Fabricating Microgeometry for Custom Surface Reflectance

Tim Weyrich

University College London

Pieter Peers

University of Southern California,
Institute for Creative Technologies

Wojciech Matusik

Adobe Systems, Inc.

Szymon Rusinkiewicz
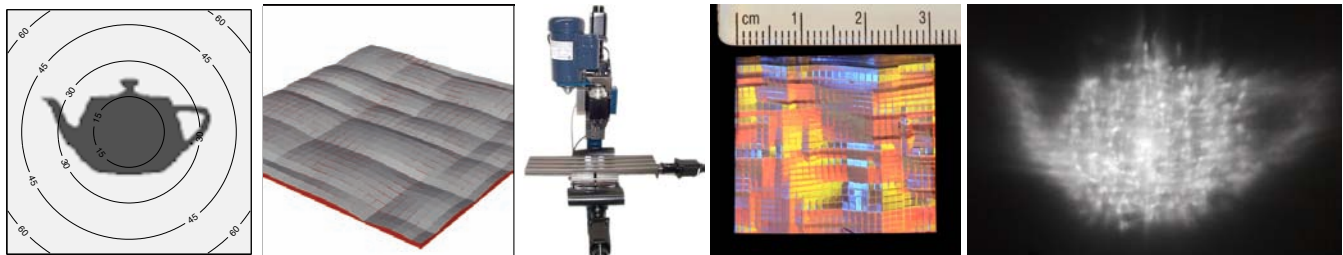
Princeton University,
Adobe Systems, Inc.

Figure 1: *From left: a user-designed highlight is converted to an optimized microfacet height field. A computer-controlled milling machine is used to manufacture the surface (30 × 30 facets, each approximately 1 mm × 1 mm), which exhibits the desired reflectance.*

Abstract

We propose a system for manufacturing physical surfaces that, in aggregate, exhibit a desired surface appearance. Our system begins with a user specification of a BRDF, or simply a highlight shape, and infers the required distribution of surface slopes. We sample this distribution, optimize for a maximally-continuous and valley-minimizing height field, and finally mill the surface using a computer-controlled machine tool. We demonstrate a variety of surfaces, ranging from reproductions of measured BRDFs to materials with unconventional highlights.

1 Introduction

The modeling of surface reflectance is a key area of study within computer graphics, since it is one of the main components that determines appearance. Although simple analytic reflectance models continue to be used in rendering applications, more complex and realistic material models are becoming practical with advances in measurement devices and methods, compact and accurate representations for such measured data, and flexible processing and rendering algorithms. This has created an opportunity to "close the loop" in this avenue of research, by considering the output stage.

We leverage research positing that complex appearance is often determined by small-scale surface structure, or *microgeometry*. Because these geometric features, typically called *microfacets*, are too small to be seen individually, the surface appearance is dictated by their aggregate distribution, in combination with the behavior of the underlying material (metal, dielectric, translucent / scattering, etc.). Many analytic reflectance models used in computer graphics are, in fact, statistical microfacet models [Torrance and Sparrow 1967; Blinn 1977; Cook and Torrance 1982; Ashikhmin et al. 2000].

This paper considers a first step towards reflectance output by creating such microgeometry directly. We compute the desired surface shape algorithmically, then bring it into the real world with a computer-controlled milling machine. In doing so, we take advan-

tage of the increased availability of such devices: consumer-grade milling machines are becoming common, and are rapidly approaching affordability (a few thousand USD). They are marketed to home users, for hobbyist projects such as woodworking.

The ability to manufacture surfaces with user-specified reflectance red may have applications in a variety of fields, especially if it is possible to carefully manufacture dies and then mass-produce materials via stamping. Architects and industrial and product designers could make material *design*, as opposed to purely material *selection*, part of their creative process. They could include milled surfaces, ranging from subtle to unusual and creative, as panels within products or as architectural elements. A famous example of (unintended) shaped highlight creation in architecture is the Berlin television tower, whose highlight is known as the *Pope's Revenge*. Interior designers could specify materials having particular reflectance to control the lighting within a space. In military applications, reflectance could be controlled as a form of camouflage, much as radar reflectance is already controlled in stealth aircraft. Smaller milled surfaces could be used as a form of security marker, much as white-light holograms are used in credit cards today. Finally, companies and advertising agencies may be interested in materials displaying product or corporate logos. In all cases, there is a need to balance predictability and fidelity to the design against practicality and constraints of the manufacturing process. Therefore, in this paper we present examples that may be practical with current consumer-grade devices, which are not able to make the microfacets so small that they become invisible at small viewing distances. As such, we concentrate on applications such as architectural elements and moderate-scale industrial design, for which there may be a greater demand for unusual materials (Figure 1) instead of exact matching of subtle material variations.

