

Surface roughness estimation using structured light projection

Marjan Shahpaski¹, Luis Ricardo Sapaico² and Sabine Süsstrunk¹

¹School of Computer and Communication Sciences, EPFL

²DxO Labs, France

{marjan.shahpaski, sabine.sustrunk}@epfl.ch, rsapaico@dxo.com

Abstract

How we visually perceive non-emissive objects in our surrounding depends on the interaction of light with the optical characteristics of the materials that comprise them. The macroscopic surface roughness can also influence the appearance through shadowing and interreflections. In this work, we use a structured light scanner to estimate the surface structure of near-planar surfaces, namely of printing textiles. We compare our scans, both qualitatively and quantitatively, to those from a commercial high-grade profilometer based on the confocal principle. We achieve comparable results to the profilometer on samples with moderately complex surfaces. We discuss the possible reasons for errors in the scans of complex surfaces, thus providing guidelines for robust depth estimation. This comparison can help other researchers build more robust acquisition setups by understanding and minimizing the errors inherent to the reconstruction methods.

Introduction

Printing on textile is finding wider use, driven by the fast production and cost-effective print runs of digital printing. However, the difficulty of obtaining a product appearance which is close to the expected one is a limiting factor for the utilization of digital printers. Some reasons include lack of knowledge of the materials used, and inefficient and inaccurate color management, especially because textiles exhibit rough surfaces.

Material appearance is related to our visual perception. We can observe that an object is a given color, whether it is matte or glossy; opaque, translucent or transparent, etc. The material appearance depends on the optical properties of the materials, but also on the illuminating and viewing directions. In addition, macroscopic surface structures can influence the appearance through shadowing and interreflections.

Surface roughness is used as a proxy for describing the macroscopic surface texture. Because it has a strong influence on material appearance, its accurate characterization is important for creating realistic simulations or reconstructions of objects. Recent color prediction models for printing account for the roughness of the inks [1], but do not take into account the roughness of the substrate, which is about two orders of magnitude greater for textile substrates. Knowledge of the surface structure could allow for devising color prediction models for macroscopic-level features that can potentially incorporate shadowing and interreflections. Work in this direction has already been performed, for the more specialized case of diffuse [2] and specular V-cavities [3, 4].

We are interested in recovering the surface structure of coated printing textiles, where the weave threads form macroscopic structures. An example surface of a printing textile that we are interested in measuring can be seen in Figure 1.



Figure 1: An image of the surface of Canon IJM 618, a coated polyester textile used as a print media. We are interested in measuring the surface roughness of printing textiles, where the weave threads form macroscopic structures.

Profilometers have traditionally been used for measuring surface roughness. These instruments, although very accurate, are not easily affordable, and have a fairly limited scanning area.

In an effort to overcome these restrictions, we design and evaluate a surface scanning system based on structured light (SL) projection. This method provides non-contact means of depth estimation by using a calibrated camera-projector pair. The simple SL setup that we constructed is composed of off-the-shelf components, and it can be seen in Figure 2.

The goal is to test the feasibility of acquiring an accurate and high-resolution depth map of the surface structure by means of an affordable structured light scanner. The depth map can then be used to compute the surface roughness of the material. We believe that access to accurate surface scans from affordable devices can open the door to more accurate color prediction models for rough substrates, and more accurate computer simulations.

The benefits of increased scanning speed and decreased cost come with a trade off in decreased spatial resolution. With the devices and measurement parameters used in our tests, the spatial resolution was approximately 5 fold higher for the profilometer per dimension. However, we are interested in measuring the macroscopic surface roughness formed by the textile threads. In our set of textiles, the smallest periods are around $420 \mu\text{m}$, which is within the spatial resolution of the SL setup.

Our work focuses on discussing thoroughly the pros and cons of the SL setup for measuring surface roughness. We show that by using state-of-the-art SL based 3D reconstruction methods we are able to achieve comparable results to the profilometer on samples with moderately complex surfaces. Moreover, we include a transparent analysis of the accuracy-related problems common to all optical-based reconstruction methods.



Figure 2: Our simple structured light scanner used in the experiments. It is composed of a projector-camera pair, and a flat articulated stage where the samples are placed. A sinusoidal encoding pattern is being projected onto the measured fabric.

