

Freeform irradiance tailoring for light fields

Brand, M.E.; Birch, D.A.

TR2019-019 May 04, 2019

Abstract

We show how to correct an optical surface to transform an arbitrary incident light field into a desired irradiance pattern on a projection surface. Beam dilation errors and optical surface corrections are derived from the pullbacks of the actual and desired irradiances. Étendue effects — the principal obstacle to extended-source tailoring - are factored out by solving a sparse linear system. The method accommodates nontrivial projection surfaces, transport phenomena, and incident wavefronts, including those from multiple extended light sources. Numerical experiments achieve high fidelity and contrast ratios in as little as $O(N \log N)$ -time for a surface represented by N height values.

Optics Express

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Freeform irradiance tailoring for light fields

MATT BRAND* AND DANIEL A. BIRCH

Mitsubishi Electric Research Laboratories, 201 Broadway, Cambridge, MA 02139, USA

**brand@merl.com*

Abstract: We show how to correct an optical surface to transform an arbitrary incident light field into a desired irradiance pattern on a projection surface. Beam dilation errors and optical surface corrections are derived from the pullbacks of the actual and desired irradiances. Étendue effects — the principal obstacle to extended-source tailoring — are factored out by solving a sparse linear system. The method accommodates nontrivial projection surfaces, transport phenomena, and incident wavefronts, including those from multiple extended light sources. Numerical experiments achieve high fidelity and contrast ratios in as little as $O(N \log N)$ -time for a surface represented by N height values.

1. Introduction

It is well known that a freeform optical surface can be tailored to convert a zero-étendue beam into almost any desired irradiance field on a projection plane. The problem has attracted diverse modeling and numerical solution strategies, exemplified by, but not limited to, use of PDEs [1], piecewise conic [2] and oval surfaces [3], Monge-Ampère equations [4], shape-from-transport [5], and paraxial caustics [6]. However, modeling assumptions — usually some combination of collimated light, shallow surface slopes, and long effective focal lengths — do not comport with the realities of extended light sources and miniaturized optics. Treatments of the more realistic non-zero étendue problem generally focus on the modest goal of obtaining a uniform irradiance field from an idealized LED, e.g., see [7, 8] and all examples in the recent state-of-the-art review by Wu et al. [9]. As noted by Sorgato et al. [10], “The 3D prescribed intensity problem for extended sources with no symmetry has no universal solution yet.”

We present a general solution for the irradiance tailoring problem for light fields, a considerably broader class of sources. This enables the design of freeform optics that produce highly structured illumination patterns from high-étendue light sources. As demonstrated in §3.3, the method can produce steeply sloped high-power optics that work at very short effective focal lengths, opening the door to high-efficiency high-brightness systems with very compact dimensions.

Starting with the zero-étendue case, we derive corrections to the optical surface from beam dilation errors, as revealed by comparing the target irradiance with the actual irradiance. In practice, this reduces to a ray-casting and a discrete Fourier transform.

We then generalize to an arbitrary incident light field (plenoptic function). This includes multiple extended light sources of arbitrary shape and directionality. The historical difficulty is that the irradiance is a highly nonstationary convolution of the light field and the optical surface ray-mapping function; this convolution has resisted mathematical characterization that would enable one to invert it and thereby isolate the effects (and thus defects) of the optical surface. Attempts to approximate the deconvolution as a stationary deblurring problem have only shown success in systems with negligible étendue (e.g., [11]). We solve the deconvolution problem by developing a decomposition of the flux through the optic that connects irradiance errors to surface corrections through a system of linear equations. The system is often sparse and rapidly solvable. Where étendue or optical path physics do not allow an exact solution, the method finds a surface that approximates the target irradiance. The resulting method is simple, fast, and versatile — accommodating nontrivial wavefronts, optical paths, and transport phenomena.

2. The model

We use italics for *scalars*, boldface for **vectors**, and $\|\cdot\|$ for the Euclidean norm. Most quantities can be interchangeably treated as points, fields, or functions. All variates are tabulated in Appendix 4, and their relationships depicted.

