Optics and Lasers in Engineering 91 (2017) 169–177

Contents lists available at ScienceDirect

Optics and Lasers in Engineering

journal homepage: www.elsevier.com/locate/optlaseng

Novel metrics and methodology for the characterisation of 3D imaging systems

John R. Hodgson*, Peter Kinnell, Laura Justham, Niels Lohse, Michael R. Jackson

EPSRC Centre for Innovative Manufacturing in Intelligent Automation, Wolfson School of Mechanical Electrical and Manufacturing Engineering, Loughborough University, LE113QZ, United Kingdom

ARTICLE INFO

Keywords:
3D imaging
Scanner
Evaluation
Performance
Surface
Roughness

ABSTRACT

The modelling, benchmarking and selection process for non-contact 3D imaging systems relies on the ability to characterise their performance. Characterisation methods that require optically compliant artefacts such as matt white spheres or planes, fail to reveal the performance limitations of a 3D sensor as would be encountered when measuring a real world object with problematic surface finish. This paper reports a method of evaluating the performance of 3D imaging systems on surfaces of arbitrary isotropic surface finish, position and orientation. The method involves capturing point clouds from a set of samples in a range of surface orientations and distances from the sensor. Point clouds are processed to create a single performance chart per surface finish, which shows both if a point is likely to be recovered, and the expected point noise as a function of surface orientation and distance from the sensor. In this paper, the method is demonstrated by utilising a low cost pan-tilt table and an active stereo 3D camera. Its performance is characterised by the fraction and quality of recovered data points on aluminium isotropic surfaces ranging in roughness average (*Ra)* from 0.09 to 0.46 μm at angles of up to 55° relative to the sensor over a distances from 400 to 800 mm to the scanner. Results from a matt white surface similar to those used in previous characterisation methods contrast drastically with results from even the dullest aluminium sample tested, demonstrating the need to characterise sensors by their limitations, not just best case performance.

1. Introduction

The process of selecting the optimal 3D imaging system for a particular industrial application is a challenging one [1,2]. This is because of the range of variables that have to be considered. Parameters such as acquisition time, acquisition rate, scanning volume, physical size, weight and cost are straightforward to use as selection criteria; they are typically the first things to be constrained by project specifications and budget. What is more challenging to understand is the performance that can be expected from a particular imaging system. The project may require specific performance parameters such as point accuracy, resolution and repeatability, which are often available on manufacturer data sheets. The problem arises that these values are usually best case parameters and do not reflect the real-world performance of a system when utilised in one of the wide array of industrial applications for 3D imaging systems [3–7]. This makes comparisons between competing devices very challenging.

The parameters in data sheets are usually derived from tests on idealised metrological artefacts or are limited to discussions of the theoretical maximum resolution based on the number of pixels in the imaging system. For instance, the VDI/VDE 2634 standard [8] recommends using matt textured spheres, planes and ball-bars to assess a variety of metrological parameters. Such artefacts are completely unrepresentative of objects encountered in most industrial applications in terms of surface finish, and therefore cannot provide accurate predictions of scanner performance. The reason for this is that most modern 3D vision systems are active, and hence rely on the return of projected light from a surface to measure it. The amount of light returned, and hence the signal to noise ratio of the signal and quality of the measurement is determined by the Bi-directional Reflectance Distribution Function (BRDF) [9,10], which depends, amongst other factors, on surface finish.

Whilst the theoretical limits of sensor performance are developed from fundamental laws of physics [11,12], understanding their real-life performance has been an active area of research. Guidi [13] has presented a thorough review of developments in the field of 3D imaging

* Correspondence to: EPSRC Centre for Innovative Manufacturing in Intelligent Automation, Wolfson School of Mechanical Electrical and Manufacturing Engineering, Loughborough University, Holywell Building, Holywell Way, Loughborough LE11 3QZ, United Kingdom.

E-mail addresses: j.r.hodgson@lboro.ac.uk (J.R. Hodgson), p.kinnell@lboro.ac.uk (P. Kinnell), l.justham@lboro.ac.uk (L. Justham), n.lohse@lboro.ac.uk (N. Lohse), m.r.jackson@lboro.ac.uk (M.R. Jackson).

http://dx.doi.org/10.1016/j.optlaseng.2016.11.007
Received 8 August 2016; Received in revised form 31 October 2016; Accepted 8 November 2016
0143-8166/ © 2016 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

system evaluation. The primary focus in literature is on achieving traceable measurements of metrological parameters such as accuracy, precision and repeatability. A few studies have dealt with the issue of surface inclination on performance [14–16], but only with regard to surfaces of optically compliant finish or varying colour. The National Physical Laboratory (NPL) offer a 3D sensor characterisation service which includes the evaluation of scanner performance on a selection of material coupons at different orientations relative to the sensor [17]. NPL also produce a freeform artefact [18] for the evaluation of shape reproduction under different lighting conditions. These services are useful to industry, particularly manufacturers of 3D sensors as a benchmarking service. However, the expense of the freeform artefact limits its use more generally and the limited set of orientations that are possible with a set of coupons inherently limits the evaluation of dimensional sensitivity to surface finish without an excessively large experimental set. Despite the lack of published investigations into characterising the effect of surface finish on general sensor performance, its importance is clearly appreciated, otherwise evaluation methodologies would not recommend the use of vapour blasted, or matt painted surfaces as test artefacts.

A further issue is the limited set of standards for scanner evaluation. Two standards are of particular relevance: VDI/VDE 2634 [8] and ASTM E2919-14 [19]. VDI/VDE 2634 is primarily concerned with determining errors by the measurement of three standard artefacts: a sphere, ball-bar and plane, which should first be vapour blasted to produce optically diffuse surfaces for optimal measurement. ASTM E2919-14 specifies a test method for evaluating systems that measure pose (position and orientation) of a rigid test object. There are no limitations placed on the test object itself, in fact, it recommends using one that is representative of the final application in terms of geometry and material. This is useful for assessing performance, but it is only valid for the test object chosen and as there is no specification for the object, the replication and comparison of results for different systems by third parties is difficult.

In previous work [20], the authors presented a methodology for collecting point cloud data from a sensor for samples of varying surface finish and inclination only. The work is extended here to incorporate samples at varying distances and tolerating small deviations of the sample from the centre of the field of view. The main focus however has been improvements of the data processing techniques and performance metrics to allow straightforward comparison of sensors in real world conditions.

It is envisaged that if a standard methodology for the collection of this information were conceived, it would allow manufacturers to provide their customers with significantly improved levels of information to make scanner selection considerably more straightforward. It would also allow third party organisations to be able to collect comparable performance evaluation data.

Section 2 gives details on the data collection methodology including sample preparation, validation, test apparatus and the calculation of performance metrics from the data. Section 3 details the presentation of results into a format that allows easy comparison of sensor performance on different surfaces.

2. Methodology for 3D imaging system evaluation

This section describes the methodology for evaluating the performance of a 3D imaging system. The process begins with preparing a selection of flat samples with different surface finishes. These samples are then placed on a pan tilt table and point clouds are collected at as many surface orientations and distances from the scanner as practical. Finally, the data is processed to calculate the performance of the scanner. It is important to note that the data processing method is based on point cloud data only. This is to ensure a third party can evaluate any scanner that produces point cloud output.

