Optimizing Environment Maps for Material Depiction

Adrien Bousseau[1,2] Emmanuelle Chapoulie[1] Ravi Ramamoorthi[2] Maneesh Agrawala[2]

[1] REVES / INRIA Sophia Antipolis [2] University of California, Berkeley



(a) Our optimized lighting emphasizes materials (b) Poor lighting de-emphasizes materials

Figure 1: *Our method (a) automatically optimizes the lighting to enhance material-specific visual features. The lighting reveals the thin and thick parts of the subsurface scattering wax candle, it accentuates the Fresnel reflections along the side of the porcelain vase and it adds strong specular highlights to emphasize the shiny chrome metal of the sculpture. Poorly designed lighting (b) diminishes these characteristic visual features of the materials. The candle appears more like solid plastic, the vase appears to be made of diffuse clay and the sculpture no longer looks like it is made of chrome.*

Abstract

This technical report is an extended version of our EGSR 2011 paper. *We present an automated system for optimizing and synthesizing environment maps that enhance the appearance of materials in a scene. We first identify a set of lighting design principles for material depiction. Each principle specifies the distinctive visual features of a material and describes how environment maps can emphasize those features. We express these principles as linear or quadratic image quality metrics, and present a general optimization framework to solve for the environment map that maximizes these metrics. We accelerate metric evaluation using an approach dual to precomputed radiance transfer (PRT). In contrast to standard PRT that integrates light transport over the lighting domain to generate an image, we pre-integrate light transport over the image domain to optimize for lighting. Finally we present two techniques for transforming existing photographic environment maps to better emphasize materials. We demonstrate the effectiveness of our approach by generating environment maps that enhance the depiction of a variety of materials including glass, metal, plastic, marble and velvet.*

1. Introduction

Shading, reflections and refractions are important visual features for understanding the shapes and materials in an image. While well designed lighting configurations can enhance these features and facilitate image perception, poor lighting design can lead to misinterpretation of image content. In Figure 1a for example, the lighting enhances the thickness variations in the subsurface scattering wax candle, it accentuates the Fresnel reflections at grazing angles of the porcelain vase and it adds strong edges in the specular highlights of the chrome sculpture. In Figure 1b however, poor lighting greatly reduces these visual features and makes it more difficult to correctly identify the materials.

In applications such as product design and advertising, lighting is often carefully configured to best depict the material properties of the objects. Such manual *lighting design* is difficult and requires a deep understanding of the way light interacts with shapes and materials. Manual lighting design is also prohibitive for large databases of 3D models. As an example, a search for "car" in *Google Warehouse* leads to more than 38000 models, and most of them are lit by a single canonical light source. While automatic lighting design techniques can place a few point lights to enhance shape depiction [SL01, Gum02, LHV06], the resulting lighting lacks the visual richness of natural lighting environments that have been shown to improve materials depiction [FDA03].

In this paper we present the first automated system for optimizing and synthesizing environment maps that emphasize distinctive visual features of the *materials* in a scene. Our system focuses on material depiction and does not address other goals of lighting design such as overall aesthetic and composition that represent challenging directions for further research. Our work makes three primary contributions:

1. Lighting design principles for material depiction. We have analyzed books on studio photography [Mon03, HFB07] and studies of material perception [PFG00, FDA03, KP03, FJB04, FTA04] to identify distinctive visual features due to lighting, for a variety of materials including glass, metal, plastic, marble and velvet. We derive a set of design principles that an environment map should fulfill in order to emphasize these visual features and effectively depict the materials in the scene (Section 3). We express each design principle in the form of a simple *image quality metric* that quantifies how well a given rendering of the scene conveys the desired visual features (Sections 4.2 and 4.4).

2. Optimization framework. We introduce a general optimization framework to solve for the environment map that maximizes our image quality metrics (see Figures 3, 4 and 5). We express the metrics within this framework as either linear or quadratic combinations of the image radiance values (Section 4). The key idea of our framework is to pre-integrate the light transport over the *image domain* in order to significantly accelerate the optimization. Our approach can be seen as a dual to precomputed radiance transfer methods [SKS02, NRH03] that integrate the light transport over the *lighting domain* to efficiently render an image.

3. Environment lighting synthesis. While our method can directly generate the optimal lighting for a given set of image metrics, this result may not match the appearance and spatial statistics of a natural environment. We suggest two methods to transform any existing photographic environment map (Section 5). In the first approach (Figures 7-9) we search for the optimal orientation of the map and show how to accelerate this search using spherical harmonic rotations [MD03, RH04]. In the second approach, we use constrained texture synthesis [Ash01] to combine the color and details of the photographic environment map with the luminance variation of the optimal lighting (Figures 11).

2. Related Work

Automatic lighting design. Researchers have developed a variety of techniques for optimizing lighting to enhance shape [SL01, Gum02, LHV06, RBD06, VPB*09]. While our system also optimizes lighting, we focus on enhancing the visual features that emphasize materials rather than shape. In addition, our system generates complex environment lighting rather than a small set of point lights. In the context of material depiction, Gooch et al. [GGSC98] describe a non-photorealistic shading technique for depicting metallic objects. In this paper we compile and automate similar guidelines to depict a variety of materials in photorealistic images.

User interfaces for lighting design. Kerr and Pellacini [KP09] present an excellent survey and evaluation of lighting design interfaces. Researchers have investigated two main interaction paradigms; *indirect interfaces* in which users adjust lighting features in the image such as shading, shadows and highlights [PF92, PTG02], and *painting interfaces* in which users paint the distribution of light over the scene and an optimizer solves for the lighting setup that best reproduces this distribution [SDS*93, AD04, PBMF07]. Okabe et al. [OMSI07] and Pellacini [Pel10] extend painting and indirect interfaces to the design of environment maps. However, these interfaces do not incorporate the guidelines expert photographers and lighting designers commonly use to enhance the appearance of materials. Users must rely on their own training and experience to set up the lighting. In contrast, we present an automated system that optimizes the lighting based on design principles for material depiction.

PRT for illumination and material editing. Precomputed radiance transfer (PRT) [SKS02, NRH03] significantly accelerates rendering under complex lighting and provides real-time feedback for illumination [OMSI07, Pel10] and material editing [BAOR06]. We share many mathematical ideas with these methods such as the use of spherical harmonics and rotational convolution [RH04]. However, while traditional PRT integrates light transport over the lighting domain to generate the image, we integrate light transport over the image domain to optimize the environment map.

3. Design Principles for Material Depiction

Books on studio photography [Mon03, HFB07] and studies of material perception [PFG00, FDA03, KP03, FJB04, FTA04] suggest a variety of techniques for enhancing the appearance of materials. We have compiled these techniques into a set of design principles for depicting many common types of materials. Each principle articulates the distinctive visual features of a class of materials and describes how environment lighting can emphasize those features. Note that some materials can benefit from several principles, like glass which is transparent and reflective.

Transparent materials (e.g. glass, ice). The reflections and refractions along contours should contrast sharply with the refractions in the interior parts of the object to depict transparent materials.


Wine Glass

We base this principle on the *bright-field* lighting technique used in studio photography [HFB07] to delineate the contours of transparent objects (see inset). Photographers place the object in front of a bright background and position dark plates (called *gobos*) around the object, outside the field of view, to produce dark reflections and refractions along contours.

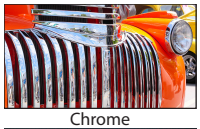
Dark-field lighting produces the complementary effect using a dark background and bright gobos.

Subsurface scattering materials (e.g. wax, marble). The shading should be darker in thick parts and brighter in thin parts to show how the amount of light passing through the object depends on its thickness when it is made of a subsurface scattering material.


Plant Leaves

Fleming et al. [FJB04] show that subsurface scattering is better perceived when objects are illuminated from behind, since backlighting provides more visual cues of light scattered through the thin parts of the object (see inset). Subsurface scattering materials appear to have low contrast and look flatter when illuminated from the front as the variations in shading are diffused by multiple scattering.