2.1. Givens

We are given an initial base surface described as a height field $z(x, y)$ and a target irradiance field I mapped onto a projection surface. Typically $z = 0$ and I is a bitmapped intensity image. On the radiant side, we assume a wavefront function $W(x, y, z)$ which provides the direction ∇W and divergence $\nabla^2 W$ of the incident rays at any optical surface point $(x, y, z(x, y))$. Slightly abusing terminology, we will refer to $\nabla^2 f$ as the local curvature of the field f . For notational convenience we also define $w(x, y)$ as the isosurface of W in the neighborhood of $(x, y, z(x, y))$. Finally, we are given a flux density $\bar{s}(x, y)$ along direction ∇W at optical surface point $(x, y, z(x, y))$, or equivalently, an incident flux density $s = \bar{s} \cos \theta_A$, where θ_A is the angle of the incident ray with the z -axis. The incident wavefront plays a role in derivations but is cancelled out in the solutions, so it need not be mathematically characterized for the reduction to practice. We assume a transport simulator that follows each ray along ∇W through optical surface z to the projection surface, where it samples the irradiant flux density. A field of these samples, parameterized by where the rays depart the optical surface, is known as the pullback of the irradiance. We will denote $u(x, y)$ as the pullback of the desired irradiance I and $\hat{u}(x, y)$ as the pullback of the actual irradiance \hat{I} produced by z . For each ray the simulator also provides the z -axis-parallel propagation distance r from the optical surface to the projection surface and the angle of incidence θ_B at the projection surface.

2.2. Transport from surface gradients

We recall a simple geometric fact governing ray propagation at an interface between two homogeneous isotropic media: Given incident ray $\mathbf{r}_i \in \mathbb{R}^3$, exitent ray $\mathbf{r}_e \in \mathbb{R}^3$, and refractive index ratio $n \doteq n_i/n_e$ ($n = -1$ for mirrors), the interface normal \mathbf{n} and tangent \mathbf{t} satisfy

$$\mathbf{n} \propto n \frac{\mathbf{r}_i}{\|\mathbf{r}_i\|} - \frac{\mathbf{r}_e}{\|\mathbf{r}_e\|}, \quad \text{equivalently,} \quad n \langle \mathbf{t}, \mathbf{r}_i \rangle \|\mathbf{r}_e\| = \langle \mathbf{t}, \mathbf{r}_e \rangle \|\mathbf{r}_i\| \quad \text{since} \quad \langle \mathbf{t}, \mathbf{n} \rangle = 0. \quad (1)$$

In our model the surface normal is $\mathbf{n} = (-\nabla z, 1) = (-\partial_x z, -\partial_y z, 1)$, the incident direction is $\mathbf{r}_i = \nabla W \propto \text{sign}(n)(-\nabla w, 1)$, and an exit ray rooted at interface point (x, y, z) meets the projection surface at ray terminus $(x, y, z) + r\mathbf{r}_e$. Expanding the tangent form of Eq. (1) in these terms and rotating coordinates to the plane of refraction reveals that higher order terms only appear in the Euclidean norms. Dropping those higher order terms is equivalent to assuming that $\|\mathbf{r}_i\| = \|\mathbf{r}_e\|$, in which case Eq. (1) is satisfied by exitent direction

$$\mathbf{r}_e \approx n\mathbf{r}_i + (1 - n)\mathbf{n} = ((n - 1)\nabla z - n\nabla w, 1). \quad (2)$$

2.3. Dilations from surface Laplacians

We are interested in the bundle of rays that travel from area element dA on the optical surface to area element dB on the projection surface (Fig. 1). The goal is to tailor the optic such that the actual flux through dA matches the desired flux through dB : $\int s dA = \int u dB$. In any area small enough that s and u are essentially constant, we can write the flux constraint as $s dA = u dB$. Therefore to achieve a desired intensity dilution we must have an equal geometric dilation:

$$s/u = dB/dA. \quad (3)$$

This can be reduced to plane geometry in a local Cartesian coordinate frame, where $dA \doteq dx dy$ is a square area element with average incident intensity s and dB is a quadrilateral area with average