By using flat samples, the number of measurements that must be taken to rigorously sample the gradient space is large; 1008 measurements taking approximately one hour per sample and distance were typical in our tests. Other sample shapes, such as hemispheres and cylinders were considered instead of flat planes, which could potentially yield information for many sample orientations in a single scan. Such a shape would have significant drawbacks however. Firstly, the cost and difficulty of producing and validating a set of artefacts with different, consistent, isotropic surface finishes is far greater than for flat plates. Secondly, the quantity of data representing a particular surface normal on a curved surface is technically infinitely small. A point grouping technique would therefore be required to select points covering a range of similar gradients, limiting the amount that can be collected and the ability to assess its quality.

The choice of sample surface finish is arbitrary, however it is best to match it as closely as possible to the types of object the scanner will be used on. The methodology and data processing steps described rely on the assumption that the samples are isotropic, so it is most important to select an appropriate finishing process, such as shot blasting, barrel finishing or random action abrasive sanding.

When deciding on the set of surface orientations to test, more orientations should be taken about the direction where self-blinding is expected to occur, as this is where the quality of scan is most sensitive to changes in surface orientation. The sample preparation and validation, test apparatus and setup and data processing steps are explained in Sections 2.1, 2.2 and 2.3 respectively.

2.1. Sample preparation

Four samples were prepared on which to evaluate the performance of the scanner. However, if the sample exhibits periodic texture, say from a turning or milling process, it will generate a directional diffraction grating effect and a non-isotropic BRDF [21]. This would introduce sample rotation and the nature of the periodicity as additional experiment variables. In this investigation, this degree of complexity was removed by considering samples with isotropic surface finish only.

Samples were manufactured from 60×60×2 mm aluminium sheet. The selection of sample size depends on many factors, including the scanner field of view, resolution, distance and the range of surface normals to be tested. Through these factors, sample size affects the number of data points that can be recovered in each scan. More data points improve the confidence of the performance metrics, especially at orientations where the sample is viewed from highly oblique angles. However, if the sample is too large relative to the sensor field of view then incidence angle will vary significantly across the sample surface. Size selection is therefore a compromise between the number of points on the surface and the variation of the viewing angle over the sample, as shown in Fig. 1. A large sample also requires a large pan-tilt table to orient it, which may be limiting. The criteria for selecting a 60 mm square plate for this evaluation is that the relative surface angle varies by no more than 5° over the sample surface at the minimum distance scanned (400 mm), and more than 500 points are still collected on a matt white surface at the maximum angle and distance tested.

The data processing step involves fitting a plane to point clouds of the sample. As such, the plate should be approximately an order of magnitude flatter than the possible resolution of the scanner in order to prevent errors of form in the sample being misinterpreted as measurement noise. At 400 mm, the Ensenso is quoted as having a depth resolution of 0.34 mm. Therefore, the flatness of the samples should ideally be less than 34 μm.

A random action orbital abrasive process using various grades of wet-dry sandpaper was used to create a range of surface finishes. Fig. 2 shows the manufactured samples. A matt white sample, sample 4, was prepared to act as a benchmark, optically compliant, surface akin to characterisation artefacts prescribed in other methods. Table 1 details the surface roughness parameters of the samples, as measured in the X

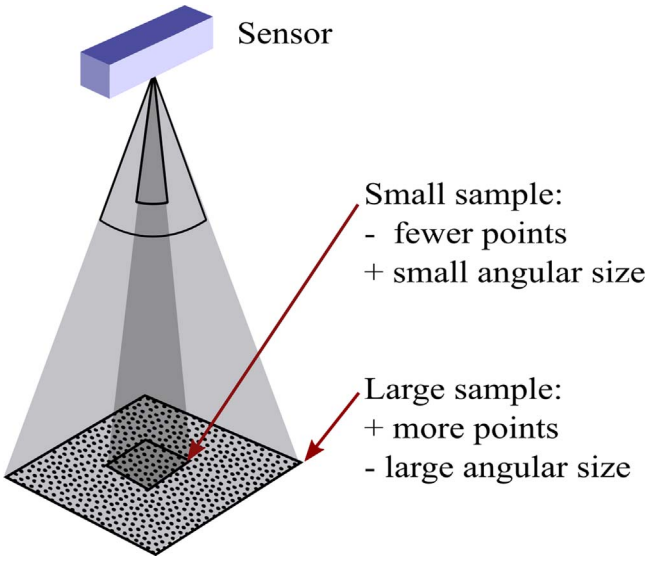

Fig. 1. The compromise of sample size on the number of points acquired and the angular size of the sample.

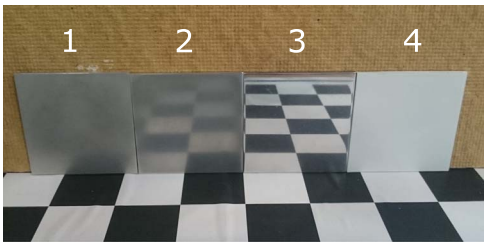

Fig. 2. Photograph of samples. The reflection of the checkerboard pattern on the samples demonstrates their relative surface finish.

Table 1
Sample surface roughness and flatness parameters.

		Sample			
		1	2	3	4
Ra (μm)	X	0.46	0.39	0.09	0.82
	Y	0.46	0.34	0.09	0.66
Rq (μm)	X	0.59	0.54	0.13	1.04
	Y	0.59	0.45	0.13	0.83
Flatness (μm)	X	15.5	20.2	11.9	24.6
	Y	39.9	13.9	30.8	19.7

and Y directions using five equally spaced profiles 55 mm long using a Talysurf CLI 2000 profilometer. To calculate *Ra* and *Rq*, a cut off wavelength of 0.8 mm was used according to EN ISO 4288. The flatness was measured by taking the maximum range of heights from the five profiles in each direction. The flatness of all four samples is acceptably close to the 34 μm required by the depth resolution of the scanner. The range of surface roughness was chosen to transition between the expected specular and diffuse behaviour of the sample in response to the Ensenso pattern projector. Sample three has an $Rq \ll \lambda$ and is therefore predominantly specular, whilst sample four has an $Rq \approx \lambda$ and is therefore diffuse.

2.2. Apparatus

The sensor selected to demonstrate the evaluation method is an Ensenso N10-304-18. The Ensenso is an active stereo vision camera that uses a pattern projector that operates in the infrared. The pattern

Table 2
Ensenso N10-304-18 specifications from manufacturer's datasheet.

General specifications	
Working distance	170–2000 mm
Optimum working distance	500 mm
Image resolution	752 × 480
Baseline	100 mm
Performance at optimum working distance (500 mm)	
Z Resolution	0.523
View Field Size	569 × 401 mm

projector augments stereo matching performance on surfaces with little texture of their own. The illuminant is not coherent, however the overall intensity of a returned coherent pattern such as one produced by a laser projection system is governed by the surface BRDF in the same way as a non coherent pattern. The only difference being the intensity of the return is modulated by the phases of photons arriving at the pixel to produce a speckle pattern. As the speckle pattern itself is unpredictable unless *a priori* knowledge of the surface texture is known, the method proposed should adequately allow the comparison of both coherent and non-coherent 3D measurement sensors. Hardware specifications of the Ensenso based on the datasheet values [22] are given in Table 2. The datasheet does not specify what surface finish the sensor will function on, nor what surface any performance evaluation has been conducted on. Stereo vision is a mature technology and as such details of the operation of the Ensenso will not be entered into here. An interested reader can refer to [23] for further details.

