Microstructure Control in 3D Printing with Digital Light Processing

A. Luongo, V. Falster, M. B. Doest, M. M. Ribo, E. R. Eiriksson, D. B. Pedersen, and J. R. Frisvad

Technical University of Denmark

Figure 1: *Hemispheres and bunnies with smooth and rough surfaces, and flat samples (smileys and QR code) with spatially varying anisotropic reflectance. The scene is observed from two different directions to exhibit the anisotropy. The sun is used as a directional light source. Each item was printed in a one-step process using the presented technique.*

Abstract

Digital light processing stereolithography is a promising technique for 3D printing. However, it offers little control over the surface appearance of the printed object. The printing process is typically layered, which leads to aliasing artifacts that affect surface appearance. An antialiasing option is to use grayscale pixel values in the layer images that we supply to the printer. This enables a kind of subvoxel growth control. We explore this concept and use it for editing surface microstructure. In other words, we modify the surface appearance of a printed object by applying a grayscale pattern to the surface voxels before sending the cross-sectional layer images to the printer. We find that a smooth noise function is an excellent tool for varying surface roughness and for breaking the regularities that lead to aliasing. Conversely, we also present examples that introduce regularities to produce controlled anisotropic surface appearance. Our hope is that subvoxel growth control in stereolithography can lead 3D printing toward customizable surface appearance. The printing process adds what we call ground noise to the printed result. We suggest a way of modeling this ground noise to provide users with a tool for estimating a printer's ability to control surface reflectance.

Keywords: 3D printing, additive manufacturing, appearance, BRDF, fabrication, reflectance, surface roughness

CCS Concepts
•*Computing methodologies* → *Reflectance modeling;*

1. Introduction

While 3D printers can often print geometric features in high quality, they lack the ability to control surface appearance by modifying roughness and reflectance properties. The ability to produce models with region-specific surface properties is crucial for artists and developers to properly design the appearance of a part. In the prototyping stage of product development, additive manufacturing is commonly used to produce parts in order to evaluate the final aesthetics of a product. For a part to look like a designed digital model, however, additional surface processing is often required. We propose a method for better control of printed surface properties to enable customization of the final appearance of a printed part.

The printing technology we work with is based on photopolymerization, which refers to the curing of liquid photo-reactive resins (photopolymers) using light. The light is usually in the ultraviolet range of wavelengths. This process is used for 3D printing with stereolithography, where a light source selectively illuminates a photopolymer to produce a solid object with a user-defined shape. If a digital light processing (DLP) projector is used as the source, the technique is referred to as DLP printing. In this case, we can specify the user-defined shape as a volume. The photopolymer is contained inside a vat and at each step a building platform is raised or lowered, depending on the setup of the DLP printer, in order to expose only a thin layer of liquid photopolymer to the projector. Each slice of the volume is then projected onto the photopolymer to produce a layer of the 3D print consisting of solidified polymer in all the pixels of the slice with value one (white voxels). In the context of DLP printing, we provide an investigation of the use of grayscale voxel values to control surface microstructure. Figure 1 displays some of our results.

Commercial 3D printers improve continually in terms of the resolution and the complexity of the geometries that can be printed. Nevertheless, the final surface appearance is typically controlled through the use of different print materials, deposition of different inks, and postprocessing of the surface. Samples with different reflectance properties can be printed directly in a one-step process, but the microstructure of the surface is then defined by the employed 3D printing technique. For example, in a material-extrusion based printer, the sample surface will exhibit layering artifacts, while a powder-based print will have a grainy surface. A DLP printer can produce smooth flat surfaces, but on vertical and curved surfaces it will produce staircase artifacts. Even if the layers are so thin that we cannot see them with our naked eyes, the layered structure still produces moiré patterns and reflects light with a glean at certain angles. To get a different appearance, such as reflective or matte, the printer must produce a more detailed geometry with smaller features. The resolution of the 3D printer typically sets the limitation and prevents us from obtaining the desired result.

In this work, we show how the use of grayscale patterns greatly increases the capabilities of a DLP printer, and how it enables us to print microfeatures and patterns on the surface of a sample in a one step process without changing the macroscopic geometry of the printed part. By using this technique, we can modify the roughness and surface appearance of a print without changing materials or applying postprocessing to the sample.

2. Related Work

Fabrication of microgeometry to obtain custom surface reflectance was pioneered by Weyrich et al. [WPMR09]. They point at many interesting applications and fabricate custom microgeometry using a micro milling approach. In a 3D printing context, a 5-axis micro milling machine can produce free-form surfaces with fairly small features. However, due to the kinematics of the milling process, it is difficult to control the surface roughness [ABRK17]. In another early technique, Matusik et al. [MAG*09] use different inks in different halftoning patterns to print a surface with spatially-varying reflectance properties. This technique is however restricted to printing on planar surfaces, and the microstructure that can be printed depends on the reflectance properties of the employed inks.

Different ways of extending these early techniques have been tested. Malzbender et al. [MSS*12] print on a paper with a static microstructure and let the selective depositing of ink control the surface reflectance. More generally, Baar et al. [BBS15] study the link between variation of print parameters and local control of the gloss appearance in a printout. However, they only consider printing of flat images. Lan et al. [LDPT13] use a 3D printer based on material jetting to produce patches with oriented facets and then coat them with glossy inks using a flatbed UV printer. However, the facets in the patches are visible to the naked eye (140 μm by 140 μm) and the fabrication process requires two steps. The use of the flatbed printer puts a constraint on the curvature of the surface that the inks can be applied to. Thus, when applying this method to a 3D surface, the object is divided into several parts that are stitched together in a post-process after inks have been deposited using the flatbed printer. Another approach requiring two steps is by Rouiller et al. [RBK*13]. They use another 3D printer based on material jetting to print microfacetted transparent domes that they stick onto a colored model, which was 3D printed using a powder bed printer. In this way, each dome modifies the reflectance in the local area where it is attached. As opposed to these techniques, we present a one-step approach where the fabrication of surface microstructure is integrated into the 3D printing process. The material jetting printers (PolyJet technology) employed in this previous work can only print binary voxels (material or not). Consequently, they do not support the grayscale voxel values that we can use when employing a vat polymerization based DLP printer.

Levin et al. [LGX*13] present a technique for printing microstructure small enough to create reflectance functions based on wave interference effects. Their technique is based on photolithography, which is a very precise but also very costly process that requires a special wafer coated by photoresist. Photolithography is currently not available as a 3D printing technique.

Pereira et al. [PLMR17] propose an entirely different approach, where magnetic microflakes are embedded into a photopolymer and controlled during printing using electromagnets. While they obtain interesting results, the magnetic flakes are significantly harder to control than our surface microstructure based on grayscale values in the projected cross-sectional images.

Use of grayscale values in DLP printing is not entirely new. Mostafa et al. [MQM17] explore to what extent grayscale values can improve the dimensional accuracy of an Autodesk Ember printer. This use case has also been investigated internally

at Autodesk [Gre16], where they improve printing fidelity using grayscale values computed with antialiasing techniques. The work presented by Greene [Gre16] is the work most closely related to ours. Greene even mentions in passing that random noise can be used to break moiré patterns and to produce a matte surface. However, to the best of our knowledge, we are the first to more carefully modify surface roughness and reflectance properties of 3D printed objects by applying grayscale patterns across surface voxels.

Some work has been done to control the subsurface scattering and absorption properties of fabricated objects [DWP*10, HFM*10,PRJ*13,ESZ*17]. In our case, these properties are determined by the photopolymer selected for the print job. We consider it an interesting challenge for future work to investigate ways of controlling the scattering properties of a photopolymer.

3. Subvoxel Growth

The resolution of DLP printing is typically in the range from 15 to 100 μm [LCR*17]. It depends on the quality and pixel resolution of the digital micromirror device (DMD) chip of the projector and on the step-precision of the building plate. It is possible to use grayscale images as input for the projector to obtain subvoxel accuracy [Gre16,MQM17]. The principle behind this idea is that the solidification process of the resin depends on the amount of UV light received, and this amount can be changed by varying either the period of time for which an image is projected (exposure time) or the intensity of the light. With grayscale values as input for the projector, we vary the intensity and thus control the growth of each voxel. This approach can be used to produce very small features and patterns on the surface of a 3D printed sample. If applied properly, the grayscale values modify the microscopic surface properties of an object without modifying its macroscopic geometrical structure.

3.1. Subvoxel control

The relation between grayscale values and voxel growth is crucial if we are to print an arbitrary microscopic pattern with high accuracy. If we project an even slope of all the grayscale values (pixel intensity values from black to white), we would ideally see the same even slope being printed. If this were the case, voxels would grow proportionally with the grayscale values.