Our algorithm begins with a description of the desired material, and finds a probability distribution over the necessary slopes of the surface. We draw samples from this distribution, using a low-discrepancy technique inspired by a stippling algorithm (Section 3.2). We then construct a surface from the individual microfacets, using a simulated-annealing optimization to minimize discontinuities and valleys (Section 3.3). We preview the result of the optimization using a path tracer (Section 3.4), verifying that our assumptions hold and that the effect of the tool shape on the milled surface is tolerable. Finally, we pass the height field to the milling machine (Section 4); a variety of results is shown in Section 5.

2 Related Work

Material Design and Editing: The reflectance of an opaque, homogeneous surface is represented by the Bidirectional Reflectance

Distribution Function or BRDF, which encodes the ratio of reflected radiance in each direction to incident irradiance from each direction [Nicodemus et al. 1977]. Though early BRDF models were simply analytic formulas, Westin et al. [1992] used a "virtual gonioreflectometer" to determine the aggregate BRDF due to arbitrary synthetic microgeometry. Ashikhmin et al. [2000] incorporated an arbitrary statistical distribution of microfacets with a physically plausible shadowing term, yielding an energy-conserving BRDF that could be used in rendering. These papers confirm that microgeometry can give rise to a rich variety of BRDF behaviors, and motivate our use of a microfacet-based optimization pipeline. Nevertheless, they focus on analysis, and do not address synthesis of microfacet distributions.

Colbert et al. [2006] present an interface for BRDF editing that involves drawing highlights directly on a sphere. This motivates our choice to let the user draw an easily understandable 2D distribution as a goal, instead of having to specify a full 4D BRDF. In our system, we do not restrict the user to a lit sphere, but rather provide several possible mappings to unfold a hemisphere into 2D.

Reflector Design: There are similarities between our work and reflector design for luminaires (see [Patow et al. 2007], and [Patow and Pueyo 2005] for an overview). The goal of reflector design is to produce a desired hemispherical light distribution given a fixed (local) position of a point-light source. By moving the local light to infinity, the goal of reflector design would be similar to ours. Nevertheless, there are two important consequences of this restriction to *non-local light* that do not apply to general reflector design: we can make our surfaces tileable and near-planar, both of which are important for practical large-scale applicability.

Milling as an Output Device: Although automatic milling machines are commonly used in computer-aided manufacturing (CAM), both the shapes to be milled and the motion path of the tool are typically specified manually. Automatic tool-path optimization remains an area of active research, but is of limited necessity when milling height fields of relatively low depth range.

Recent projects have considered the generation and output of artist-generated [Sourin 2001] and computer-assisted [Cignoni et al. 1997; Weyrich et al. 2007; Song et al. 2007; Kerber et al. 2007] bas-reliefs. The optimization performed in this paper is unlike those appropriate for bas-reliefs, since our goal is to specify the aggregate effect of an entire patch of surface rather than the spatially-dependent details. Nevertheless, some of the constraints considered by bas-relief synthesis systems (e.g., continuity of height) are due to the physical milling process, and hence are also appropriate constraints in our optimization.

3 Designing Height Fields

We now describe our pipeline for designing height fields for milling, given a target BRDF as input. We begin by stating the assumptions necessary to ensure that the problem is tractable (Subsection 3.1), then describe the microfacet design (Subsection 3.2), sampling, and optimization algorithms (Subsection 3.3). We verify the results of these algorithms using a simulator (virtual gonioreflectometer) based on Monte-Carlo path tracing (Subsection 3.4).

3.1 Assumptions

The range of achievable aggregate BRDFs is a function of the BRDF of the underlying material (which we will call the "base BRDF") and the height field that is milled.

BRDF: We assume a spatially homogeneous base BRDF. This excludes both spatial variation across the surface and in depth (i.e., different layered base materials). Furthermore, we assume that the base material is opaque. While, lifting each of these restrictions would enable exciting possibilities with a greater flexibility in the range of achievable aggregate BRDFs, we focus in this paper on

creating the basis tools necessary to enable these possibilities in future work.

Height Field: Due to physical constraints of the milling hardware and available tools (see Section 4), we restrict the slope of the microfacets to 65 degrees. As a result, shadowing, masking and interreflections are negligible (see Section 3.4 for a verification). Consequently, we will ignore these throughout this paper. Additionally, in order to be able to create large continuous surfaces, the height fields need to be tileable, and no significant height discontinuities should be introduced at the tile edges.

Under the above assumptions, the aggregate BRDF is simply the sum of many copies of the base BRDF, each rotated according to the normal of one of the microfacets. To derive the net effect of this, it is easiest to work in the half-angle / difference-angle parameterization proposed by Rusinkiewicz [1998]. As shown by Han et al. [2007], in this parameterization, the half-angle component of the aggregate BRDF is the spherical convolution of the half-angle component of the base BRDF with the microfacet normal distribution function, while the difference-angle component is simply carried over from the base BRDF.

