantly smaller than the scanner accuracy, the true surface can be effectively considered to be planar over regions with diameter smaller than the sampling length used to define R_a and R_z .

The sampling length is 0.8mm. Thus, for each measured point on the surface, we fitted a local plane to the measured data within a circular neighbourhood of diameter 0.8mm. The measured 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. However, because the surface fitted to the noisy measured points has a roughness exceeding the range designated by R_a and R_z , it is not a good approximation to the true surface. Thus, we

Fig. 5. Extracted noise for N1 sample given Assumption A3, as per Figure 1.

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 (this required 323 iterations). The final approximate points can be considered to be good approximations to the real surface points: the difference between each measured point and the corresponding approximate point then gives an estimate of the noise. Note that because of the iteration process, the surface may exhibit a global shift—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. 5 shows the extracted noise surface using the above method. The noise pattern is clearly very similar to that estimated using Assumption A2, shown in Fig. 4, although 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, and that this noise

Fig. 6. Differences between Fig. 4 and Fig. 5, in mm.

really does contain the structures visible in Figs. 4 and 5.

We also determined the standard deviation of the differences between noise values estimated using A2 and A3 to be $4.9\mu\text{m}$, which is only 10% of the specified accuracy of the scanner. While these differences are not negligible, the difference surface is slowly varying over a long distance with low amplitude (see Fig. 6). 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 simpler computation, so in the rest of the paper, our analysis is based on Assumption A2.

We now make perhaps the most important point in the whole paper, which is clear from a visual inspection of Figures 4 and 5: **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 mesh denoising, which often assume scanner noise to be independent Gaussian noise at each measurement point [1,2,5,8,18,19].

3 Noise Analysis

We now analyse the properties of the estimated noise. We carry out both a statistical analysis, and a Fourier analysis of the noise. Because we only have an estimate of the scanner noise (based on assumption A2), we call our analysis a *quasi-statistical* analysis of the real scanner noise.

Fig. 7. Normal-optical axis angle θ

3.1 Quasi-Statistical Analysis

3.1.1 Noise Variation under Different Scanning Conditions

We start our quasi-statistical analysis with a discussion of how the noise varies with different factors, including ambient light, distance of specimen to the camera, scanner laser intensity, and orientation of the specimen (defined by θ , the normal-optical axis angle: see Fig. 7). We assess this by computing the standard deviation of the noise at all sample points, as these various factors change. Many other environmental factors, such as temperature, may also influence the level of noise, but because the user typically does not have control over such factors, we do not analyse variation with respect to them in this paper.

Firstly, we tested whether typical changes in ambient light have any significant effect on the standard deviation of noise. (Again, we did not consider subtle changes under controlled lighting, but rather the gross effects of large changes in lighting.) To do so, we aligned the sample with 24 different choices of θ , and for each we scanned the sample first 5 times with the fluorescent ceiling lamp on (giving normal ‘office’ lighting conditions), and then 5 times with it off (giving a more-or-less darkened room). We calculated the standard deviation in noise for each such set of scanned data, denoted by $\sigma_{i,j}$, where i corresponds to a different orientation and j numbers the datasets at a given orientation. For each orientation i , $i = 1, \dots, 24$, we calculated the mean μ'_i and standard deviation σ'_i of the above standard deviations themselves $\sigma_{i,j}$, $j = 1, \dots, 10$, and then calculated the relative variation: σ'_i/μ'_i . The result shows that the relative variation for $i = 1, \dots, 24$ lies between 0.46% and 2.05%. Thus, for a given orientation, the variation in level of noise is very small, irrespective of the ambient light conditions. This result also implies that the magnitude of

Fig. 8. Standard deviation of noise: variation with surface distance

Fig. 9. Standard deviation of noise: variation with scanner laser intensity

scanner noise has a high repeatability for a given set of conditions.

