

Goal-based Caustics

Marios Papas¹ Wojciech Jarosz¹ Wenzel Jakob¹ Szymon Rusinkiewicz² Wojciech Matusik¹ Tim Weyrich³

¹ Disney Research Zürich ² Princeton University ³ University College London

Abstract

We propose a novel system for designing and manufacturing surfaces that produce desired caustic images when illuminated by a light source. Our system is based on a nonnegative image decomposition using a set of possibly overlapping anisotropic Gaussian kernels. We utilize this decomposition to construct an array of continuous surface patches, each of which focuses light onto one of the Gaussian kernels, either through refraction or reflection. We show how to derive the shape of each continuous patch and arrange them by performing a discrete assignment of patches to kernels in the desired caustic. Our decomposition provides for high fidelity reconstruction of natural images using a small collection of patches. We demonstrate our approach on a wide variety of caustic images by manufacturing physical surfaces with a small number of patches.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Line and curve generation

1. Introduction

Recently, increased attention has been paid to the design of physical materials with desired reflectance or refraction properties [WDB*07, WPMR09, MAG*09, DWP*10, FHM*10, FDL10]. In this paper, we propose a novel system for designing and manufacturing surfaces that produce desired caustic images when illuminated by a light source. This allows, for instance, to create a translucent object that, when lit by direct sunlight, projects a custom image on a projection plane behind it. (This is different from projecting light through a slide, as light is redirected rather than absorbed, maintaining higher light levels and contrast.)

A traditional application of reflective caustic design is luminaire design: the creation of lamp reflectors that create a desired radiance distribution for a given light source position. Other applications could be for instance in architecture (pattern glass that casts attractive caustics into a room) or security features, creating a hard-to-forge artefact that can be “read out” under normal daylight.

Previous work on physical surface reflectance or refraction design that would be capable of creating caustic patterns either uses a bottom-up approach, constructing geometry from discrete building blocks [WPMR09], or uses a global analysis-by-synthesis approach [PP05, PPV07, FDL10]. The bottom-up approach suffers from quantization artifacts due to their choice of primitives, as we will show, while global,



Figure 1: We design and mill piecewise-smooth arrays of micropatches which reshape incident light onto a collection of Gaussian caustics, to approximate arbitrary images at a projection plane. Above, we show a photograph of a physical sample projecting the *Lena* image.

continuous optimizations cannot guarantee convergence.

In this paper, we propose a novel, discrete (bottom-up) approach for surface generation that band-limits the constructed appearance and thus avoids quantization artifacts while guaranteeing convergence. In addition, we show that our surfaces are more amenable to milling than the current state of the art in caustic generation [FDL10], at the same

time providing higher contrast.

At the core of our system is a nonnegative image decomposition into (possibly overlapping) anisotropic Gaussian kernels, the superposition of which yields an aliasing-free reconstruction of an input image. From this decomposition, we create a surface consisting of an array of continuous surface elements, each of which focuses light onto one of the kernels, forming a Gaussian caustic through either refraction or reflection. In analogy to the concept of microfacets, we call these smooth surface elements *micropatches*. Together, these micropatches create a joint caustic that approximates the input image.

We show how to derive the shape of each micropatch and how to near-optimally arrange such patches so that they project onto the kernels in the desired caustic. Our decomposition provides for high fidelity reconstruction of natural images even for small collections of patches. We compare the results of our algorithm to previous work and demonstrate our approach on a wide variety of caustic images by manufacturing physical surfaces with a small number of patches.


2. Related Work

The process of finding the underlying structure that matches the high-level reflection or refraction specification have been explored by a number of researchers. The computer-based optimization approaches for designing custom reflectors have been investigated since the early 1990s (see, for example, Neubauer [Neu94]). The most common approach uses analysis-by-synthesis, where the simulation framework evaluates the light distribution for a given state and compares it to the desired distribution. Typically the reflector surface is assumed to be continuous and therefore continuous optimization approaches can be used. In this context, the work by Patow et al. [PP05, PPV07] use SAN optimization to obtain the reflector surface for a given desired distribution and a fixed light position. As the search space is inherently large and difficult to optimize researchers have used GPUs to decrease the simulation time [MMP09, FDL10], have used more efficient surface representations (e.g., B-Splines [FDL10], NURBS [MMP09, ASG08]), have reduced the search space to radially symmetric reflectors [ASG08], and they have used more efficient optimization strategies (for example, SPSA [FDL10]).