Unfortunately, the photopolymerization is initiated only when a critical energy level is reached, and the cure depth then follows a logarithmic curve with increasing energy [Jac92, LPA01, Ben17]. Thus, we can determine the relationship between pixel intensity and voxel growth. With τ denoting the thickness of a print layer, the cure depth and thus the voxel growth height is

$$\tau f(I) = \begin{cases} \alpha + \beta \log(I - \gamma), & \text{for } I > e^{-\alpha/\beta} + \gamma, \\ 0 & , & \text{for } I \leq e^{-\alpha/\beta} + \gamma, \end{cases} \quad (1)$$

where I is the pixel intensity, and α, β, and γ are parameters that need to be fitted for a particular photopolymer.

Through inversion of the function f, we obtain a mapping to the proportionality relation, which significantly eases control of the
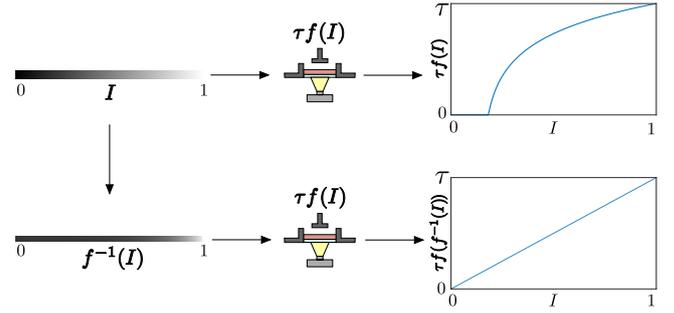

Figure 2: *Inversion of non-linear voxel growth to have printed voxel height proportional to grayscale pixel intensity, I.*
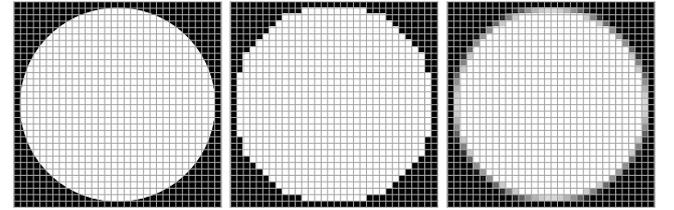

Figure 3: *A desired circular print layer geometry (left), its rasterization according to the resolution of the projector (middle), and the same layer with grayscale values for antialiasing (right).*

voxel growth. We have

$$f^{-1}(I) = \begin{cases} e^{\frac{\tau I - \alpha}{\beta}} + \gamma, & \text{for } I > 0, \\ 0 & , & \text{for } I = 0, \end{cases} \quad (2)$$

and using $f^{-1}(I)$ as the grayscale values of the pixels in a projection, the printer prints voxels of height τI. This is illustrated in Figure 2. Greene [Gre16] presented a similar result, but they used a quadratic f function while suggesting that a logarithmic function seems a better choice. We found the right f function by considering the photopolymerization cure depth.

3.2. Grayscale patterns

The ability to control voxel growth using a linear scale of grayscale values enables us not only to improve fidelity and reduce aliasing artifacts, as demonstrated by Greene [Gre16], it also enables us to print smooth microfeatures in a single layer and thereby modify the reflectance properties of the surface.

3.2.1. Antialiasing

When printing an object, we have to slice the geometry to generate an image for each layer. Slices are obtained by rasterizing the geometry, and if no measures are taken, spatial aliasing will be present along edges of the layers in the form of pixelated boundaries, see Figure 3. Grayscale values based on supersampling (in all three dimensions) can be used to counteract this effect and produce a smoother surface [Gre16]. However, this is not enough to completely remove staircase artifacts in a surface. These artifacts lead to visible reflectance anisotropy and moiré patterns.
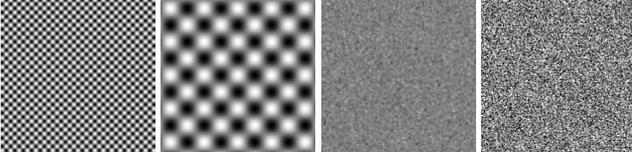
Figure 4: *Sinusoidal patterns with different wavelengths (leftmost with $\lambda_u = \lambda_v = 100\ \mu m$ and middle left with $\lambda_u = \lambda_v = 400\ \mu m$). Sparse convolution noise with different amplitude and frequency factors (middle left with $A = 0.625$ and $B = 16$ and rightmost with $A = 3$ and $B = 32$). Both types of patterns are useful for controlling roughness. Due to its irregularities, the noise function is also useful for antialiasing.*

Figure 5: *Sinusoidal patterns with different wavelengths along the two axes (left with $\lambda_u = 50\ \mu m$ and $\lambda_v = 200\ \mu m$ and middle with $\lambda_u = 50\ \mu m$ and $\lambda_v = 400\ \mu m$) and sequences of parallel ridges (right). These 2D patterns are useful for printing anisotropic surface roughness and reflectance contrast.*

Greene [Gre16] suggests the use of Gaussian smoothing that produces grayscale values in a thick band around the edges to further reduce aliasing. A broad Gaussian smoothing is however likely to also smoothen the macroscopic geometry of the object if the surface is not spherical. This would compromise object fidelity. Another suggestion by Autodesk [Gre16] is to add random noise to all the grayscale values. This breaks the moiré patterns, but it also leads to a matte surface. In other words, when printing in 3D, existing work leaves us with the choice of an aliased or a matte surface appearance. In the following, we demonstrate how a smooth low-amplitude solid noise function can be used to break moiré patterns while retaining surface smoothness. In addition, we explore the use of procedural methods for inserting grayscale values in surface voxels to control the surface microstructure.

3.2.2. Reflectance properties

The roughness of a surface is given by its microstructure. The features are so small that they are only individually visible at the microscale, but they affect the macroscopic surface appearance. Our goal is to apply grayscale patterns along the surface of an object to print surfaces with different roughnesses, going from smooth to almost diffuse, and also to print spatially varying anisotropic reflectance properties.

As rough surfaces are characterized by having a distribution of microfacet normals pointing in various directions, one way to obtain isotropic roughness is to use a curved surface [TR75]. We therefore test a grayscale pattern with surface voxel values set according to a (2D) sinusoidal function running along the surface. The function is

$$I(u,v) = \frac{1}{2}\sin\left(\frac{2\pi}{\lambda_u}u\right)\sin\left(\frac{2\pi}{\lambda_v}v\right) + \frac{1}{2}, \qquad (3)$$

where u and v are parameters measuring physical length in a uniform parametrization of the surface, so that λ_u and λ_v represent the wavelengths along these two dimensions. The wavelengths of the sinusoid then control the roughness of the surface, see Figure 4. This kind of grayscale pattern will generate a periodic sequence of micro-cavities and micro-bumps on the 3D printed object, and this structure will produce a rough surface when the frequency of the sinusoid is high (more bumps and cavities), and a smooth surface when the frequency is low.

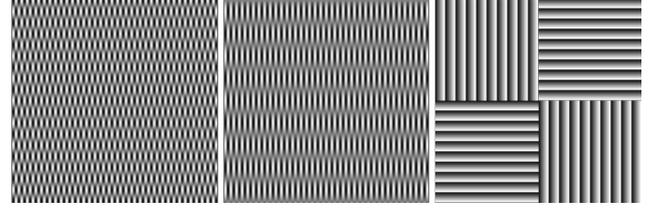An issue with the sinusoidal surface is its regularity. Since the

function is regular, it does not entirely prevent the aliasing problems due to layered printing. We therefore decided to also use a smooth noise function, as it is irregular but produces a similar effect in terms of the microfacet normal distribution. To avoid the grid-aligned artifacts seen in Perlin noises [Per85, Per02, MSRG12], we employ a solid sparse convolution noise (Appendix A). The difference between sinusoidal patterns and noise slices is illustrated in Figure 4. By controlling the frequency and amplitude of the noise function, we are able to obtain smooth and rough surfaces with very few staircase artifacts (hemispheres and bunnies in Figure 1).

We print anisotropic reflectance properties using a 2D sinusoidal function with different frequencies along the two axes, or a sequence of parallel ridges, as described by Luongo et al. [LFD*17], see Figure 5. These patterns are useful for producing anisotropic reflectance contrast (smileys and QR code in Figure 1). While we only test these patterns on a 2D surface, they could be texture mapped onto a curved surface to obtain a 3D surface with anisotropic reflectance. Texture coordinates for a given model can be generated using a 3D modeling tool such as Maya or Blender. If we want to avoid this task, a solid noise function (Appendix A) can be stretched along the tangent space of a 3D surface using line integral convolution [BSH97]. To obtain a consistently oriented tangent space without use of texture coordinates, we can use the function for building an orthonormal basis by Frisvad [Fri12].