Secondly, we adjusted the distance between the scanner and the test piece, to assess its influence on the standard deviation of noise; we attempted to keep the the normal-optical axis angle almost constant at zero degrees. (Because we did not have a precision instrument to fix the angle, we could only keep the angle approximately unchanged and check it was satisfactory after scanning). Fig. 8 shows the variation of standard deviation of noise with respect to distance: the standard deviation increases with distance.

Thirdly, we considered the influence of scanner laser intensity on the standard deviation of noise, for the sample in a fixed position. See Fig. 9. Provided that the intensity is kept within a suitable range, the variation in noise level is very

Fig. 10. Standard deviation of noise: variation with normal-optical angle. Dashed line: original values. Solid line: after removal of outliers.

small—the relative variation over the intensity range from 45 to 85 (using the manufacturer’s units) is only 5.2%. However, if the intensity is inappropriately chosen, the standard deviation can become much larger. In this example, when the intensity is set to 175, the standard deviation reaches $345\mu\text{m}$. Looking in detail at the data captured under these conditions, we found it to contain many outliers, and such data would not be useful for any practical application. This emphasises the need to correctly set up the laser scanner to obtain satisfactory results, following the manufacturer’s procedures. Note that correct choice of intensity depends on the color and reflectance of the surface being measured.

Fourthly, we considered the influence of the normal-optical axis angle θ on the standard deviation of noise. See Fig. 10. The dashed line shows the noise estimated with assumption A2. Unfortunately, it was not possible to pick a single laser intensity which would both capture a full set of data, and avoid outliers, over this complete range of angles. We thus processed the raw scans to remove the outliers, and the solid line in Fig. 10 shows the standard deviation of noise after outliers in the data have been discarded. To remove outliers, we iteratively removed those points with noise greater than three times the standard deviation of the whole data set. After doing so, the variation of standard deviation of noise with θ becomes much smaller. In particular for angles of less than -15 degrees, which has fewer outliers, the noise varies little with angle (the relative variation is 8%). Overall, there is a trend to increasing noise with increasing normal-optical axis angle, although the variation is not smooth.

As a final comment, we note that the measurements above show that the scanner does indeed meet the manufacturer’s claimed accuracy of $100\mu\text{m}$.

Fig. 11. Histogram of estimated noise magnitude. Red line: best fit Gaussian distribution.

3.1.2 Is the Noise Gaussian?

We next discuss the distribution of the estimated noise for a particular scan. Fig. 11 shows a histogram of the magnitude of the estimated noise taken at all points of the test surface, 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 with the Gaussian distribution, but to decide this issue more definitely, a statistical test is necessary.

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 a Gaussian function of appropriate standard deviation. The standard deviation of the estimated noise is $\sigma = 16.2\mu\text{m}$. For this we use Pearson's chi-square test [27]. 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 (measurement points: $N = 9375$), m is the number of bins in the histogram (discussed 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 close enough to $\chi^2(m-1)$ provided that we ensure that the expected numbers of events per bin, Np_i , is greater than five in each case [27]. If the probability $P(\chi^2(m-1) > K) < \alpha$, the null hypothesis is said to be rejected at the significance level α . Commonly used α values are 5%, 1%, and 0.1%. The smaller the α , the stronger the conclusion. We use the $\alpha = 5\%$ significance level.

We generated histograms using $m_0 = 3, \dots, 100$ equal-width bins, except that we combined neighbouring bins with small expected numbers of events into larger bins so that each bin was expected to have more than five events. m was the number of bins after merging small bins.

We then used Eq. (5) to compute K , and computed the probability that $\chi^2(m-1) > K$. 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 $\alpha = 5\%$ level of significance. In simple words, **we should reject the hypothesis that the estimated noise distribution is Gaussian**. The test does not tell us what the true distribution is. Nevertheless, Fig. 11 makes it clear that the distribution of the estimated noise is quite Gaussian-like.

We must caution that, as well as the direct 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.

3.1.3 Noise Correlation

We now consider noise autocorrelation in the x - and y -directions. 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, and $e(i, j)$ to denote the estimated noise at this point. $N_x = 125$ and $N_y = 75$ for our data.

