

# Dual-Color Mixing for Fused Deposition Modeling Printers

Tim Reiner<sup>1,2</sup>

Nathan Carr<sup>1</sup>

Radomír Měch<sup>1</sup>

Ondřej Štáva<sup>1</sup>

Carsten Dachsbacher<sup>2</sup>

Gavin Miller<sup>1</sup>

<sup>1</sup> Adobe Systems Incorporated

<sup>2</sup> Karlsruhe Institute of Technology



**Figure 1:** Left: Lena image 3D printed on a MakerBot Replicator 2X. This plate is 3.5 cm tall and about 2 mm thick. We achieve approx. 50 dpi vertically with a 60×71 texture modulated on 142 alternating black and white layers. Center: 3D models printed with continuous tone. Corresponding texture maps: an ambient occlusion texture gives the GARGOYLE (10.2 cm tall) an aged appearance while we fit LUCYCAT (5.4 cm tall) with a calico coat. Right: the Lena texture remains sharp after projection.

**Acknowledgements:** We thank Zoltan Bourne (Lucy the cat) and Julie Dorsey (Gargoyle) for permission to use their models.

## Abstract

In this work we detail a method that leverages the two color heads of recent low-end fused deposition modeling (FDM) 3D printers to produce continuous tone imagery. The challenge behind producing such two-tone imagery is how to finely interleave the two colors while minimizing the switching between print heads, making each color printed span as long and continuous as possible to avoid artifacts associated with printing short segments. The key insight behind our work is that by applying small geometric offsets, tone can be varied without the need to switch color print heads within a single layer. We can now effectively print (two-tone) texture mapped models capturing both geometric and color information in our output 3D prints.

Categories and Subject Descriptors: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling

## 1. Introduction

Recent technological advances have made 3D printers affordable for home users and their increased use is foreseeable. Unsurprisingly, 3D printing has gained a lot of attention in the computer graphics community, who is contributing to this new output medium in the areas of geometric modeling, rendering, and perception. Obviously it is very desirable to print 3D objects in full color, however, low-cost entry-level printers are restricted to print using only monochrome filaments. In this paper we present a dithering technique for printers with at least two nozzles. We produce interleaved color patterns, which convey an impression of

smooth color blending, which ultimately provides a broader color (or grayscale) spectrum. We used the MakerBot Replicator 2X desktop printer for all our experiments with black and white filament for maximum contrast.

We make the following contributions: we show how to construct a basic checkerboard pattern which gradually dissolves into an in-between tone, while minimizing the switching of nozzles at the same time; we propose a simple scheme to project this pattern onto arbitrary shapes; we demonstrate how to texture map these objects efficiently; finally, we reveal a technique to avoid artifacts associated with printing using two nozzles, significantly improving print quality.

## 2. Previous Work

Recent advances in 3D printers along with their increased availability have led to numerous research efforts. We refer to a recent state of the art report from Hullin et al. [HIH\*13], which, among others, gives a broad overview of recent fabrication methods. There also exist numerous investigations on improving the structural integrity of 3D printed forms [TJ11, SVB\*12]. For large objects, methods have been developed to effectively partition the model into pieces that can be fabricated within a fixed print volume [LBRM12]. Methods for printing both articulated and deformable structures have been proposed [BBO\*10, BBJP12, CCA\*12, CTN\*13]. The technique of Wang et al. [WWY\*13] focuses on reducing the cost of printing, while other approaches exist for optimizing quality [HBA13]. All of these techniques, however, focus on the geometric form of the printed shape.

Our work is more closely aligned to that of OpenFab and Spec2Fab [VWRKM13, CLD\*13] in that we focus on the appearance of the printed shape and not just its geometric form. The general idea of printing color is not a new one. For example, the work of Levin et al. [LGX\*13] showed that full BRDFs can be fabricated, and Hašan et al. [HFM\*10] use multi-material 3D printers to approximate a given or measured heterogeneous subsurface scattering function. These works make use of high-end printers which support color printing, such as printers from 3D Systems (powder-based ProJet series) and Mcor Technologies (paper-based). Although these printers can produce high-quality color prints, it is unclear whether their costs can come down to reach the consumer space any time soon. A new company, BotObjects, is offering an FDM-based printer where the filaments are mixed to produce a variation in color. It will be possible to create slow gradations in color throughout the object, but not sharp transitions between colors without time and material costly purging and refilling of the mixing chamber. Our work, by contrast, is aimed at leveraging low-cost consumer printers based on FDM. We restrict our research to focus on the reproduction of continuous tone, contrary to a full BRDF. Our work is specifically targeted at the broad hobbyist and consumer 3D printing communities providing an effective practical solution whose quality scales as the printers increase in both accuracy and precision.