Our system is the most similar to Weyrich et al.'s approach for the fabrication of microgeometry [WPMR09], in that we use a regular assembly of discrete surface elements that redirect light. We also use a similar optimization framework to find a near-continuous arrangement of these elements. In contrast, however, we use higher-order elements (reflectors/refractors with a Gaussian caustic), which leads to better tonal reproduction, and our optimization has to consider local light sources and local target points when orienting these elements, as we will discuss.

3. Overview

The final output of our system is a regular array of (reflective or refractive) micropatches, where, under fixed illumination conditions, each micropatch generates a Gaussian caustic and where all elementary caustics sum up to form a desired input image



in a known projection plane. Figure 1 shows a photograph of this and we illustrate this setup in the inset above.

In the remainder, we present the components of this system. Section 4 describes our nonnegative image decomposition into Gaussian kernels. Section 5 derives the construction of an individual micropatch that generates a desired Gaussian caustic. Finally, in Section 6 we present a global optimization scheme to implement the Gaussian image decomposition with micropatches, while considering physical constraints of a fabrication process.

The system described here works equally well for reflective and refractive surfaces. For simplicity, we will focus on the more complex effect of refraction, while the reflective case follows by analogy.

4. Gaussian Image Decomposition

We approximate images using nonnegative linear combinations of anisotropic Gaussian kernel functions. In the context of normalized statistical distributions, such representations are usually referred to as Gaussian mixture models (GMMs).

The parameters to be computed in our application consist of the weights w_i , means μ_i and covariance matrices Σ_i of a 2-dimensional m -term GMM of the following form:

$$f(\mathbf{x}) = \sum_{i=0}^m \frac{w_i}{2\pi\sqrt{\det \Sigma_i}} \exp\left(-\frac{1}{2}(\mathbf{x} - \mu_i)^T \Sigma_i^{-1} (\mathbf{x} - \mu_i)\right) \quad (1)$$

Rather than finding the parameters of this model using general nonlinear optimization techniques (as done by Tsai and Shih [TS06] for approximation environment maps using spherical radial basis functions), we use expectation maximization (EM) [DLR*77] specialized to GMMs, which robustly converges to high-quality approximations using a comparably small number of Gaussians terms.

Each of the summands in Equation (1) eventually leads to a micropatch on the milled surface. To facilitate the manufacturing process, it will be useful to introduce the requirement that all of the patches collect an equal amount of energy, allowing them to be manufactured with identical sizes. In the above expression, this constraint translates into setting $w_i := \frac{1}{m}$ ($i = 1, \dots, m$).

Note that instead of directly operating on the input image, the algorithms in this section instead make use of a dense




Figure 2: We approximate images using a linear combination of anisotropic Gaussian kernels. An initial estimate of their parameters is computed and iteratively improved using a sparse version of expectation maximization.




Figure 3: Our image representation produces high-quality approximations using comparatively few Gaussian terms. The obtained quality continuously improves as more terms are added. In this example, the number of samples \mathbf{x}_i is held constant, so that sampling noise becomes visible for large numbers of Gaussians (d).

point representation, which is obtained by sampling a large number of points proportional to the pixel luminances (we use $n = 4,000,000$ in our examples, in order for it to be orders of magnitudes larger than m , the number of patches). In the following, these will be referred to as \mathbf{x}_i ($i = 1 \dots n$).

4.1. Initialization using CCVT

It is important to note that while EM is guaranteed to converge, it will generally only find a locally optimal solution to Equation (1). Hence it is important to supply it with a reasonable initial parameter estimate. We experimented with several strategies for computing such an estimate, and obtained the most reliable results using regular point sets that adapt to the image intensity.