Fig. 1. Dilation of a ray bundle from optical surface element dA to projection surface element dB (derived in §2.2–§2.3; see §2.1 for the variable meanings).

irradiant intensity u : The area of quadrilateral dB is half the cross product of its vector diagonals, $\frac{1}{2}(\mathbf{d}_1 \times \mathbf{d}_2)$. These vectors are the change in the ray terminus as we move the ray root across the two diagonals of dA : specifically $\mathbf{d}_1, \mathbf{d}_2 \approx dx(\mathbf{e}_x + r\partial_x \mathbf{r}_e) \mp dy(\mathbf{e}_y + r\partial_y \mathbf{r}_e)$ where \mathbf{e}_x and \mathbf{e}_y are the unit vectors in the x - and y -directions. It is more convenient to calculate the area dQ of a quadrilateral formed by taking a constant- z cut through the ray bundle where it meets the projection surface, in which case the diagonals simplify to 2D vectors $\mathbf{d}_1, \mathbf{d}_2 \approx dx(\mathbf{e}_1 + r\partial_x \nabla q) \mp dy(\mathbf{e}_2 + r\partial_y \nabla q)$, where $\nabla q \doteq \nabla((n-1)z - nw)$ is the exitent ray's rate of lateral displacement as per Eq. (2). For small dA , the two quadrilateral areas are related by $dQ \cos \theta_Q = dB \cos \theta_B$ where θ_B (resp. θ_Q) is the angle between the local normal of the projection surface (resp. constant- z plane) and the central ray \mathbf{r}_e of the bundle. Putting this all together yields the projective area relation

$$dB \cos \theta_B = (1 + r\nabla^2 q + r^2((\partial_{xx}q)(\partial_{yy}q) - (\partial_{xy}q)^2)) \cos \theta_Q dA. \quad (4)$$

If we assume the surface exhibits little curvature inside dA , the coefficient to r^2 in Eq. (4) is negligible, so we drop the r^2 term and divide both sides by $\cos \theta_B dA$ to obtain the area dilation

$$\frac{dB}{dA} = \frac{1 + r\nabla^2 q}{\cos \theta_B / \cos \theta_Q} = (1 + r\nabla^2((n-1)z - nw))/o, \quad (5)$$

where the second equality uses $q = ((n-1)z - nw)$ from Eq. (2) and the cosine terms are absorbed into the obliquity term $o \doteq \cos \theta_B / \cos \theta_Q$.

3. Tailoring

We now develop Eq. (5) into three tailoring fixpoints: In §3.1 we solve directly for the surface z ; in §3.2 we eliminate the wavefront w and solve for surface corrections c ; in §3.3 we leverage the correction approach into a solution for light-field tailoring. The information flows of these algorithms are summarized in Fig. 2.

3.1. As a Poisson problem

Recalling that the intensity dilution (l.h.s. Eq. (3)) must equal the beam dilation (r.h.s. Eq. (5)), we equate them and algebraically isolate $\nabla^2 z$ to reveal the Poisson problem

$$\nabla^2 z = \frac{1}{n-1} \left(\frac{m}{r} + n\nabla^2 w \right), \quad (6)$$




Fig. 2. Algorithmic summary: In the zero-étendue case (left and §3.1–3.2), rays from the incident wavefront propagate through the optic to determine sampling locations on the projection surface. In §3.1, sampled target irradiance intensities are pulled back along the existent rays to the optical surface and compared with incident intensities from the wavefront to estimate local curvatures for the optic via Eq. (6). In §3.2, both target and actual irradiances are pulled back and compared to obtain a field of curvature corrections via Eq. (8). For positive-étendue light-fields (right and §3.3), the physics are reversed: Pseudo-rays from test points on the target determine sampling locations and directions on a radiant surface. The pushforward of energy flows along these rays forms a system of linear equations that connects irradiance errors to curvature corrections (Eq. (12)).

where $m \doteq \bar{s}o/u - 1$ is the mismatch between the light incident on the optic from the source and the light pulled back to the optic from the target. In preparation for future results, we have switched from incident flux density s at dA to directional flux density $\bar{s} = s/\cos\theta_A$, consequently the obliquity term $o \doteq \cos\theta_A \cos\theta_B \sec\theta_Q$ gains an additional projective cosine.