The connection between microfacet normals and half-angle distributions is intuitive, since the half-angle represents the normal of a microfacet that, under ideal specular reflection, would reflect the incident ray into the exitant direction. Therefore, the probability density of microfacets with such a normal represents the relative brightness of the reflection, for the given light and view directions.

Given the spherical-convolution formulation, we are now able to state some limitations on the range of target BRDFs we can produce. First, the user only has control over the half-angle component of the target; its difference-angle dependency is fixed. Second, the target BRDF must be band-limited according to the non-zero frequency components of the base BRDF. For the near-specular base BRDFs typical of polished metals, this is not a significant limitation. Finally, it is not possible to change the albedo (net hemispherical reflectance) of the base BRDF: a material that reflects most of its incident light cannot absorb most of it after milling (unless the milling introduces significant interreflection). Related to this latter restriction is that the user-specified microfacet distribution must be a true probability distribution: it must be normalized to unit area.

3.2 From Highlight Shape to Microfacet Distribution

Inspired by Colbert et al. [2006], we propose to design a BRDF by simply painting the desired specular highlights via an intuitive design interface. This can be performed either on an orthographic view of a sphere assumed to have the desired material and illuminated by a headlight, or directly on the hemisphere of exitant directions parameterized, for example, using the parabolic projection [Heidrich and Seidel 1998]. A distribution over half-angles can then be trivially constructed from the user-specified distribution of reflected intensities. Alternatively, we can also extract half-angle distributions from measured materials, such as brushed metals, paints, and cloth (silk/satin/velvet).

The next step is to convert the half-angle distribution into the desired microfacet (normal) distribution, by accounting for the effect of the base BRDF. As noted in Section 3.1 the effect of the base BRDF can be seen as a convolution. Compensating for this convolution effectively requires solving a deconvolution problem. We use the iterative Lucy-Richardson deconvolution algorithm [Richardson 1972; Lucy 1974] using the half-angle component of the base BRDF as the convolution kernel. Because the kernel is strictly positive, no ringing artifacts occur. Figure 2 shows an example of a user-provided target highlight shape, and the deconvolved microfacet distribution, all displayed using a parabolic mapping of the
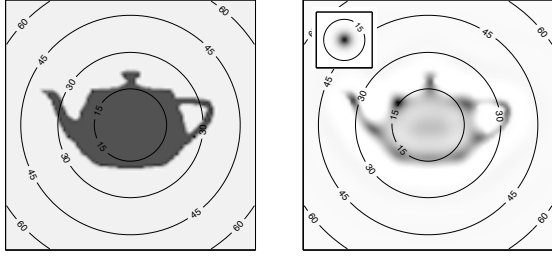
Figure 2: Left: *User-provided target highlight.* **Right:** *Deconvolved microfacet distribution with the base BRDF (inset).*

hemisphere[1]. The deconvolution kernel (i.e., half-angle component of the base BRDF) is shown in the inset.

3.3 From Microfacet Distribution to Height Field

While in principle there may exist an infinite number of height fields producing the same microfacet distribution, not every facet distribution corresponds to a continuous tileable surface. This is similar in spirit to the question whether a gradient field is integrable, except that the spatial distribution of the gradients (facet normals) is not known. It can be shown that a necessary condition to obtain a continuous tileable height field is that any integral over a straight line through the origin of the microfacet distribution should equal zero. This constraint is trivially fulfilled if the distribution is symmetric around the surface normal. Meeting the more general precondition, however, would rule out many desired highlight shapes, such as in Figure 2. Luckily, we find that it is sufficient to enforce a weaker constraint, by rotating the microfacet distribution such that its mean is perpendicular to the surface, and still creates microfacet distributions that can be translated into a height field at reasonable precision.

Given a microfacet distribution we create a height field as follows. First we sample the microfacet distribution. Since the order of the microfacets does not impact the distribution, we can then optimize their order to maximize tileability, and minimize discontinuities. Finally, we solve for the optimal height of each facet.

Sampling: In a first step we sample the microfacet distribution to transform it to a discrete set of microfacet orientations.

A possible sampling strategy for microfacet distributions could employ importance sampling. While in the limit this generates a height field with the desired distribution, it is possible to achieve greater fidelity and lower noise for a limited number of samples by using low-discrepancy sampling techniques. A particular challenge is that we cannot control the "brightness" of of each microfacet's reflection, i.e., light is reflected according to the base BRDF, which is assumed to spatially homogeneous. A very similar problem to microfacet distribution sampling is image stippling, where there is also no control over the "brightness" of the stipples. Similar to [Secord 2002], we will employ a centroidal Voronoi tessellation technique to place the samples proportional to the local density, while maintaining a good global distribution of the samples. Figure 3 compares random importance sampling versus the low-discrepancy sampling for height fields for 20×20 samples drawn from the target microfacet distribution depicted on the left.