We first discuss the linear autocorrelation. Linear autocorrelation coefficients are defined by

$$\rho_x(k) = \frac{\sum_{j=1}^{N_y} \sum_{i=1}^{N_x-k} e(i, j)e(i+k, j)}{\sqrt{\sum_{j=1}^{N_y} \sum_{i=1}^{N_x-k} e^2(i, j)} \sqrt{\sum_{j=1}^{N_y} \sum_{i=1}^{N_x-k} e^2(i+k, j)}}, \quad (6)$$

$$\rho_y(k) = \frac{\sum_{i=1}^{N_x} \sum_{j=1}^{N_y-k} e(i, j)e(i, j+k)}{\sqrt{\sum_{i=1}^{N_x} \sum_{j=1}^{N_y-k} e^2(i, j)} \sqrt{\sum_{i=1}^{N_x} \sum_{j=1}^{N_y-k} e^2(i, j+k)}}, \quad (7)$$

Having obtained the autocorrelation coefficients $\{\rho_x(k), k = 1, \dots, N_x - 1\}$ and $\{\rho_y(k), k = 1, \dots, N_y - 1\}$, we consider the null hypothesis that the noise at each measurement location is uncorrelated, corresponding to

$$H_0 : \rho_x(i) = 0, \quad \{i = 1, \dots, N_x - 1\},$$

and similarly for y . The alternative hypothesis is that the noise at each measurement location is correlated, corresponding to

$$H_1 : \rho_x(i) \neq 0,$$

and similarly for y .

If the noise distribution were Gaussian, we could use the t -test to assess the correlation [28]. However, as we have shown above, the noise distribution is not strictly Gaussian, and we can only conservatively borrow the conclusions from performing a t -test. Nevertheless, the noise distribution is quite close to Gaussian—at least visually—from the histogram shown in Fig. 11, which suggests that the t -test results will be useful.

Given a correlation coefficient ρ , (ρ could be either $\rho_x(k)$ or $\rho_y(k)$), the statistic

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

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

Fig. 12 shows how the probability $P(t(M-2) > |T|)$ varies with k ; the green line is set 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 46 in the y -direction, respectively; other correlation lengths greater than these values also have $P(t(M-2) > |T|) < 5\%$. Thus we conclude that the noise is autocorrelated in both x - and y -directions, at various length scales.

Because the above hypothesis test is based on an assumption of Gaussian statistics, which the estimated noise does not exactly exhibit, it is more appropriate to use a non-parametric test. We used Spearman rank-order correlation

Fig. 12. Linear autocorrelation showing probability $P(t(M-2) > |T|)$ in x - and y -directions.

Fig. 13. Spearman autocorrelation showing probability $P(t(M-2) > |T|)$ in x - and y -directions.

coefficients, which do not require knowledge of the probability distribution of the noise [28].

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 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

$$\rho_{Sx}(k) = \frac{\sum_{j=1}^{N_y} \sum_{i=1}^{N_x-k} \tilde{R}_{x1k}(i, j) \tilde{R}_{x2k}(i + k, j)}{\sqrt{\sum_{j=1}^{N_y} \sum_{i=1}^{N_x-k} \tilde{R}_{x1k}^2(i, j)} \sqrt{\sum_{j=1}^{N_y} \sum_{i=1}^{N_x-k} \tilde{R}_{x2k}^2(i + k, j)}}. \quad (9)$$

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

$$\rho_{Sy}(k) = \frac{\sum_{i=1}^{N_x} \sum_{j=1}^{N_y-k} \tilde{R}_{y1k}(i, j) \tilde{R}_{y2k}(i, j + k)}{\sqrt{\sum_{i=1}^{N_x} \sum_{j=1}^{N_y-k} \tilde{R}_{y1k}^2(i, j)} \sqrt{\sum_{i=1}^{N_x} \sum_{j=1}^{N_y-k} \tilde{R}_{y2k}^2(i, j + k)}}. \quad (10)$$