3.3. Assessing reflectance controllability

We assess how well our method controls the reflectance properties of a printed surface using two different approaches. For anisotropic microstructure, we predict the expected contrasts in light reflection when the surface is illuminated from different directions. We do this by rendering the surface appearance due to the varying microstructures using analytic BRDF models derived for those specific microstructures. For the ridged structure in Figure 5, we use the model presented by Luongo et al. [LFD*17]. For the anisotropic sinusoidal patterns, we derived a new model, which is described in Appendix B. We then qualitatively compare the rendered images with photographs of printed samples. The comparison is not in terms of photorealism, but in terms of contrast in light reflection. For irregular noise-based microstructure, such as the patterns generated using sparse convolution noise (Figure 4), we compute the corresponding bidirectional reflectance distribution function (BRDF) using a path tracer. We path trace a representative patch
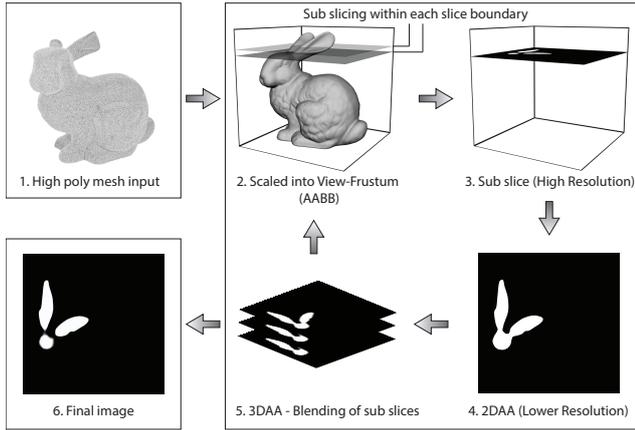
Figure 6: *Mesh slicing pipeline based on rasterization. Used for generating cross-sectional layer images for the DLP projector.*

Figure 7: *Schematic of the homebuilt DLP printer.*

of the noise used as grayscale input for the printer. Measuring the printed microstructure using a microscope, we can then compare the BRDF of the desired microstructure with the BRDF of a corresponding printed microstructure.

Interestingly, Ribardière et al. [RBSM19] provide an algorithm for generating height fields with microstructure corresponding to the normal distributions used in popular analytic microfacet BRDF models [WMLT07]. These height fields can be used as grayscale maps in our printing process and would allow for assessments similar to ours but with the commonly used BRDFs. We leave a complementary investigation of this kind for future work.

3.4. Mesh slicing

To generate antialiased cross-sectional layer images for the DLP projector, we have tested two different approaches: one based on rasterization and one based on ray tracing, both running on the graphics processing unit (GPU). Our rasterization procedure is illustrated in Figure 6, and the different steps are described in the following paragraphs.

In both approaches, a closed triangle mesh is provided as input (step 1) and the print volume is represented by the view frustum of an orthographic camera placed above the mesh looking downwards. The background color is set to black and the frame buffer resolution is set to the projector resolution. The latter ensures that each pixel of a generated layer image corresponds to a voxel with physical dimensions as described in Section 4. To determine the number of slices that we need, we calculate the object height in number of voxels using the desired physical height of the printed object.

In rasterization, we slice the mesh by moving the near cutting plane of the camera through the print volume in steps of the print layer thickness (step 2). The far cutting plane is placed at the end of the print volume and depth testing is enabled. For frontfacing triangles, the color is based on a procedural texture (sinusoid or noise) but the fragment is only rendered to the color buffer if it is within the current layer. Frontfacing triangles behind the current layer are only rendered to the depth buffer. Backfacing triangles passing the
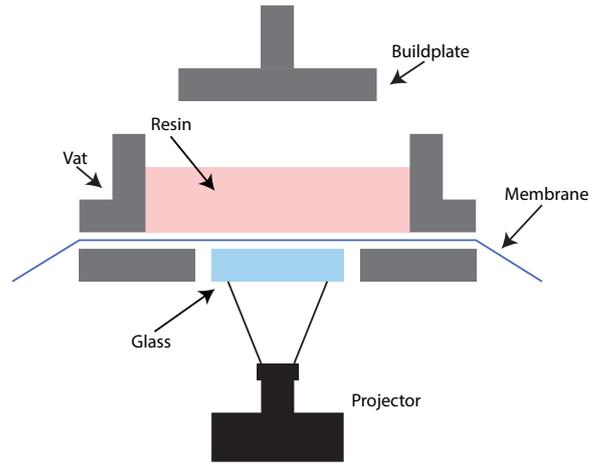
depth test are rendered with a flat white color. For each slice, we generate a number of subslices (step 3) to include supersampling in the depth dimension.

In ray tracing, we trace a ray from the image plane through all surfaces until it reaches the front surface of the current layer. The ray keeps a counter for each intersection, so that the counter is even when the ray is outside the object, odd when inside. A ray is then traced in the same direction from the front to the back of the layer. The fraction of the distance traveled by this ray that was also inside the object provides a grayscale value for antialiasing in the depth direction. Combining this with jitter sampling of the ray origin in the camera pixel, we obtain grayscale values incorporating full 3D antialiasing. As in rasterization, the grayscale value is modulated by a procedural texture when the ray going through the layer intersects a frontfacing triangle.

In rasterization, antialiasing requires more passes. To have 2D antialiasing in each slice, we use hardware supported full screen antialiasing with four samples in each pixel (4xFSAA). This is done in eight times higher resolution and downsampled to the projector resolution (step 4). The subslices are then blended into the same frame buffer (step 5) to produce one antialiased cross-sectional layer image for the printer (step 6).

4. Experiments

We run our experiments on a homebuilt bottom-up DLP printer, which is based on the work of Jørgensen [Jør15]. A schematic of the printer is in Figure 7. The photopolymer resin is inside the vat. The building platform starts at the bottom of the vat and moves upwards during the printing process. The step precision of the building platform is 1 μm, which enables us to print very thin layers. A transparent membrane is placed at the bottom of the vat in order to separate the photopolymer from the glass. This is done to facilitate the peeling effect and the release of the sample from the vat when the platform is raised [PZNH16].

The DLP projector we use is a LUXBEAM Rapid System by Visitech equipped with a DMD chipset of the DLP9000 family by

Texas Instruments. It has an array of 2560×1600 micro-mirrors and pixel pitch of 7.54 μm. The projector is placed underneath the vat and can be raised and lowered to focus it. We use a projection lens from Visitech with a magnification factor of $1.0\times$, yielding an image pixel pitch of 7.54 μm, or alternatively a lens with a factor of $2.0\times$ and pixel pitch of 15.08 μm.

According to the manufacturer, the projector is more stable for high values of the UV LED amplitude, but even low values of UV LED amplitude can overcure the photopolymer in our setup. This would ruin the quality of the prints, so we equipped the projector with two absorptive neutral density filters from Thorlabs. Each filter transmits 10% of the incoming light, so that the amount of light reaching the photopolymer is 1% of the light emitted by the projector. In this way, we can use higher values of UV LED amplitude for our prints, which means that we get a more stable behavior from the projector (less flickering, for example).

The photopolymer we use is Industrial Blend (red) resin from Fun To Do. In order to inspect and measure the properties of our prints we used an optical measuring device based on focus-variation, Infinite Focus by Alicona, which can produce high-quality 3D measurements of the surfaces and measure the surface roughness with nanometer precision.

After the printing process, the sample is cleaned with iso-propanol in an ultrasonic cleaner in order to remove any residual resin from the surface. We then do additional curing in a UV curing box to ensure that the sample has solidified properly, and to remove the risk of contamination when touching the sample.

Our setup enables us to print high resolution samples. However, the presence of the membrane, which mitigates peeling forces, is a source of some defects: when the membrane is installed on the glass, some wrinkles may be present and air can be trapped between the membrane and the glass causing the formation of bubbles. Such issues affect the final quality of the sample, where we sometimes observe bumps and scratches on the surface. Scratches start appearing as the membrane gets worn.

4.1. Parameter calibration

The photopolymer curing process is determined by the intensity of the projected UV light, by the exposure time, and by the amount of resin that we want to cure (layer thickness). All these parameters vary for different materials, and a calibration operation is required in order to find the optimal configuration for a certain setup.