In the low-end space, some attempts have been made to print continuous tones. Hobbyists started to feed multiple filaments into a single nozzle and try to mix the color at the head of the printer. While this type of approach may have promise, it does not allow the rapid color changes necessary to print high-frequency detailed imagery on the surface of 3D shapes. We refer to this approach as active color mixing, while we consider our work as a passive mixing technique. Cho et al. [CSPT03] used dithering to produce tone variation. From our experiments we found this type of approach impractical on low-end FDM printers; they cannot print short spans or small dots without producing significant

geometric artifacts due to the lack of precise flow rate control of filament from the heads of the nozzle. The insight behind our approach is to enable the appearance of such dithering while maintaining long contiguous spans of single color.


## 3. Algorithm

To print continuous variation of tone using two colors one could adapt a halftone technique to 3D printers: very small regions of the model would be assigned different colors and long continuous print spans produced by the slicer would be divided up into very short spans (or dots) of interleaved color. While at first glance this straightforward approach may sound promising, FDM printers produce significant artifacts when filament streams are not long and continuous; this is due to the rate at which the printer can start and stop the flow of melted filament. Furthermore, such an approach would require incessant switching between print heads, increasing the total printing time significantly. Thus our challenge is to derive a scheme that minimizes switching between print heads while maximizing the length of continuous spans of filament emitted from the nozzle.

### 3.1. Tone Variation using Geometric Offsetting

An initial experiment involved printing a test cylinder where every other layer from the slicer was printed in alternating black and white color. Such a print when viewed at a distance should produce a constant gray appearance. To our surprise our test print produced more variation in tone than expected. Upon close examination we noticed that subtle misalignment of the color print heads lead to slightly offset circles alternating every other layer. Thus one side of the cylinder appears darker where the black filament protruded, and the other side lighter due to more exposed white filament. In between the color went to a smooth tone of gray where the two filaments were aligned and equally exposed. This very simple test print leads to the observation that subtle geometric offsetting between layers enables the control of variation of tone due to occlusion and gravity: overhanging filament sags a little, closing the gap by hiding the recessed filament color. Therefore, one can easily modulate the tone by offsetting or inseting each layer slice producing the desired tonal variation along the vertical print axis.

In Fig. 2 we show how to exploit this observation to produce interleaving patterns on a test cylinder. We apply a high frequency sine pattern to the circular outlines (visualized as a circle on the build plate) during slicing and print the first layer in black. By phase shifting its sine pattern by  $\pi$ , the second layer printed in white overlays into the first, producing an interleaved chain of very short horizontally alternating black and white strips. We retain this phase and print the third layer in black directly atop the second layer. The fourth layer is again printed in white, reverting to the original phase of the first layer. The outcome is a basic checkerboard building block which is simulated and visualized in the closeups.



**Figure 2:** Printing a basic checkerboard pattern. The layers overlay into each other, resulting in an interleaved effect.

We use this technique to construct the cylinder in Fig. 3(a): always after printing a pair of two layers we shift the phase by  $\pi$ . This leads to filament extruded into mid-air, but due to gravity the layers of the real print interleave with each other. Comparing (a) with (b) demonstrates that the appearance of a printed result can be simulated well before starting an actual printing process. This allows for faster modeling iterations and avoids costly trial-and-error prints.


In Fig. 3(c) we again alternate between black and white for every layer, but do not shift the phase between pairs of layers; instead, we investigate the effect of phase shifting the white layers against the black layers for varying amounts. Starting at the bottom, layers are not shifted against each other. Note how horizontal stripes can be seen. Upward, the shift is  $1/4\pi$ ,  $1/2\pi$ ,  $3/4\pi$ , and  $\pi$  in the center. Note how this allows printing vertical stripes while switching colors on a horizontal basis. Further upward, the shift is  $5/4\pi$ ,  $3/2\pi$ ,  $7/4\pi$ , until there is eventually no more shifting again.