Given the rank-order correlation coefficients, we can again use the statistic T defined in (8) (ρ now is either $\rho_{Sx}(k)$ or $\rho_{Sy}(k)$), to evaluate the autocorrelation of the estimated noise. Fig. 13 shows how the probability $P(t(M - 2) > |T|)$ varies with k . From the figure it can be seen that this test yields very similar results to those obtained with the linear autocorrelation coefficients: the probability that $P(t(M - 2) > |T|)$ in this case is less than 5% for all correlation lengths less than 11 in the x -direction and 47 in the y -direction, respectively. We can rely on the robustness of the non-parametric test [28] to conclude from the above results that **the estimated noise shows significant autocorrelation**.

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 analysis. In this section, we perform an alternative Fourier analysis of the *measurements themselves*, which does not require any assumption concerning specimen shape.

We first performed 1D Fourier analysis on the *measurement data* for a particular scan (not the extracted *noise*). Fig. 14 shows a log-log scale plot of the power spectrum for the x - and y -directions respectively. 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 from the surface signal itself of the specimen. Considering that the specimen surface is quite smooth and its fluctuations are quite small, such

Fig. 14. Power spectrum along x - (top) and y - (bottom) directions; the green vertical lines are at a frequency $1/0.8\text{mm}$.

a 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. 14 is at the frequency $f = 1/0.8\text{mm}$, corresponding to the roughness sampling length. Above this frequency, surface shape effects are negligible, and we can assume the surface is flat. Thus, the spectral power for frequencies $f > 1/0.8\text{mm}$ (to the right of each green line) is certainly due to scanner noise. (It may also be mainly due to scanner noise below this cut-off, but we cannot guarantee this). Since the power spectra are not constant above this cut-off, the measurement noise is not *white noise*.

To consider the noise as a 2D process, we also performed a 2D Fourier analysis of the measurement data. Fig. 15 shows its 2D discrete Fourier transform. As for the 1D Fourier transforms, the magnitude of the 2D Fourier transform decreases with frequency, varying less at high frequencies. 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 again tells us that the measurement data are correlated, and furthermore, the measurement noise is not white noise.

4 Noise Synthesis

We now know that the noise is not white and its distribution is not quite Gaussian, so we cannot generate realistic synthetic noise using straightforward

Fig. 15. 2D discrete Fourier transform (zero-frequency at the centre). Top: logarithmic magnitude of the spectrum. Bottom: phase of the spectrum, in degrees.

pseudo-random noise sequences. However, using the Fourier analysis results from Section 3.2, we may generate synthetic noise by using the inverse Fourier transforms.

The Fourier transform of the measured data represents both the noise, and an underlying signal (due to non-planarity of the test specimen). Because the *signal* is potentially significant mainly at low frequencies, *noise* represents most of the amplitude of the transform at high frequencies, so we may simply set the lowest frequency components of the Fourier transform to zero, and use the remaining components to approximately model the noise. While a sharp cut-off filter can introduce ringing artifacts, these were negligible in practice.

The top of Fig. 16 shows the inverse Fourier transform of the original scan data with the lowest 5 frequency 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 seen in Figs. 1, 4 and 5, while removing the obvious curvature of the specimen surface visible in Fig. 1.

Having obtained the Fourier transform of the measurement data, we can use

Fig. 16. Generating noise by inverse Fourier transform with low frequency components zeroed. Top: original magnitude and phase except for low frequencies; Middle: synthetic noise with fitted magnitude and original phase except for low frequencies; Bottom: synthetic noise with fitted magnitude and random phase except for low frequencies.

a simple model to fit the magnitudes with respect to frequency. We tried several models, and selected that particular model which gave minimum fitting residual among those models. The model we use is

$$Z(i, j)^{\frac{1}{4}} = \begin{cases} a_0 + b_0i + c_0j + d_0i^2 + e_0ij + f_0j^2 & i \neq 0, j \neq 0, \\ a_1 + b_1i + d_1i^2 & j = 0, \\ a_2 + b_2j + f_2j^2 & i = 0, \end{cases} \quad (11)$$