Shiny materials (e.g. metal, plastic, chrome). The specular reflections should contain high contrast edges to convey shiny materials. The edges should be sharp for mirror-like materials and blurry for rough materials like brushed metal.


Chrome


Brushed Metal

Pellacini et al. [PFG00] show that under the Ward reflectance model, contrast and sharpness of highlights are the two main dimensions in the perception of gloss. Fleming et al. [FDA03] demonstrate that real-world illumination contributes to accurate perception of shiny materials because it contains edges and bright light sources. Fleming et al. [FTA04] also show that distortions of the reflected environment structures contribute to the perception of shape. Finally, Gooch et al. [GGSC98] show how metal can be conveyed in non photorealistic rendering by alternating sharp dark and light bands oriented in the direction of minimum curvature.

Fresnel materials (e.g. glass, plastic, varnish). The specular reflections at grazing angles should maintain high contrast with respect to the diffuse color of the object to reveal the strong view dependent reflection that is characteristic of Fresnel materials.


Car Paint

The intensity of specular reflections varies with respect to the angle of incidence according to the Fresnel equations. While the strength of reflections is almost constant for conductor materials like metal, it is much stronger at grazing angles for dielectric materials like plastic. In the inset, Fresnel reflections in the car paint delineate the silhouette of the car. Illustrators often exaggerate the Fresnel reflections when drawing cars [Rob03].

Asperity scattering materials (e.g. short fur, velvet). The highlights at grazing angles should maintain high contrast with respect to the diffuse color of the object to reveal the scattering properties of materials covered with a thin layer of asperities like velvet, short fur or dust.


Velvet

Koenderink and Pont [KP03] show that asperity scattering materials produce strong highlights when viewed or lit at grazing angles. In extreme cases like black velvet, strong highlights along occluding contours are the only visible features.

4. Optimization Framework

Our optimization framework automatically generates lighting environments that best fulfill the material specific design principles. Our key insight is that we can pre-integrate the light transport over the image domain and then efficiently compute the optimal lighting for any material design principle we can express as either a linear or quadratic *image quality metric*. While we instantiate our system to tackle the specific problem of material depiction, our optimization framework is general and should be of interest for many other lighting design applications such as shape depiction and environment map painting.

Consider a static scene lit by an environment map L. The radiance at each point $\mathbf{x}$ in the image is given by

$$B(\mathbf{x};L(.)) = \int_{S^2} T(\mathbf{x},\boldsymbol{\omega})L(\boldsymbol{\omega})d\boldsymbol{\omega}, \quad (1)$$

where $\boldsymbol{\omega}$ is the incoming light direction and T is the light transport operator that includes the effects of reflectance, visibility, the cosine term, and interreflections. In matrix form we obtain

$$\mathbf{B} = \mathbf{TL}, \quad (2)$$

where $\mathbf{T}$ is the light transport matrix and $\mathbf{B}$ and $\mathbf{L}$ are vectors of the image and lighting environment. Each column of $\mathbf{T}$ contains an image of the scene as lit from the direction $\boldsymbol{\omega}$, and each row of $\mathbf{T}$ contains the values of pixel $\mathbf{x}$ when lit from all directions. Traditional PRT methods use Equation 2 to compute the image $\mathbf{B}$ under arbitrary lighting $\mathbf{L}$ [NRH03]. Our goal instead is to find the optimal lighting $\mathbf{L}_{opt}$ according to a given material-dependent image quality metric.

4.1. Linear Metric

In the linear case, we define the image quality metric C as

$$C(L(.)) = \int_{\mathbf{x}} f(\mathbf{x})B(\mathbf{x})d\mathbf{x}, \quad (3)$$

where $f(\mathbf{x})$ is a per-pixel linear weighting function designed to enhance material-specific visual features in the image B.

Combining Equations 1 and 3 gives

$$C(L(.)) = \int_{\mathbf{x}} \int_{S^2} f(\mathbf{x}) T(\mathbf{x}, \boldsymbol{\omega}) L(\boldsymbol{\omega}) d\mathbf{x} d\boldsymbol{\omega} \quad (4)$$

$$C = \mathbf{f}^t \mathbf{B} = \mathbf{f}^t \mathbf{T} \mathbf{L}, \quad (5)$$

where $\mathbf{B}$ is a vector of image pixels and $\mathbf{f}$ is a vector of pixel weights. Our goal is to maximize $|C|$.

Pre-integration. The image quality metric C is a double integral over the image and the lighting. By inverting the order of the integrals, we can pre-integrate the metric over the image domain

$$F(\boldsymbol{\omega}) = \int_{\mathbf{x}} f(\mathbf{x}) T(\mathbf{x}, \boldsymbol{\omega}) d\mathbf{x} \quad (6)$$

$$\mathbf{F}^t = \mathbf{f}^t \mathbf{T}. \quad (7)$$

The linear metric then becomes a single integral over the lighting domain which we can evaluate as a dot product

$$C(L(.)) = \int_{S^2} F(\boldsymbol{\omega}) L(\boldsymbol{\omega}) d\boldsymbol{\omega} \quad (8)$$

$$C = \mathbf{F}^t \mathbf{L}. \quad (9)$$

$\mathbf{F}$ is the vector of weights transferred into the lighting domain. By precomputing $\mathbf{F}$ we avoid explicit storage of the transport matrix $\mathbf{T}$ and allow the fast evaluation of C as a dot product between $\mathbf{F}$ and any lighting $\mathbf{L}$. In Section 4.6 we discuss the precomputation of $\mathbf{F}$ and evaluation of $\mathbf{F}^t \mathbf{L}$.

Optimal Lighting. Since Equation 9 defines C as a dot product between $\mathbf{L}$ and $\mathbf{F}$, $|C|$ is maximized when $\mathbf{L}$ and $\mathbf{F}$ are aligned, i.e. when $\mathbf{L}$ is proportional to $\mathbf{F}$. Thus,

$$\mathbf{L}_{opt} = \frac{\mathbf{F}}{\sqrt{\mathbf{F}^t \mathbf{F}}}. \quad (10)$$

The normalization ensures that the magnitude of $\mathbf{L}_{opt}$ is 1 and prevents infinite values in the solution. However, negative values in $\mathbf{F}$ lead to negative values in the lighting, which violate the physics of light. We enforce $\mathbf{L}_{opt} \geq 0$ by clamping $\mathbf{F}$ to positive values

$$\mathbf{L}_{opt} = \frac{\hat{\mathbf{F}}}{\sqrt{\hat{\mathbf{F}}^t \hat{\mathbf{F}}}} \qquad \hat{\mathbf{F}} = \max(\mathbf{F}, 0). \quad (11)$$

The resulting lighting is one solution maximizing $|C|$ under the constraint $L(\boldsymbol{\omega}) \geq 0$. We obtain a second when $\hat{\mathbf{F}} = \max(-\mathbf{F}, 0)$. It should be noted that $\mathbf{L}_{opt}$ is only optimal with respect to our principles for material depiction and may be sub-optimal with respect to other criteria such as composition and aesthetic.

4.2. Material-Dependent Linear Metrics

Transparent materials. For transparent objects we obtain either dark-field or bright-field lighting by maximizing the contrast between the pixels along object contours and the surrounding pixels. To differentiate these two groups of pixels we define indicator functions $w(\mathbf{x})$ and $\overline{w}(\mathbf{x}) = 1 - w(\mathbf{x})$, where $w(\mathbf{x})$ is close to 1 along contours and close to 0 away from contours. We set $w(\mathbf{x}) = (1 - (\mathbf{n}_\mathbf{x}.\mathbf{v}_\mathbf{x}))^\gamma$ where $\mathbf{n}_\mathbf{x}$ and $\mathbf{v}_\mathbf{x}$ are the surface normal and view vector at point $\mathbf{x}$ respectively. The γ parameter controls the decay of the function and we have found that setting $\gamma = 3$ delineates thin contours. To ensure good contrast against the background, we set $w(\mathbf{x}) = 0$ for the background pixels immediately surrounding the object. Then we set

$$f(\mathbf{x}) = \frac{w(\mathbf{x})}{\int_\mathbf{x} w(\mathbf{x})} - \frac{\overline{w}(\mathbf{x})}{\int_\mathbf{x} \overline{w}(\mathbf{x})}, \quad (12)$$

where the first term corresponds to the contour pixels and the second term corresponds to the surrounding pixels. The normalization terms divide by the number of pixels in each group. They ensure that both groups have equal influence on the overall contrast measure when it is integrated over the image according to Equation 3.