Based on previous experiments performed on the same printer [Rib17], we decided to use a value of $\tau = 18$ μm for the layer thickness. This value is small enough to give us microfeatures, which can affect the reflectance properties of an object without being visible to the naked eye, and it is thick enough so that the features created with grayscale images are not overexposed.

To calibrate the projector intensity and exposure time, we created a calibration sample with the same pattern repeated 36 times on the top surface, see Figure 8. For each of these 36 patterns, we use a different value of intensity or exposure time. One out of the 36 patterns has a physical size of 1920×1920 μm^2 and consists of four black-and-white checkerboards with different scales for the
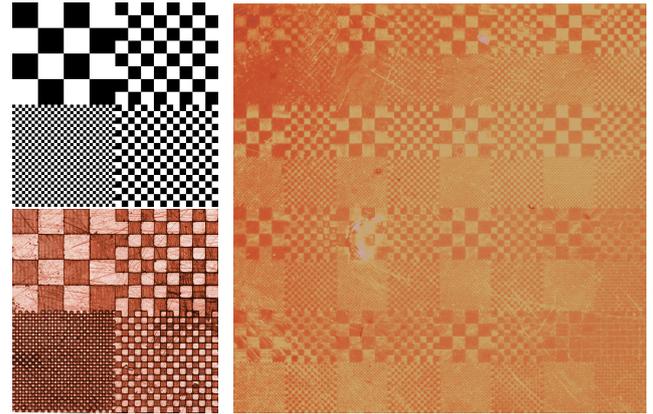

Figure 8: *Pattern used to calibrate projector parameters (top left) and microscope image of a printed pattern (bottom left). The pattern is composed of four black-and-white checkerboards at different scales, and it is repeated 36 times in a calibration sample. On the right, a microscope image with 16 of the 36 checkerboard pattern repetitions in a calibration sample.*

size of the squares. We first print a calibration sample with increasing UV LED amplitude for each pattern repetition while keeping the exposure time constant. The same experiment is then repeated with increasing exposure time while keeping the UV LED amplitude constant. A good combination of parameters is found when a pattern shows sharp features which are neither underexposed nor overcured. With this experiment, we found that for a layer thickness of $\tau = 18$ μm the optimal parameters of our setup are an UV LED amplitude of 230 and an exposure time of 3 seconds.

4.2. Voxel height measurements

As mentioned in Section 3.1, the relation between pixel intensity and growth of the corresponding voxel is logarithmic, Eq. 1. In order to apply our correction, Eq. 2, we need to find the values of the parameters α, β, and γ.

We printed several samples with a repeated linear grayscale gradient containing all the values from black to white, the upper left part of Figure 9 shows two examples. We then examined the samples with the Infinite Focus microscope and measured the surface with a vertical resolution of 0.4 μm. The collected data were used to find a fit for Eq. 1, see the lower left part of Figure 9, and we estimated the parameter values to be $\alpha = 17.71$ μm, $\beta = 10.24$ μm, and $\gamma = -0.01$. By having the same pattern repeated multiple times we got a better estimate and were able deal with some of the noise introduced by the printing process.

The corrected grayscale pattern and the corresponding printed samples are shown in the upper right part of Figure 9. The surface of the sample now looks more smooth and the resin solidifies everywhere on the surface. The blue plot in the lower right part of Figure 9 is a measurement of the surface height, while the red plot is the ideal linear behavior that we would like to have when printing with grayscale images. Even though the blue plot shows some irregularities, it proves that by applying Eq. 2 to our patterns we
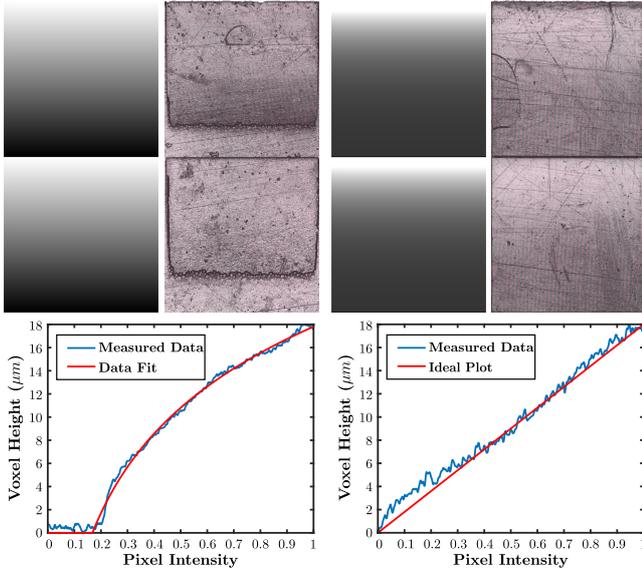
Figure 9: *Grayscale layer images and microscope images of printed results used for estimating α, β, and γ to control voxel growth (two repetitions). The linear gradient (left) is used for fitting to Eq. 1. The logarithmic gradient (right) is used for testing the linearity of the printed gradient after correction with Eq. 2.*

Table 1: *Average roughness measured as S_a and S_q for samples with sparse convolution noise applied using different amplitudes A and frequencies B.*

	$A = 0.625$	$A = 2$	$A = 3$	
S_a (μm)	2.21	3.30	5.53	$B = 16$
S_q (μm)	2.82	4.20	6.94	
S_a (μm)	2.90	4.49	7.95	$B = 32$
S_q (μm)	3.64	5.71	10.10	

obtain the desired geometry, and we therefore have the ability to control subvoxel-sized surface microstructure.

4.3. Roughness measurements

To verify that we can print surfaces with different roughnesses by applying sparse convolution noise with varying amplitude and frequency parameters (Appendix A), we printed several samples and measured their surface roughness with the microscope. The parameters used in this experiment and the corresponding results are in Table 1. These results show quantitatively that by increasing the amplitude *A* and the frequency *B* of the noise function the area roughness parameters S_a (arithmetic average height) and S_q (root mean square roughness) increase as well. Thus, we obtain a smoother surface if we apply a grayscale pattern with sparse convolution noise using lower values of *A* and *B*, and more diffuse-like surfaces if we use higher values of these two parameters.
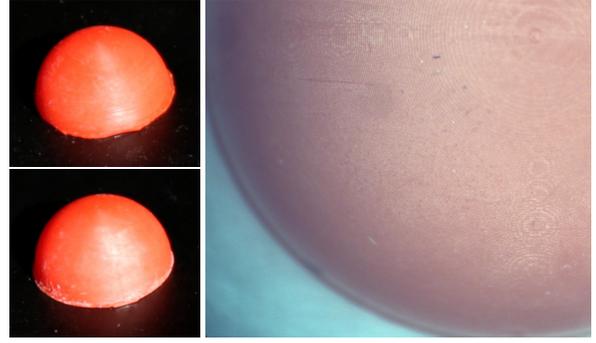

Figure 10: *Hemispheres printed with grayscale values calculated using supersampling. On the left, the hemispheres were printed using a 2× magnifying lens: one with supersampling only (top left) and one with both supersampling and sparse convolution noise (bottom left, parameters A = 0.625 and B = 32). On the right, the hemisphere was printed with supersampling and 1× magnifying lens. Even at a scale this small, moiré patterns are still visible when the surface is observed in a microscope.*

4.4. Antialiasing abilities of supersampling

As discussed by Greene [Gre16] and in Section 3.2.1, we can use supersampling to calculate grayscale values for spatial antialiasing during the slicing process. However, we find (as did Greene [Gre16]) that the surface still exhibits reflectance anisotropy and moiré patterns. The hemisphere in Figure 10 (top left) was printed using 2× magnifying lens and supersampling for antialiasing. Nevertheless, it still has an elongated highlight that we would only expect to see when the surface exhibits anisotropic reflectance [AS00]. Even if printed with 1× magnifying lens and supersampling, we still see staircases and moiré patterns when looking through a microscope (Figure 10, right). On the other hand, we find a smooth irregular noise function (like the one presented in Appendix A) useful for obtaining improved antialiasing and more isotropic reflectance properties. The hemisphere in Figure 10 (bottom left) includes sparse convolution noise with parameters $A = 0.625$ and $B = 32$. While this sample is not completely free of aliasing artifacts, it does exhibits a more rounded highlight and, thus, more isotropic reflectance properties. The same hemisphere is illuminated by a more directional source in Figure 1.