Optimization: Once we have the desired set of microfacets, we need to arrange them in the *xy*-plane in a tileable and manufacturable configuration. This problem is closely related to solving the inverse patch transfer problem [Cho et al. 2008], which is shown to be NP-complete. There exist numerous methods for finding approximate solutions to such combinatorial optimization problems. We

---

[1] We tabulate half-way instead of reflectance vectors, and thus the parabolic map corresponds to the full hemisphere of reflection directions. All plots show only a central region, specified by lines of equal reflectance angle.

opt for a simulated-annealing optimization, since it allows flexibly incorporating non-local constraints, as we will describe below.

As, at this point, the facets' heights $z(x, y)$ are not yet known, this optimization considers first derivatives (facet slopes) only, trying to rearrange discrete slopes to meet necessary preconditions for a smooth, manufacturable surface. This goal is formulated as a minimization of three heuristical energy functions. The first energy function, $C = C_x + C_y$, penalizes slope incompatibility between adjacent facets in x and y direction, respectively:

$$C_x = \frac{1}{4} \sum_y \sum_x \left\| \frac{\mathrm{d}z(x+1, y) - \mathrm{d}z(x, y)}{\mathrm{d}y} \right\|^2, \qquad (1)$$

and C_y analogously. This enforces that neighboring facets have a similar slope along their common edge, which makes parallel rows of the height field slope more synchronously. Note that this term only considers directional derivatives orthogonal to those that are part of the final height optimization discussed below. This reduces the chance of both optimizations to interfere adversely. We do, however, enforce integrability along entire rows and columns by simultaneously minimizing a second energy function $I = I_x + I_y$ ensuring that in a cyclic arrangement, the derivatives along each row or column sum to zero,

$$I_x = \sum_y \left\| \sum_x \frac{\mathrm{d}z(x+1, y) - \mathrm{d}z(x, y)}{\mathrm{d}x} \right\|^2, \qquad (2)$$

and I_y correspondingly.

A third constraint is to compensate for the shape of the milling bits. The extent of these bits is nonzero, and as such we cannot manufacture arbitrary concave shapes. Where neighboring microfacets form a valley, the bit that we use leaves narrow horizontal strips that introduces erroneous angles of reflection. To minimize this effect, we add an additional term $V = V_x + V_y$ to the energy function that adds a constant penalty whenever two neighboring facets' slopes build a concave angle:

$$V_x = w_V \left| \left\{ (x, y) \mid \frac{\mathrm{d}z(x+1, y)}{\mathrm{d}x} > \frac{\mathrm{d}z(x, y)}{\mathrm{d}x} \right\} \right|, \qquad (3)$$

and V_y analogously, with w_V a weighting term balancing the influence of the valley minimization relative to the other penalty terms.

Starting from a random $w \times h$ assembly of our microfacets we use a simple simulated-annealing process to minimize the joint energy function

$$E = \frac{C + V}{(wh)^2} + \frac{I}{(w + h)^4}. \qquad (4)$$

The weights are chosen so that the effect of C and I approximates the squared surface area of verticals due to depth discontinuities. Furthermore, we have empirically determined that a w_V of $1/2000$ reduces the number of valleys by 50 to 60 percent, without noticeable increase in C and I. Higher weights decrease the number of valleys only slightly, at the cost of a dramatic increase in discontinuities in the resulting microfacet field.

Integration: Using the optimal 2D arrangement of the microfacets from the previous optimization, we subsequently optimize for maximal continuity of the resulting height field by shifting the
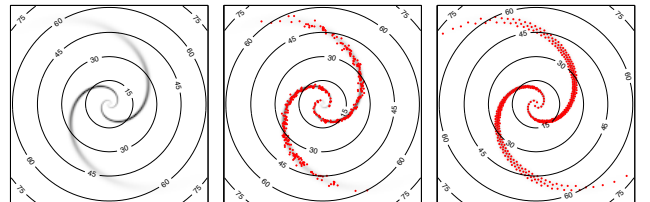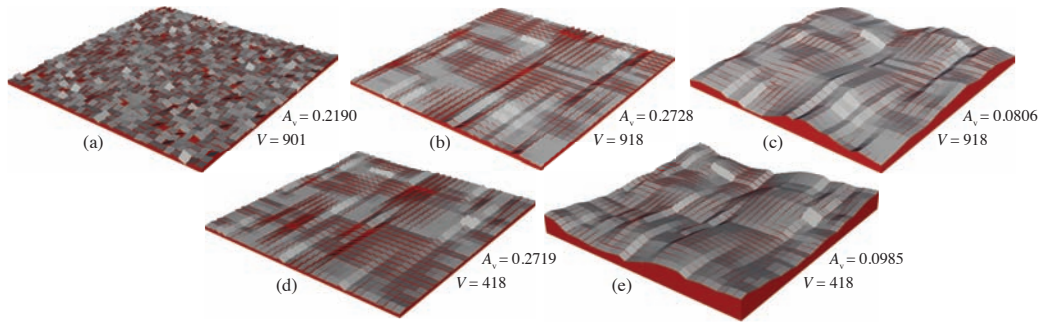