Modulating the amplitudes of the sine waves locally then can control the tonal variation. To make the tone appear brighter, we attenuate the amplitude of black layers. Vice versa, attenuating white layers makes the tone appear darker. In Fig. 3(d), we can perceive a smooth gradient on the cylinder. This modulation technique is also used to apply texture mapping to our prints, as demonstrated in the teaser (Fig. 1).

Note that this approach avoids the rapid switching of color print heads. Within each layer our algorithm produces the same number of long continuous spans that would have been emitted from the slicer for single color printing. Thus our method only incurs the overhead required to switch color print heads between layers, which is minimal, and the artifacts associated with starting and stopping filament streams are no worse than that of regular single color printing.

### 3.2. Arbitrary Shapes

The checkerboard pattern nicely dissolves into an in-between tone when viewed from a distance and is the foundation for texture mapping our models. We can map this reg-



**Figure 3:** (a) Checkered cylinder. Top and bottom: three out of four layers are black or white, respectively. Rendered with mental ray. (b) Photo of actual print. (c) Gradually shifting phase between layers. (d) Gradually attenuating amplitudes.

ular pattern perfectly onto the previous cylinder examples or other extruded shapes. For arbitrary shapes, however, we either have to alter the sine pattern frequency or allow for errors in phase between levels. Unfortunately, even small deviations in this pattern lead to salient artifacts in the prints, as the human eye instantly recognizes subtle pattern misalignments. Thus the challenge becomes bounding distortions in frequency and phase between layers.

While one could perform some global optimization using all layer slices in parallel, our solution operates locally on the current layer slice and requires only the knowledge of the past  $k \geq 1$  sliced layers. We iteratively perform a projection and relaxation process, moving from bottom to top, which integrates seamlessly into the workflow of existing slicers with minimal memory and computational overhead.

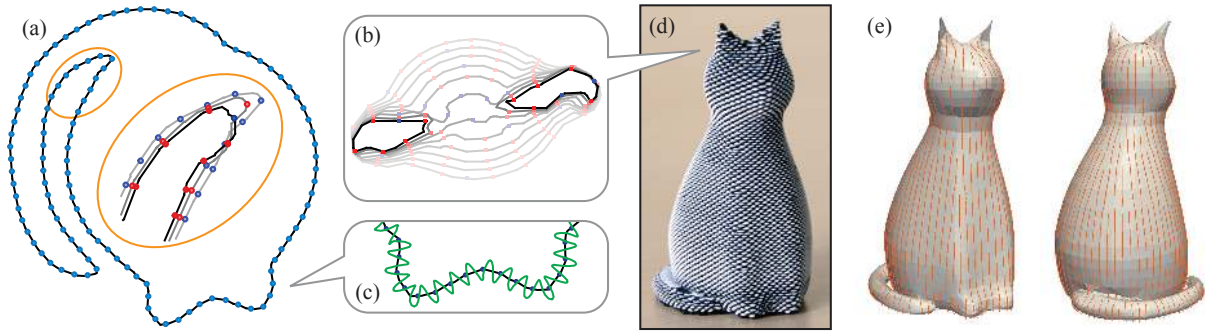
Our slicer transforms a given 3D object into  $n$  contiguous slices. Each slice is a set of polygons  $P_\ell$  where  $\ell \in [0, n-1]$  denotes the  $\ell$ -th slicing layer starting at the bottom of the model and proceeding upward. Our goal is to convert the straight line segments of each of these polylines into high frequency sine waves with near constant wavelength  $\lambda$ . Always after two layers, sine waves get phase shifted by  $\pi$ . To obtain the desired checkerboard pattern, waves must be aligned as well as possible between layers, i.e., positions of peaks and valleys of adjacent layers have to match.

Our algorithm works by placing *samples* on the polygons in each layer  $P_\ell$ . Each sample  $s$  corresponds to a peak or valley of a wave on odd or even pairs of layers, respectively. By controlling the spacing and number of samples within a layer  $\ell$  we can optimize for maintaining a near constant wavelength  $\lambda$ . By aligning samples between successive layers we can optimize the phase alignment. We employ a projection and relaxation strategy to map samples from layer to layer. Starting with the base layer  $P_0$  we compute an optimal number of samples for all polygons  $p_0^i \in P_0$  as follows:

$$k^i = \left\lfloor \frac{A(p_0^i)}{\lambda} \right\rfloor, \text{ where } A(\cdot) \text{ is the arc length of } p_0^i.$$

These samples are then equally spaced along  $p_0^i$  by an arc length distance of  $A(p_0^i)/k^i$ .