Figure 2 shows the weighting vector $\mathbf{f}$ in the image domain and the corresponding vector $\mathbf{F}$ in the lighting domain for a wine glass. Figure 3a shows the two optimal lighting solutions for this model. When $\mathbf{L}_{opt}$ is proportional to $\mathbf{F}$ it enhances the bright contours of the wine glass with the dark-field effect. The second solution, proportional to $-\mathbf{F}$, produces the bright-field effect with dark contours.

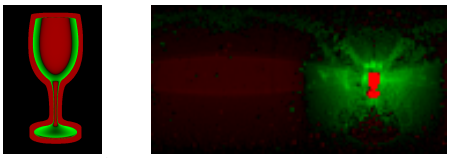


Figure 2: *For glass, the image domain weight vector* $\mathbf{f}$ *is positive (green) along contours and negative (red) in the surrounding regions. In the corresponding lighting domain vector* $\mathbf{F}$ *the green pixels spread out over a wide area indicating that the contours are lit by a relatively wide region of the environment map.*

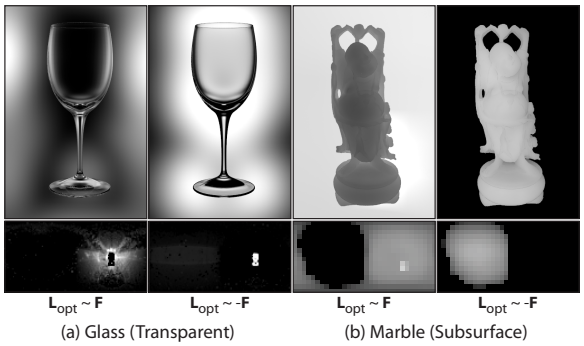


Figure 3: *The optimal lighting delineates the contours of the wine glass (a) and produces either a dark-field or a bright-field effect. For marble (b), when the optimal lighting is proportional to* $\mathbf{F}$*, backlighting reveals subsurface scattering through thin parts of the object. When it is proportional to* $-\mathbf{F}$*, thin parts are darker yielding a less effective depiction.*

Subsurface scattering materials. For subsurface scattering materials we enhance the contrast between thin and thick parts of the object. We first use a raytracer to compute the thickness of the object $t(\mathbf{x})$ along each viewing ray. For simplicity we approximate $t(\mathbf{x})$ as the distance between the first and the last intersection of a ray with the object. To estimate the transmittance through a slab of thickness t we use the indicator function $w(\mathbf{x}) = e^{-\frac{t(\mathbf{x})}{l}}$ where l is the *mean free path* in the material [JMLH01]. In our examples we manually set the mean free path high enough to allow light to travel through the thinner parts of the object. Finally, we use this new indicator function in Equation 12 to compute the contrast between thick and thin parts of the object.

Figure 3b shows that setting $\mathbf{L}_{opt}$ proportional to $\mathbf{F}$ produces the desired effect of brightening thin parts of the object and darkening thicker parts. However, setting $\mathbf{L}_{opt}$ proportional to $-\mathbf{F}$ brightens thick parts and diminishes the characteristic visual features of subsurface materials. Subsurface scattering is a special case of the linear metric where we must optimize for C rather than $|C|$.

4.3. Quadratic Metric

Our optimization framework also applies to quadratic image quality metrics expressed as

$$C(L(.)) = \int_{\mathbf{x}}\int_{\mathbf{y}} Q(\mathbf{x},\mathbf{y})B(\mathbf{x})B(\mathbf{y})d\mathbf{x}d\mathbf{y} \quad (13)$$

$$C = \mathbf{B}^t\mathbf{Q}\mathbf{B} = \mathbf{L}^t(\mathbf{T}^t\mathbf{Q}\mathbf{T})\mathbf{L} \quad (14)$$

where $Q(\mathbf{x},\mathbf{y})$ is a quadratic weighting function between pixels $\mathbf{x}$ and $\mathbf{y}$.

Pre-integration. As in the linear case, we accelerate the evaluation of the quadratic metric by pre-integrating components of the metric over the image domain

$$\mathbf{Q}' = \mathbf{T}^t\mathbf{Q}\mathbf{T} \quad (15)$$

$$C = \mathbf{L}^t\mathbf{Q}'\mathbf{L}. \quad (16)$$

In Section 4.6 we provide algorithmic details explaining how to precompute $\mathbf{Q}'$ and evaluate $\mathbf{L}^t\mathbf{Q}'\mathbf{L}$.

Optimal lighting. The optimal lighting $\mathbf{L}_{opt}$ is the solution to

$$\underset{\mathbf{L}}{\text{argmax}}(C = \mathbf{L}^t\mathbf{Q}'\mathbf{L}) \qquad \mathbf{L}^t\mathbf{L} = 1 \qquad \forall\boldsymbol{\omega},\ L(\boldsymbol{\omega}) \geq 0 \quad (17)$$

where the unit norm constraint $\mathbf{L}^t\mathbf{L} = 1$ prevents the degenerate solution of an infinite vector. We turn this problem into a minimization by defining $\hat{C} = -C$ and we treat the unit norm constraint as a quadratic penalty to obtain

$$\underset{\mathbf{L}}{\text{argmin}}(\hat{C} = -\mathbf{L}^t\mathbf{Q}'\mathbf{L} + \alpha(\mathbf{L}^t\mathbf{L} - 1)^2) \quad \forall\boldsymbol{\omega},\ L(\boldsymbol{\omega}) \geq 0 \quad (18)$$

where α controls the strength of the unit norm constraint. We set $\alpha = 0.5$ for our results. The gradient of $\hat{C}$ is given by

$$\nabla\hat{C} = -2\mathbf{Q}'\mathbf{L} + 4\alpha\mathbf{L}(\mathbf{L}^t\mathbf{L} - 1). \quad (19)$$

We solve for $\mathbf{L}_{opt}$ using the quasi-Newton L-BFGS-B algorithm [BLNZ95]. While this algorithm does not guarantee a global minimum, we have found that the resulting environment maps effectively depict the desired material properties.

An alternative approach is to compute $\mathbf{L}_{opt}$ using power iteration to maximize the Rayleigh Ritz ratio $\frac{\mathbf{L}^t\mathbf{Q}'\mathbf{L}}{\mathbf{L}^t\mathbf{L}}$, while clamping negative values at each iteration. We have applied this approach to a few test cases and found that it produces visually identical results to the quasi-Newton optimization.

4.4. Material-Dependent Quadratic Metrics

Shiny materials. Shiny materials are best depicted when the reflections contain high contrast edges. To encourage the presence of such reflection edges we maximize the metric

$$C(L(.)) = \int_{\mathbf{x}} \left(h \otimes B(\mathbf{x})\right)^2 d\mathbf{x} \quad (20)$$

where h is a high pass filter and $\otimes$ is the convolution operator. In matrix form we obtain

$$\mathbf{Q}' = (\mathbf{H}\mathbf{T})^t(\mathbf{H}\mathbf{T}), \quad (21)$$

$\mathbf{H}$ is formed using a Laplacian kernel along the diagonal

$$\mathbf{H}_{i,j} = \begin{cases} 4 & \text{if } i = j \\ -1 & \text{if } i \neq j \text{ and pixel } i \text{ is adjacent to pixel } j \\ 0 & \text{otherwise.} \end{cases}$$

In multiplying $\mathbf{H}$ and $\mathbf{T}$ we are applying the high-pass filter to the columns of the light transport matrix. Note that to prevent the object silhouette from contributing to the metric, we initially render an object mask, and only apply the filter to image pixels covered by the object. We illustrate the benefit of this metric in Figure 4a, where the bright reflected highlights contain high frequency variations.