Such a point set can easily be created by computing a Capacity-Constrained Voronoi Tessellation (CCVT, [BSD09]). Following that, the centroid of each Voronoi region is chosen as the center of a Gaussian in the initial parameter set. (This bears similarity to the stippling approach by Weyrich et al. [WPMR09]. Note, however, that in our approach, stipple point locations are only a starting guess to guide the optimization of a more complex image approximation.)

Initially, Gaussian kernels are centered at the stipple points and initialized to an isotropic covariance $\Sigma_i = \text{diag}(\sigma^2, \sigma^2)$, with σ set to the radius of a k -nearest neighbor query across sample points \mathbf{x}_i near its center, and



Figure 4: Starting from the initial estimate, our optimization can compute either isotropic or anisotropic approximations. Due to the significant qualitative improvement, we always use anisotropic Gaussians for our surfaces.

with k equal to the capacity of the associated Voronoi region. This leads to a similar weighted contribution of samples to each Gaussian's region.

4.2. Refinement using EM

The EM algorithm improves the initial estimate using a sequence of alternating E and M steps; the E step computes the probabilities of each sample conditioned on the current set of parameters, whereas the M step maximizes the expected log likelihood of the parameters subject to the previously computed probabilities.

The standard formulation of EM considers all pairs of Gaussians and sample points. In our context, this would lead to prohibitive running times and storage requirements, given the large numbers of Gaussians and samples involved. In-

stead, we disregard pairs whose distance exceeds a certain threshold, on the grounds that our Gaussians are comparatively narrow and approach zero for larger distances. This effectively turns each E and M step into a sparse problem. By using a kd-tree to look up relevant pairs, we speed up the EM convergence from $O(n^2)$ to $O(n \log n)$, if we assume m and n to be proportional to each other. In addition, we use a sparse matrix representation to pass probabilities into the M step, to avoid $O(n^2)$ memory usage.

For the number of sample points and Gaussians used in this paper, each pair of EM-iterations took approximately 17 seconds on a 3Ghz Intel Core 2 Duo machine. In all cases, we ran between 50 and 100 iterations. The aforementioned sparsity optimizations were crucial to make the fitting process practical, since the $O(n^2)$ memory costs of the unmodified method quickly exhausted our machine's 8 GiB of memory.

5. Micropatch Design

We wish to compute the shape of a square surface patch which produces a 2D Gaussian irradiance footprint on a projection surface when illuminated with incident lighting. We call such a surface patch a micropatch. In the reflective case, this will be a minute concave reflector, in the refractive setting a plano-convex microlens. We start out by determining the normal field of the desired patch, obtaining the final shape through subsequent integration. As stated before, we are confining the description to the refractive case.

Assumptions: We assume the configuration illustrated in Section 3, showing incident lighting, the milled surface, and the projection plane. Without loss of generality we describe the case where the projection surface is a plane parallel to the micropatch (at $z = -D$), but the approach is the same given an appropriate 2D domain on an arbitrary curved projection surface. We orient our surfaces to have the milled side facing the projection plane and the flat side facing the light. For simplicity of exposition, we assume distant illumination coming from the z^+ direction and can therefore ignore the first refraction (on the un-milled side of the surface). This is not a restriction of our method, and we can in fact account for local light sources and refractions at both interfaces (by solving for the first refraction point as done by Walter et al [WZHB09]), with the remaining procedure largely unchanged.

Approach: To compute the surface normal of the facet we first need a mapping between points on the micropatch, \mathbf{g} , and the corresponding points, \mathbf{p} , on the projection plane. To enforce that the resulting irradiance profile is a 2D Gaussian, the determinant of the Jacobian of the mapping needs to be a 2D Gaussian. Given this continuous correspondence, we analytically solve for the normal direction producing the required refraction direction at any point of the micropatch, and we then integrate a discretized version of this normal