5. Results

Let us compare printed surface microstructure with the surface microstructure given as input grayscale values for the printing process. The first column of Figure 11 is examples of input noise at amplitudes $A = 0.625, 2, 3$ and the third column is examples of printed surface microstructure for input noise at the same amplitudes. It is clear that the printing process introduces additional noise, let us call it ground noise, caused by the membrane and the cleaning process. We can now use path tracing of a specular surface patch with geometry given by these height maps to calculate a corresponding BRDF lobe (second column of Figure 11). The input noises produce a highly specular lobe, so we also draw these using a logarithmic scale in Figure 12 to make their differences more easily
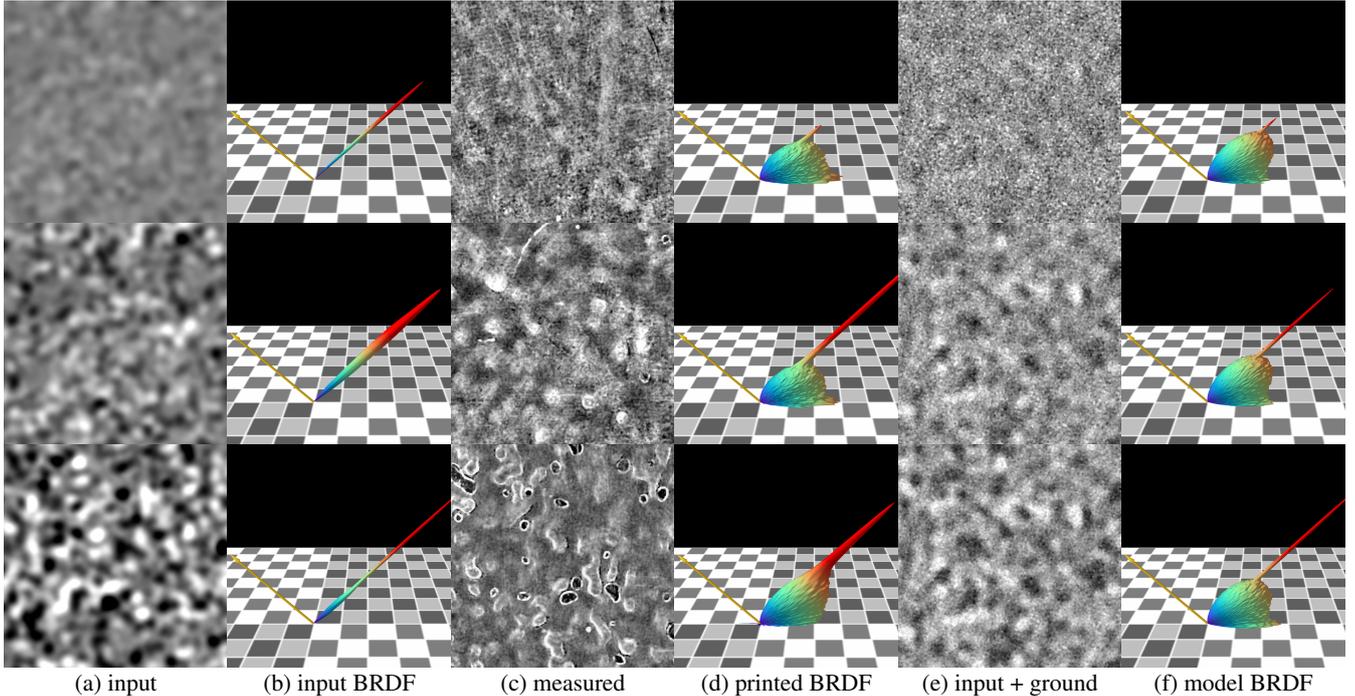
| (a) input | (b) input BRDF | (c) measured | (d) printed BRDF | (e) input + ground | (f) model BRDF |

Figure 11: *(a) Input grayscale noise values of amplitudes $A = 0.625, 2, 3$, (c) surface microstructure printed using input of the same amplitudes and measured using a microscope, (e) ground noise added to the input noise. (b, d, f) Lobe images showing the BRDF values for a 45 degrees angle of incidence. The lobes were computed using path tracing.*
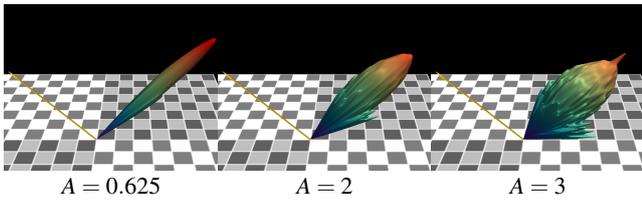

| $A = 0.625$ | $A = 2$ | $A = 3$ |

Figure 12: *Log transformed versions of the BRDF lobes based on the input noise values alone (second column of Figure 11).*

observable. We observe that the shape of the lobe broadens with increasing amplitude. The height maps obtained by imaging printed surfaces using the Infinite Focus microscope result in a much more broadly scattering lobe that we visualize in the fourth column of Figure 11. The reflectance properties of the input surfaces and the printed surfaces are so different that they are hard to compare. However, the results are important as we can use them to build a model of the printer's added ground noise.

Through inspection of the measured height maps and using the noise function in Appendix A, we manually found that the following function is a good model for our printer's ground noise:

$$\mathrm{ground}(\boldsymbol{x})$$
$$= \frac{2}{3}\mathrm{noise}\left(\frac{\boldsymbol{x}}{50\,\mu\mathrm{m}}\right) + \frac{1}{9}\mathrm{noise}\left(\frac{\boldsymbol{x}}{25\,\mu\mathrm{m}}\right) + \frac{1}{12}\mathrm{noise}\left(\frac{\boldsymbol{x}}{2\,\mu\mathrm{m}}\right).$$

We believe this is useful as an example if one were to build a similar model for the ground noise of another printer. Finding an expression for the ground noise of a printer is important as it models

the imprecision of the printing process. Since the printer adds noise similar to the ground noise to the input grayscale values, the ground noise function provides us with an outline of the printer's limitations in terms of reflectance control. If the printer is improved, we can repeat the experiment and see if the ground noise has diminished. To model the BRDF output of the printer, we add the ground noise to the input grayscale values and flatten the result a bit by clamping to include the membrane in the model. The fifth column of Figure 11 is examples of the surface microstructure estimated by this model, and the sixth column indicates that the resulting BRDF lobes come fairly close to the printed BRDF in the fourth column.

Figure 1 displays some of the visual effects enabled by our technique. It is remarkable that the rather small difference in the BRDFs that we estimated (Figure 11) produces a fairly obvious visual difference. In the following, we explore different techniques for printing surfaces with anisotropic reflection, and we demonstrate why the irregular noise function is important when printing 3D surfaces. Regarding the quality of antialiasing and the rate at which slices are generated, both techniques introduced in Section 3.4 perform similarly, and either one can be used to obtain the following results.

Figure 13 (top row) shows the grayscale patterns used for printing the smiley sample displayed in Figure 1. The figure also shows microscope images of the printed result (bottom row). We printed this sample with the $1\times$ lens to test how well we can print surfaces with anisotropic reflectance properties. In this example, we used the grayscale pattern for the last layer of the printing process only. We used the 2D sinusoid to generate the patterns in the main diagonal of Figure 13, with parameters $\lambda_u = 150\,\mu m$ and $\lambda_v = 50\,\mu m$
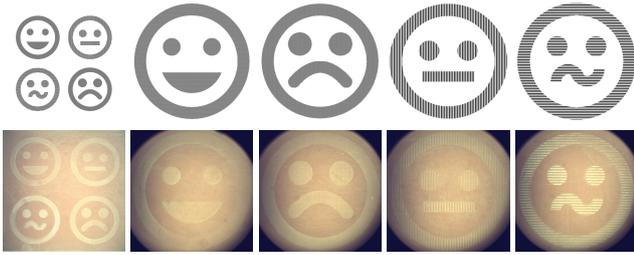
Figure 13: *Sample generated using two different anisotropic patterns with orthogonal orientation. The first two smileys have been printed with anisotropic sinusoidal patterns but with two different orientations. The last two have been printed with a ridged pattern with two different orientations.*
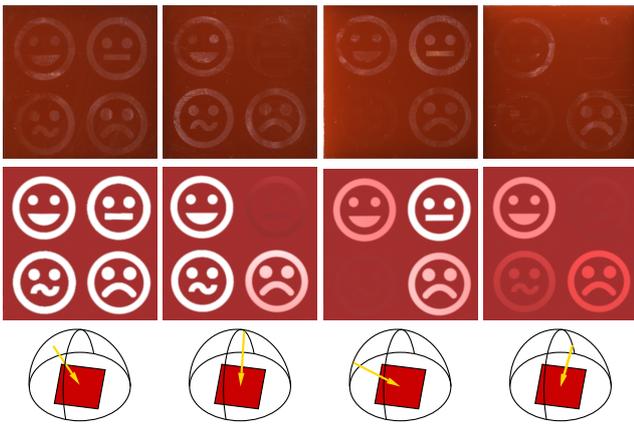

Figure 15: *Hemisphere printed without applying a grayscale pattern to the surface, leftmost, and from second to rightmost when using noise with amplitude $A = 0.625, 2, 3$, respectively, and frequency $B = 32$. Mesh slicing was done with ray tracing. The light-view configuration is the same within each row.*

Figure 14: *Photos of the anisotropic smiley samples of Figure 13 (top row) with light incident from the directions shown in the bottom row. Reflection contrast predictions based on our analytic BRDF models are in the middle row. While the contrast seen in the printed samples is not as clear as in the predictions, the intensity variations are qualitatively similar.*

respectively $\lambda_u = 50 \ \mu$m and $\lambda_v = 150 \ \mu$m. In the antidiagonal, we used ridged patterns [LFD*17] with an inclination of $10°$ and pitch length of $100 \ \mu$m. The ridges of the patterns in these two smileys have orthogonal orientations. The QR code in Figure 1 is another example of a surface with orthogonal ridged structures, but this was printed using the $2\times$ magnifying lens.