Any method is appropriate to control the sample orientation, providing it allows sufficient repeatability over a requisite range of angles. The angle range of the table must be adequate to expose the performance limitations of the sensor on the sample surface finishes. From previous experience of characterising sensor performance, diffuse surfaces require large changes of surface orientation to noticeably change scanner performance parameters. Shiny surfaces however have much higher rates of change. On the shiniest sample tested (a near mirror finish), the transition between maximum and minimum performance occurs over a range of approximately 20° of sample tilt. If we assume we require at least 10 points to adequately describe this transition, this places a modest limit on tilt table resolution of 2°. As such, low cost pan-tilt tables can be used in this characterisation method. The table may be manually or computer controlled, although the speed benefits of an automatable system cannot be overstated. Regardless of the orientation method, it must be possible to define a surface normal with respect to the camera co-ordinate system. This requires knowledge of the transformation between the tilt table and camera coordinate frames.

In this evaluation, a simple pan-tilt table, constructed using Lego®, was used to orient the samples as shown in Fig. 3. The table is controlled using the RWTH - Mindstorms NXT Toolbox for MATLAB® [24]. The toolbox provides control over motor movement and access to encoder positions. Functions were written to control of the sample normal, $\boldsymbol{n}$, by specifying polar co-ordinates azimuth, θ, and polar angle, Φ, up to a maximum of 55°. The table has a repeatability of ± 1.5°. The co-ordinate systems of the pan-tilt table and the Ensenso camera are shown in Fig. 4. The transformation between the coordinate systems C and T consists of a translation, $\boldsymbol{V}$, and a rotation of 180° about the y-axis. The exact value of $\boldsymbol{V}$ is determined during the data processing stage, but the rotation is fixed using an alignment jig on the table top, positioned carefully with reference to the camera mounting frame to ensure that y_T and y_c are parallel. This jig also coarsely locates O_T, the origin of the tilt table, along the axis z_C. The relative angle between the sample normal, $\boldsymbol{n}$, and $\boldsymbol{V}$ is Φ_R. The angle between $\boldsymbol{V}$ and z_c is β.

The sensor mounting frame performs two functions, the first is to maintain geometry; axis z_C remains perpendicular to the table top and

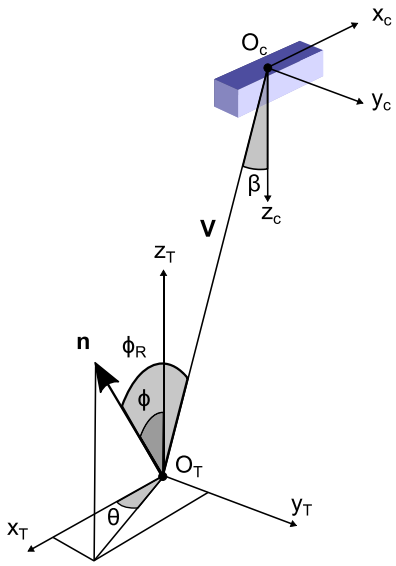

Fig. 3. The camera and tilt table co-ordinate systems, denoted by subscript C and T respectively.

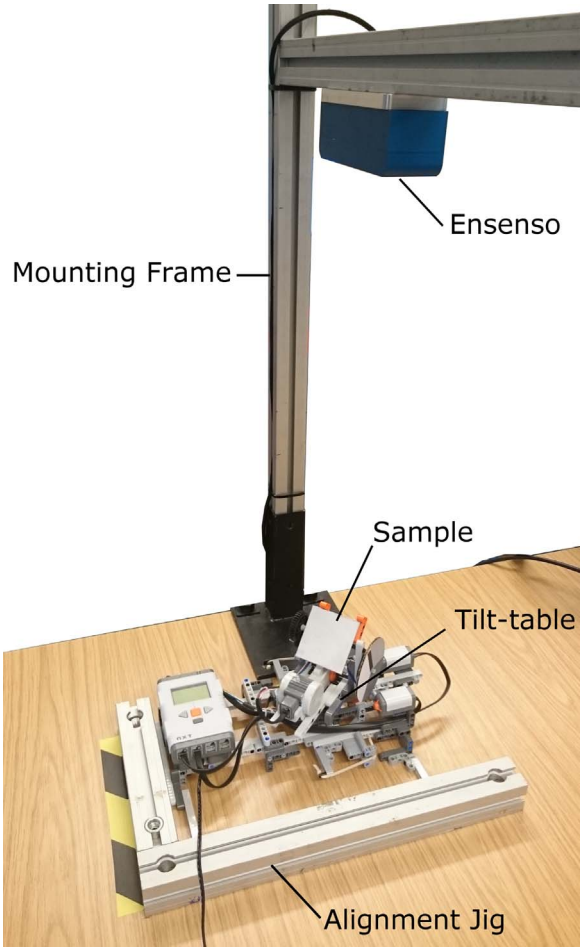

Fig. 4. Photograph of the experimental setup.

axis y_C parallel to y_T. The second is to allow the translation of the sensor along the z_C, to change the distance between the sensor and sample. For each of the four samples, sets of point clouds were recorded at distances of 400–800 mm in increments of 100 mm. Each set consists of point clouds measured at azimuths of 0–350° in steps of 10° and polar angles of 1° to +55° in increments of 2°. Point clouds were captured in synchrony with the MATLAB control script using the Ensenso SDK [25] and stored in text files. The Ensenso is capable of capturing at 30 Hz. The rate determining step in the experimental process is the movement speed of the pan-tilt table, which was able to capture an image on average every four seconds. A set of scans for a given surface and distance therefore took approximately one hour. A more consistent pan-tilt table would reduce this significantly however, as the table used had to undergo a recalibration procedure every 50 scans to compensate for drift in positioning accuracy.

2.3. Point cloud processing

The raw point clouds require processing to extract parameters describing the quality of the data measured from the sample surface at each surface normal. This is achieved in three steps. First, the points acquired from the sample surface must be segmented from the rest of the scene. Second, a plane is fitted to the remaining points. Finally, the performance metrics are calculated based on the number of points acquired and point noise. All processing was performed in MATLAB.

2.3.1. Point cloud segmentation

For each point cloud the origin of the tilt table, $\boldsymbol{O_T}$, must be located in order to reliably segment the point cloud. This is the centre of rotation of the sample, and hence remains the same for every point cloud for a particular sample and distance experiment. The sample surface itself lies 4 mm above the axis of rotation due to the design of the tilt table. As such, the sample both translates and rotates as it sweeps through polar angle. The centre of the sample, $\boldsymbol{S}$, can therefore be calculated as $\boldsymbol{S}=\boldsymbol{O_T}+\boldsymbol{n}d$, where $d=4$ *mm*. A point is segmented from the cloud if it lies within a distance of $r=22$ *mm* from $\boldsymbol{S}$, as shown in Fig. 5. The origin was selected manually for each sample and distance combination, such that the point $\boldsymbol{S}$ consistently lies on the sample surface for all orientations.