Figure 3: Left: *Target microfacet distribution.* **Center:** *Random sampling with 400 samples.* **Right:** *Low-discrepancy sampling with the same number of samples.*

Figure 4: *Microfacet array for the "spiral" MFD in Figure 3.* (a) *Initial arrangement.* (b,c) *Without valley penalty, after position optimization (b) and after height optimization (c).* (d,e) *With valley penalty, after position optimization (d) and after height optimization (e).*

individual facets along the z direction to their optimal height. This can be expressed as a discrete Poisson equation,

$$\nabla^2 z = \operatorname{div} \boldsymbol{g} , \qquad (5)$$

with a discrete gradient field $\boldsymbol{g} = (\boldsymbol{g}_x, \boldsymbol{g}_y)$ that expresses the condition of seamless facet connectivity in the x direction as,

$$\boldsymbol{g}_x(x,y) = \frac{1}{4}\frac{\mathrm{d}z(x-1,y)}{\mathrm{d}x} + \frac{1}{2}\frac{\mathrm{d}z(x,y)}{\mathrm{d}x} + \frac{1}{4}\frac{\mathrm{d}z(x+1,y)}{\mathrm{d}x} , \qquad (6)$$

and analogously in y. As we target a tileable arrangement (i.e., the system is cyclic), a Dirichlet boundary condition at a single point is sufficient to obtain a solution.

After solving the Poisson problem based on the output of the annealing simulation, we finally obtain a field of microfacets that approximates a C^0 surface. Figure 4 shows five 30×30 height fields, with microfacet orientations given by the sampling of Figure 3. From left to right, we show the initial (random) arrangement, the results of position optimization, and of the height optimization. The top row has been optimized without the valley penalty, the bottom row with. Shown next to each arrangement is its vertical surface area A_v (a measure for discontinuities) and the valley term V. Note that without position optimization, the result of the Poisson reconstruction would yield a much larger A_v of 0.1782 and $V = 901$.

3.4 Verification

In order to verify the effectiveness of our sampling pipeline, we simulate reflection from the resulting height fields using a virtual gonioreflectometer. In contrast with the rest of our processing pipeline, this stage uses a physically-correct ray tracer that incorporates all effects such as shadowing, masking, and interreflection. Figure 5 presents such renderings for the low-discrepancy sampling shown in Figure 3, comparing a material simulation with a photograph of a reflection off an actual milled material.

3.5 Discussion

Separating the optimization into two discrete stages may appear unnecessary. Jointly optimizing facet ordering and individual heights, however, would increase dimensionality of the search space and
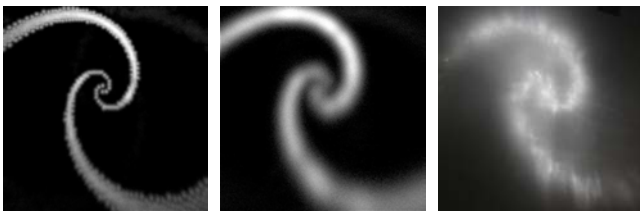
with it the number of local minima. We argue that by separating the optimization into its inherently NP-hard part and into a linear problem, the optimization becomes more tractable.

Also, as noted earlier, there is a similarity between our system and reflector design: they both optimize for surfae shape, given certain objectives and constraints. The key difference is that we take the angular distribution of surface slopes as a hard constraint and optimize for smoothness and manufacturability, while reflectance design takes smoothness as a hard constraint and optimizes for matching the desired angular distribution. This allows us to perform low-discrepancy sampling, to tolerate small surface discontinuities that do not adversely affect manufacurability, and to incorporate additional cost functions such as tileability.

4 Milling Height Fields

Milling Machine: For our initial experiments we have used a MicroMill 2000 by MicroProto Systems, a desktop Computer Numerical Control (CNC) milling machine with 3 axes of motion, which is shown in Figure 1. The maximum operating range is 9, 5.75, and 6 inches in the x, y, and z directions, respectively, which is much greater than the size of the height fields we have milled. The stepper motor has a resolution of 0.000125 inch, and its repeatability is 0.0005 inch. We use the Mach3 control software to execute our GCode (a programming language for CNC machine tools).