Treating the Poisson equation (Eq. (6)) as a fixpoint yields a simple tailoring algorithm: (1) Use a provisional surface estimate z to calculate (via simulated transport) the ray-dependent r.h.s. terms in Eq. (6) (target irradiance pullback u , propagation depths r , obliquities o); (2) solve the Poisson problem in Eq. (6) via FFT to update z ; (3) repeat until convergence. The fixpoint works well for substantially linear problems where ray bends are modest, and converges quickly if the rays through the initial surface provide a good sampling of the target irradiance pattern. In the essentially linear case of uniform collimated light ($\bar{s} = o = 1, \nabla^2 w = 0$), a long projection distance ($r \gg 0$), and a nearly flat optical surface ($z \approx 0$), Eq. (6) contains as special cases freeform methods that view the irradiance as the Laplacian of z , e.g., [6, 12].

3.2. Via curvature corrections

A more effective fixpoint is obtained by comparing the desired irradiance u with the irradiance \hat{u} actually produced by a provisional surface z . We will seek a field c of corrections to surface z that satisfies the linear dilation-dilution model,

$$u = s \frac{dA}{dB_c} = \frac{\bar{s}o}{1 + r\nabla^2((n-1)(c+z) - nw)}, \quad (7)$$

in which dB_c is simply dB in Eq. (5) with c added to z . We write Eq. (7) for u and \hat{u} (with and without c), difference them, and solve for the Laplacian of c to obtain a curvature correction

$$\nabla^2 c = \frac{\bar{s}o}{(n-1)r} \left(\frac{1}{u} - \frac{1}{\hat{u}} \right) = \frac{1}{u} \frac{(\hat{u} - u)}{(n-1)r} \frac{\overline{dB}}{dA}, \quad (8)$$

where \overline{dB}/dA is the dilation dB/dA in Eq. (5) without the obliquity terms. Note that once we have u and \hat{u} via simulation, we only need the incident intensities (first form) or dilations (second form) to compute the correction; the wavefront w has been algebraically eliminated.

Algorithmically, we proceed as in §3.1 above, but add calculation of the provisional irradiance \hat{u} to the simulator's duties, and update the provisional surface $z \leftarrow z + c$. Compute time for




Fig. 3. Two mirrors estimated via Eq. (8) to produce the same irradiance, but from different base surfaces, as per schematic ray transport diagrams. Grayscale shading on the mirror surfaces indicates local surface curvature.

each correction scales log-linearly with the number of surface height values, since the dominant calculation is the DFT in the Poisson solver. We typically start with the simplest surface that provides good coverage of the target; convergence is fastest when most of the rays sample most of the target density. For example, tailoring a mirror represented by 128×64 height values to accurately reproduce a detailed high-contrast image of a human eye (Fig. 3) took less than one second on a laptop computer. See Appendix 4 for numerical considerations.

The correction method has some advantages: Most of the nonlinearities of the optical path are captured on the r.h.s. of Eq. (8) and can be calculated exactly in the simulation of transport, while the linearizations of §2.2–§2.3 are at play only in the relatively small curvature corrections, where they are accurate. Our experience is that Eq. (8) converges faster than Eq. (6) for "easy" problems and yields superior results for harder problems where steep slopes, high curvatures, short effective focal lengths, Fresnel losses, or high contrast ratios are present. An added utility is that the tailored optic can inherit some desired properties from a suitably designed base surface, e.g., demagnifying to increase brightness or crossing all rays to obtain globally convex or concave surfaces (e.g., the r.h.s. of Fig. 3).