**Figure 4:** LUCYCAT model. (a) Base layer 0 with initial placement of samples (blue dots). Inset: projection of samples around the tail onto layers 1 and 2. Projected samples are shown in red, newly inserted samples again in blue. (b) Projection of samples to start forming the ears. (c) Applying sine wave pattern to the outline. (d) Resulting 3D print. (e) Traceline visualization.

Given this initial sample assignment for the base layer, we perform both projection and relaxation on all subsequent layers until we reach the top of the model. Specifically, given an already sampled layer  $\ell$ , we take each sample  $s \in s_\ell^i$  and project it as  $s'$  onto the next layer up, unless the distance between  $s$  and  $s'$  would be greater than a certain threshold. This avoids the projection of samples in areas where slices are changing rapidly; such areas will have discontinuities anyway. The projection forms new samples  $s_{\ell+1}^j$  with  $j \leq i$ . We then relax their locations, spacing them more evenly apart while avoiding excessive phase distortion with respect to the layer(s) below. Next, we prune any samples that are spaced less than  $\lambda$  in arc length apart. We detect empty spans between samples longer than  $2\lambda$  and insert new evenly spaced samples to fill that gap such that each new sample is spaced no closer than  $\lambda$  from its adjacent samples. Once new samples have been inserted and deleted we reapply relaxation to better space the samples, and then proceed to the next layer. The high level approach is summarized in Algorithm 1.

Fig. 4 shows how this technique processes the LUCYCAT model. A great advantage of the projection scheme is that it is completely agnostic about topology changes—see for instance 4(b)—and naturally handles any change of number in polygons between successive layers. After all samples have been placed on all layers, we can easily tessellate the outline

---

#### Algorithm 1 Projection-Relaxation Scheme

---

Input:  $n$  sliced layers, i.e.,  $n$  sets of polygons  
Output: locations on slices to place peaks or valleys of sine curves  
assign initial sample placement on base layer  $\ell = 0$   
**for** each layer  $\ell \in [1, n - 1]$  **do**  
  project samples from layer  $\ell - 1$  onto layer  $\ell$   
  relax sample locations on layer  $\ell$  (and optionally  $\ell - 1, \dots$ )  
  remove samples closer than  $\lambda$  on layer  $\ell$   
  insert new samples on layer  $\ell$  in gaps larger than  $2\lambda$   
  relax sample locations on layer  $\ell$  (and optionally  $\ell - 1, \dots$ )  
**end for**

---

in-between samples with a sine pattern (c), which eventually produces the G-Code to 3D print the model (d).


By connecting projected samples, we can visualize the entire process as a series of tracelines moving up the side of a model (Fig. 4(e)). When the model narrows, some of the tracelines will dead end, and in widening regions new tracelines form. The tracelines, i.e., samples, are provably spaced between  $\lambda$  and  $2\lambda$  apart: any samples spaced less than  $\lambda$  will get pruned, and any samples spaced greater than  $2\lambda$  will not exist since a new sample can be inserted in between. Our relaxation maintains the constraint that any two samples are bounded in distance  $d \in [\lambda, 2\lambda)$  from one another, i.e., we never move samples to locations that violate this criteria.

## 4. Implementation

We implemented our dual-color mixing technique within our own slicer, however, this would be straightforward to integrate into other existing slicers: they could iteratively project and relax samples while processing layer by layer and then modulate the outlines. Our slicer is further parallelized to process the input model using multiple threads.

### 4.1. Texture Mapping

For textured input models, our slicer evaluates and stores texture coordinates for every vertex forming the polylines of the dissected model. While tessellating the outlines these coordinates are interpolated for texture lookup. The result is then used to attenuate the corresponding amplitudes. Since we have parametrized our layers with a target fixed wavelength, according to the Nyquist limit we are not able to capture textures at all frequencies. Bilinear interpolation produces textures aliasing on the prints if the texture resolution is too high; thus we apply trilinear interpolation (mip-mapping) to properly prefilter textures based on our target sample spacing. For our teaser (Fig. 1) we prefiltered the Lena image using Lanczos resampling to produce a  $60 \times 71$  texture which creates a perfect mapping of texels to 71 pairs of layers.