Fresnel materials. To enhance the grazing-angle reflections of Fresnel materials we maximize the contrast between the specular and diffuse components of the shading near object contours using

$$C(L(.)) = \int_{\mathbf{x}} w(\mathbf{x})(B_s(\mathbf{x}) - B_d(\mathbf{x}))^2 d\mathbf{x} \quad (22)$$

$$C = (\mathbf{W}(\mathbf{B}_s - \mathbf{B}_d))^t(\mathbf{B}_s - \mathbf{B}_d)$$

$$\mathbf{Q}' = (\mathbf{T}_s - \mathbf{T}_d)^t\mathbf{W}(\mathbf{T}_s - \mathbf{T}_d) \quad (23)$$

where B_s is the specular component of the image, B_d is the diffuse component and $w(\mathbf{x}) = (1 - (\mathbf{n}_\mathbf{x}.\mathbf{v}_\mathbf{x}))^\gamma$ is an indicator function that forces the contrast to be strongest near the contours. $\mathbf{W}$ is a diagonal matrix with $\mathbf{W}(\mathbf{x},\mathbf{x}) = w(\mathbf{x})$.

Figure 4b shows how the optimal lighting enhances Fresnel reflections on a car. Because the optimal lighting places all of the light at grazing angles, parts of the car that are oriented toward the viewer appear under-exposed. In Section 5 we describe methods for generating environment maps with more natural lighting distributions.

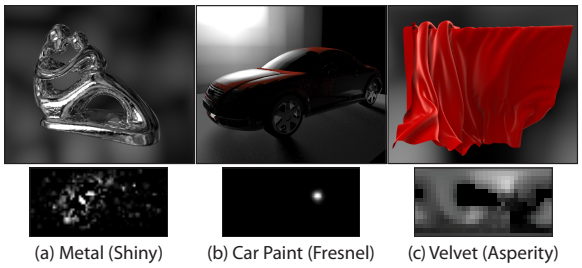


Figure 4: *For shiny metal (a) the optimal lighting adds high frequency edges to the reflection. Because of the complex shape of this object, the reflections come from many different lighting directions. For the car paint (b) the optimal lighting positions most of the light at grazing angles creating strong Fresnel reflection. For the velvet (c) the optimal light creates strong highlights along the grazing angles of the folds.*

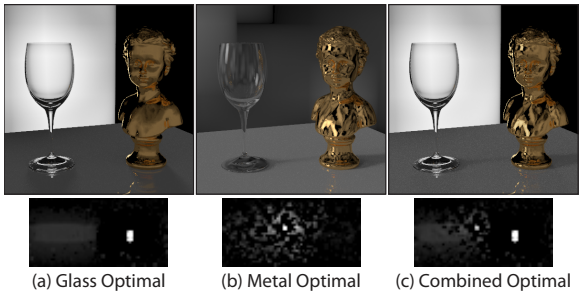


Figure 5: *Combining multiple metrics produces an environment map that enhances both the reflections along countours of the glass and the sharp reflections of the shiny sculpture.*

Asperity scattering materials. Like Fresnel materials, asperity scattering materials are enhanced when the lighting produces high contrast reflections along contours. We use the same image quality metric for asperity scattering materials as we introduced for Fresnel materials (Equation 23).

Figure 4c shows the effect of optimizing this metric to light a piece of velvet fabric. We use the isotropic microfacet model of Ashikhmin et al. [APS00] to render this velvet BRDF. The optimal lighting emphasizes the specular highlights at grazing angles along the folds of the cloth to enhance the depiction of the velvet material.

4.5. Combining Multiple Objectives

Scenes usually contain multiple materials. We can optimize the lighting for all materials by combining the linear and quadratic quality metrics into a single quadratic form. We first transform linear metrics into quadratic metrics by setting $\mathbf{Q}' = \mathbf{F}\mathbf{F}^t$ and then combine the metrics as a weighted sum of the individual terms

$$C = \mathbf{L}^t\mathbf{Q}'\mathbf{L} \qquad \mathbf{Q}' = \sum_i w_i'\mathbf{Q}'_i. \tag{24}$$

We define each individual weight w_i' as a combination of a user-specified weight w_i and a normalization factor that

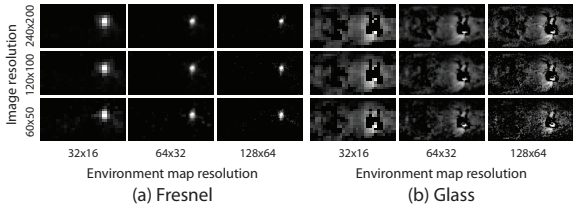


Figure 6: *Optimal lighting for Fresnel (a) and glass (b) materials. As the angular resolution of the lighting increases (left to right) the optimal lighting gains detail, especially for the sharp lighting in (a). As the spatial resolution of the image increases (bottom to top) the optimal lighting remains largely unchanged.*

normalizes each metric by its maximum value, given by the optimal lighting $\mathbf{L}_{opt_i}$

$$w_i' = \frac{w_i}{\mathbf{L}_{opt_i}^t\mathbf{Q}'_i\mathbf{L}_{opt_i}}. \tag{25}$$

Setting $w_i = 1$ gives equal importance to each individual metric in the combination. When C_i is a subsurface scattering metric, we only consider the solutions where $\mathbf{F}_i^t\mathbf{L} \geq 0$.

Figure 5 illustrates such a combination with a scene that contains a wine glass and a shiny sculpture. The combined metric produces an optimal environment map that combines the features of the optimal lighting for each material.

4.6. Implementation

The input to our system is a description of the scene and its materials. We precompute the light transport $\mathbf{T}$ using a custom renderer in the PBRT raytracer [PH10] that generates each column of $\mathbf{T}(.,\boldsymbol{\omega})$ using a black environment map where only the pixel $\boldsymbol{\omega}$ is set to 1; this is essentially the same approach as in standard PRT. For a 100^2 image and a 64×32 sampling of the environment, computing $\mathbf{T}$ takes about an hour. Note that for the Fresnel and asperity scattering metrics we generate two transport matrices, one for the diffuse and one for the specular term of the BRDF. This roughly doubles computation time.

Choosing the appropriate spatial (image) and angular (lighting) resolution is a practical issue in applying our approach. Figure 6 investigates how the computed optimal lighting varies with spatial and angular resolution for two different materials. When the optimal lighting is sharp, as in Figure 6a, increasing angular resolution (from left to right) produces a more accurate result while, as expected, the lighting is somewhat blurred at low angular resolutions. For lower frequency optimal lighting, as in Figure 6b, there is less difference with angular resolution. The relationship with image resolution is more interesting. Unlike standard PRT, our metric is pre-integrated over the image, rather than visualized directly. Therefore, while the integral (and optimal lighting) does become slightly smoother and more accurate as image resolution increases (bottom to top in Fig-

ure 6), low image resolutions are usually adequate to capture the main features of the optimal lighting. In practice, we vary the image dimension between 100^2 and 300^2 and the environment map dimension between 32×16 and 128×64. We emphasize that once the optimal lighting is found, the final image is then rendered at full resolution; the lower spatial resolution is used only for pre-integrating the image quality metric.