Each height field tile is a 1.2-inch square (x/y), and the maximum height (z) varies between 0.04 and 0.12 inches. Each height field contains $30 \times 30 = 900$ microfacets. Therefore, the size of each microfacet is 0.04×0.04 inches, or roughly a square millimeter. Each height field is milled in two passes, first in x- then in y-scanline order. Two passes are necessary to reduce the effect of the tool path on the shape of the resulting height field. We mill the surface at 0.001 inch resolution in both x and y directions (i.e., we upsample each microfacet to a plane, using samples with 0.001 inch spacing). Therefore, in one pass we execute about 144,000 lines of GCode. A typical milling speed is 9 inches per minute, and the total milling time for two passes is about 5.5 hours.

We have experimented with a number of different substrates for milling, mainly metals and plastics. The results in this paper are all milled in aluminum, because it is soft and therefore easy to mill. Moreover, it is not harmful to humans, in contrast with some alternative metals (including most colored ones). Once the surface has been milled, we use a standard polishing compound to remove any small imperfections and to make the microfacets near-specular.

Figure 5: Left: *Simulated reflectance off the material against a wall, using a point-light source and assuming a near mirror-like base BRDF.* **Center and Right:** *Verification of the simulation by simulating a more realistic base BRDF (center) to match a photograph of a reflection from a physical sample (right).*

We have also experimented with a variety of different drill bits, ultimately choosing an inexpensive bit that has high durability and small size: a high-grade steel bit by Dremel (125 High Speed Cutter), shown on the right. This bit has a conical shape, with a sharp yet durable tip and a small opening angle (about 38°). In order to minimize breakage, we use coolant on the milling surface during machining.
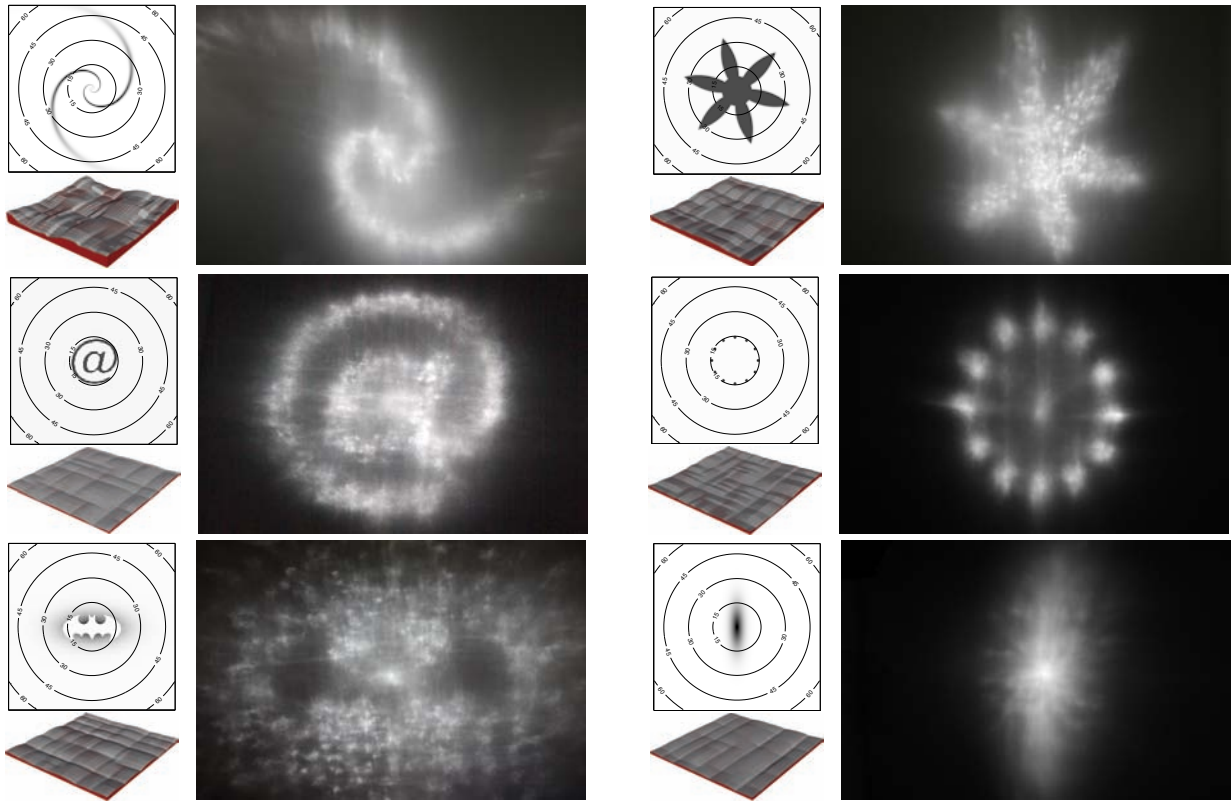
Figure 6: Results of surface manufacturing. The bottom-right sample is a measured brushed aluminum BRDF, while the remaining examples were user created. Within each example, we show the following: Top-left: The target half-angle distribution specified by the user. The iso-lines denote equal reflection (not half-way) angles. Bottom-left: The height field output by our algorithm. Right: Photograph of the reflection produced by the milled surface, under directional illumination. The imaging configuration is shown at right: a projector illuminates the sample, and the reflection is visible on a wall parallel to the sample. Note that the large size of the projector's aperture introduces additional blurring in the results.