**Figure 5:** The cleaning tower (far right) primes nozzles and prevents salient artifacts such as superfluous filament strips.

#### 4.2. Cleaning Tower

There is at most one active nozzle at a time during the printing process, and due to the heat the inactive nozzle slowly starts extracting unwanted strips of filament. They smear the actual print, leading to salient artifacts in form of small strips sticking out of the model. Delayed filament extraction after switching nozzles is another problem. To avoid these issues, we print a cleaning tower in parallel. Its W-shaped form primes nozzles and improves its own stability (Fig. 5).

#### 5. Results

All results were printed using a MakerBot Replicator 2X desktop printer. The layer thickness is set to 0.25 mm for high-quality prints. If two layers are shifted against each other, they transform into an interleaved chain with 0.5 mm thickness. The cylinders in Fig. 3 are rather small and when viewed from a distance their color is perceived as a nice in-between tone. This is the key achievement of our work. Moreover, these cylinders now feature an interesting touch and feel and provide more grip. Holding these prints in the hands, their surfaces appear like woven fabric. Turning them a little exposes them to a viewing angle dependency of perceived shades, as some geometric offsets start hiding others.

LUCYCAT has a very printer-friendly shape (Fig. 4) which does not require any support structures. Much more complex and involved is the GARGOYLE model, portrayed in Fig. 6. We texture mapped our model into a texture atlas, and used a 3D painting system to paint the cracks of this figurine to accentuate shading, giving it an aged look, as dirt collects in the cracks over time.

We determined an optimal wavelength  $\lambda$  of 1200 microns to space our samples apart. We keep below a projection threshold of 600 microns to project samples up onto subsequent layers. During a relaxation phase we correct 5% of the misalignment for each sample from its perfectly centered location between neighboring samples. When modulating outlines with sine patterns, the maximum amplitude is 750 microns, which can be attenuated down to 375 microns.



**Figure 6:** GARGOYLE model during printing process with multiple support structures and final texture mapped result.

Listed below is more information about our models, including timings how long it takes to print them using a single color, our dual-color mixing technique, and with enabled cleaning tower:

	LUCYCAT	GARGOYLE
Layers	216	408
Height	5.4 cm / 2.1 in	10.2 cm / 4.0 in
Single Color	0 h 45 m	5 h 48 m
Dual-Color	0 h 58 m	6 h 56 m
w/ Cleaning Tower	1 h 42 m	8 h 28 m


The additional processing time *for the slicer* to incorporate the dithering technique is negligible, particularly in relation to the actual printing time. Whether a model is textured or not does not have any impact on the duration of the prints. Our method also produces more geometry, but even for GARGOYLE the uncompressed G-Code remains below 40 MB, which nowadays is far from being a constraint.

#### 6. Analysis

A key property of our approach is that we trade geometric detail for continuous tone imagery. We found this loss of precision acceptable for our models, but assume our input polygons have been filtered to remove high frequency geometry detail which is clearly higher than the frequency of the applied sine waves. This limits the effective curvature of our models; otherwise outlines would start to overlap with too much filament extracted at these regions, resulting in severe artifacts. We now characterize the impact of our geometric processing on the quality of the appearance of our prints.

##### 6.1. View-Dependent Appearance / Gamut Specification

By observing the resulting prints we realize that their appearance is highly anisotropic. In order to investigate this issue in greater depth and in an efficient and flexible way, we set up a ray tracing environment which allows us to



**Figure 7:** Set of gamut maps obtained from a synthetic calibration object which was generated for various surface slopes. For varying viewing angles  $\varphi$  and  $\theta$  the intensity response of the gradient is encoded into the corresponding column of such a map.

generate synthetic renderings of our prints. In practice we found that by modeling the geometric detail of each print down to the G-Code level, a standard global illumination engine produced results that were nearly indiscernible from the real-world prints. They adequately capture all of the larger shading impacts resulting from chosen interleaving patterns. More specifically, we found that any perceived impacts resulting from small changes to the local dielectric microfacet shading model were minor in comparison to the more macroscopic anisotropic shading impacts resulting from the geometric patterning of the filaments. Thus we found that any reasonable BRDF approximation of plastic ABS filament sufficed, whereas accurately modeling each strand of geometry was critical for achieving correct synthetic reproductions.