2.3.2. Measurement noise

Following segmentation, a plane, W, is fitted to the data points in the least squares sense as shown in Fig. 6a. The perpendicular distance, D, from each point to the plane is calculated as follows:

$$D_k = \hat{\boldsymbol{n}}_{\boldsymbol{W}} \bullet (\boldsymbol{P}_k - \boldsymbol{W}_0)$$

Where $\boldsymbol{P_k}$ are the co-ordinates of a point in the point cloud with index k, $\boldsymbol{n_W}$ is the normal of the plane W and $\boldsymbol{W_O}$ is an arbitrary point on the plane.

Point standard deviation, σ, is used as a measure of point noise. This is calculated as the standard deviation of the perpendicular distances from each point to the plane, where N is the number of points in the segmented point cloud:

$$\sigma = \sqrt{\frac{1}{N-1}\sum_{k=1}^{N} |D_k - \overline{D_k}|^2}$$

Where $\overline{D_k}$ is the mean distance from each point to the plane. As the plane was fitted to the points in the least squares sense, the value of $\overline{D_k}$ is zero. A histogram showing the distribution of perpendicular distances from each point to the plane is shown in Fig. 6b. A Gaussian probability density function (pdf) with a mean of zero and standard deviation σ is overlaid. The pdf of D is well represented by the Gaussian for this particular case, however on some surface and scanner combinations it may differ and a large number of tests are required to determine the underlying pdf. Due to the lack of a general noise model for 3D sensors, standard deviation is taken to be the measurement noise metric.

2.3.3. Fraction of recovered points

The measurement noise alone is not sufficient to characterise the

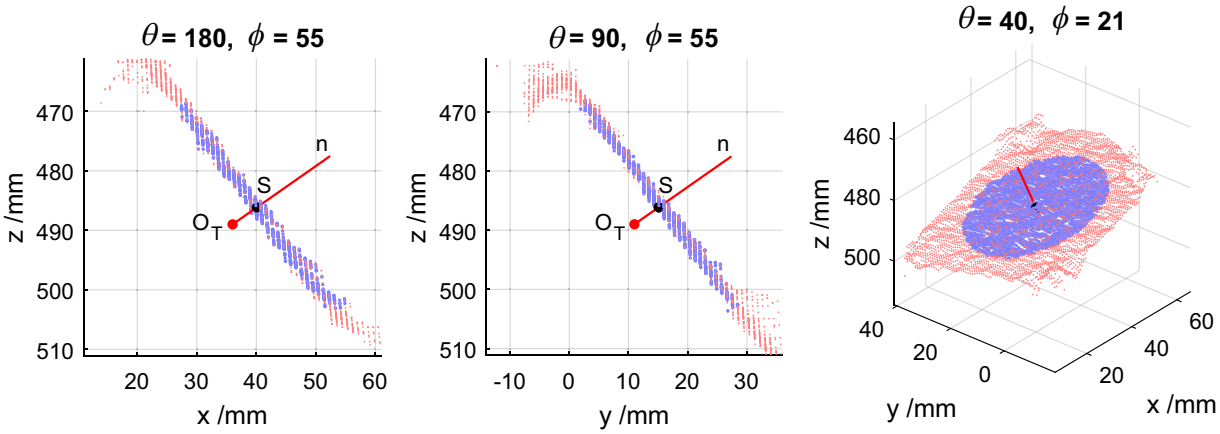

Fig. 5. Method for determining $\mathbf{O_T}$ and **S**. Point clouds showing a segmented region (blue) for various sample orientations. Data is of sample 4 at 500 mm. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

performance of a 3D sensor. It is equally important to know the probability of actually acquiring a point on a particular surface. A simple measurement of this may be to calculate the point density, ρ, with units of points/mm^2, by counting the number of points recovered and dividing it by the area over which they were measured. This value could then be used to predict the number of points it is possible to measure on a given surface at a given distance and orientation. However, this parameter cannot be used to compare relative performance over different variables, as it says nothing about the number of points it is actually possible for the scanner to measure. For instance, a sample at zero inclination may yield 0.5 points/mm^2 at 800 mm distance, and 2 points/mm^2 at 400 mm distance. The scanner does not necessarily perform 4 times better at 400 mm. If ρ_{max} is the maximum density of points possible and we assume that at 800 mm ρ_{max}=1 point/mm^2 and at 400 mm ρ_{max}=2 point/mm^2, then our scanner has recovered 50% of possible points at 800 mm and 100% at 400 mm, so in fact only performs twice as well at 400 mm. This normalised point density is referred to as the fraction of recovered data points, and is calculated as $F=\rho/\rho_{max}$. If the point density is not normalised in this way it masks where the sensor actually reaches its performance limits and starts to recover less data than expected. Provided ρ_{max} can be calculated, F is independent of both sample orientation and distance.

To calculate ρ_{max}, the Ensenso is modelled as a pinhole camera to determine the area imaged by a pixel at a given distance, d, and angle β from the sensor. Fig. 7 shows the geometry of a pinhole camera imaging a small square sample area, Δ^2 on a pixel with real size, s. The camera focal length is f and the angle subtended by a pixel on the sample is γ. For the Ensenso camera, f=3.6 mm and s=6 μm from the manufacturers datasheet. From the cosine rule, we can calculate the angle γ:

$$\gamma = cos^{-1}\left(\frac{b^2 + c^2 - s^2}{2bc}\right)$$

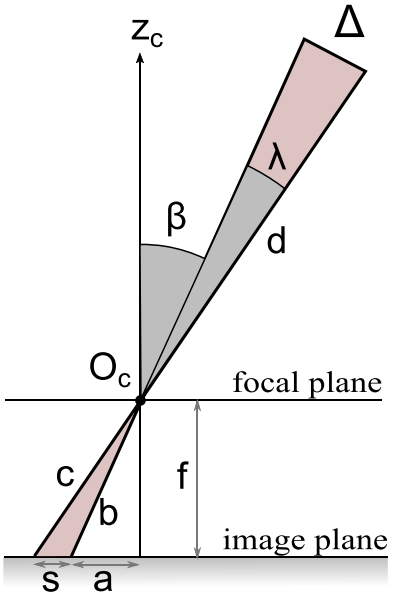

Fig. 7. Pinhole camera geometry imaging a small square area Δ^2.

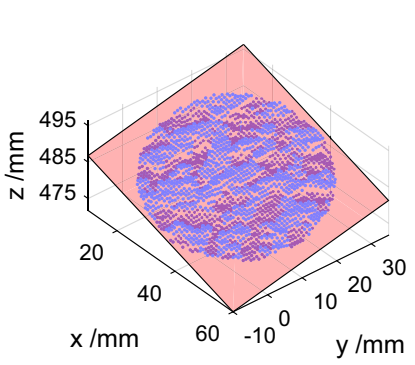

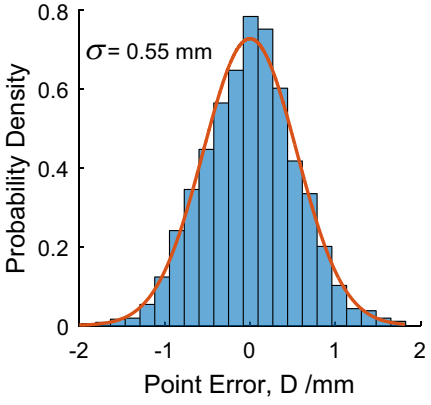

Fig. 6. Data for sample 1 at 500 mm, θ=40°, Φ=21° showing (a) A segmented point cloud with a fitted plane and (b) a histogram showing the distribution of perpendicular distances from each point to the plane.