Printing these anisotropic patterns, we obtain a sample with spatially varying reflectance properties without adding any extra step to the DLP printing process. Figure 14 exemplifies how the different parts of the sample reflect light differently under different lighting conditions. The ridged structure generates different contrasts as the light rotates around the sample. The sinusoid structure also results in anisotropic properties, but the difference in contrast between the two different pattern orientations is not as strong as for the ridged pattern. On the other hand, the 2D sinusoid structure results in a more diffuse-like anisotropic effect. We validated these results by comparing the photographs in the top row of Figure 14 with images rendered using the corresponding BRDF models (as explained in Section 3). The printed samples present light reflec-
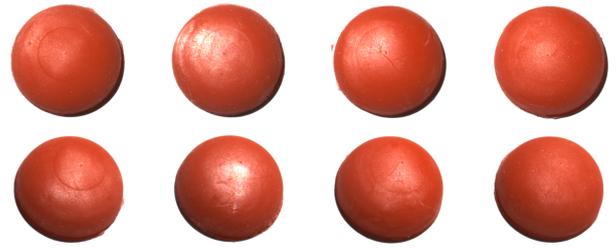
tion contrast that is qualitatively similar to the rendered images. The difference in reflection contrast between the printed samples and the rendered images are mainly due to our choice of using BRDF models (no subsurface scattering), and due to the ground noise introduced by the printing process.

In Figure 15, we compare a hemisphere printed without applying any grayscale pattern to the surface (leftmost) with samples where we applied sparse convolution noise of different amplitudes (*A*). The presence of a grayscale pattern produced by a smooth irregular noise function with low amplitude makes the surface smoother and removes the majority of the staircase aliasing artifacts introduced by the layered printing process. As the amplitude increases, the specular highlight becomes less visible and the surface appears to be more diffuse. This is a visual indication that the noise function enables us to control roughness not only in flat samples (as measured in Section 4.3) but also in curved 3D printed surfaces.

Finally, we applied grayscale patterns to a more complex geometry, namely the Stanford Bunny. The results are in Figure 1 and in Figure 16. In the leftmost column of Figure 16, the bunny was printed without applying a pattern to the surface. It exhibits an anisotropic specular highlight which is caused by the staircase that is a by-product of the layered printing. In the middle left column, we tried to remove the anisotropy and smoothen the printed surface by applying a low-frequency 2D sinusoid. While this approach to some extent reduces staircase artifacts in the highlights, a line-like reflection is still visible across the back of the bunny (bottom image). In addition, the regularity of the sinusoid pattern makes it visible on the back and the ears of the bunny (top image). A better result was achieved by using sparse convolution noise (middle right column and rightmost column). With a value of $A = 0.625$, we obtained a smoother surface with highlights similar to the ones obtained with the sinusoid pattern but without introducing visible sinusoidal features. With $A = 3$, the bunny is more rough and the appearance is more diffuse-like. In Figure 1, we used the sun as the light source. This somewhat resembles a directional light and makes the difference between the rough and the smooth bunny stand out clearly.

We observed that the effect of our technique is less visible at the bottom of the ears of the Stanford Bunny. This is the case for surface voxels that are backfacing as seen from the projector. Here, the slicer applies antialiasing but no greyscale pattern. In addition, the
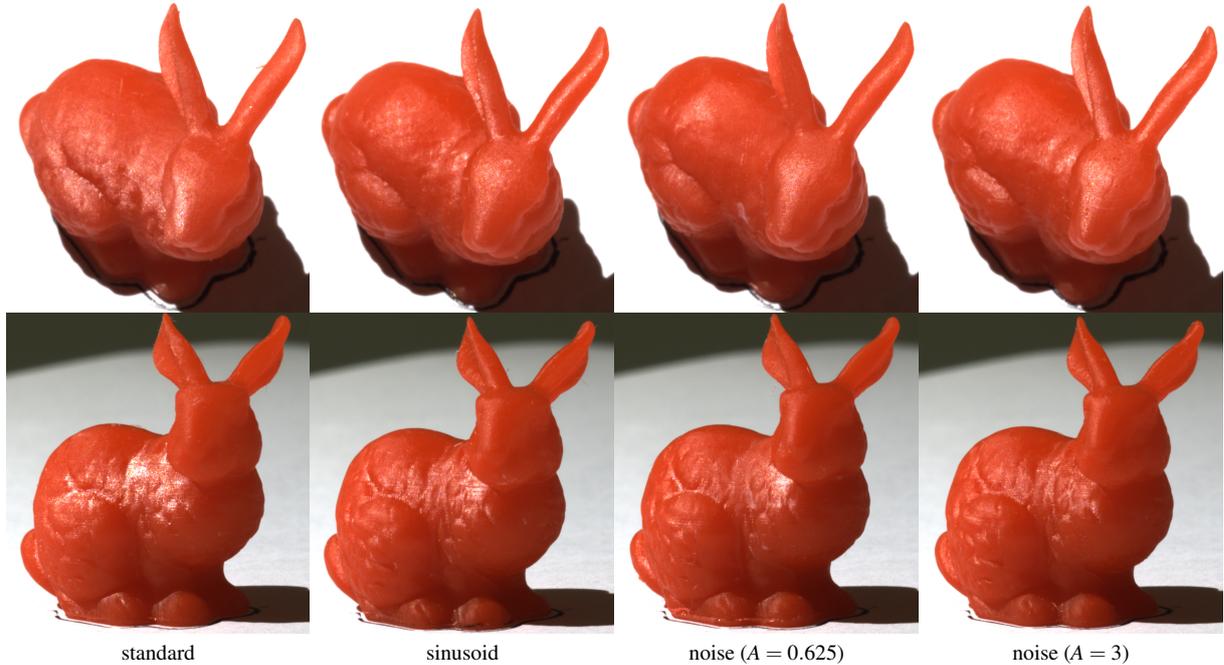
standard   sinusoid   noise ($A = 0.625$)   noise ($A = 3$)

Figure 16: *Stanford Bunny printed and photographed in two different light-view configurations (rows). The bunny was printed without any grayscale pattern applied (*standard*), with an isotropic 2D sinusoid function applied (*sinusoid*, $\lambda_u = \lambda_v = 400 \ \mu m$), and using sparse convolution noise with low and high amplitudes (A) and frequency $B = 64$ (*noise*). The glean due to anisotropic reflection caused by layering artifacts is clearly observable for the standard technique. The sinusoid pattern reduces the problem but introduces regularity artifacts. The noise function more effectively reduces the problem. As compared with the rough bunny (A = 3), the smooth bunny (A = 0.625) is brighter in the highlight regions and darker outside those regions as expected. Mesh slicing was done with rasterization.*

ground noise is probably different due to curing without adhesion (interlaminar bonding) to an existing solidified layer. A technique such as monitoring the photopolymerization process using a photorheometer [HOBS18] might be used to improve the precision of a 3D printer for backfacing surface voxels.

6. Conclusion

In this work, we presented a one-step technique for controlling surface appearance in DLP printing. Our technique is based on projection of grayscale images to control the voxel growth and enable printing of subvoxel sized microstructure. We provided a procedure for correcting the nonlinearity of the photopolymerization process, and the validity of this procedure was experimentally verified. We also demonstrated that application of different grayscale patterns to surface voxels is useful for modifying the microstructure of a surface and for printing spatially varying anisotropic reflectance properties. An important discovery in our work is that a smooth irregular noise function (sparse convolution noise, in our case) is useful both for antialiasing to obtaining a smooth surfaces without staircase artifacts and for controlling surface roughness. We have described a pipeline for applying grayscale patterns to surface voxels during the slicing of mesh geometry. Finally, we included a procedure for calibrating the parameters of a DLP printer and for estimating the ground noise added to the surface by the printing process. Our results demonstrate that by modulating the UV light intensity of a DLP projector with grayscale images we can print samples with

spatially varying reflectance properties, such as anisotropic effects and surface roughness.