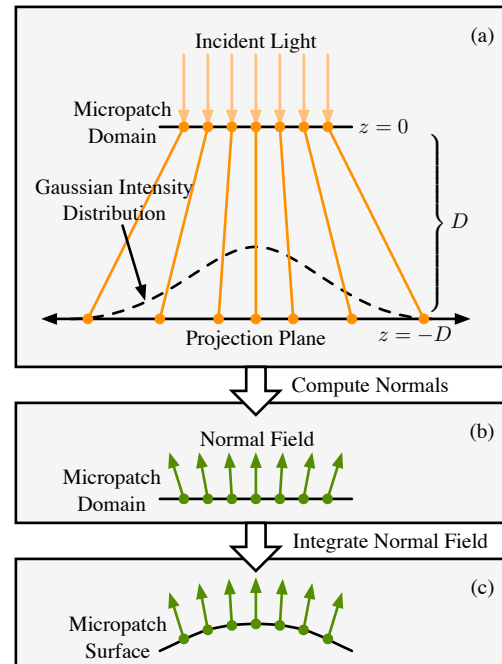


Figure 5: We (a) define a bijective mapping between points in the micropatch domain and points on the projection plane, such that the irradiance distribution follows a Gaussian profile. We then (b) analytically compute the surface normals that refract/reflect the light in this way. We integrate this normal field to arrive at the required micropatch surface (c).

field to obtain the elevated micropatch surface. We illustrate this process in Figure 5.

5.1. Irradiance Warp

We first find a mapping from any point in the micropatch domain to a zero-mean unit Gaussian at the projection plane \mathbf{q} . We subsequently explain how we modify this mapping to account for anisotropy and arbitrary means.

Since our micropatch is square, we parametrize positions as $\mathbf{g}(s, t) = [\mathbf{g}_x, \mathbf{g}_y, \mathbf{g}_z]^T$ where $(s, t) \in [0, 1]^2$. We therefore need a transformation from the uniform square to a standard 2D normal distribution. We could use the Box-Muller transformation [BM58] to map points on the patch to points on the projection plane. However, we wish this mapping to be “well behaved” (i.e. continuous, low-distortion, and without singularities) to ensure the resulting normal field is integrable. Unfortunately, the original Box-Muller transform does not satisfy these requirements. In its standard form it transforms points on a disk to points on a Gaussian, and moreover this mapping contains a singularity in the polar radius (by mapping the boundary of the disk to the origin, and the center of the disk to infinity).

To overcome this we use a two-stage transformation. We first apply Shirley’s low-distortion concentric mapping [SC97] to map points (s, t) in the square micropatch




Figure 6: Our multi-stage process of mapping points in the micropatch domain (s, t) to points on a Gaussian caustic.

domain to points on a unit disk (u, v) . We then apply a modified Box-Muller transform, which non-uniformly scales the radius without singularities. Given the point (u, v) , we compute the new radius as:

$$R = \sqrt{\frac{-2 \log(1 - u^2 - v^2)}{u^2 + v^2}}, \quad (2)$$

and the resulting point on a unit Gaussian is $\mathbf{q} = [uR, vR]^T$. We illustrate this two-stage mapping in Figure 6.

5.1.1. General Gaussians

We now need to transform our unit Gaussian position \mathbf{q} to obey a specific mean μ and covariance matrix Σ . Standard approaches for doing this are typically only concerned with properly transforming the distribution, by e.g. interpreting the coordinates as offsets along the scaled eigenvectors of Σ . This results in a sequence of scale, rotation and translation operations. Unfortunately, using this approach can introduce large distortions in the mapping from patch coordinates to the final projection plane positions, affecting the integrability of the normal field. To avoid this, we eliminate rotation from our transformation, and instead directly scale the point \mathbf{q} along the directions of the eigenvectors of Σ and then translate the point by the mean μ . This is equivalent to transforming \mathbf{q} by the square root of the covariance matrix (as defined by the matrix logarithm):

$$\mathbf{p} = -\sqrt{\Sigma} \mathbf{q} + \mu. \quad (3)$$

We complete our mapping by placing this 2D point onto the projection plane, which involves appending a z -component of $\mathbf{p}_z = -D$ to the above result.