Sources of error in structured light scanning

To overcome the high cost and the limited scanning area of the profilometers, we turned to depth estimation based on structured-light projection. The goal is to compute an accurate and high resolution depth map of the surface structure, which can then be used to compute surface roughness. The high resolution is required to capture the surface features formed by the fabric weave, and the accuracy of the reconstruction could allow for devising accurate color prediction models that incorporate shadowing and local interreflections.

This field of research has been quite active in the past four decades. A myriad of SL encoding techniques is now available, all offering different trade-offs that are suited for specific purposes. Salvi et al. in [5] present a comprehensive overview and hierarchical classification of more than forty SL encoding techniques. The aim of SL is to encode the position of each pixel (or a region) in the projected pattern in a unique way, such that a correct correspondence can be established between the projected and the captured patterns.

Care has to be taken when using structured light scanning techniques, because the measured objects can be translucent and experience significant amounts of indirect illumination [6, 7, 8]. Furthermore, the scenes can be dynamic due to object movement and vibrations [9, 10]. Given these optical and mechanical noises, many coded structured light techniques fail to decode the patterns correctly, which hinders the accuracy of the reconstruction. Note that indirect illumination can also degrade the scans of the confocal microscope, as we will see in the Results section.

Indirect or global illumination is the set of effects that include interreflections, subsurface scattering and defocus. These effects alter the directly projected patterns. Interreflections are considered long-range effects, since they generally appear as low-frequency waves across the captured scene. Subsurface scattering and defocus, on the other hand, are considered short-range effects, as their effect is local and manifests in a local low-pass filtering on the projected patterns.

Interreflections are mainly caused by concavities in the measured objects' shape, where the directly projected light gets reflected from the object surface onto other portions of the object. It is generally assumed to be a reflection with low spatial frequency, however, smooth metallic or plastic surfaces can produce high-frequency mirror reflections.

Subsurface scattering happens with translucent objects, where the light entering the object at one point experiences multiple scattering events and internal reflections, and exits the object over a small area around the entry point. Subsurface scattering together with camera/projector defocus act as low-pass filters upon the projected patterns.

In addition to degrading the projected patterns, subsurface scattering can introduce measurement bias. During the scanning of Michelangelo's David, Levoy et al. [11] measured a depth bias of approximately 40 μm for light at normal incidence for the Carrara Statuario marble. This happened because the incident light from the laser beam got scattered and formed a volume below the surface of the marble. It created a brighter light centroid that was detected by the camera, instead of detecting the direct reflection from the object's surface.

Scene dynamics are another major source of artifacts that mainly affect the techniques that project multiple patterns when encoding the measured scene. For example, when measuring a part of the human body (e.g., face, eye curvature, etc.) it is highly unlikely that the subject will stay still during all the projections. Furthermore, when scanning a scene with high precision even the slightest vibrations can introduce measurement noise [10].

We are interested in estimating the surface profile of printing textiles. They have near planar surfaces and are white, since they are used as print media. An example patch of a Canon IJM 618, white, polyester textile coated with a special polymer coating, is shown in Figure 1.

We should consider the indirect illumination effects when scanning these textiles to choose the best encoding strategy. The local interreflections should not cause major issues when decoding the projected patterns, since the surface features are relatively small and cannot reflect considerable amounts of light towards other parts of the fabric. Subsurface scattering is, however, present with these materials. This will result in smoothing of the projected patterns, and can cause problems with high-frequency binary patterns. Finally, these are static scenes, and allow us to use temporal encoding techniques that project multiple patterns.

Taking these considerations into account, we decided to use a continuous phase shifting encoding. Unlike the discrete coding methods, these method allow us to use the full resolution of the camera, which is generally higher than that of the projector. The slight blurring of the patterns due to sub-surface scattering does not affect these patterns adversely. Furthermore, it is often suggested to slightly defocus the projector, such that the pixel grid will not be visible in the camera image.