As an addendum, Mark Wheadon has presented a webpage that describes an interesting experimental technique called velocity painting (www.velocitypainting.xyz). This technique enables use of grayscale values in fused deposition modeling (FDM) printing. The grayscale input images modify and control the print speed of an FDM 3D printer. This enables printing of patterns on the sample surface without modifying the filament or using multiple extruders. We leave investigation of the microstructure controllabilities of such a technique to future work. Nevertheless, we find it exciting that our calibration and grayscale microstructure control techniques can perhaps be transferred to the more commonly available nozzle-based 3D printers.

Appendix A: Sparse Convolution Noise

We use sparse convolution noise [Lew84, Lew89] in the version presented by Frisvad and Wyvill [FW07], but implemented as a closed function. This is a solid noise function in the classical sense [Per85], but without the grid-aligned regularity artifacts seen

in Perlin noise and with no need for tabulated data. The noise function uses a simple linear congruential pseudo-random number generator:

$$t_{n+1} = (bt_n + c) \bmod m$$
$$\text{rnd}(t_n) = t_{n+1}/m,$$

where we use $b = 3125$, $c = 49$, and $m = 65536$, and a cubic filter kernel function

$$\text{cubic}(\mathbf{v}) = \begin{cases} (1 - 4\mathbf{v} \cdot \mathbf{v})^3 & \text{for } \mathbf{v} \cdot \mathbf{v} < \frac{1}{4}, \\ 0 & \text{otherwise}. \end{cases}$$

A sparse distribution of randomly placed random impulses are then blended using this cubic filter to obtain the noise function. As the filter radius is $\frac{1}{2}$, we can use a regular grid offset by half a unit, so that we only need to consider the impulses in the eight nearest grid cells. Suppose i is the neighbor index of the grid cell, j is the impulse index, and N is the number of impulses per cell. We let $\alpha_{i,j}$ denote the value of the impulse, $\boldsymbol{\xi}_{i,j}$ the local position of the impulse in its grid cell, and $n_{i,j}$ the seed of the pseudo-random number generator for an impulse. The noise function is then

$$\text{noise}(\mathbf{p}) = \frac{4}{5\sqrt[3]{N}} \sum_{i=0}^{7} \sum_{j=1}^{N} \alpha_{i,j} \, \text{cubic}(\mathbf{x}_{i,j} - \mathbf{p}),$$

$$\mathbf{x}_{i,j} = \mathbf{q}_i + \boldsymbol{\xi}_{i,j}$$
$$\alpha_{i,j} = \text{rnd}(t_{n_{i,j}})(1 - 2(j \bmod 2))$$
$$\boldsymbol{\xi}_{i,j} = \left(\text{rnd}(t_{n_{i,j}+1}), \text{rnd}(t_{n_{i,j}+2}), \text{rnd}(t_{n_{i,j}+3}) \right)$$
$$n_{i,j} = 4(N\mathbf{q}_i \cdot \mathbf{a} + j)$$
$$\mathbf{q}_i = \left\lfloor \mathbf{p} - \left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2} \right) \right\rfloor + \left(i \bmod 2, \left\lfloor \frac{i}{2} \right\rfloor \bmod 2, \left\lfloor \frac{i}{4} \right\rfloor \bmod 2 \right),$$

where N should be an even number to avoid a bias toward negative impulse values. We use $N = 30$ and $\mathbf{a} = (1, 1000, 576)$.

To generate a noise function for procedural texturing with values in $[0,1]$, we use

$$\min \left(\max \left(0, \frac{A}{2} \text{noise}(B\mathbf{p}) + \frac{1}{2} \right), 1 \right),$$

where the parameters A and B control the amplitude and the frequency (also called the scale) of the noise, respectively.

Appendix B: Masking and Shadowing for a Sinusoidal Structure

This appendix briefly describes the BRDF model that we used to predict the reflection contrast produced by a 2D sinusoidal microstructure. The microstructure is described by Equation 3. The model that we used is similar to the one presented by Luongo et al. [LFD*17] for the ridged surface microstructure, with the main difference that we here use a different geometrical attenuation function, G. We still use the separation [WMLT07]

$$G(\omega_i, \omega_o, \mathbf{n}) = G_1(\omega_i, \mathbf{n}) G_1(\omega_o, \mathbf{n}),$$

where ω_i and ω_o are incoming and outgoing light directions and $\mathbf{n}$ is the surface normal.

We consider a generic 2D sinusoidal function

$$f(x,y) = A \sin \left(\frac{2\pi}{\lambda_x} x \right) \sin \left(\frac{2\pi}{\lambda_y} y \right),$$
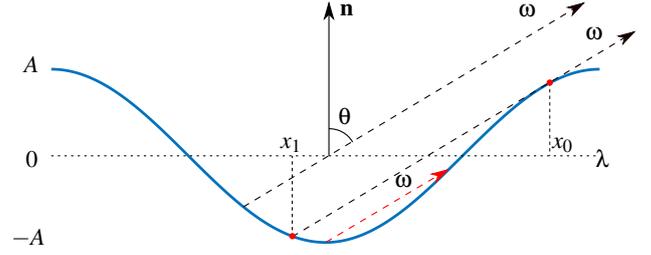

Figure 17: *The surface fraction masked by the sinusoidal structure is given by the ratio between $|x_1 - x_0|$ and λ.*
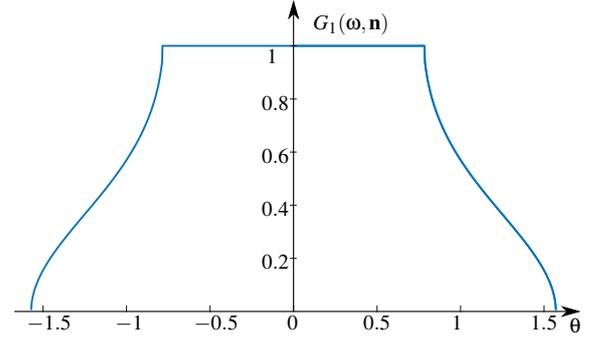

Figure 18: *Plot of the masking function G_1 for $A = 1$ and $\lambda = 2\pi$.*

where A represents the amplitude, and λ_x and λ_y are the wavelengths along the x and y axes. For simplicity, we derive the geometrical attenuation function for the 1D function

$$f(x) = A \cos(kx)$$

with $k = \frac{2\pi}{\lambda}$, and we then extend it to the 2D case.

For a given direction ω forming an angle θ with the surface normal $\mathbf{n}$, as shown in Figure 17, we would like to determine if this direction is tangent to $f(x)$. This is determined by solving

$$f'(x) = -Ak \sin(kx) = m \tag{4}$$

with $m = \tan \left(\frac{\pi}{2} - \theta \right)$. Equation 4 admits

$$x_0 = \arcsin \left(-\frac{m}{Ak} \right) \frac{1}{k}$$

as solution only if $\left| \frac{m}{Ak} \right| < 1$. We can now define the function G_1 by

$$G_1(\omega, \mathbf{n}) = \begin{cases} 1 - \frac{|x_1 - x_0|}{\lambda}, & \left| \frac{m}{Ak} \right| < 1, \\ 1, & \left| \frac{m}{Ak} \right| > 1, \end{cases} \tag{5}$$

where x_1 is the intersection point between $f(x)$ and the tangent line

$$f_t(x) = f(x_0) + m(x - x_0),$$

as shown in Figure 17, and this is found by numerically solving the equation $f_t(x) - f(x) = 0$.