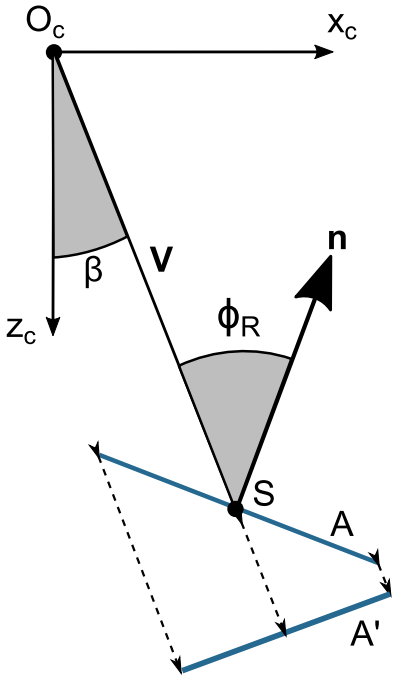

Fig. 8. The projection of area A onto area A′ along the direction of **V.**

Where by Pythagoras, $b^2=a^2+f^2$ and $c^2=(a+s)^2+f^2$. The angle $\beta = \cos^{-1}(\hat{V}\bullet\boldsymbol{k})$ and $a = f\tan\beta$. $\boldsymbol{k}$ is the unit vector along the z_C axis; [0 0 1].

Using the small angle approximation, the size of the surface element $\Delta=\gamma d = \gamma \mid V \mid$. It follows that surface points will be recovered on the sample in a grid with a spacing of Δ, therefore the maximum point density, $\rho_{max}=1/\Delta^2$. However, this point density is only correct for surfaces which are perpendicular to the vector $\boldsymbol{V}$. To account for this, the calculation of ρ is simply the number of points in the segmented cloud, N, divided by the projected sample area, A', as shown in Fig. 8. In the case of this segmentation method, $\rho = N\pi r^2\cos\beta$.

There are two disadvantages with this model. The first is that it does not take into account radial distortion of the camera optics, and hence should only be used for objects close to the centre of the field of view. To correct for this, it would be necessary to perform an intrinsic camera calibration. Whilst possible with the Ensenso, it would make the method impossible to implement on a 3D scanner that does not allow the capture of raw images from the camera. The second is that it requires knowledge of the focal length and pixel size of the camera, which is not always available in a 3D scanner's datasheet. An alternative approach would be to take the point density from a matt white sample as ρ_{max}. Doing so removes the need for *a priori* knowledge of the camera, but increases the number of tests required to characterise a sensor.

3. Analysis of characterisation data and presentation of results

During the experimental phase of the sensor evaluation, a large volume data is collected. In our evaluation, with only four samples and five sample distances, over twenty thousand point clouds were captured. Each point cloud was processed to extract the parameters described in Section 2.3. Careful consideration must be given to present the results in a way that allows the meaningful comparison of different scanner systems. This section describes the methodology and reasoning to arrive at such results. The point clouds from this evaluation are available with the http://dx.doi.org/10.17028/rd.lboro.4258274.

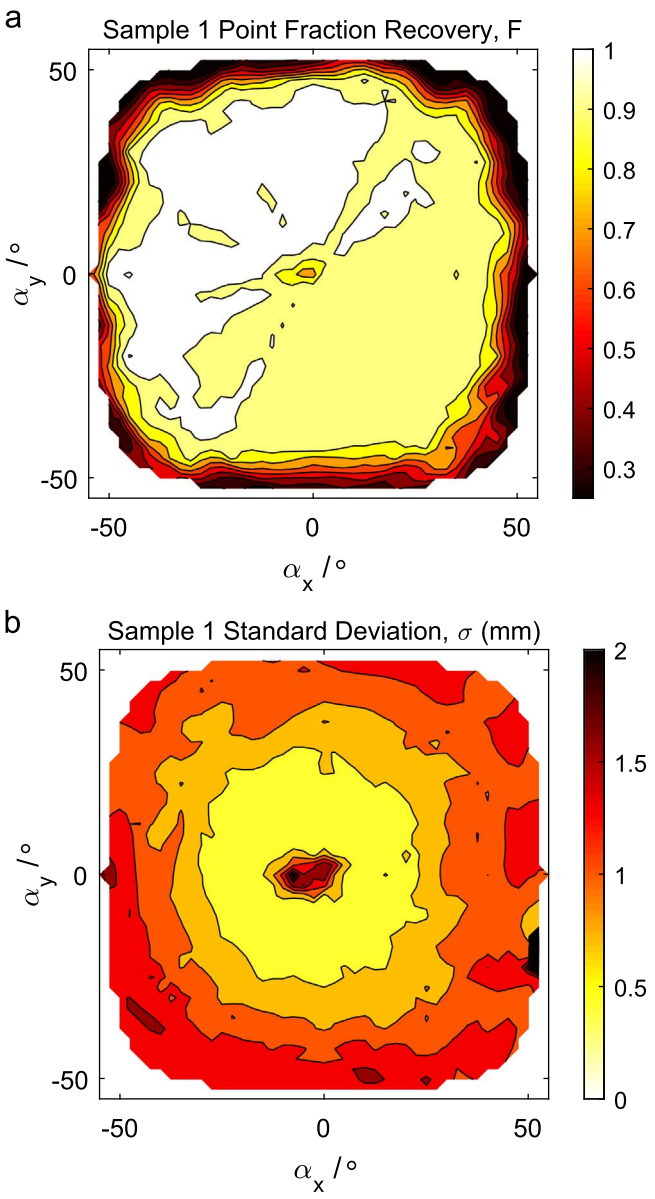

Fig. 9. Contour maps for sample 1 at 500 mm for (a) point fraction recovery, F and (b) point standard deviation, σ.

Fig. 9 shows contour plots of results for F and σ for sample 1 at a distance of 500 mm. The results are linearly interpolated onto a grid with a 2.5° spacing. The graph is plotted on axes of X and Y angle, where if **n** is defined as:

$$n=[n_x n_y n_z]^T$$

The x and y angles for this normal are therefore:

$$\alpha_x=tan^{-1}\left(\frac{n_x}{n_z}\right)\alpha_y=tan^{-1}\left(\frac{n_y}{n_z}\right)$$

Of particular interest is the central region of self-blinding resulting in significant point uncertainty, as indicated by the high standard deviation. This is the region of angles where the sample is reflecting the light from the projector directly back into one of the two cameras, resulting in image saturation and/or poor contrast of the projected pattern. A drop in the fraction of recovered points at high inclinations is visible, dropping to 0.3 at angles of 50°, due to a poor return of the projected light pattern from the projector back to the camera. The point uncertainty is seen to degrade far more gradually over the same range.