Measuring Base BRDF: To measure the "base BRDF" of our milled aluminum, we first mill a constant-height surface using the techniques described above. Because the surface is close to ideal-specular, we adopt the approach of shining a laser onto the surface (at near-normal incidence) and observing its reflection on a plane parallel to the surface. We found that the reflection varied slightly from location to location on the surface, but the width at half-maximum intensity was approximately 2 ± 1 inches, at a distance of 96 inches. This corresponds to a spread of 0.02 radians (approximately 1 degree), which is smaller than the typical size of elements in our target half-angle distributions. For this reason, the results in this paper did not require deconvolution, though surfaces milled using different techniques or different materials might do so.

We also verified that after two passes (in the *x* and *y* directions) the base BRDF was isotropic, to within measurement precision. This is in contrast to the results we obtained with one milling pass, which did exhibit anisotropy (approximately 2:1 elongation of the highlight). Although our processing pipeline can incorporate anisotropic base materials, we chose to perform two-pass milling for the results presented in this paper.

5 Results and Discussion

Figure 6 shows results for a number of user-created distributions, in addition to a brushed-aluminum BRDF, fit to measured data. For each example, we show the target distribution, optimized height field, and a photograph of the reflection. As shown at lower right,

we use a configuration in which a projector illuminates the milled portion of the surface, causing a reflection to appear on the wall.

Discussion of Results: The results exhibit a number of imperfections, which are due to the limitations of the milling process in reproducing the desired height field. First, several of the results display a strong central highlight that is not present in the desired BRDF. By examining the milled surfaces in detail, we have determined that these highlights are due to the nonzero extent of the milling bit. Where a pair of adjacent microfacets forms a valley, the bit leaves a strip of surface that is horizontal (parallel to the *xy*-plane) instead of milling a perfect, sharp concavity. Although we attempt to minimize this effect by including an additional term in our optimization, as described in Section 4, we are unable to eliminate the effect completely. (We have verified that *not* including the valley-minimization term in the optimization results in a significantly stronger central highlight.) We believe that surfaces manufactured on more accurate, commercial-grade milling machines would exhibit a greatly reduced central highlight. A second artifact is the presence of spurious cross-shaped (horizontal and vertical) highlights, which are most visible on the twelve-star configuration (middle right). We believe that these are due to a combination of the remaining vertical discontinuities in the height field with the inability to mill such vertical surfaces.

Limitations: A major limitation on the highlight shapes we are able to produce is due to the resolution at which we sample the

microfacet distribution. All of the height fields shown in this paper included only 30×30 facets, and large areas of uniform intensity are sampled poorly with so few points. This effect is most visible in the bat logo at lower left and the teapot in Figure 1. The situation is analogous to stippling of an image where obtaining an accurate approximation of large, uniform areas requires more samples.

Further limitations on the achievable range of BRDFs may be due to the base BRDF or, more significantly, on the restriction to microfacet distributions with zero line integrals through the origin. We found, however, that slight deviations from this condition, such as in the example in Figure 1, still lead to good reconstructions in practice; the major downside is the slight increase in the vertical discontinuities in the facet field.

6 Conclusions and Future Work

Although this represents only a first investigation into the possibilities of material output, our results suggest that it is possible to manufacture surfaces with user-specified reflectance, using off-the-shelf equipment. With higher-precision tooling, higher fidelity should be achievable by reducing the size of microfacets, and increasing their number and precision. Our algorithms for producing height fields from target BRDFs, encompassing low-discrepancy sampling, simulated-annealing optimization, and discontinuity minimization via the Poisson equation, are efficient and should scale to the design of larger and more complex surfaces.

One immediate avenue of future work is to manufacture larger patches of microfacetted surface. By simply tiling one of our height fields many times, it should be possible to produce surfaces that, at sufficient viewing distances, allow shaped highlights to be observed directly on the surface (as opposed to being observable only in the reflection onto another surface). This is because variation in viewing direction (for a finite observer) would cause different angles of the BRDF to be observed at different surface locations, even for a single observer. This effect would, in fact, be facilitated by milling a curved surface, as shown at right. This would provide more rapid variation in the part of the BRDF that is observed, thus allowing shaped highlights to be observed more easily.