Measurement method

The measurement principle of our setup is shown in Figure 3. The measured scene is illuminated with structured light by the projector. Since we are using a continuous coding technique (phase shifting), and a non-linear projector, the values of the projected patterns need to be distorted by an equal amount, but in the opposite direction. This is where the radiometric calibration of the projector is used. When projected, the patterns become linear, which is what the decoding algorithm expects in the captured images. The projected patterns are captured by the camera, and a correspondence is established between the projector and the camera pixels. With a continuous coding techniques, we can leverage the full resolution of the camera, because we can establish sub-

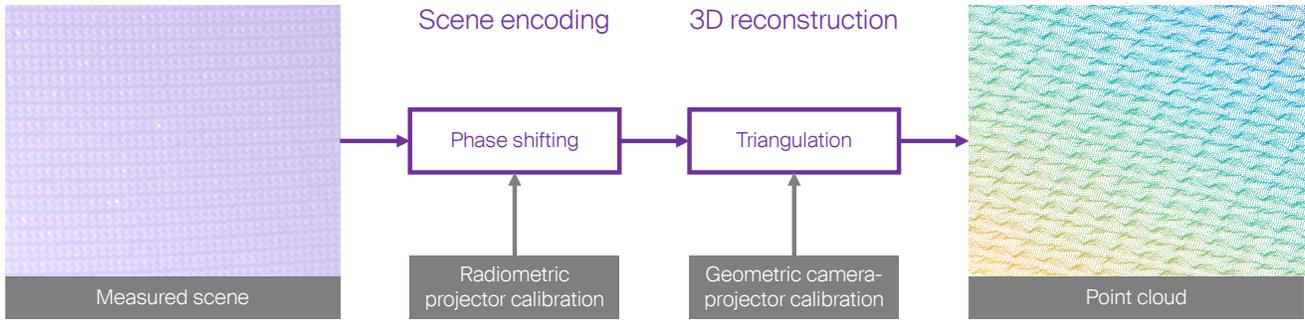


Figure 3: The measurement principle of our structured light setup. The measured scene is illuminated with structured light by the projector. The projected patterns are captured by the camera, and a correspondence is established between the projector and camera pixels. A triangulation procedure then computes a point cloud with a 3D position for each camera pixel.

pixel correspondences with the projector image. A triangulation procedure then computes the 3D positions of each camera pixel, for which the geometric calibration of the projector-camera pair is required. The output of this process is a point cloud, with as many points as there are camera pixels.

Measurement setup

We used a very simple structured light setup composed of a single camera and a single projector, which is shown in Figure 2. The projector that we used was an Acer H6502BD (color, DLP, 1920 px × 1080 px) and the camera was a Canon 5D Mark II (color, CMOS, 5616 px × 3744 px) with a Canon EF 50 mm f/2.5 Macro lens. The samples were affixed to the flat aluminum composite board shown in the background. The samples were at a distance of approximately 30 cm from the projector-camera pair. It was not possible to bring the samples closer, because we were limited by the minimal focal distances and the fields of view of the projector and camera lenses.

To calibrate the camera and the projector geometrically we used the simultaneous geometric-radiometric calibration method by [12]. The camera is capable of capturing linear images, however, the projector applies non-linear processing to the input signal. Since we are using a continuous coding technique, we performed a radiometric calibration of the projector. Instead of capturing a separate set of images for the radiometric calibration, we received it for free by using the above calibration method.

Roughness metric

Roughness can be defined in several ways, and here we will use the root mean square height of the surface points with respect to the mean plane of the surface they comprise [13]:

$$S_q = \sqrt{\frac{1}{A} \int \int z^2(x, y) dx dy},$$

where A is the area of the measured surface, z is the height from the mean plane of a data point with x and y its horizontal and vertical coordinates, respectively. In the case of a uniformly discretized surface, as in our case, we can write:

$$S_q = \sqrt{\frac{1}{N} \sum_i \sum_j z^2(x_i, y_j)}.$$

To test the geometric calibration of the projector-camera pair and the radiometric calibration of the projector, we reconstructed a portion of the flat board that holds the samples. The reconstruction of the board after a calibration with 15 patterns is shown in Figure 4a. The figure shows a 5 mm × 5 mm patch, corresponding to the size of the profilometer scan. We can see that the reconstruction is fairly accurate, with a roughness of only 3.03 μm. The visible waves in the height profile that propagate along the x-axis are caused by small errors in the radiometric calibration.

Error compensation

The fabrics are flexible, and might bend slightly even when attached to the flat supporting surface. Therefore, we performed a global rectification of all samples, measured with both the profilometer and with our setup. We divided the sample's surface into a coarse 3 × 3 grid, and calculated the average height from the best fitting plane for each cell. We then applied an interpolation to compute a compensation for every point of the sample's point cloud. The coarse grid should capture the global distortion of the fabric from the best fitting plane, if there is any. The expectation is that the coarse grid will capture an area where the average deviation from the best fitting plane should be zero, since we expect the samples to be near-planar, and with repetitive surface features. This compensation would not be necessary if the samples were rigid.