In Equation (3), we actually scale by the negative square root, which introduces a reflection about the origin prior to translation. This results in micropatches which always focus light (between the micropatch and the projection plane) regardless of the size of the Gaussian caustic. This gives our refractive micropatches more curvature, which is better for milling.

5.2. Normal Computation


Given this mapping, we need to compute the surface normal at each location \mathbf{g} which would refract light towards \mathbf{p} .

Since the refracted direction, surface normal, and incident direction must lie in the same plane, we view this as a 2D problem in this coordinate system, which we illustrate on the right.

The angle between the outgoing direction and the incident direction is $\theta_b = \arctan(\beta)$ where $\beta = \frac{d}{D}$. By using the fact that $\theta_o = \theta_i - \theta_b$ and combining with Snell's Law, we can solve for the incident angle as

$$\theta_i = \frac{\arctan(\eta_o \sin(\theta_b))}{\eta_o \cos(\theta_b) - \eta_i}, \quad (4)$$

where η_i and η_o are the incident and outgoing indices of refraction. We obtain the surface normal vector by rotating the incident vector towards the outgoing vector by θ_i radians.



5.3. Normal Field Integration

The remaining step follows the height field optimization of Weyrich et al. [WPMR09]: similarly, we model the micropatch as a rectilinear grid of elementary microfacets with the desired normal orientations, solving a Poisson problem to lift each facet onto a level that leads to a near-continuous surface. In contrast to the previous work, we do not optimize for tileability, as only an individual micropatch is modeled. Note that we also skip the “shuffling” stage, as, by design, our mapping already leads to a continuous and optimal facet layout.

In our experiments, we found that for a wide range of target Gaussians, the residual of the optimization to be diminishingly small, suggesting that the input normal fields are very close to integrable. In fact, the error is far below the precision of our manufacturing tools, and subsequent polishing of our samples leads to smooth patches with near-Gaussian caustics. We show visualizations of individual micropatches in Figure 7.

6. Micropatch Assembly

Given a Gaussian image decomposition, we aim at finding an assembly of micropatches that generate the respective Gaussian caustics, while building a surface that is as continuous




Figure 7: Refractive micropatches producing (left) an isotropic Gaussian caustic, and (right) a highly anisotropic Gaussian caustic.

as possible, in order to facilitate fabrication and to minimize stray light at boundaries between patches.

A naïve approach would be to pre-generate micropatches for all Gaussian kernels in the image, and to treat their layout as an independent problem. Unfortunately, since we account for local lights and projection distances, the relative position of each micropatch with respect to its target Gaussian affects its shape, rendering an independent layout optimization impossible. Conversely, including a recomputation of each micropatch’s shape into a global optimization is costly at best, and would negatively affect scalability to larger problems.

Instead, we propose an optimization scheme that approximates each patch by a planar proxy geometry that refracts a central ray toward the center of the target Gaussian. We found that the orientation of such a proxy is a good approximation of the overall “tilt” of a Gaussian micropatch that projects toward the same center. We use this planar approximation during optimization to gauge *compatibility* of neighboring patches, in the sense of Equation (1) in Weyrich et al. [WPMR09]. This approximation puts us once again into the same domain as their microfacet optimization, this time each microfacet acting as the proxy for one micropatch.

Accordingly, we use a similar optimization scheme, starting with a permutation optimization, followed by a Poisson integration to solve for facet heights. After the last step, we replace each microfacet by a micropatch that aims at the corresponding Gaussian kernel.

6.1. Facet Permutation

Due to the similarity of this pipeline, we only point out differences to this previous work, using the same notation as in Weyrich et al. [WPMR09], referring to the previous paper for more detail.

Permutation Loop. We, too, deploy simulated annealing (SA) as a very simple algorithm to optimize the permutation of facets. Finckh et al. [FDL10] have pointed out the superiority of simulated perturbation stochastic approximation (SPSA) over SA, but we found the simpler SA iteration to produce satisfactory results nevertheless. The central transition operation is still the swap of two facets, but in contrast to the case of BRDF design, where only the net effect of a microfacet assembly under infinite lighting and viewing is of interest, a swap alters the orientation of the two facets: swapping two facets basically corresponds to swapping the assignment of two kernel centers to their facets, triggering the recomputation of each facet’s normal.