We developed a virtual gonioreflectometer on top of our custom rendering system and were able to rapidly analyze the anisotropic impacts of the filament patterns on the perceived shading. In Fig. 7 we examine a synthetic calibration object, textured with a gradient map, with its plastic ABS filament approximated with a standard Phong-like shading model. The printer paths were modeled as sphere-capped cylinders, where warping was performed to distort them to have oval cross-sections to better simulate the impacts of gravity. We produce slanted variants of this object in steps of 10 degrees until the surface is almost horizontal, to show how the gamut behaves for varying surface slopes.

The gonioreflectometer views each slanted tile face on, i.e., the viewing direction is perpendicular to the surface normal, and then locally averages sets of pixels. This results in an intensity response of the tonal gradient. For all slanted objects we analyze the gamut for varying viewing angles, and rotate the gonioreflectometer first in a vertical fashion (varying  $\varphi$ ,  $\theta = 0$ ), and then horizontally (varying  $\theta$ ,  $\varphi = 0$ ), i.e., remaining parallel to the individual layers of the object.

Each obtained *gamut map* contains one column per viewing angle encoding the measured intensity response of the tonal gradient. As expected, the high anisotropy of the perceived shade is reflected in these maps. Viewing the surface directly on (cf. center column in a gamut map) produced the largest gamut, however, the range of the gamut is limited in that it does not reach full black or white. As the material is


viewed more edge on, the gamut starts to collapse and the gray level intensities move towards complete black or white. This effect also occurs when focusing on curved parts, e.g., gamut degrades when looking more toward the silhouettes of a model. The gamut dissolves completely for extremely slanted objects. Consequently, a wider gamut is only visible under a limited range of viewing angles, but it might be interesting to inspect the anisotropic appearance as future work.

## 6.2. Calibration of Tone

Our analysis shows that tones in a given model map to different tones after interleaved printing, which is due to small-scale surface details and more importantly due to their mutual occlusion. By analogy, for display systems such as monitors and projectors, this nonlinearity is called *gamma* and its compensation is known as *gamma correction*.


The virtual gonioreflectometer enables us to measure the appearance of a calibration object over a broad range of viewing angles, which we used to calibrate and optimize our prints. Given both the slope and orientation of the surface with respect to the viewer, we can calibrate our prints using a corresponding gamut map. Rather than optimizing for a specific view location (possible future work), we assume we are always locally viewing the geometry perpendicular to the surface and employ the center column of a gamut map.

This column shows how a gradient modeled from black to white is perceived after printing. For these intensities we cre-



**Figure 8:** The calibration improved contrast on this print.





**Figure 9:** The effective texture resolution degrades for more horizontal surfaces. We increase the slant of an  $8 \times 8$  cm test object (far left) in steps of 10 degrees. First row: the viewing direction is fixed. Second row: views perpendicular to the surface normal.

ate an inverse map, so that for a computed gray value we find an index in this column for which this value is closest. We precompute such an inverse intensity table for fast look-ups and use it to obtain calibrated prints. Fig. 8 demonstrates the calibration effect resulting in increased perceptible detail.

### 6.3. Resolution on Slanted Surfaces

Our approach works by alternating filament colors every other layer enabling long continuous print paths. The obvious limitation with this approach is that under extreme slopes the tonal resolution degrades. We investigate this effect in Fig. 9, where we again gradually increase the slant of a synthetic test object. In the first row, we retain a fixed camera position with a fixed viewing direction parallel to the individual layers of the object. In the second row, the viewing direction is perpendicular to the surface normal. Note that the resolution does not visibly degrade until thirty degrees from flat, and even at twenty degrees from flat the resolution print pattern is still prominently visible.