All projected light systems must cope with self-blinding and adverse

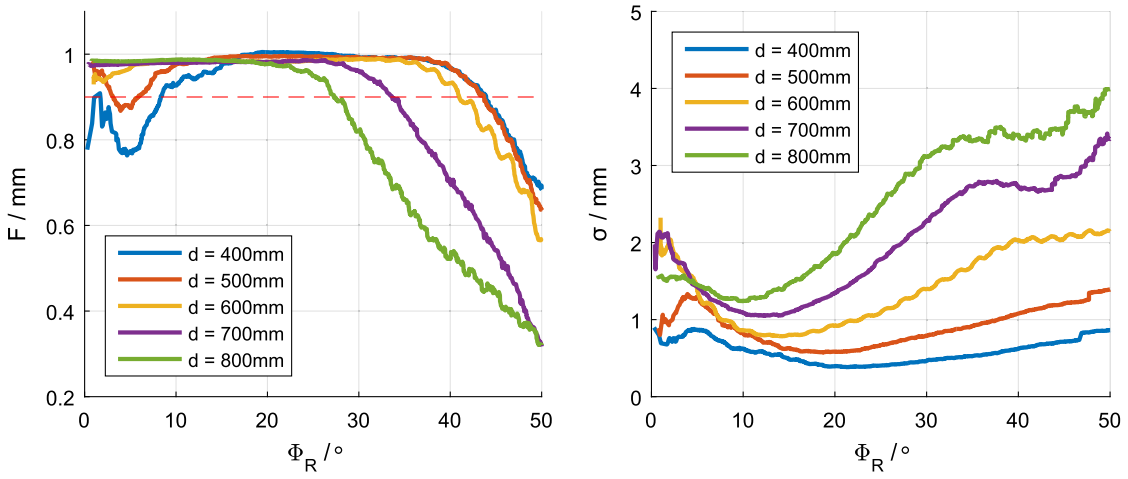

Fig. 10. Sample 1 results for (a) Fraction of recovered points, F, showing the 90% level as the dashed line and (b) standard deviation, σ.

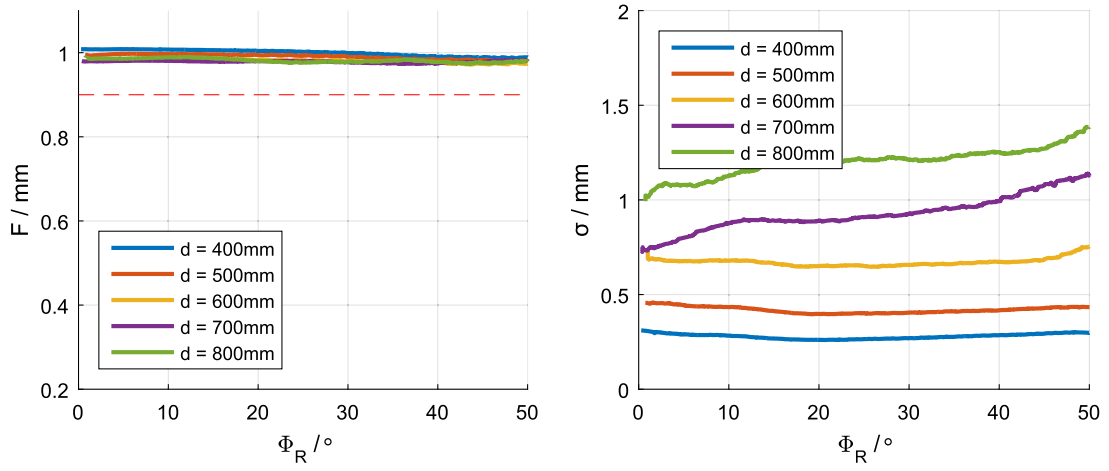

Fig. 11. Sample 4 (matt white) results for (a) Fraction of recovered points, F, showing the 90% level as the dashed line and (b) standard deviation, σ.

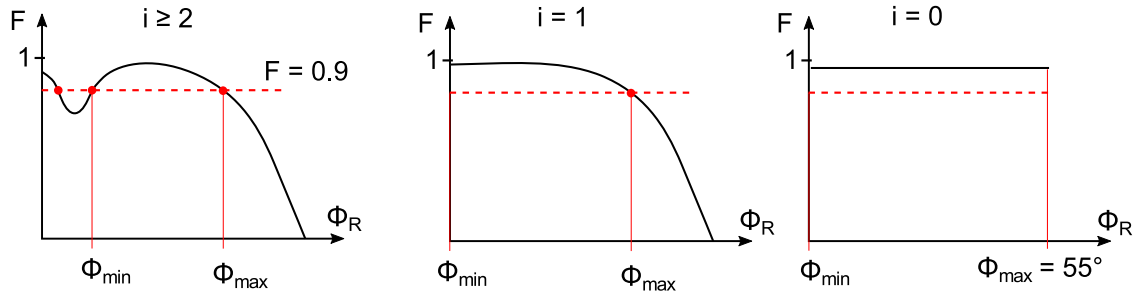

Fig. 12. The selection of Φ_{max} and Φ_{min} for different numbers of intersections.

scattering. The variation in measurement systems in terms of lighting and imaging strategies and processing methods mean that systems will vary in performance; for example some high-end industrial systems make use of multiple exposure imaging, and use multiple cameras to extend dynamic range and reduce sensitivity to surface texture and form. However, the functionality offered by such systems usually comes at significant extra cost and without a method to directly compare like for like performance, there is no way for a user to assess if the extra cost is warranted, or indeed what the limits of any technology are.

Ideally, the contour plots should be perfectly symmetrical. However, the experiments were performed in a laboratory with no controls over ambient light, as this is the condition the sensor is used in on a day to day basis. As such, the uneven features are due to windows and overhead lights reflecting on the sample and different orientations and reducing the signal to noise ratio of the images.

Whilst the contour plots are useful for analysing results at a particular sample at a given distance, 40 charts (5 distances, 4 samples, 2 metrics) are required to fully display the data from all the characterisation experiments. For ease of use and efficient comparisons, it is therefore necessary to reduce the dimensionality of the data, with the aim of reducing the results to a single performance chart per surface type, incorporating both F and σ.

The first step to achieve this is to plot F and σ versus relative surface angle, Φ_R, therefore reducing the need for 2 angles, α_x and α_y, to describe a surface orientation. This is possible as the samples are isotropic and hence have a BRDF that is independent of the sample rotation about **n**. This is exploiting the axial symmetry present in Fig. 9. Figs. 10 and 11 show the results of doing this for samples 1 and 4 respectively. Each line represents how the value of a performance parameter, F or σ, changes as a function of sample angle, Φ_R, for a