To eliminate the waviness due to the errors in the radiometric calibration (visible in Figure 4a), we computed the average height of each column along the y-axis, and compensated the height profile accordingly. This is possible because our projected patterns' encoding is along the x-axis, thus the error patterns propagate exclusively along the x-axis. This compensation assumes that the mean deviation from the best fitting plane of a column of points along the y-axis is zero, and a possible height bias might exist due to the errors in the radiometric calibration. More concretely for our measured samples, the compensation assumes that the weaving patterns are periodic, and that there are enough periods along the y-axis in the fabric that the mean deviation is zero. This claim appeared to be correct, because the computed compensations, a signal formed by the mean value along the y-axis for each x-coordinate, of all samples had very similar frequencies, and they were also very close to the one for the flat surface.

Figure 4b, shows the effects when both compensations are applied. The surface roughness of 1.61 μm may be due to the surface not being perfectly flat, and/or having some roughness.

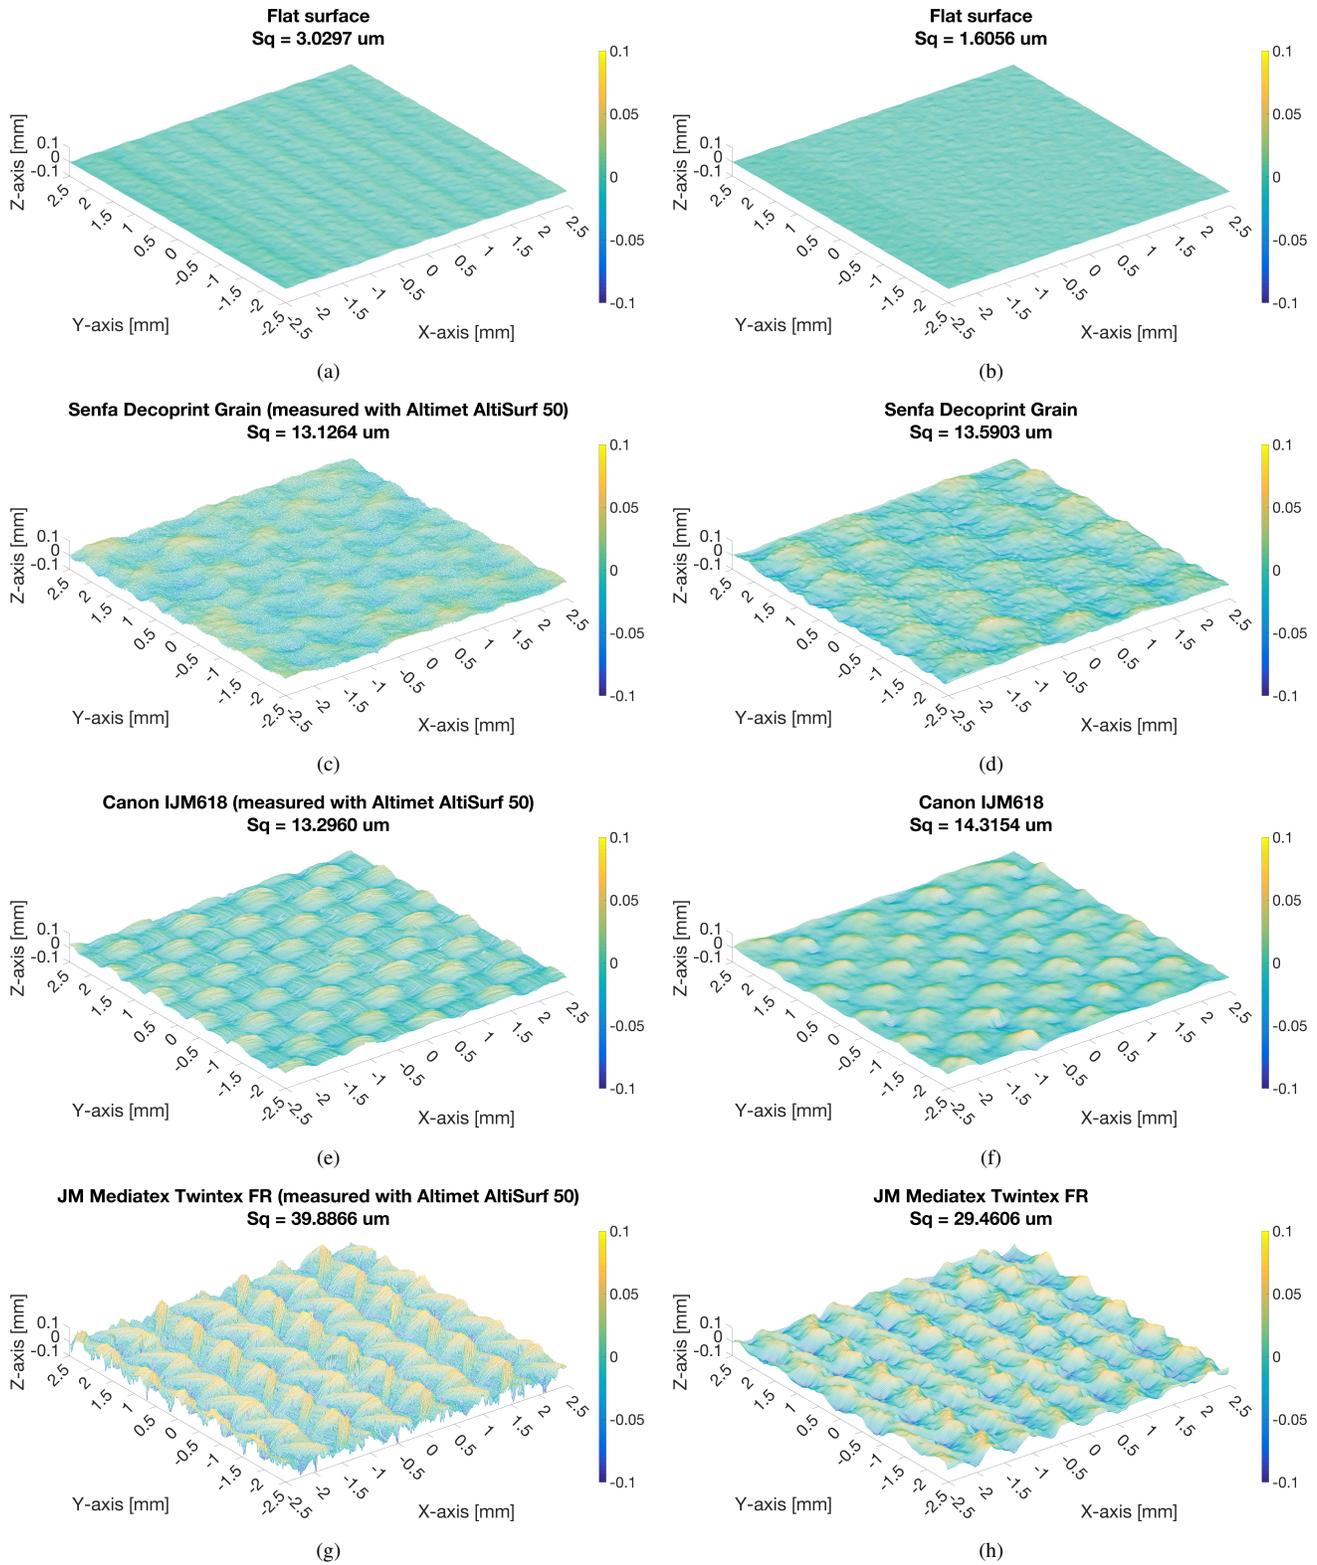


Figure 4: Surface structure of scanned materials. The colormap on the right side color codes the height of the surface points. The height is represented as the deviation from the best fitting plane to the scanned point cloud. All units are mm, except in the title, where they are μm . From the second row, the left subfigure shows the point cloud computed by the Altimet AltiSurf 50 profilometer, after we apply our simple compensation for the non-planarity of the sample. The right subfigure shows the point cloud computed with our structured light setup, after we apply the compensation for non-planarity of the sample and for the errors due to the radiometric calibration.

Experiments

To test the accuracy of our structured light scanner, we re-constructed the surface structure of printing textiles. The measured fabrics have near planar surfaces, and are white since they are used as print media. The representative set of five printing textiles was characterized by different weave frequencies and feature heights. The surface features have periods ranging from 0.42 mm to 1 mm, and maximum heights of the surface from 81 μm to 247 μm . Table 1 summarizes the main characteristics of the measured textiles. We compared our surface reconstructions qualitatively and quantitatively with those from a profilometer. We compared scans of 5 mm \times 5 mm patches.