Energy Function. The previous approach minimizes the following constraints: facet compatibility C , tileability I , and the avoidance of valleys V . In our setting, tileability is not required, and valley avoidance is of reduced value, as the (refractive) micropatches create slight ridges at all times. In

order to bias the solution toward more low-frequency solutions, however, we add a smoothness term $S = S_x + S_y$ with

$$S_x = \sum_y \sum_x \left\| \frac{dz(x, y)}{dx} \right\|^2, \quad (5)$$

with dz defined as in Weyrich et al. [WPMR09], and similarly for S_y . This leads to surfaces that are easier to polish after the milling process (see Section 7). The final energy function is then:

$$E = \frac{C + S}{m^2}, \quad (6)$$

with m the total number of micropatches. Figure 8 shows the energy convergence plots for the samples used in this paper.




Figure 8: Convergence plots of the energy minimization using simulated annealing.

6.2. Height Optimization

The height field integration to determine facet placements in z direction follows the previous approach, except that we once more ignore tileability. In our setting, however, changing the height level of a microfacet requires re-aiming it at its caustic center. We consequently run three iterations of height optimization and normal recomputation, which typically leads to reliable convergence after the second iteration. After this, we replace each facet by the corresponding micropatch, leading to near-continuous surfaces (see Figure 9).

7. Results

We evaluate our refractive designs both through digital renderings (Figures 10 and 14), as well as by milling physical samples from PMMA (trade name Plexiglas).

Milling Procedure. We use a Roland DGA EGX-600 Engraver with a maximum cutting area of $61 \times 40 \times 4.2$ cm, maximum software spatial resolution 0.01 mm, and a repeat accuracy of 0.05 mm. The engraver uses a C2-171-010K Carbide Tipped Engraving Cutter (0.13 mm Tip and 4.36 mm Shank).

For smoothest surface results, we use one draft and one fine milling pass, each time along the same axis (x or y). We run the draft mode with a 0.16 mm path spacing, which cuts off layers of maximally 2 mm thickness per pass. This helps to prevent damaging the milling bit. This pass leaves




Figure 9: Visualizations of the micropatch assemblies used to produce the results in this paper.

0.2 mm of material for the final pass at 0.02 mm path spacing. For our surfaces (10 × 10 cm), the draft pass takes between 12 and 24 hrs, while the fine mode: requires another 24 to 48 hrs, depending on the depth variation of the surface.

Afterwards, we (manually) polish the surface using Novus Plastic Clean & Shine No. 2, with a final pass with No. 1. Polishing takes about 1 hr.

Physical Samples. Figure 12 shows two sample datasets milled according to this procedure. Each of the figures shows the target image, a decomposition into 1024 Gaussian kernels, a ray-traced simulation of the resulting micropatch assembly, and a photograph of the caustic generated by the real sample. The samples are 9 × 9 cm in size, comprising 32 × 32 micropatches of 2.8 mm edge length. We image the caustics by using a white-light projector as an approximation of a distant light source. In the examples, the physical projection surface was at about 25 cm behind the sample.

As the examples show, our technique translates well to physical constraints, albeit not without artifacts. Differences between the simulation and the captured results are due to a number of error sources, of which the largest is arguably




Figure 10: Application to an image of Venus de Milo.

the polishing procedure. Great care has to be taken not to damage the material during manual polishing; especially the ridges between micropatches are hard to reach and may lead to spurious caustic patterns as particularly visible in Figure 12. But also the milling procedure may lead to imperfections. Systematic error sources, such as stepper quantization, play a role, and minute shifting of the sample over the long milling time and local overheating of the material tend to impair the quality of the results. Given that, we believe that our results suggest that our approach is robust enough for physical applications.