Given the light transport matrix, we precompute the linear term $\mathbf{F}$ using Equation 7. The complexity is $O(p^2n^2)$, where p^2 and n^2 are the number of pixels in the image and the environment map respectively. The computation of the quadratic term $\mathbf{Q}'$ using Equation 15 is more expensive, with complexity $O(p^2n^4 + p^4n^2)$. However the matrix $\mathbf{Q}$ is usually sparse, which significantly reduces the cost (Equation 21 and 23). As reported in Table 1, $\mathbf{F}$ can be calculated in a few seconds, while precomputing the quadratic term $\mathbf{Q}'$ can take a few hours, on the same order as determining the light transport $\mathbf{T}$. Since $\mathbf{F}$ and $\mathbf{Q}'$ can be expressed as integrals of products of functions, an interesting future work is to explore the use of importance sampling over the image domain, dual to the acceleration of lighting integrals in Monte Carlo rendering. In addition, most of the computation can be accelerated by a parallel implementation on a cluster or on the GPU. We report in Table 1 the timings with a single CPU and with a cluster of 100 nodes for the most expensive computations.

The optimization computes the optimal lighting $\mathbf{L}_{opt}$ as described earlier in this section. For the linear metric, this is a direct calculation, while an iterative optimization is needed in the quadratic case, which takes a few seconds.

5. Synthesizing Natural Environment Maps

While the optimal lighting $\mathbf{L}_{opt}$ enhances material depiction, it lacks the color and often the spatial statistics of natural environment maps. To give lighting designers more flexibility, we propose to rotate a given photographic environment map $\mathbf{L}$ so that it maximizes the desired image quality metric. Formally, we seek to find the rotation $\mathbf{R}$ that optimizes

$$C(\mathbf{R}) = \mathbf{F}^t\,\mathbf{R}(\mathbf{L}) \quad \text{or} \quad C(\mathbf{R}) = \mathbf{R}(\mathbf{L})^t\,\mathbf{Q}'\,\mathbf{R}(\mathbf{L}) \qquad (26)$$

in the linear and quadratic cases respectively. Since there are only 3 parameters (such as Euler angles) that determine the rotation, we directly search across the 3D space of all orientations. Advanced numerical optimization techniques could be used to accelerate this search. In some cases, physical considerations, such as maintaining a vertical orientation, may restrict the search space to a 1D or 2D set of rotations.

We pre-integrate the metric C over the image domain to obtain $\mathbf{F}$ or $\mathbf{Q}'$, as described in Section 4.6. To evaluate the metric we first rotate the environment map to determine $\mathbf{R}(\mathbf{L})$ and then compute a single dot product in the linear case, or a matrix multiply and a dot product in the quadratic case. We repeat this procedure for a set of regularly sampled rotations $\mathbf{R}$, and we pick the one that maximizes C.

Computation	Complexity	CPU	Cluster
Precomputation			
$\mathbf{T}$ transport	Scene dependent	$\sim$ 1 h.	7 min.
$\mathbf{F}$ linear	$O(p^2n^2)$	12 sec.	
$\mathbf{Q}'$ quadratic	$O(p^2n^4 + p^4n^2)$	4 : 40 h.	32 min.
Metric evaluation			
$C = \mathbf{F}^t\mathbf{L}$ linear	$O(n^2)$	1 milisec.	
$C = \mathbf{L}^t\mathbf{Q}'\mathbf{L}$ quadr.	$O(n^4)$	2.25 sec.	
Optimization			
Linear		0.8 milisec.	
Quadratic		2.65 sec.	
Best rotation			
$C(\mathbf{R})$ linear direct	$O(m^3n^2)$	13 min.	2 min.
$C(\mathbf{R})$ linear SH	$O(n^4)$	6.25 sec.	
$C(\mathbf{R})$ quadr. direct	$O(m^3n^4)$	78 h.	2 : 40 h.
$C(\mathbf{R})$ quadr. SH	$O(n^6)$	1 : 45 h.	14 min.

Table 1: *Complexity and timing of our method, where p^2 and n^2 are the spatial and angular resolution of the light transport matrix and m^3 is the number of sampled rotations of the lighting. The timings are for $p^2 = 100^2$, $n^2 = 50^2$ and $m^3 = 50^3$, for a single CPU and a cluster of 100 nodes.*

Speedup using spherical harmonic rotations. For linear metrics, the direct search over a regular sampling of all possible 3D rotations gives a complexity of $O(m^3n^2)$ where m is the number of samples in each rotation dimension and n^2 is the number of pixels in the environment map. In the quadratic case, the direct search complexity is $O(m^3n^4)$. To speed up this search in the linear case, we observe that the computation of Equation 8 for all rotated lighting $L(\mathbf{R}\boldsymbol{\omega})$ is a rotational convolution over the lighting domain. This enables fast computation of $C(\mathbf{R})$ using spherical harmonic rotations. Mathematical details of the derivation, and a fast multistage algorithm are given in the Appendix. The complexity is reduced by at least a factor of n, which can lead to a speedup of one to two orders of magnitude. While the quadratic case in Equation 16 does not have the same convolution form, we can still use spherical harmonic rotations to speed up the computation of the inner integral.

As shown in Table 1, spherical harmonic rotations yield a significant speedup. In the linear case spherical harmonics reduce running time by over two orders of magnitude and can find the optimal rotation in seconds. Although the computation is more expensive in the quadratic case, spherical harmonics reduce the search time from days to hours.

Compositing multiple rotations. When a scene contains multiple materials, we can optimize the rotation of the environment map for each object individually. However, the optimal rotation for one material is unlikely to be optimal for another material. A better approach is to find the optimal rotation after combining the image quality metric for all the materials (Section 4.5). However, we have found that for some photographic environment maps, a single rotation cannot satisfy the requirements imposed by all materials. Instead, we propose creating a *composite* result in which each



Figure 7: *Worst, median and best orientations of photographic environment maps for transparent glass (a) and subsurface scattering marble (b). In the final renderings we blurred the environment maps to mimic shallow depth of field. The best orientations enhance the contours of the glass and the thickness of the marble, while the median and worst orientations diminish these material-specific features. The best orientation of the photographic environment map aligns strong light sources with the brightest pixels of the optimal lighting. Compare these renderings to those in Figure 3.*