Ground truth profilometer measurements

Ground truth measurements were performed with an Altimet AltiSurf 50 profilometer. It is based on the confocal principle. The x-, y- and z-axis steps were 10 μm , 10 μm and 0.156 μm , respectively. The measured area was 5 mm \times 5 mm, yielding 501 \times 501 data points, or 251'001 data points in total.

The surface structure of the Canon IJM 618 sample from Figure 1 can be seen in Figure 4e. The resolution in every axis is very fine, such that we can see the fibers that make up the threads. That is already a very fine scale, and we are interested in the macroscopic roughness of the surface.

The surface structure of the JM Mediatex Twintex FR is shown in Figure 4g. Here we can see that although the scan is fairly consistent, anomalies are visible as isolated peaks or ridges. These are errors in the scan, and are mostly likely the result of indirect illumination effects, or specular reflections. This shows that indirect illumination is a general problem, which affects both measurement techniques.

Structured light measurements

The resolution of our scans is much lower in all 3 dimensions. The point clouds contain around 11'450 data points for a grid of approximately 107 \times 107 data points, which yields a step of approximately 46.7 μm in the x- and y-axes. The theoretical step in the z-axis is 51.5 μm . We computed the roughness values on the native point clouds, and we then upsampled our point clouds to the same 501 \times 501 point grid for display.

To lower the impact of surface variability, we scanned the same physical samples with both devices, at approximately the same locations. Since the scans with the SL setup are wider, we manually cropped the point clouds to further align them.

Results

The surface scans for three samples from Table 1 are shown in Figure 4. Starting from the second row, the scans from the profilometer are shown on the left-hand side, and from our SL

Table 1: The characteristics of the textiles that we scanned. All textiles are woven, except for Mediatex Twintex, which is knitted.

Supplier	Product name	Weight [g/m ²]	Thickness [mm]	Period [mm]
Senfa	Decoprint Grain	375	0.45	1
Canon	IJM 617	310	0.39	0.55
Canon	IJM 618	200	0.27	0.6
Senfa	Decoprint Night	320	0.33	0.42
JM	Mediatex Twintex FR	380	0.61	0.625

setup on the right-hand side. The height values are computed from the best fitting plane, and are therefore both positive and negative.

In the second row of Figure 4 we can see the surface of the Senfa Decoprint Grain. Both scans reveal the structure of the weave with distinct zig-zag peaks. Our scan is rougher and somewhat noisier, and registers a slightly higher roughness value of 13.59 μm , as compared to 13.13 μm . The difference can be caused by several factors, including slight warping of the sample, non-equal surface structure due to the weave or coating, different number of weave periods due to limited scan areas, or simply errors in the measurement. Nevertheless, the roughness values are very similar, with a difference of 3.5 %.

The Canon IJM 618 sample shows a very regular weave, which can be seen in the third row of Figure 4. Here we can clearly see the higher resolution of the profilometer. The structure is correct in our scan, as both the higher and the lower threads can be distinguished as yellow and faint green peaks, respectively. We also see some anomalies in the valleys where four threads meet, likely caused by shadowing, specular reflections or interreflections. The roughness that we compute from our scan is 14.31 μm , and from the AltiSurf scan it is 13.3 μm . This difference is likely caused by the anomalies in our scan, and it amounts to 7.59 %.

The final row of Figure 4 shows the surface structure of the JM Mediatex Twintex FR. This textile is knitted, and features the roughest surface, with the greatest height range. The two scans now show obvious differences. Our scan manages to recover the correct global structure, and the correct structure frequency, but fails to recover the finer structures that comprise each loop. The deviation between the roughness values amounts to -26.15 %.

Discussion

A summary of the roughness results can be seen in Table 2. After a precise calibration of our structured light setup, we managed to achieve comparable roughness values to those of the profilometer, with 3.5 % to 8.26 % difference for the materials that we could scan reliably. Such errors could arise from global warping in the shape of the samples, and from uneven weaving and coating. The materials that we could scan reliably were the Senfa Decoprint Grain, Canon IJM 617, and Canon IJM 618, which had weave periods of 1 mm, 0.55 mm, and 0.6 mm, and a maximum surface heights of 122 μm , 144 μm , and 81 μm , respectively.