7.1. Comparison to Continuous Optimization

Finckh et al. recently presented an algorithm to optimize a single, continuous surface to exhibit a custom caustic [FDL10]. The authors kindly agreed to run our “Lena” test specification through their algorithm to allow for a direct comparison to our approach. While their system readily converged, the result does not span the full contrast range of the rather complex target image (see Figure 11). The surface furthermore turned out extremely flat (0.07 mm depth range over a 100 × 100 mm area), which made it very susceptible to stepper quantization of our mill. Figure 11c shows a simulation of the quantization introduced by the mill’s z resolution. We hence decided to abort the experiment of creating a physical surface from this data set.

While there may be alternative target specifications (shorter throw distance or larger projection) that would still lead to millable surfaces, we cautiously conclude that our algorithm outperforms Finckh et al.’s global optimization in terms of physical viability of the results. We hope to be able to investigate this further to achieve a fair comparison.

8. Discussion and Future Work

Despite other potential applications, our evaluation focuses on image projections onto a plane, which is arguably the most convenient domain to gauge the quality of the method’s output. It, however, invites comparison with a much easier method for image generation—that of slide projection. In contrast to slide projection, our method does not modulate,




Figure 11: We visualize the surface (a) generated using the method by Finckh et al. [FDL10], as well as a simulation of the resulting caustic (b). Unfortunately this surface poses a challenge for physical reproduction due to its small vertical extent, which results in only a small number of quantized isocontours (c).




Figure 12: Our method applied to the *Lena*, *Einstein*, and *Venus* images. The Gaussian decomposition is shown in (b) The resulting surfaces are visualized in Figure 9, and we show a PBRT simulation of the caustic in (c). A photograph of the resulting caustic from a milled reproduction of the surface is shown in (d).

but only redirect light, avoiding light loss due to absorption. This is similar in spirit to the light reallocating projection engine of Hoskinson et al. [HSHF10] and meets the hard requirement in luminaire design, that all energy should be used for lighting.

One limitation of our current approach is the restriction to equal-sized surface patches, a choice made to simplify the layout of micropatches. This implies roughly equal energy for all Gaussians (ignoring Fresnel effects), which leads to the side effect that brighter regions receive more, and higher-frequency splats than dark regions, which means that bright regions show more detail than darker regions (visible, for instance, in the *Venus* dataset). This effect could be avoided by allowing for micropatches of different sizes, at the cost of having to solve the packing problem of seamlessly and optimally arranging patches of various sizes.

Another limitation is the fact that the generated surface focuses the caustic on a projection plane with a fixed distance.

As shown in Figure 13 the generated caustic is warped and distorted when the actual distance deviates from the optimal.

It is further conceivable to extend the caustic generation from a single plane to a two-plane target specification. This would enter the realm of light field construction and has been demonstrated as an initial prototype by Fuchs et al. [FRSL08]. Similarly, it is worth investigating whether the refractive domain allows for lighting-dependent image formation, as shown by Alexa et al. for diffuse reflectors [AM10].

9. Conclusion

In this paper we proposed a novel system for designing and manufacturing surfaces that produce desired caustic images when illuminated by a light source. The system produces results superior to previous work on goal-based caustic design and proves to be robust even for complex caustics. Our nonnegative image decomposition scales well to very large




Figure 13: Positioning the sample at different distances to the projection plane shows the effective depth of field of the caustic.




Figure 14: Application to the woodcut “The Great Wave off Kanagawa” by Hokusai. We show the input image (a), the Gaussian decompositions (b), and the resulting caustic images simulated using PBRT for (c) a refractive and (d) reflective surface.

numbers of Gaussian kernels, paving the way for highly intricate caustic designs.

At the same time, we show that our anisotropic Gaussian primitives are flexible enough to describe a wide range of natural image content even with only moderate numbers of kernels, making it possible to fabricate functioning prototypes with off-the-shelf hardware, as we showed with our refractive designs.

We hope that our approach of using continuous functions as discrete building blocks for physical reflectance design will allow for future developments in creating more complex lighting and reflectance phenomena.