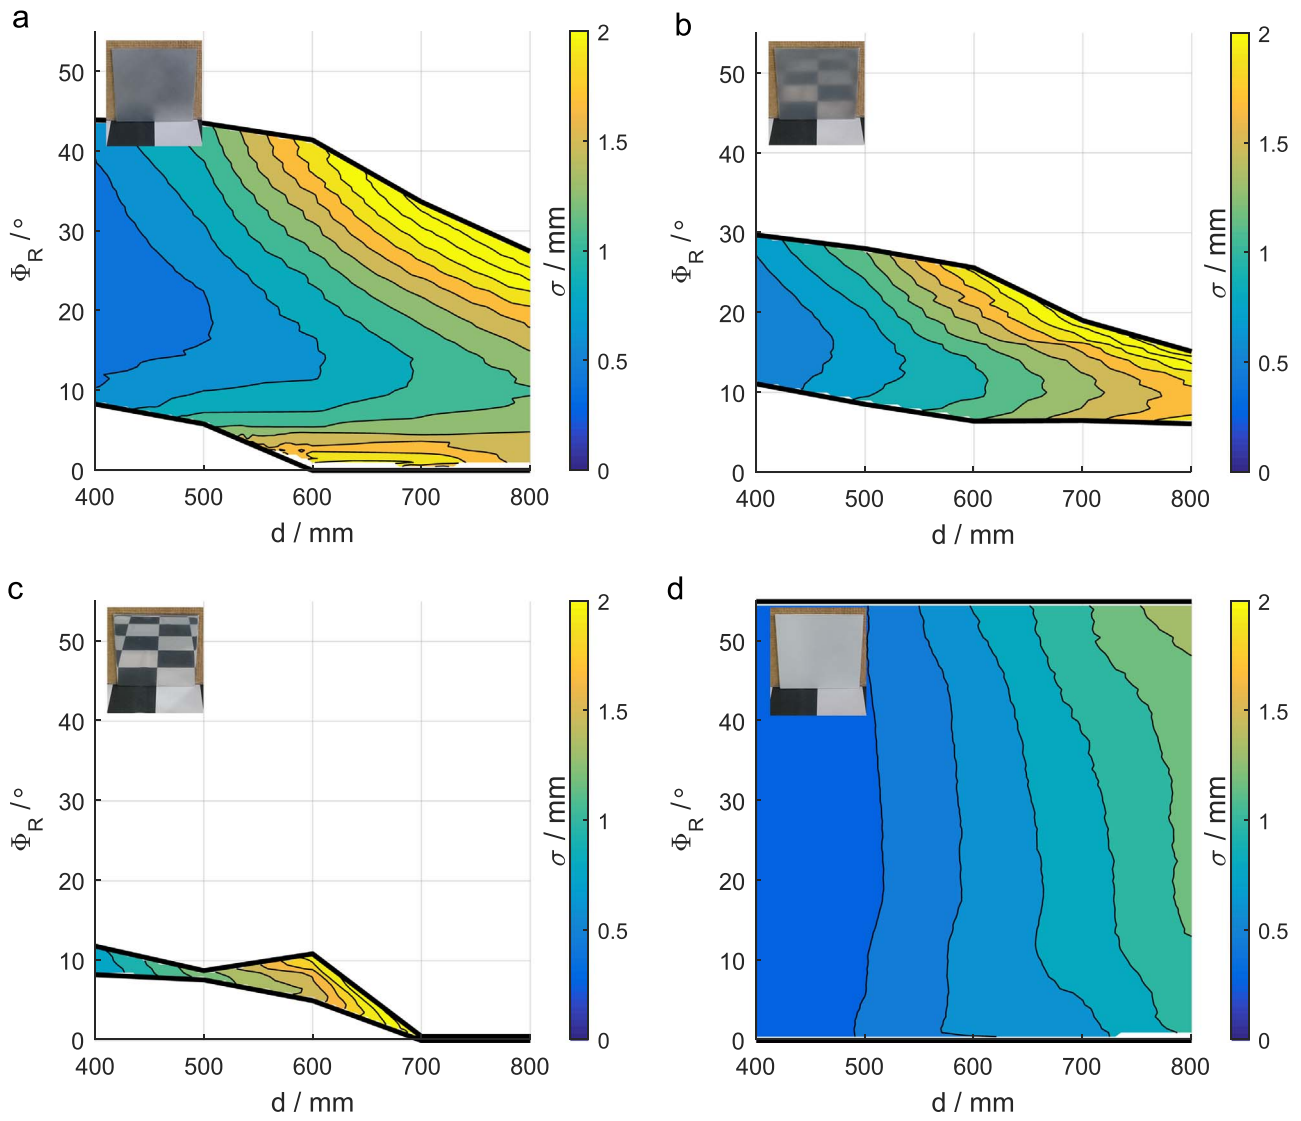

Fig. 13. Performance charts for F≥0.9 for samples (a)1 (b) 2 (c) 3 and (d) 4, coloured by point standard deviation. Inset photographs are of a checkerboard reflecting in the corresponding sample to illustrate relative shininess. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

particular distance, d, and sample. Each data series is filtered by a 20 point moving average. For sample 1, this reveals an increasing drop off in F as a function of both distance and sample angle, which is accompanied by an increase in point noise. In addition, the effect of self-blinding is seen to be small after a distance of 600 mm. The matt white surface, sample 4, shows nearly 100% point recovery over the range of distances tested, and a point standard deviation below 1.5 mm for all measurements, compared to 4 mm for sample 1. No self-blinding occurs on the matt white sample.

To further reduce the number of graphs required to describe the sensor performance, it is assumed that they need not show the probability of recovering a point at an arbitrary surface angle, but rather show where there is a probability above a certain threshold of recovering a point. As such, the parameters Φ_{max} and Φ_{min} are determined for each distance curve at the intersection of the line $F=F_{lim}$. The selection of the cut off F_{lim} is somewhat arbitrary and can be chosen to reflect the performance requirements for a particular application. In this characterisation, it is taken as 0.9. Fig. 12 shows the selection process of Φ_{max} and Φ_{min} for different numbers of intersections, *i*.

Finally, Φ_{max} and Φ_{min} can be plotted for each sample as a function of distance. The region bounded by Φ_{max} and Φ_{min} represents the range of surface angles where fractions of points greater than F_{lim} are expected to be recovered. Each point in this region has co-ordinates (d, Φ_R) and therefore has a standard deviation associated with it, which can be calculated by interpolating between the curves for σ *vs* Φ_R at the corresponding Φ_R coordinate. Once the region is mapped by standard deviation, it can be colour mapped and displayed as seen in Fig. 13. The graph therefore describes the expected standard deviation on any surface orientation where more than F_{lim} points are expected to be recovered. For example, to plot the standard deviation at a distance of 550 mm and a sample angle of 20°, the value of σ is calculated by interpolating between the 500 and 600 mm curves on the plot of σ *vs* Φ_R.

As is to be expected, the self-blinding at low values of Φ_R becomes more severe on shinier samples, and at shorter distances. Similarly, shiny surfaces cease to yield a useful number of points at shallower inclinations than dull ones. This is not surprising to anyone who has even a modest experience with 3D scanners. However, outcomes of this methodology enables a user to easily identify the optimum scanner orientation for a given surface, distance and scanner combination, or indeed determine without trial and error if a particular scan will be possible. For a particularly challenging surface, such as sample 3, it identifies the narrow range of conditions at which it is possible to get useful information. It is envisaged that this data could be used to predict the statistical properties of a point cloud if the surface finish of

the subject were known, and subsequently predict the optimum position to scan an object from.

Crucially, the method also shows the contrast in performance between even the dullest metallic sample and the matt white sample representative of typical characterisation artefacts. Performance degrades gradually across the relatively large range of surface roughness tested as the surface transitions from diffuse to specular behaviour. In these results, performance similar to that on the ideal sample is only achieved over a very narrow band of surface orientations for sample 1, and never for samples 2 and 3. Therefore, it is essential to perform any characterisation on surfaces similar to those to be used in the final application. In addition, any performance metric should always be quoted with details of the surface finish of any artefact used to measure it.