We also showed the limitations of our structured-light setup, which appeared when scanning the Senfa Decoprint Night, and the JM Mediatex Twintex FR textiles. The Senfa Decoprint Night is a textile that has a high frequency weave (period of 0.42 mm),

Table 2: A summary of the roughness results. The roughness computed from the 5 mm \times 5 mm point clouds produced with the Altimet AltiSurf 50 and with our structured light setup are compared, and the difference is shown in the *AltiSurf diff.* column. The units of the surface roughness, S_q , and the maximal height of the surface, S_z , are μm .

Product name	Altimet AltiSurf 50		SL setup	
	S_z	S_q	S_q	AltiSurf diff.
Senfa Decoprint Grain	122	13.13	13.59	3.5 %
Canon IJM 617	144	14.65	13.44	-8.26 %
Canon IJM 618	81	13.3	14.31	7.59 %
Senfa Decoprint Night	100	15.98	12.63	-20.96 %
JM Mediatex Twintex FR	247	39.89	29.46	-26.15 %

with deep crevasses between the threads (maximal surface height of 100 μm). The JM Mediatex Twintex FR textile on the other hand has a lower frequency knitted pattern (period of 0.625 mm in the x-axis, and 0.9 mm in the y-axis), but very prominent surface features (maximal surface height of 247 μm).

For these materials we generally need higher resolution scans in order to capture the finer surface details. Although in theory we have adequate resolution, the effective resolution limit was imposed by the camera lens. Furthermore, the main issue was not in the planar XY-axes resolution, rather in the depth Z-axis resolution. To increase the effective resolution, we will need a higher quality and/or a higher focal length lens for the camera, and, optionally, for the projector to gain better signal-to-noise levels. This will increase the cost of the SL setup only slightly relative to the cost of the profilometer. Additionally, increasing the baseline between the projector and the camera will increase the depth resolution, however, it might also cause issues given the limited depth of field of the devices.

With the current SL setup we can scan areas of up to 150 mm \times 100 mm. This is very advantageous, as it allows us to scan multiple materials at once, which dramatically improves acquisition speeds. It also allows us to scan surfaces with larger periods, whereas with the profilometer we are constrained to surfaces with relatively small periods (maximum of 2.5 mm if the scanning area is 5 mm \times 5 mm). Both systems have a similar acquisition time for a 5 mm \times 5 mm patch, however, the SL setup can scan the much larger areas without increasing the acquisition time. To significantly extend the scanned area with the profilometer we have to resort to stitching, which is prone to errors.

Conclusion

The accurate characterization of surface roughness is important for creating realistic reconstructions of physical objects and for simulating real materials in virtual environments. While surface roughness is generally measured by optical profilometers, they are expensive instruments and have a fairly limited scanning area. In this work, we evaluated the feasibility of an inexpensive SL setup to be used for computing a reliable surface depth map and roughness. Our setup offered wider scanning area and/or faster scanning than the profilometer, at an order of magnitude lower cost, and it was built by readily available off-the-shelf components. However, it comes at the expense of lower resolution.

The SL scanner achieved comparable results to the profilometer on samples with moderately complex surfaces. Higher resolution scans to capture finer surface details can be acquired by using a higher-quality lens and/or camera, at a small increase in the price. Although this analysis was done with textiles, the practical considerations we propose and evaluate are easily extendable to scanning any kind of near-planar surface.

This exploration was to show that printing textile's surface roughness can be acquired with a SL scanner, where the actual system implementation will depend on the desired accuracy, speed, and price. Knowledge of the surface structure could allow for devising color prediction models for macroscopic-level features that can incorporate shadowing and local interreflections.

Acknowledgments

We would like to thank Marine Page for providing us with the Altimet AltiSurf 50 measurements of our samples.