There are also alternative microfacet manufacturing technologies that could be explored, possibly including laser-engraving, chemical etching, and water-jet milling. For mass production, we anticipate the possibility of milling a master die, then stamping many replicas (both tiled on a single surface and allowing for rapid production of many exemplars). Some of these technologies would allow for easier production of surfaces with significant slopes, offering the hope of a greater range of achievable appearances, at the cost of a more complex optimization (which would have to account for non-negligible shadowing/masking/interreflection — a full inverse-rendering problem). Finally, many of these technologies would allow for a wider range of base materials, further increasing the range of appearances available to a designer.

Acknowledgments

References

ASHIKHMIN, M., PREMOŽE, S., AND SHIRLEY, P. 2000. A Microfacet-Based BRDF Generator. In *Proc. of ACM SIGGRAPH*, 65–74.

BLINN, J. 1977. Models of Light Reflection for Computer Synthesized Pictures. In *Proc. of ACM SIGGRAPH*, 192–198.

CHO, T. S., BUTMAN, M., AVIDAN, S., AND FREEMAN, W. T. 2008. The Patch Transform and its Applications to Image Editing. In *Proc. IEEE CVPR*, 1–8.

CIGNONI, P., MONTANI, C., AND SCOPIGNO, R. 1997. Automatic Generation of Bas- and High-Reliefs. *Journal of Graphics Tools, Vol. 2, No. 3*, 15–28.

COLBERT, M., PATTANAIK, S., AND KŘIVÁNEK, J. 2006. BRDF-Shop: An Artistic Tool for Creating Physically Correct BRDFs. *IEEE CG&A, Vol. 26, No. 1*, 30–36.

COOK, R. L., AND TORRANCE, K. E. 1982. A Reflectance Model for Computer Graphics. *ACM Trans. Graph., Vol. 1, No. 1*, 307–316.

HAN, C., SUN, B., RAMAMOORTHI, R., AND GRINSPUN, E. 2007. Frequency Domain Normal Map Filtering. *ACM Trans. Graphics (Proc. ACM SIGGRAPH), Vol. 26, No. 3*.

HEIDRICH, W., AND SEIDEL, H.-P. 1998. View-Independent Environment Maps. In *Proc. Graphics Hardware*, 39–45.

KERBER, J., BELYAEV, A., AND SEIDEL, H.-P. 2007. Feature Preserving Depth Compression of Range Images. In *Proc. of the 23rd Spring Conference on Computer Graphics*, Comenius University, Slovakia, 110–114.

LUCY, L. B. 1974. An iterative technique for the rectification of observed distributions. *Astron. Journal, Vol. 79, No. 6*, 745.

NICODEMUS, F. E., RICHMOND, J. C., HSIA, J. J., GINSBERG, I. W., AND LIMPERIS, T. 1977. *Geometrical Considerations and Nomenclature for Reflectance*. National Bureau of Standards, U. S. Department of Commerce.

PATOW, G., AND PUEYO, X. 2005. A Survey of Inverse Surface Design From Light Transport Behavior Specification. *Computer Graphics Forum, Vol. 24, No. 4* (Dec.), 773–789.

PATOW, G., PUEYO, X., AND VINACUA, A. 2007. User-Guided Inverse Reflector Design. *Computers & Graphics, Vol. 31, No. 3*, 501–515.

RICHARDSON, W. H. 1972. Bayesian-based iterative method of image restoration. *JOSA, Vol. 62, No. 1*, 55–59.

RUSINKIEWICZ, S. 1998. A New Change of Variables for Efficient BRDF Representation. In *Proc. EGRW*, 11–22.

SECORD, A. 2002. Weighted Voronoi Stippling. In *Proc. NPAR*, 37–43.

SONG, W., BELYAEV, A., AND SEIDEL, H.-P. 2007. Automatic Generation of Bas-Reliefs from 3D Shapes. In *Proc. SMI*, 211–214.

SOURIN, A. 2001. Functionally Based Virtual Computer Art. In *Proc. I3D*, 77–84.

TORRANCE, K. E., AND SPARROW, E. M. 1967. Theory for Off-Specular Reflection from Roughened Surfaces. *JOSA, Vol. 57, No. 9*, 1104–1114.

WESTIN, S., ARVO, J., AND TORRANCE, K. 1992. Predicting Reflectance Functions from Complex Surfaces. In *Proc. ACM SIGGRAPH*, 255–264.

WEYRICH, T., DENG, J., BARNES, C., RUSINKIEWICZ, S., AND FINKELSTEIN, A. 2007. Digital Bas-Relief from 3D Scenes. *ACM Trans. Graphics (Proc. ACM SIGGRAPH), Vol. 26, No. 3*.