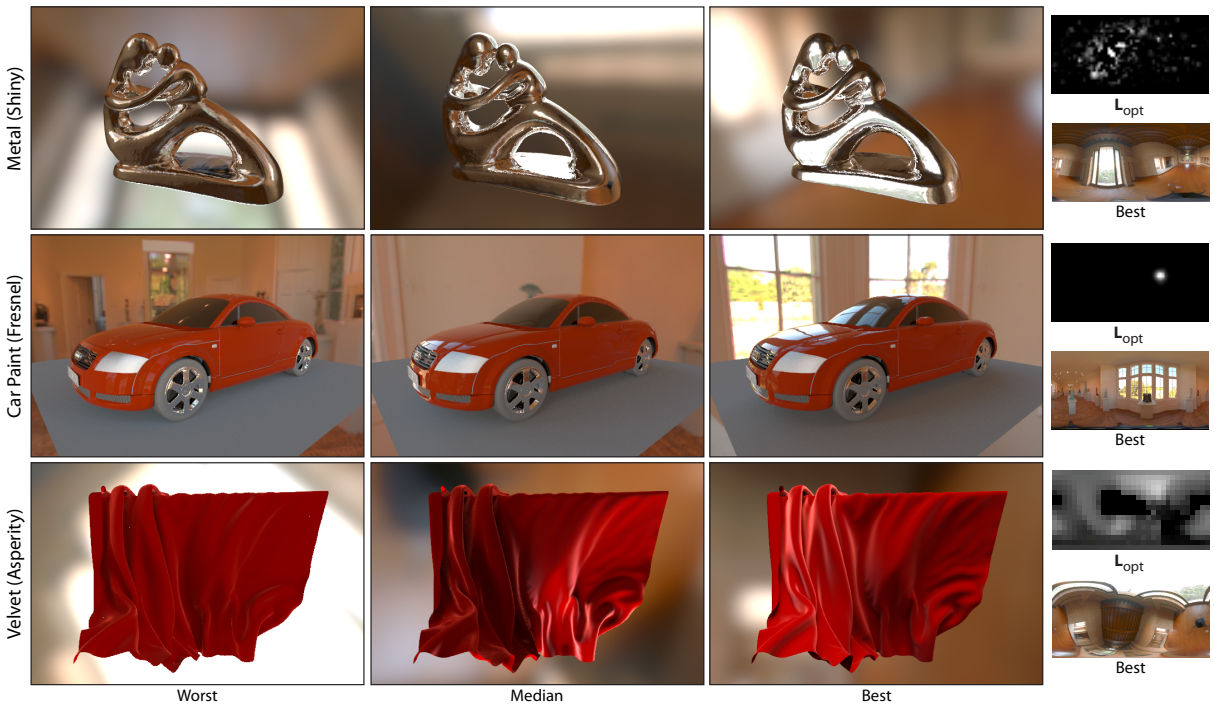


Figure 8: *Worst, median and best orientations of photographic environment maps for shiny metal (top row), car paint (middle row) and velvet (bottom row). The best orientations produce high frequency edges in specular highlights of the metal sculpture, and strong grazing angle reflections for both the Fresnel car paint and the asperity scattering velvet. The median and worst orientations diminish these material-specific features. The best orientation of the photographic environment map aligns strong light sources with the brightest pixels of the optimal lighting. Note that in the car example we restricted the orientation search space to a 1D set of rotations in order to preserve the vertical orientation of the environment. We blurred the environment map for the metal and velvet example to mimic shallow depth of field. Compare these renderings to those in Figure 4.*

material is lit by its own rotated environment map. Although the lighting in the resulting image is physically incorrect, human observers are remarkably tolerant to inconsistencies in reflections [KRFB06]. For transparent materials however, inconsistencies in the refractions may be noticeable. Thus, if the scene contains a glass object we set the rotation of the background environment map to the optimal rotation for the glass. Otherwise when the scene only contains opaque materials, we can choose any orientation for the background. Blurring the background environment or using a shallow depth of field can improve the composite further.

Lighting texture transfer. As an alternative to simply rotating a photographic environment map we can use constrained texture synthesis [Ash01] to generate a new environment map that includes the color and spatial detail of the photographic environment and the intensity distribution of the optimal lighting. Such texture synthesis is especially

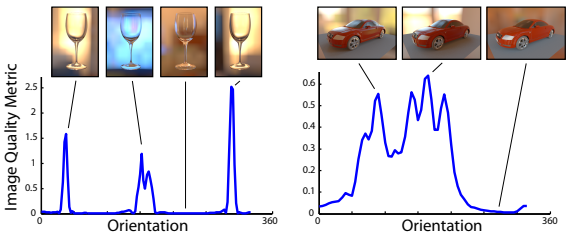


Figure 9: *The image quality metric for materials varies as we rotate a photographic environment. For glass (left) the peaks produce the bright-field effect. For Fresnel car paint (right) there is a large region of suitable orientations because windows provide light at grazing angles. For illustration purposes we only show rotations along a single axis.*

useful when the distribution of light in the photographic environment is such that no single rotation will emphasize the materials in the scene.

We apply the fast texture-by-numbers method of Busto et al. [PELS10] treating the input photographic environment map as the *texture exemplar*, its luminance channel as the *source label map*, and the optimal lighting as the *target label map*. Since this randomized algorithm is prone to local minima, we rotate the photographic environment map to its best orientation to provide a good initial solution.

Studio Lighting. While overall exposure falls outside the scope of our optimization criteria, our approach is complementary to common studio lighting design practice where a few lights (e.g. *key light*, *fill light*, etc.) are placed in canonical positions to ensure that objects are well exposed. In Figure 8 we add a fill light to the optimized environment map for the car to prevent underexposure.

6. Results

Optimal environment maps. The framework of Section 4 can automatically generate environment maps that enhance material depiction, as already shown in Figures 3 and 4. Metrics for multiple materials can be combined to create an overall optimal lighting (Figure 5).

Optimizing lighting orientation. Figures 7 and 8 demonstrate how our system can optimize the orientation of a photographic environment map to maximize a material-specific image quality metric. In all of these examples the *best* orientation aligns the strong light sources in the photographic environment with the brightest regions of the optimal lighting. When the objects are lit with the best orientation, the material-specific visual features are emphasized just as they are in our direct optimization results (Figures 3 and 4). For comparison we include images resulting from the median and worst orientations of the environment map. In these cases, the characteristic features of each material are greatly diminished. Note that for some of our results we blurred the



Figure 10: *The optimal environment map rotations for the wine glass (a) and the metallic sculpture (b) differ considerably. With composite image (c) we separately render each object with its own optimal lighting and then composite them together. The result enhances the depiction of both materials.*

environment map to mimic shallow depth of field and focus attention on the objects.

Figure 9 shows how our image quality metrics for transparent and Fresnel materials evolve as we rotate the photographic environment map. For the wine glass there are three strong peaks in this plot that maximize the metric and produce the bright-field effect with dark contours. For this environment most orientations do not produce the desired contrast along contours and generate a very low image quality score. For the car, a window in the environment provides light at a grazing angle for a large set of orientations, but when the window rotates behind the viewer both the image quality score and the Fresnel reflections are reduced.

Figure 1a is an example of a composite result combining the optimal rotations for the wax candle (subsurface scattering), porcelain vase (Fresnel) and chrome sculpture (shiny metal). Figure 1b shows the objects lit consistently using the same environment but with a poorly chosen rotation; this lighting de-emphasizes the material-specific characteristics and the objects appear to be made of different materials than in Figure 1a. Figure 10 shows another example of compositing multiple environment rotations. In this case the metric for the wine glass encourages solid backlighting to obtain the bright-field effect while the metric for the metal sculpture encourages the presence of light sources that produce strong edges in the specular reflections. In the composite result each material is lit by its own optimal environment map. To preserve consistency in the refractions, we use the optimal lighting for the wine glass as the background environment. While the lighting on the metal sculpture is inconsistent with the lighting on the glass, the composite result does not appear to be inconsistent.

Lighting texture transfer. Figure 11 illustrates the use of constrained texture synthesis to produce the bright-field effect on a complex glass dinosaur. The best orientation for the exemplar photographic environment is cluttered with high-frequency edges inside the body of the dinosaur. In contrast the synthesized environment map properly delineates the contour and creates a high-contrast bright region in the surrounding pixels to enhance the depiction of glass. How-

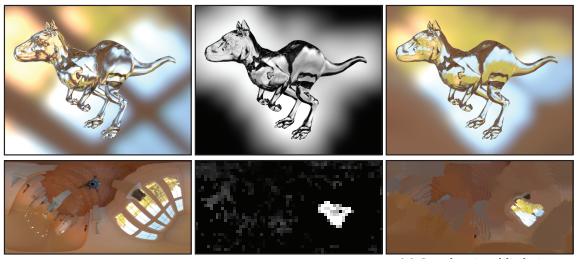


Figure 11: *We use constrained texture synthesis to produce an environment map that has the local appearance of a photographic environment map (a) and the global structure of the optimal lighting (b). The synthesized result produces a bright-field effect that strongly delineates contours.*

ever, the synthesized environment map does introduce some unnatural blockiness in the background image. Compositing the object over another background or blurring the background environment can mitigate such issues.

Sorting environment maps. Because our system is automatic it can be used to automatically process large collections of 3D models and environment maps. Figure 12 illustrates such application where we find the best orientation for a collection of environment maps for three objects made of glass, metal and plastic. In addition, we can use our approach to rank order the most effective environment maps to depict the materials of the object.