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 giving a best fit are $a_0 = 1.5600$, $b_0 = -0.0185$, $c_0 = -0.0176$, $d_0 = 0.0001$, $e_0 = 0.0003$, $f_0 = 0.0000$, $a_1 = 1.9134$, $b_1 = -0.0417$, $d_1 = 0.0005$, $a_2 = 1.8352$, $b_2 = -0.0530$, $f_2 = 0.0008$; with these values, the squared sum of the residual errors is 2.2% of the squared sum of the noise magnitudes. Fig. 16, middle, shows the result of computing an inverse Fourier transform using the modelled magnitude together with the original phase data. This generated noise has a structure quite close to the original structure, verifying the utility of this model.

Because the phase spectrum is rather unstructured—see Fig. 15—we cannot easily fit a model to it, but we do not wish to store it as a table of values, either. Instead, we suggest the use of random phases to generate synthetic noise. Fig. 16, bottom, shows noise synthesised by inverse Fourier transform of the modelled amplitudes and random phases. The synthetic noise in this case has a slightly different structure to the original one, which is to be expected as there is *some* structure in the phase of the noise—again see Fig. 15. Nevertheless, the structure of this synthetic noise is quite similar to that of 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 acceptable. It is certainly closer to real scanner noise than white Gaussian noise.

5 Denoising experiments

As demonstrated, scanner noise is neither strictly Gaussian, nor independent at each location or white in nature, contrary to the most common assumptions adopted by many papers considering denoising or other processing of scanner data. We compare denoising results by various algorithms on three sets of data: real scanner data, data generated by adding synthetic noise described in Section 4 to the best quadratic surface fitted to the scanning data, and

Fig. 17. Denoising using Laplacian smoothing [14]: real scanner data (left), data with synthetic noise (middle), and data with Gaussian white noise (right). Top: original data. Other rows: denoising results after 5, 20, and 50 iterations.

data obtained by adding Gaussian white noise with the same standard deviation as that of the scanned data, again to the best fitted quadratic surface. Gaussian noise has often been used in the past to assess denoising algorithms. Three representative denoising algorithms were chosen: the original Laplacian smoothing algorithm [14], the bilateral filtering algorithm [6], and our own feature-preserving denoising algorithm [12].

Fig. 17 shows the denoising results obtained for each kind of noise using the original Laplacian smoothing algorithm [14] with 5, 20, and 50 iterations, as commonly used in the literature. The Figure shows that fewer iterations are needed to remove Gaussian white noise than to remove the synthetic noise generated by our method, and real scanner noise requires yet more iterations. Furthermore, although the original Laplacian algorithm can effectively remove Gaussian noise, it is less successful at completely removing the bumpy structures present in real scanner noise, and to a lesser degree, in the synthetic noise generated by our method.

Figs. 18 and 19 show comparable denoising results for the bilateral filter-

Fig. 18. Denoising using bilateral filtering [6]: real scanner data (left), data with synthetic noise (middle), and data with Gaussian white noise (right). Rows: results after 5, 20, and 50 iterations.

Fig. 19. Denoising using feature-preserving denoising [12]: real scanner data (left), data with synthetic noise (middle), and data with Gaussian white noise (right). Rows: denoising results after 5, 20, and 50 iterations.

ing algorithm [6] and our feature-preserving denoising algorithm [12]. Both algorithms show similar results to those obtained for Laplacian smoothing:

Gaussian noise is easiest to remove, then our synthetic noise, and finally real scanner noise. However, because these two algorithms are designed to preserve features, unlike Laplacian smoothing [14], these algorithms have more difficulty than Laplacian smoothing in removing the structured features (bumps) due to correlations in real scanner noise and our synthetic noise. Compared to our algorithm, bilateral filtering is slightly more efficient in removing such structural features of the noise.