References

- [1] Phan Van Song, Théo and Andraud, Christine and Ortiz Segovia, Maria V. Spectral predictions of rough ink layers using a four-flux model. *Color and Imaging Conference (CIC)*, pg. 251 (2017).
- [2] Saint-Pierre, Dorian and Deeb, Rada and Muselet, Damien and Simonot, Lionel and Hébert, Mathieu. Light interreflections and shadowing effects in a Lambertian V-cavity under diffuse illumination. *Electronic Imaging (EI)*, pg. 166 (2018).
- [3] Saint-Pierre, Dorian and Chavel, Pierre and Simonot, Lionel and Hébert, Mathieu. Angular reflectance model for ridged specular surfaces, with comprehensive calculation of inter-reflections and polarization. *Journal of the Optical Society of America A*, 36, 11 (2019).
- [4] Saint-Pierre, Dorian and Simonot, Lionel and Hébert, Mathieu. Reflectance computation for a specular only V-cavity. *International Workshop on Computational Color Imaging*, pg. 289 (2019).
- [5] Salvi, Joaquim and Fernandez, Sergio and Pribanic, Tomislav and Llado, Xavier. A state of the art in structured light patterns for surface profilometry. *Pattern Recognition*, 43, 8 (2010).
- [6] Gupta, Mohit and Agrawal, Amit and Veeraraghavan, Ashok and Narasimhan, Srinivasa G. A practical approach to 3D scanning in the presence of interreflections, subsurface scattering and defocus. *International Journal of Computer Vision*, 102, 1-3 (2013).
- [7] Gupta, Mohit and Nayar, Shree K. Micro phase shifting. *IEEE Conference on Computer Vision and Pattern Recognition*, pg. 813 (2012).
- [8] Chen, Tongbo and Seidel, Hans-Peter and Lensch, Hendrik PA. Modulated phase-shifting for 3D scanning. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pg. 1 (2008).
- [9] Su, Xianyu and Zhang, Qican. Dynamic 3-D shape measurement method: a review. *Optics and Lasers in Engineering*, 48, 2 (2010).
- [10] Blais, François and Picard, Michel and Godin, Guy. Accurate 3D acquisition of freely moving objects. *Proc. 2nd International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT)*, pg. 422 (2004).
- [11] Levoy, Marc et al. The digital Michelangelo project: 3D scanning of large statues. *Proc. of the 27th annual conference on Computer graphics and interactive techniques*, pg. 131 (2000).
- [12] Shahpaski, Marjan, Luis Ricardo Sapaico, Gaspard Chevassus, and Sabine Süssstrunk. Simultaneous geometric and radiometric calibration of a projector-camera pair. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pg. 4885 (2017).
- [13] ISO25178-2:2012(en), Geometrical product specifications (GPS) – Surface texture: Areal – Part 2: Terms, definitions and surface texture parameters, International Organization for Standardization (2012).

Author Biography

Marjan Shahpaski obtained his PhD in Computer Science from the École Polytechnique Fédérale de Lausanne (EPFL, Switzerland) in 2020. His research activity is focused on 3D reconstruction with structured-light, and the measurement of BTDF with computational methods.

Luis Ricardo Sapaico obtained his Ph.D (Computer Science) from the Tokyo Institute of Technology in 2011. He currently works as a Senior Imaging Processing Engineer at DxO Labs in France. His main research interests are in color science, computer vision and image processing.

Sabine Süssstrunk is a full professor and Director of the Images and Visual Representation Laboratory (IVRL) in the School of Computer and Communication Sciences at EPFL. Her main research areas are in computational imaging, computational photography, color image processing and computer vision, multimedia, and computational image quality and aesthetics. She is a Fellow of IEEE and IS&T.

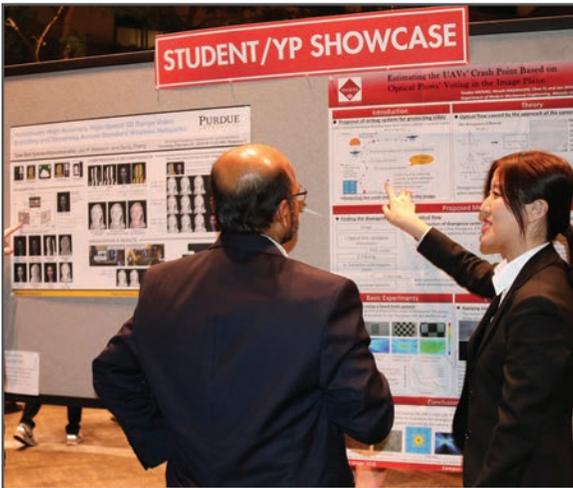
JOIN US AT THE NEXT EI!

IS&T International Symposium on

Electronic Imaging

SCIENCE AND TECHNOLOGY

Imaging across applications . . . Where industry and academia meet!



- **SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •**
- **INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •**

www.electronicimaging.org