Real scenes. An interesting future direction is to apply our approach to design lighting for real-world scenes. Debevec et al. [DHT*00] captured the light transport of a knight dressed in armor using their LightStage. Figure 13 shows our initial attempt at using our system to compute the optimal lighting for depicting the rough metal material of the armor. One challenge for such real-world scenes is to estimate the geometric information required by some of our image quality metrics (e.g. contour locations, thickness, etc.). In this case we manually created a mask for the chest plate and ran our shiny image metric within this masked region.

7. Limitations

Our precomputation framework relies on the linearity of light transport and does not account for the effect of non-linear tone mapping. Only a few papers have studied the influence of non-linear tone mapping on material perception [PFS09], and it is unclear how to model these effects in our context.

While we have demonstrated that design principles for many different classes of materials can be expressed using simple image quality metrics, these principles may not capture all of the subtle features that distinguish one material from another. We believe that further perceptual research is



Figure 13: *Our system emphasizes the shiny metal of the knight's armor. This real-world light transport was captured by Debevec et al. [DHT*00] using their LightStage. We manually created a mask for the knight's chest plate and ran our shiny image quality metric within this masked region. The optimal lighting orientation favors the presence of highlights on the chest plate while these highlights are diminished in the worst and median orientations.*

necessary to determine how people recognize different materials. Nevertheless our optimization framework applies to any linear or quadratic metric and could directly benefit from advances in perceptual understanding of materials.

8. Conclusion and Future Work

Automatic lighting design is a very challenging problem and in its most general form it requires balancing a variety of different criteria, many of which are application specific. In this paper we demonstrate initial steps toward automatic design of environment maps, focusing on the specific goal of enhancing material appearance. Our system is based on material-specific lighting design principles and we show how to quantify these principles as image quality metrics. Extending our system to account for principles on other goals such as composition or portrait photography represents an exciting direction for future work.

Our main technical contribution is a general optimization framework for computing the environment map that maximizes any linear or quadratic image metric. We have also demonstrated how this framework can be used to transform existing photographic environments maps to enhance material depiction while retaining the spatial statistics of natural environments. We plan to further explore the duality between traditional PRT which integrates over the lighting domain and our strategy of pre-integrating over the image domain. Our mathematical framework and pre-integration techniques may also benefit other lighting design techniques such as environment map painting [OMSI07, Pel10].

Acknowledgments. We thank the anonymous reviewers and George Drettakis for their constructive comments. Many thanks to Wesley Willet for the submission video and to Bruno Levy and Manmohan Chandraker for help on optimization. The 3D models used in this paper come from the PBRT book [PH10], the AIM@SHAPE project and the Stanford 3D scanning repository. The environment maps come

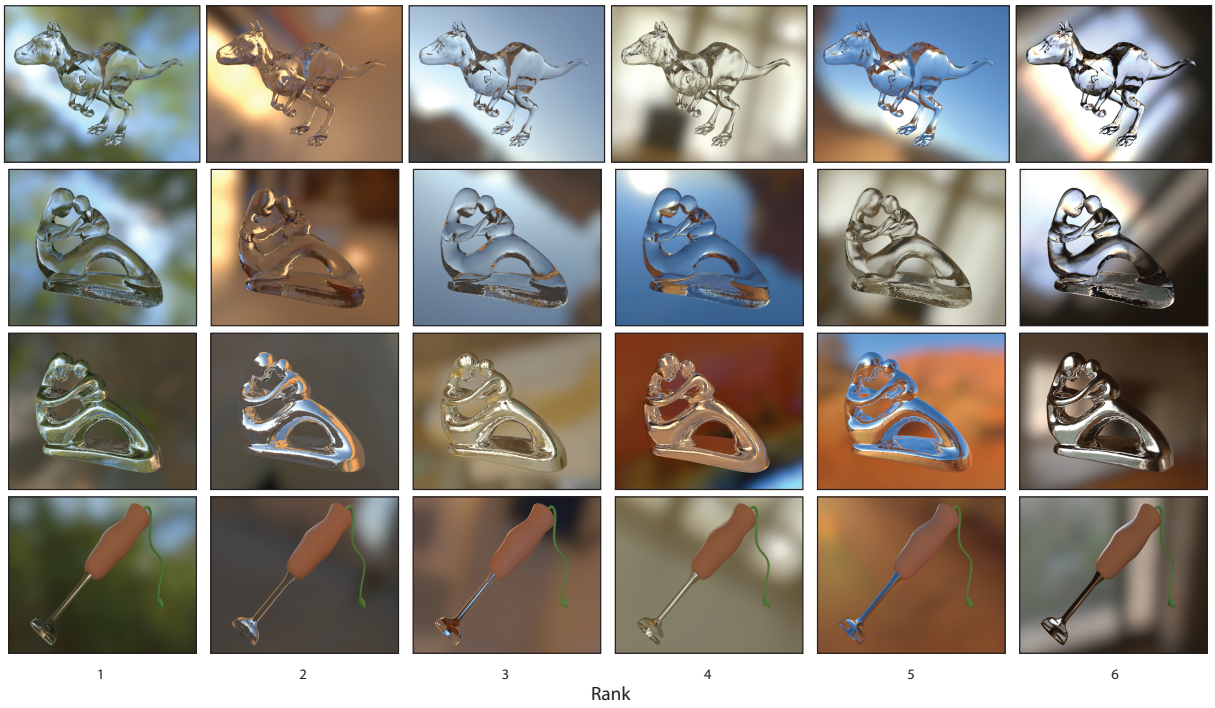


Figure 12: *We used our automatic approach to find the best best orientation of 6 environment maps for different objects and materials — glass (first and second row), shiny metal (third and fourth row) and glossy plastic (fourth row). We additionally rank order the environment maps according to their score (1 for low score, 6 for high score).*

from the ICT Graphics lab, the sIBL Archive and Bernhard Vogl. This work was funded in part by NSF grants CCF-0924968, IIS-1011832, CCF-0643552. UC Berkeley thanks Intel, NVIDIA, Adobe Systems and Pixar for additional support through equipment and funding. INRIA acknowledges the generous support of NVIDIA (Professor Partnership Program), Autodesk (Software donation of Maya and 3DSMax) and Adobe Systems (research donation).

References

[AD04] ANRYS F., DUTRÉ P.: Image-based lighting design. In *Proc. of the 4th IASTED International Conference on Visualization, Imaging, and Image Processing* (2004).

[APS00] ASHIKMIN M., PREMOŽE S., SHIRLEY P.: A microfacet-based BRDF generator. *SIGGRAPH '00* (2000).

[Ash01] ASHIKHMIN M.: Synthesizing natural textures. In *Proc. of Interactive 3D Graphics* (2001).

[BAOR06] BEN-ARTZI A., OVERBECK R., RAMAMOORTHI R.: Real-time brdf editing in complex lighting. *ACM TOG (Proc. of SIGGRAPH) 25* (2006).

[BLNZ95] BYRD R. H., LU P., NOCEDAL J., ZHU C.: A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific and Statistical Computing 16* (1995).

[DHT*00] DEBEVEC P., HAWKINS T., TCHOU C., DUIKER H.-P., SAROKIN W., SAGAR M.: Acquiring the reflectance field of a human face. *SIGGRAPH '00* (2000).

[FDA03] FLEMING R. W., DROR R. O., ADELSON E. H.: Real-world illumination and the perception of surface reflectance properties. *Journal of Vision 3*, 5 (2003).

[FJB04] FLEMING R. W., JENSEN H. W., BÜLTHOFF H. H.: Perceiving translucent materials. In *APGV '04* (2004).

[FTA04] FLEMING R. W., TORRALBA A., ADELSON E. H.: Specular reflections and the perception of shape. *Journal of Vision 4*, 9 (2004).

[GGSC98] GOOCH A., GOOCH B., SHIRLEY P., COHEN E.: A non-photorealistic lighting model for automatic technical illustration. *SIGGRAPH '98* (1998).

[Gum02] GUMHOLD S.: Maximum entropy light source placement. In *Proc. IEEE Visualization* (2002).