References

- [AM10] ALEXA M., MATUSIK W.: Reliefs as images. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 29 (2010), 60:1–7. 8
- [ASG08] ANSON O., SERON F. J., GUTIERREZ D.: Nurbs-based inverse reflector design. In *In: Proceedings of CEIG 2008* (2008), pp. 65–74. 2
- [BM58] BOX G. E. P., MULLER M. E.: A note on the generation of random normal deviates. *The Annals of Mathematical Statistics* 29, 2 (1958), 610–611. 4
- [BSD09] BALZER M., SCHLÖMER T., DEUSSEN O.: Capacity-constrained point distributions: A variant of Lloyd’s method. *ACM Trans. Graph. (Proc. SIGGRAPH)* 28, 3 (2009), 86:1–8. 3
- [DLR*77] DEMPSTER A., LAIRD N., RUBIN D., ET AL.: Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* 39, 1 (1977), 1–38. 2
- [DWP*10] DONG Y., WANG J., PELLACINI F., TONG X., GUO B.: Fabricating spatially-varying subsurface scattering. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 29, 4 (2010), 1–10. 1
- [FDL10] FINCKH M., DAMMERTZ H., LENSCH H.: Geometry construction from caustic images. In *Proc. of European Conference on Computer Vision (ECCV)* (Sep 2010). 1, 2, 6, 7
- [FRSL08] FUCHS M., RASKAR R., SEIDEL H.-P., LENSCH H. P. A.: Towards passive 6d reflectance field displays. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 27, 3 (2008), 58:1–8. 8
- [HFM*10] HAŠAN M., FUCHS M., MATUSIK W., PFISTER H., RUSINKIEWICZ S.: Physical reproduction of materials with specified subsurface scattering. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 29, 4 (2010), 61:1–10. 1
- [HSHF10] HOSKINSON R., STOEBER B., HEIDRICH W., FELS S.: Light reallocation for high contrast projection using an analog micromirror array. *To appear in ACM Trans. on Graphics (Proc. SIGGRAPH Asia)* 29, 5 (2010). 8
- [MAG*09] MATUSIK W., AJDIN B., GU J., LAWRENCE J., LENSCH H. P. A., PELLACINI F., RUSINKIEWICZ S.: Printing spatially-varying reflectance. *ACM Trans. on Graphics (Proc. SIGGRAPH Asia)* 28, 5 (2009), 128:1–9. 1
- [MMP09] MAS A., MARTÍN I., PATOW G.: Fast inverse reflector design (fird). *Computer Graphics Forum* 28, 8 (Dec. 2009), 2046–2056. 2
- [Neu94] NEUBAUER A.: The iterative solution of a nonlinear inverse problem from industry: design of reflectors. In *In: Proceedings of the International Conference on Curves and Surfaces in Geometric Design* (1994), A. K. Peters, Ltd., pp. 335–343. 2
- [PP05] PATOW G., PUEYO X.: A survey of inverse surface design from light transport behavior specification. *Computer Graphics Forum* 24, 4 (Dec. 2005), 773–789. 1, 2
- [PPV07] PATOW G., PUEYO X., VINACUA A.: User-guided inverse reflector design. *Computers & Graphics* 31, 3 (2007), 501–515. 1, 2
- [SC97] SHIRLEY P., CHIU K.: A low distortion map between disk and square. *Journal of Graphics Tools* 2, 3 (1997), 45–52. 4
- [TS06] TSAI Y., SHIH Z.: All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation. *ACM Trans. Graph.* 25, 3 (2006), 967–976. 2
- [WDB*07] WEYRICH T., DENG J., BARNES C., RUSINKIEWICZ S., FINKELSTEIN A.: Digital bas-relief from 3D scenes. *ACM Trans. Graph. (Proc. ACM SIGGRAPH)* 26, 3 (2007), 32:1–7. 1
- [WPMR09] WEYRICH T., PEERS P., MATUSIK W., RUSINKIEWICZ S.: Fabricating microgeometry for custom surface reflectance. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 28, 3 (2009), 32:1–6. 1, 2, 3, 5, 6
- [WZHB09] WALTER B., ZHAO S., HOLZSCHUCH N., BALA K.: Single scattering in refractive media with triangle mesh boundaries. *ACM Transactions on Graphics* 28, 3 (aug 2009). 4