### 6.4. Parametrization Seams

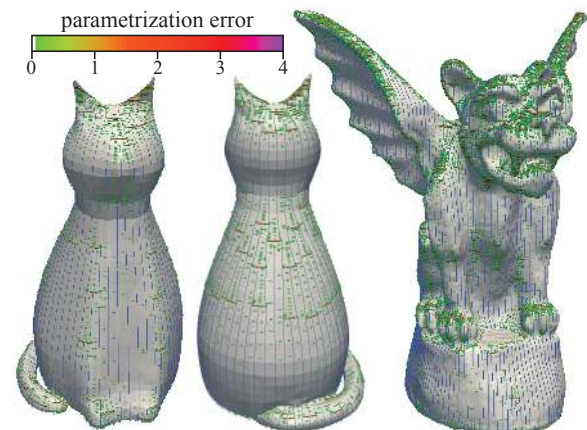
To analyze the generated parametrizations of our projection and relaxation scheme, we equidistantly sample all polylines of our models. We propagate these samples up to their layer above, and determine how well the desired checkerboard pattern is preserved. The divergence from the optimal pattern is then accumulated as error. A sample has a tessellation value (displacement along the surface normal)  $t_1$  in the range from  $-1$  to  $1$  as this is the result of the sine function. Its projected sample above has a tessellation value of  $t_2$ . (Assume the model is not textured and ignore the phase shifting after every two layers.) We then assess the parametrization error between these samples using the metric  $d(t_1, t_2) = (t_1 - t_2)^2 \in [0, 4]$ . Samples are discarded if they exceed the projection threshold.

Fig. 10 visualizes parametrization seams on our test models. A transfer function maps “no error” to transparent and fades over to opaque green for an error value of 0.05, then orange at 1.5, red at 3.0, and purple at 4.0. Parametrization seams in the green range are invisible to the naked eye, but higher error leads to noticeable discontinuities. Regions where streamlines begin and end are primarily af-

ected: there will be a horizontal seam where areas with wavelengths  $\lambda$  and  $2\lambda$  adjoin each other, visualized in red.

The surface of LUCYCAT is to more than 80% free from error (107799 out of 133539 samples in total). Another 18% of the samples fall into the interval  $(0, 0.5)$  where error is not yet perceptible. The remaining 2% of the surface area contains salient seams. This number might sound low, but it is enough to be a disturbing factor to the human eye. (Imagine, by analogy, a  $10 \times 10$  image with two corrupt pixels.) An overall analysis of the GARGOYLE model results in similar numbers (79% error-free and 19% not yet perceptible errors), but the visualization reveals that error primarily accumulates at intricate sections of the head and strongly curved parts of the wings, where most seams will occur.

Besides our projection scheme, we experimented with piecewise planar projections onto the sides of a model, which, by contrast, produces vertical seams and would not be practical for intricate shapes. In another experiment, we adapted known subdivision techniques, which smoothly insert new vertices into a mesh: in areas in-between  $\lambda$  and  $2\lambda$ , we applied different blending functions to create intermediate wave patterns. Unfortunately, they feature frequencies much higher than  $\lambda$ , which makes them impractical to print. On the positive side, they scatter parametrization issues evenly along the whole model, avoiding singularities.



**Figure 10:** Visualizing parametrization error: noticeable seams (in red) occur primarily where tracelines form or end.

## 7. Discussion

**Model orientation.** Fig. 9 shows the loss of resolution on nearly horizontal surfaces. Such limitations may be avoided by reorienting the model within the print volume. The Lena image (Fig. 1), for instance, prints upright vertically instead of lying flat on the build plate. This approach may not work for all models and can introduce the need for support structures, and problematic regions may still exist. These slopes can be improved by adjusting the layer thickness to the smallest amount supported by the print device, slowing down printing, but improving the effective printed resolution. For larger models, we could imagine adapting the approach from Hildebrand et al. [HBA13] to separate models and then aligning individual chunks to be printed vertically.

**Solid color.** Textures may not exclusively consist of regions with continuous tonal variation. If a texture contains a large part of solid color which can be perfectly represented by one of the used filaments, it would be worthwhile to instruct the slicer that the corresponding chunk of the model gets printed using only that filament. This would be compatible with our method, but requires nozzle switching within a layer.

**Noise functions.** We tried modulating outlines with 3D solid noise instead of a sine pattern. Arbitrary slices through noise space are stationary and bandwidth-limited, eliminating the need for parametrization. Instead of phase shifting a curve, we inverted the noise function. Alas, this pattern does not dissolve into an in-between tone, and individual colors remain prominent. We can imagine mapping great visual complexity onto shapes by using colors with less contrast, e.g., brown and black, to model a tree trunk, dark and light gray to model rocks, or blue and white to simulate a water surface.

## 8. Conclusions and Future Work

Recent advances in 3D printing have led to a number of low-cost printers in the consumer space. Although the geometric fidelity of these printers is quite high, we are still on the way to full color printing at the consumer level. Recently low-end printers with two color heads have arrived on the market, and we demonstrated how they can produce interleaved color patterns necessary to print two-tone images. With our passive color mixing technique we achieve a greater perceived color space and can print two-tone texture mapped objects.