As the presented methodology stands, providing care is taken to control lighting and sample position, it allows for a direct comparison of 3D imaging systems under the same circumstances. The range of surface finishes available from manufacturing process is vast however, and producing a representative set of samples for characterisation is a significant challenge. This presents a limitation for predicting performance on an arbitrary object, as a sample must either be manufactured to the same surface specification of the object or a sample with similar optical properties must be used instead. Determining surface properties which will allow either the interpolation between data sets from known samples, or the selection of similarly performing samples would therefore be a beneficial area for future work. Due to the complexity of dealing with anisotropic surfaces, the work so far has been based on isotropic surfaces only; this is in line with almost all other metrological artefacts used to assess the performance of 3D vision systems, which have isotropic surface finishes.

A potential future application for this method is the ability to predict the statistical properties of a point cloud based on knowledge of an objects surface properties and geometry. This could allow the optimisation of scanner location on production lines or in freeform assembly or reverse engineering applications, where an estimation of object position could be used to find the optimum location to perform a more detailed scan. The characterisation method presented in this paper would be completely appropriate for any object with an isotropic finish, for example, metal parts that have been cast, forged, sand-blasted, shot-peened, selective laser sintered, injection moulded, or the vast majority of moulded plastic parts or ceramic parts. Characterising and modelling the effects of anisotropic surface finish is the primary challenge to achieving sensor simulation on parts with completely arbitrary surface finish, which will be investigated in future work.

It is important to note that this paper is intended to present guidelines of a method to produce performance metrics that are generic to any 3D sensor. The Ensenso is used to demonstrate the procedure; it was not the intention to present a comparison of sensors as to do so would be cumbersome and detract from the presentation of the method itself. In future work, studies will be undertaken to evaluate multiple 3D imaging systems and technologies with the proposed methodology. The authors also invite other researchers active in the field of 3D vision system design and characterisation to consider the use of this methodology and metric.

4. Conclusions

This paper presents a methodology that fills a critical gap in the characterisation procedures for 3D imaging systems; it allows the evaluation of sensor performance in a way that is representative of real world measurements, and exposes a sensors' limitations in terms of measureable surface types and orientations. Two metrics allow a simple and pragmatic approach to sensor comparison and a convenient method for visualisation of sensor performance with respect to these metrics was defined. The only constraint on the sensor technology is that it must be possible to produce point cloud output and no intimate working knowledge of the sensor is required. Combined with the low cost of sample manufacture and apparatus, this allows manufacturers and third parties alike to characterise and compare sensors, and assess sensors capability for different applications.

Acknowledgements

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) through grant numbers EP/IO33467/1 and EP/L01498X/1. The authors would like to thank the support staff of the EPSRC Centre for Innovative Manufacturing in Intelligent Automation for providing equipment and facilities to conduct this research.

References

[1] Vezzetti E. Computer aided inspection: design of customer-oriented benchmark for noncontact 3D scanner evaluation. Int J Adv Manuf Technol 2008;41:1140–51.
[2] Beraldin J-A, Mackinnon DK, Cournoyer L. Metrological characterization of 3D imaging systems: progress report on standards developments. In: Proceedings of the 17th International Congress Metrology; 2015, p. 13003. doi:10.1051/metrology/20150013003.
[3] Besl PJ. Active, optical range imaging sensors. Mach Vis Appl 1988;1:127–52.
[4] Ikeuchi K. Generating an interpretation tree from a CAD model for 3D-object recognition in bin-picking tasks. Int J Comput Vis 1987;1:145–65.
[5] Martin RR, Varady T, Cox J. Reverse engineering geometric models-an introduction. Comput Des 1997;29:255–68.
[6] Schneider R, Schick A, Ko P, Ninomiya T. High-speed optical three-dimensional scanner for automatic solder joint inspection. Opt Eng 1997;36:2878–85.
[7] Weingarten JW, Gruener G, Siegwart R. A State-of-the-Art 3D sensor for robot navigation. In: Proceedings of the IEEE/RSJ international conference intelligent robots system, Sendai, p. 2155–2160; 2004.
[8] VDI/VDE 2634-2. Optical 3D measuring systems: optical systems based on area scanning; 2012.
[9] Torrance KE, Sparrow EM. Theory of off-specular reflection from roughened surfaces. J Opt Soc Am 1967;57:1105–14.
[10] Nicodemus FE, Richmond JC, Hsia JJ, Ginsberg IW, Limperis T. Geometrical considerations and nomenclature for reflectance, 160. Washington: U.S. Dept of Commerce, National Bureau of Standards; 1977. http://dx.doi.org/10.1109/LPT.2009.2020494.
[11] Hausler Gerd, Ettl S. Limitations of optical 3D sensors. Opt. Meas. Surf. Topogr 2011:23–48.
[12] Dresel T, Häusler G, Venzke H. Three-dimensional sensing of rough surfaces by coherence radar. Appl Opt 1992;31:919–25.
[13] Guidi G. Metrological characterization of 3D imaging devices. In: Remondino F, Shortis MR, Beyerer J, Puente León F, editors. In: Proceedings of the SPIE 8791, videometrics, range imaging, applications XII; Automated visual inspection, vol. 8791; 2013, p. 87910M. doi:10.1117/12.2021037.
[14] Paakkari J, Moring I. Method for evaluating the performance of range imaging devices. In: Proceedings of the SPIE 1821, industrial applications of optical inspection, metrology, and sensing; 1993. p. 350–356.
[15] Beraldin J-A, El-Hakim SF, Blais F. Performance evaluation of three active vision systems built at the national research council of canada. Optical 3-D Measurement Techniques III;1995. p. 352–361.
[16] Vukašinović N, Možina J, Duhovnik J. Correlation between Incident angle, measurement distance, object colour and the number of acquired points at CNC laser. Strojniški Vestn – J Mech Eng 2012;58:23–8.
[17] Dury MR, Brown S, McCarthy M, Woodward S. 3D optical scanner dimensional verification facility scanner dimensional verification facility at the NPL's "National FreeForm Centre". Laser Metrology and Machine Performance XI (LAMDAMAP 2015); 2015. p. 191–200.
[18] McCarthy MB, Brown SB, Evenden A, Robinson AD. NPL freeform artefact for verification of non-contact measuring systems. Int Soc Opt Photonics 2011;7864:78640K.
[19] ASTM E2919 – 14. Standard Test Method for evaluating the performance of systems that measure static, Six Degrees of Freedom (6DOF), Pose; 2014.
[20] Hodgson JR, Kinnell P, Justham L, Jackson MR. Characterizing the influence of surface roughness and inclination on 3D vision sensor performance. In: Proceedings of the 8th International Conference on Machine Vision (ICMV 2015), vol. 9875; 2015. p. 1–7.
[21] Stover JC. Optical scattering measurement and analysis. New York: McGraw-Hill, Inc; 1990.
[22] IDS Imaging Development Systems GmbH. N10-304-18 Specifications. n.d.
[23] Zhang S. Handbook of 3D machine vision. Boca Raton: CRC Press; 2013.
[24] Schneider D, Staas M, Atorf L, Schnitzler R, Behrens A, Knepper A, et al. RWTH – Mindstorms NXT Toolbox; 2013.
[25] *I*DS Imaging Development Systems GmbH. Ensenso SDK n.d.