From these experiments, we see that although the synthetic noise generated by our method does not behave exactly in the same way as real noise, it more closely resembles real noise than does Gaussian noise, and it is somewhat better at predicting how algorithms will perform in removing real noise.

We have shown that denoising algorithms, whether feature-preserving or not, cannot effectively deal with the structured nature of real scanner noise. We thus conclude that **many previous papers claiming good denoising results based on experiments with synthetic Gaussian noise are over-optimistic in their assessment of the ability of algorithms to remove real scanner noise**. Synthetic noise of the type generated by our method described in Section 4 would be a better basis for evaluating smoothing algorithms.

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. This does not necessarily require precise modelling of the scanner noise. Instead, by just assuming the noise to be non-Gaussian and correlated, we may be able to make better generic algorithms than if we assume uncorrelated Gaussian noise, even if the optimal filters would involve too many specific parameters dependent on the precise scanning conditions.

6 Conclusions

We have investigated the noise characteristics in measurement data obtained by a Konica Minolta Vivid 910 scanner. We have found that there is little systematic variation in noise resulting from different scanning conditions, providing that care is taken in choice of laser intensity.

Consideration of careful estimates of the noise has shown that real scanner noise is not quite Gaussian, but more importantly, neither is it 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 which applies the inverse discrete Fourier transform to a model-generated power spectrum with random phase. Such synthetic noise

has broadly the same visual structural features as real scanner noise. Experimentally, our synthetic noise behaves in a qualitatively similar way to real scanner noise in denoising algorithms, and thus is more suitable than Gaussian white noise for algorithm evaluation.

Our above analysis is based on scanning an N1 specimen from a Microsurf 315 set of test surfaces. An N8 specimen from the same set was also tested and gave qualitatively similar results, so we do not report the details here. Note, however, that the precise numerical results, and suggested synthetic noise model, only hold for this specific type of surface material and this specific laser scanner. For other types of surface material or other kinds of laser scanners, the results may be quite different.

We have also tested the effectiveness of three typical mesh denoising algorithms against real measurement data, Gaussian synthetic white noise, and our synthetic noise. The results have shown that it is more difficult to remove noise from real measurement data than from data with our synthetic noise, and it is easiest to remove Gaussian white noise.

New denoising algorithms are required which take into account the real nature of scanner noise; we intend to investigate such methods in the future. Improvement on robust estimation algorithms such as robust moving least-squares fitting [29] 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, IEEE Computer Society, Washington, DC, USA, 2002, pp. 51–57.
- [2] C.-Y. Chen, H.-Y. Cheng, A sharpness dependent filter for mesh smoothing, Computer Aided Geometric Design 22 (5) (2005) 376–391.
- [3] U. Clarenz, U. Diewald, M. Rumpf, Anisotropic geometric diffusion in surface processing, in: Proc. Conf. Visualization, IEEE Computer Society, 2000, pp. 397–405.