[HFB07] HUNTER F., FUQUA P., BIVER S.: *Light: Science and Magic: An Introduction to Photographic Lighting*. Focal Press, 2007.

[JMLH01] JENSEN H. W., MARSCHNER S. R., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. *SIGGRAPH '01* (2001), 511–518.

[KP03] KOENDERINK J., PONT S.: The secret of velvety skin. *Machine Vision and Applications 14* (2003), 260–268.

[KP09] KERR W. B., PELLACINI F.: Toward evaluating lighting design interface paradigms for novice users. *ACM TOG (Proc. of SIGGRAPH) 28* (2009).

[KRFB06] KHAN E. A., REINHARD E., FLEMING R. W., BÜLTHOFF H. H.: Image-based material editing. *ACM TOG (Proc. of SIGGRAPH) 25* (2006), 654–663.

[LHV06] LEE C. H., HAO X., VARSHNEY A.: Geometry-

dependent lighting. *IEEE Trans. on Visualization and Computer Graphics 12* (2006).

[MD03] MAKADIA A., DANIILIDIS K.: Direct 3d-rotation estimation from spherical images via a generalized shift theorem. *CVPR* (2003).

[Mon03] MONTIZAMBERT D.: *Creative Lighting Techniques for Studio Photographers, 2nd edition*. Amherst Media, 2003.

[NRH03] NG R., RAMAMOORTHI R., HANRAHAN P.: All-frequency shadows using non-linear wavelet lighting approximation. *ACM TOG (Proc. of SIGGRAPH) 22* (2003).

[OMSI07] OKABE M., MATSUSHITA Y., SHEN L., IGARASHI T.: Illumination brush: Interactive design of all-frequency lighting. In *Proc. of Pacific Graphics* (2007).

[PBMF07] PELLACINI F., BATTAGLIA F., MORLEY K., FINKELSTEIN A.: Lighting with paint. *ACM TOG 26*, 2 (2007).

[Pel10] PELLACINI F.: envylight: an interface for editing natural illumination. *ACM TOG (Proc. of SIGGRAPH) 29* (2010).

[PELS10] PANAREDA BUSTO P., EISENACHER C., LEFEBVRE S., STAMMINGER M.: Instant Texture Synthesis by Numbers. *Vision, Modeling & Visualization 2010* (2010), 81–85.

[PF92] POULIN P., FOURNIER A.: Lights from highlights and shadows. In *Proc. of Interactive 3D Graphics* (1992).

[PFG00] PELLACINI F., FERWERDA J. A., GREENBERG D. P.: Toward a psychophysically-based light reflection model for image synthesis. *SIGGRAPH '00* (2000).

[PFS09] PHILLIPS J. B., FERWERDA J. A., STEFAN L.: Effects of image dynamic range on apparent surface gloss. In *IS&T 17th Color Imaging Conference* (2009).

[PH10] PHARR M., HUMPHREYS G.: *Physically Based Rendering: From Theory to Implementation, second edition*. Morgan Kaufmann Publishers Inc., 2010.

[PTG02] PELLACINI F., TOLE P., GREENBERG D. P.: A user interface for interactive cinematic shadow design. *ACM TOG (Proc. of SIGGRAPH) 21* (2002).

[RBD06] RUSINKIEWICZ S., BURNS M., DECARLO D.: Exaggerated shading for depicting shape and detail. *ACM TOG (Proc. of SIGGRAPH) 25*, 3 (2006).

[RH04] RAMAMOORTHI R., HANRAHAN P.: A signal-processing framework for reflection. *ACM TOG 23* (2004).

[Rob03] ROBERTSON S.: *How to Draw Cars the Hot Wheels Way*. MBI, 2003.

[SDS*93] SCHOENEMAN C., DORSEY J., SMITS B., ARVO J., GREENBERG D.: Painting with light. *SIGGRAPH '93* (1993).

[SKS02] SLOAN P.-P., KAUTZ J., SNYDER J.: Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM TOG (Proc. of SIGGRAPH) 21* (2002).

[SL01] SHACKED R., LISCHINSKI D.: Automatic lighting design using a perceptual quality metric. *Computer Graphics Forum 20* (2001).

[VPB*09] VERGNE R., PACANOWSKI R., BARLA P., GRANIER X., SCHLICK C.: Light warping for enhanced surface depiction. *ACM TOG (Proc. of SIGGRAPH) 28*, 3 (2009).

Appendix: SH Computation of $C(\mathbf{R})$

The computation of Equation 8 for all rotated lighting $\mathbf{R}(L)$ is a rotational convolution, which we compute efficiently using spherical harmonic rotations [MD03, RH04]. We first expand the lighting in terms of spherical harmonics

$$L(\boldsymbol{\omega}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} L_{lm} Y_{lm}(\boldsymbol{\omega}) \tag{27}$$

where Y_{lm} is a spherical harmonic and L_{lm} are the spherical harmonic coefficients of L. We then denote the rotation as

$$\mathbf{R} = \mathbf{R}_z(\beta)\mathbf{R}_y(\alpha)\mathbf{R}_z(\gamma), \tag{28}$$

where $\mathbf{R}_z$ is a rotation about the z axis, and $\mathbf{R}_y$ is a rotation about the y axis, while (α, β, γ) are the Euler angles. The expression for a rotated lighting becomes

$$L(\mathbf{R}\boldsymbol{\omega}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{n=-l}^{l} L_{lm} D^l_{mn}(\alpha) e^{Im\beta} e^{In\gamma} Y_{ln}(\boldsymbol{\omega}) \tag{29}$$

where $I = \sqrt{-1}$ and D^l are the rotation matrices for rotation about the y axis (the rotations about z are simply complex exponentials). The matrices D are related to the matrix representations of the rotation group $SO(3)$. Finally we also expand F from Equation 6 in terms of spherical harmonics

$$F(\boldsymbol{\omega}) = \sum_{p=0}^{\infty} \sum_{q=-p}^{p} F_{pq} Y_{pq}(\boldsymbol{\omega}). \tag{30}$$

Combining Equations 29 and 30 and using orthonormality of the spherical hamonics we obtain

$$C(\mathbf{R}) = \sum_{l,m,n,p,q} F_{pq} L_{lm} D^l_{mn}(\alpha) e^{Im\beta} e^{In\gamma} \int_{S^2} Y_{pq}(\boldsymbol{\omega}) Y_{ln}(\boldsymbol{\omega})\, d\boldsymbol{\omega}$$

$$C(\alpha, \beta, \gamma) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{n=-l}^{l} (-1)^n F_{l,-n} L_{lm} D^l_{mn}(\alpha) e^{Im\beta} e^{In\gamma}, \tag{31}$$

where in the last line, we use $Y^*_{ln} = (-1)^n Y_{l,-n}$ (and therefore we are able to set $p = l, q = -n$). We can compute this efficiently in a multi-stage factored process for numerical implementation,

$$C^l_{mn} = (-1)^n F_{l,-n} L_{lm} \tag{32}$$

$$C_{mn}(\alpha) = \sum_{l=0}^{l_{\max}} C^l_{mn} D^l_{mn}(\alpha) \tag{33}$$

$$C_n(\alpha, \beta) = \sum_{m=-l_{\max}}^{l_{\max}} C_{mn}(\alpha) e^{Im\beta} \tag{34}$$

$$C(\alpha, \beta, \gamma) = \sum_{n=-l_{\max}}^{l_{\max}} C_n(\alpha, \beta) e^{In\gamma}. \tag{35}$$

The first part of computing coefficients in Equation 32 is only $O(l^3_{\max})$, while Equations 34 and 35 can be computed using fast fourier transforms. The bottleneck is equation 33 that must be computed for each m, n, and α, which has a cost $O(l^4_{\max})$ assuming an angular resolution of the order of $l_{\max}$. This is a savings of $l_{\max}$ over the direct angular computation using equation 8 that would be $O(l^5_{\max})$.