For future work, this method could be extended to exploit three (or more) nozzles. Our technique works very well for mostly vertically aligned shapes and should be broadened to account for rather horizontal surfaces, while investigating how to alter geometry in order to work against viewing angle dependency. There is a large body of research on parameterization methods which could be applied to further enhance mapping quality on arbitrary surfaces. We believe our work will find many applications in new passive color mixing techniques, and even for single color prints, it can be used to create interesting geometric patterns on the shapes of models to give them a new tactile finish.

## References

- [BBJP12] BÄCHER M., BICKEL B., JAMES D. L., PFISTER H.: Fabricating articulated characters from skinned meshes. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 31, 4 (2012), 47:1–47:9. 2
- [BBO\*10] BICKEL B., BÄCHER M., OTADUY M. A., LEE H. R., PFISTER H., GROSS M., MATUSIK W.: Design and fabrication of materials with desired deformation behavior. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 29, 4 (2010), 63:1–63:10. 2
- [CCA\*12] CALÌ J., CALIAN D. A., AMATI C., KLEINBERGER R., STEED A., KAUTZ J., WEYRICH T.: 3d-printing of non-assembly, articulated models. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 31, 6 (2012), 130:1–130:8. 2
- [CLD\*13] CHEN D., LEVIN D. I. W., DIDYK P., SITTHI-AMORN P., MATUSIK W.: Spec2Fab: A reducer-tuner model for translating specifications to 3D prints. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 32, 4 (2013), 135:1–135:10. 2
- [CSPT03] CHO W., SACHS E. M., PATRIKALAKIS N. M., TROXEL D. E.: A dithering algorithm for local composition control with three-dimensional printing. *Computer-Aided Design* 35, 9 (2003), 851–867. 2
- [CTN\*13] COROS S., THOMASZEWSKI B., NORIS G., SUEDA S., FORBERG M., SUMNER R., MATUSIK W., BICKEL B.: Computational design of mechanical characters. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 32, 4 (2013), 83:1–83:12. 2
- [HBA13] HILDEBRAND K., BICKEL B., ALEXA M.: Orthogonal slicing for additive manufacturing. *Computers & Graphics (Proc. SMI)* 37, 6 (2013), 669–675. 2, 8
- [HFM\*10] HAŠAN M., FUCHS M., MATUSIK W., PFISTER H., RUSINKIEWICZ S.: Physical reproduction of materials with specified subsurface scattering. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 29, 4 (2010), 61:1–61:10. 2
- [HIH\*13] HULLIN M. B., IHRKE I., HEIDRICH W., WEYRICH T., DAMBERG G., FUCHS M.: State of the art in computational fabrication and display of material appearance. In *Eurographics State of the Art Reports* (2013). 2
- [LBRM12] LUO L., BARAN I., RUSINKIEWICZ S., MATUSIK W.: Chopper: partitioning models into 3D-printable parts. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 31, 6 (2012), 129:1–129:9. 2
- [LGX\*13] LEVIN A., GLASNER D., XIONG Y., DURAND F., FREEMAN B., MATUSIK W., ZICKLER T.: Fabricating BRDFs at high spatial resolution using wave optics. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 32, 4 (2013), 144:1–144:14. 2
- [SVB\*12] STAVA O., VANEK J., BENES B., CARR N., MĚCH R.: Stress relief: improving structural strength of 3D printable objects. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 31, 4 (2012), 48:1–48:11. 2
- [TJ11] TELEA A., JALBA A.: Voxel-based assessment of printability of 3D shapes. In *Proc. of the international conference on Mathematical morphology and its applications to image and signal processing (ISMM'11)* (2011), pp. 393–404. 2
- [VWRKM13] VIDIMČE K., WANG S.-P., RAGAN-KELLEY J., MATUSIK W.: Openfab: A programmable pipeline for multi-material fabrication. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 32 (2013), 136:1–136:12. 2
- [WWY\*13] WANG W., WANG T., YANG Z., LIU L., TONG X., TONG W., DENG J., CHEN F., LIU X.: Cost-effective printing of 3D objects with skin-frame structures. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 32, 6 (2013), 177:1–177:10. 2