- [4] M. Desbrun, M. Meyer, P. Schröder, A. H. Barr, Implicit fairing of irregular meshes using diffusion and curvature flow, in: Proc. 26th Conf. Computer Graphics and Interactive Techniques, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 1999, pp. 317–324.
- [5] J. R. Diebel, S. Thrun, M. Brünig, A Bayesian method for probable surface reconstruction and decimation, *ACM Trans. Graphics* 25 (1) (2006) 39–59.
- [6] S. Fleishman, I. Drori, D. Cohen-Or, Bilateral mesh denoising, *ACM Trans. Graphics* 22 (3) (2003) 950–953.
- [7] K. Hildebrandt, K. Polthier, Anisotropic filtering of non-linear surface features, *Computer Graphics Forum* 23 (3) (2004) 391–400.
- [8] T. R. Jones, F. Durand, M. Desbrun, Non-iterative, feature-preserving mesh smoothing, *ACM Trans. Graphics* 22 (3) (2003) 943–949.
- [9] Y. Ohtake, A. Belyaev, H.-P. Seidel, Mesh smoothing by adaptive and anisotropic Gaussian filter applied to mesh normals, in: Vision, Modeling, and Visualization, 2002, pp. 203–210.
- [10] O. Schall, A. Belyaev, H.-P. Seidel, Feature-preserving non-local denoising of static and time-varying range data, in: Proc. ACM Symp. Solid and Physical Modeling, ACM Press, New York, NY, USA, 2007, pp. 217–222.
- [11] J. Shen, B. Maxim, K. Akingbehin, Accurate correction of surface noises of polygonal meshes, *Int. J. Numerical Methods in Engineering* 64 (12) (2005) 1678–1698.
- [12] X. Sun, P. L. Rosin, R. R. Martin, F. C. Langbein, Fast and effective feature-preserving mesh denoising, *IEEE Trans. Visualization and Computer Graphics* 13 (5) (2007) 925–938.
- [13] X. Sun, P. L. Rosin, R. R. Martin, F. C. Langbein, Random walks for feature-preserving mesh denoising, *Computer Aided Geometric Design* 25 (7) (2008) 437–456.
- [14] G. Taubin, A signal processing approach to fair surface design, in: Proc. SIGGRAPH’95, 1995, pp. 351–358.
- [15] G. Taubin, Linear anisotropic mesh filtering, IBM Research Report RC22213(W0110-051), IBM T.J. Watson Research Center (October 2001).
- [16] J. Vollmer, R. Mencl, H. Müller, Improved Laplacian smoothing of noisy surface meshes, *Computer Graphics Forum* 18 (3) (1999) 131–138.
- [17] H. Yagou, Y. Ohtake, A. G. Belyaev, Mesh smoothing via mean and median filtering applied to face normals, in: Proc. Geometric Modeling and Processing, 2002, pp. 124–131.
- [18] S. Yoshizawa, A. Belyaev, H.-P. Seidel, Smoothing by example: Mesh denoising by averaging with similarity-based weights, in: Proc. Shape Modeling and Applications, IEEE Computer Society, Washington, DC, USA, 2006, p. 9.

- [19] Y. Yu, K. Zhou, D. Xu, X. Shi, H. Bao, B. Guo, H. Shum, Mesh editing with Poisson-based gradient field manipulation, *ACM Trans. Graphics* 23 (3) (2004) 644–651.
- [20] J. Shen, Synthesization of edge noises for touch probe and laser sensor measurement, *Computer-Aided Design & Applications* 4 (1–4) (2007) 247–256.
- [21] N. Van Gestel, S. Cuypers, P. Bleys, J.-P. Kruth, A performance evaluation test for laser line scanners on cmms, *Optics and Lasers in Engineering* Doi:10.1016/j.optlaseng.2008.06.001.
- [22] Y. Hwang, J.-S. Kim, I.-S. Kweon, Sensor noise modeling using the skellam distribution: Application to the color edge detection, in: *CVPR*, 2007, pp. 1–8.
- [23] X. Sun, P. L. Rosin, R. R. Martin, F. C. Langbein, Noise in 3D laser range scanner data, in: *Proc. IEEE Int. Conf. on Shape Modeling and Applications*, IEEE, 2008, pp. 37–45.
- [24] C. Boehnen, 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*, IEEE Computer Society, Washington, DC, USA, 2005, pp. 310–317.
- [25] ISO, Surface Roughness - Terminology - Part 1: Surface and its Parameters, International Standard ISO 4287/1 (1984).
- [26] Rubert & Co Ltd, Roughness Parameters, <http://www.rubert.co.uk/Ra.htm>.
- [27] F. James, *Statistical Methods in Experimental Physics*, 2nd Edition, World Scientific Publishing Ltd, New Jersey, 2006.
- [28] W. H. Press, S. A. Teukolsky, W. T. Vetterling, B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*, 2nd Edition, Cambridge University Press, Cambridge, U.K., 1992.
- [29] S. Fleishman, D. Cohen-Or, C. T. Silva, Robust moving least-squares fitting with sharp features, *ACM Trans. Graph.* 24 (3) (2005) 544–552.