Equation 5 is plotted in Figure 18 for the parameters $A = 1$ and $\lambda = 2\pi$. The function G_1 is extended to the 2D sinusoidal case by considering projections of ω on the planes spanned by $\mathbf{n}$ and the

x-axis as well as **n** and the *y*-axis. We refer to these projections as ω_x and ω_y and define G_1 by

$$G_1(\omega, \mathbf{n}) = G_1(\omega_x, \mathbf{n})G_1(\omega_y, \mathbf{n}).$$

References

[ABRK17] AURICH J. C., BOHLEY M., REICHENBACH I. G., KIRSCH B.: Surface quality in micro milling: Influences of spindle and cutting parameters. *CIRP Annals 66*, 1 (2017), 101–104. 2

[AS00] ASHIKHMIN M., SHIRLEY P.: An anisotropic Phong BRDF model. *Journal of graphics tools 5*, 2 (2000), 25–32. 7

[BBS15] BAAR T., BRETTEL H., SEGOVIA M. V. O.: Towards gloss control in fine art reproduction. In *Measuring, Modeling, and Reproducing Material Appearance* (2015), vol. 9398 of *Proceedings of SPIE Electronic Engineering 2015*, p. 93980T. 2

[Ben17] BENNETT J.: Measuring UV curing parameters of commercial photopolymers used in additive manufacturing. *Additive Manufacturing 18* (December 2017), 203–212. 3

[BSH97] BATTKE H., STALLING D., HEGE H.-C.: Fast line integral convolution for arbitrary surfaces in 3D. In *Visualization and Mathematics*. Springer, 1997, pp. 181–195. 4

[DWP*10] DONG Y., WANG J., PELLACINI F., TONG X., GUO B.: Fabricating spatially-varying subsurface scattering. *ACM Transactions on Graphics (SIGGRAPH 2010) 29*, 4 (July 2010), 62:1–62:10. 3

[ESZ*17] ELEK O., SUMIN D., ZHANG R., WEYRICH T., MYSZKOWSKI K., BICKEL B., WILKIE A., KŘIVÁNEK J.: Scattering-aware texture reproduction for 3d printing. *ACM Transacions on Graphics (SIGGRAPH Asia 2017) 36*, 6 (November 2017), 241:1–241:15. 3

[Fri12] FRISVAD J. R.: Building an orthonormal basis from a 3d unit vector without normalization. *Journal of Graphics Tools 16*, 3 (August 2012), 151–159. 4

[FW07] FRISVAD J. R., WYVILL G.: Fast high-quality noise. In *Proceedings of GRAPHITE 2007* (December 2007), ACM, pp. 243–248. 10

[Gre16] GREENE R.: High-fidelity 3D printing techniques. Additive Manufacturing Today, Videos, April 2016. 3, 4, 7

[HFM*10] HAŠAN M., FUCHS M., MATUSIK W., PFISTER H., RUSINKIEWICZ S.: Physical reproduction of materials with specified subsurface scattering. *ACM Transactions on Graphics (SIGGRAPH 2010) 29*, 4 (July 2010), 61:1–61:10. 3

[HOBS18] HOFSTETTER C., ORMAN S., BAUDIS S., STAMPFL J.: Combining cure depth and cure degree, a new way to fully characterize novel photopolymers. *Additive Manufacturing 24* (December 2018), 166–172. 10

[Jac92] JACOBS P. F.: *Rapid prototyping & manufacturing: fundamentals of stereolithography*. Society of Manufacturing Engineers, 1992. 3

[Jør15] JØRGENSEN A. R.: *Design and development of an improved direct light processing (DLP) platform for presion and additive manufacturing*. Master's thesis, Technical University of Denmark, 2015. 5

[LCR*17] LOW Z.-X., CHUA Y. T., RAY B. M., MATTIA D., METCALFE I. S., PATTERSON D. A.: Perspective on 3D printing of separation membranes and comparison to related unconventional fabrication techniques. *Journal of Membrane Science 523* (February 2017), 596–613. 3

[LDPT13] LAN Y., DONG Y., PELLACINI F., TONG X.: Bi-scale appearance fabrication. *ACM Transactions on Graphics (SIGGRAPH 2013) 32*, 4 (July 2013), 145:1–145:11. 2

[Lew84] LEWIS J. P.: Texture synthesis for digital painting. *Computer Graphics (SIGGRAPH '84) 18*, 3 (July 1984), 245–252. 10

[Lew89] LEWIS J. P.: Algorithms for solid noise synthesis. *Computer Graphics (SIGGRAPH '89) 23*, 3 (July 1989), 263–270. 10

[LFD*17] LUONGO A., FALSTER V., DOEST M. B., LI D., REGI F., ZHANG Y., TOSELLO G., NIELSEN J. B., AANÆS H., FRISVAD J. R.: Modeling the anisotropic reflectance of a surface with microstructure engineered to obtain visible contrast after rotation. In *Proceedings of International Conference on Computer Vision Workshop (ICCVW 2017)* (October 2017), IEEE, pp. 159–165. 4, 9, 11

[LGX*13] LEVIN A., GLASNER D., XIONG Y., DURAND F., FREEMAN W., MATUSIK W., ZICKLER T.: Fabricating BRDFs at high spatial resolution using wave optics. *ACM Transactions on Graphics (SIGGRAPH 2013) 32*, 4 (July 2013), 144:1–144:13. 2

[LPA01] LEE J. H., PRUD'HOMME R. K., AKSAY I. A.: Cure depth in photopolymerization: experiments and theory. *Journal of Materials Research 16*, 12 (December 2001), 3536–3544. 3

[MAG*09] MATUSIK W., AJDIN B., GU J., LAWRENCE J., LENSCH H., PELLACINI F., RUSINKIEWICZ S.: Printing spatially-varying reflectance. *ACM Transactions on Graphics (SIGGRAPH Asia 2009) 28*, 5 (December 2009), 128:1–128:9. 2

[MQM17] MOSTAFA K., QURESHI A. J., MONTEMAGNO C.: Tolerance control using subvoxel gray-scale DLP 3D printing. In *Proceedings of ASME International Mechanical Engineering Congress and Exposition (IMECE17)* (2017), p. V002T02A035. 2, 3

[MSRG12] MCEWAN I., SHEETS D., RICHARDSON M., GUSTAVSON S.: Efficient computational noise in GLSL. *Journal of Graphics Tools 16*, 2 (2012), 85–94. 4

[MSS*12] MALZBENDER T., SAMADANI R., SCHER S., CRUME A., DUNN D., DAVIS J.: Printing reflectance functions. *ACM Transactions on Graphics 31*, 3 (May 2012), 20:1–20:11. 2

[Per85] PERLIN K.: An image synthesizer. *Computer Graphics (SIGGRAPH '85) 19*, 3 (July 1985), 287–296. 4, 10

[Per02] PERLIN K.: Improving noise. *ACM Transactions on Graphics (SIGGRAPH 2002) 21*, 3 (July 2002), 681–682. 4

[PLMR17] PEREIRA T., LEME C. L., MARSCHNER S., RUSINKIEWICZ S.: Printing anisotropic appearance with magnetic flakes. *ACM Transactions on Graphics (SIGGRAPH 2017) 36*, 4 (July 2017), 123:1–123:10. 2

[PRJ*13] PAPAS M., REGG C., JAROSZ W., BICKEL B., JACKSON P., MATUSIK W., MARSCHNER S., GROSS M.: Fabricating translucent materials using continuous pigment mixtures. *ACM Transactions on Graphics (SIGGRAPH 2013) 32*, 4 (July 2013), 146:1–146:12. 3

[PZNH16] PEDERSEN D. B., ZHANG Y., NIELSEN J. S., HANSEN H. N.: A self-peeling vat for improved release capabilities during DLP materials processing. In *Proceedings of the 2nd International Conference on Progress in Additive Manufacturing (Pro-AM 2016)* (2016), pp. 241–245. 5

[RBK*13] ROUILLER O., BICKEL B., KAUTZ J., MATUSIK W., ALEXA M.: 3D-printing spatially varying BRDFs. *Computer Graphics and Applications 33*, 6 (2013), 48–57. 2

[RBSM19] RIBARDIÈRE M., BRINGIER B., SIMONOT L., MENEVEAUX D.: Microfacet BSDFs generated from NDFs and explicit microgeometry. *ACM Transactions on Graphics 38*, 5 (2019), 143:1–143:15. 5

[Rib17] RIBO M. M.: *3D Printing of Bio-inspired Surfaces*. Master's thesis, Technical University of Denmark, 2017. 6

[TR75] TROWBRIDGE T. S., REITZ K. P.: Average irregularity representation of a rough surface for ray reflection. *Journal of the Optical Society of America 65*, 5 (1975), 531–536. 4

[WMLT07] WALTER B., MARSCHNER S., LI H., TORRANCE K.: Microfacet models for refraction through rough surfaces. In *Proceedings of Eurographics Symposium on Rendering (EGSR 2007)* (2007), The Eurographics Association, pp. 195–206. 5, 11

[WPMR09] WEYRICH T., PEERS P., MATUSIK W., RUSINKIEWICZ S.: Fabricating microgeometry for custom surface reflectance. *ACM Transactions on Graphics (SIGGRAPH 2009) 28*, 3 (August 2009), 32:1–32:6. 2