3.3. With light fields, including extended light sources

Thus far we have assumed that the wavefront W associates a single ray vector ∇W to every point in space. While more general than a collimated or point light source, this falls far short of nature's full plenoptic function. For the purposes of tailoring, the plenoptic function can be summarized as a light field of spatially- and directionally-varying radiances $\ell(\mu, \nu, \phi, \theta)$ sampled at a (possibly fictive) radiant surface, i.e., ℓ gives the radiance from surface point (μ, ν) in direction (ϕ, θ) . This contains multiple/extended/asymmetric light sources as special cases; ℓ can also incorporate the effects of other surfaces prior to the freeform in the optical path. We note that the second law of thermodynamics almost always excludes the possibility of an exact solution for a tailoring problem with a positive-étendue source, but good approximations have practical utility.

To extend the correction method to incident light fields, we consider both the radiant and irradiant surfaces, but reverse their roles. We decompose the flux through the optic into a superposition of "light-collecting" spherical wavefronts that propagate backward from the

projection surface. Adapting the dilation-dilution model (Eq. (7)) to relate light collection in each wavefront to the freeform's local curvatures, the superposition becomes a sparse system of linear equations whose solution isolates and quantifies the freeform's curvature errors.

To work out the details, it is useful to imagine a camera that views the optical surface through a small aperture of area dT located at a point $\mathbf{p} \in \mathbb{R}^3$ on the projection surface. The camera sees a distorted image of the light source(s) on a subset Ω of the optical surface. Ideally, the total observed flux matches $I_{\mathbf{p}}dT$, where $I_{\mathbf{p}}$ is the desired average irradiant intensity at \mathbf{p} . To wit, we want $I_{\mathbf{p}}dT = \int_{\Omega} s_{\mathbf{p}} dA$ for the flux density field $s_{\mathbf{p}}$ that the camera at \mathbf{p} observes on the optical surface. We can calculate this quantity from the area dilation (Eq. (5)) by reversing the roles of the radiant and irradiant surfaces: We imagine projection surface point \mathbf{p} emitting a spherical wavefront $w_{\mathbf{p}}$ back through the optical surface z , and calculate a field of dilations $\frac{dB_{\mathbf{p}}}{dT} = \frac{dA}{dT} \frac{dB_{\mathbf{p}}}{dA}$ of this wavefront as it propagates backward from the projection surface to the optical surface z and then to the radiant surface, where we take a 2D subset $\ell_{\mathbf{p}}$ of the 4D light field function ℓ containing those radiance values destined for \mathbf{p} . The radiance pushforward field $\ell_{\mathbf{p}}$ plays an analogous role to the irradiance pullback field u in previous sections and is similarly parameterized by the optical surface's coordinate system. Conservation of energy requires that each step have the same total flux:

$$I_{\mathbf{p}}dT = \int_{\Omega} s_{\mathbf{p}} dA = \int_{\Omega} \left(\ell_{\mathbf{p}} \left(\frac{\cos \theta_{Q,\mathbf{p}}}{r_{\mathbf{p}}} \right)^2 \cos \theta_{B,\mathbf{p}} dB_{\mathbf{p}} \right) \cos \theta_{A,\mathbf{p}} dA, \quad (9)$$

where the integrand on the right-hand side is the standard radiometric calculation of flux from a radiant patch $dB_{\mathbf{p}}$ through an aperture dA , for each $dB_{\mathbf{p}}$ and dA pair that deliver light to dT [13].

We now adapt the dilation-dilution model (Eq. (7)) to find the reverse dilation $dB_{\mathbf{p}}$. First, we are now tracing light backward, so we must negate $z + c$ and invert n . Second, the dilation of interest is now between the optical surface to the radiant surface; the relationship between the optical surface "facet" size dA in Eq. (7) and the projection surface "pixel" size dT is arbitrary. For simplicity of exposition we choose $dA = dT$. Applying these changes to Eq. (7), solving it for dB , and substituting the result into the irradiance integral (Eq. (9)) yields a relation between curvature corrections c and the irradiance at \mathbf{p} :

$$I_{\mathbf{p}}dT = \int_{\Omega} \left(\ell_{\mathbf{p}} \frac{o_{\mathbf{p}}}{r_{\mathbf{p}}^2} (1 - (n^{-1} - 1)r_{\mathbf{p}} \nabla^2 (c + z + \frac{n}{n-1} w_{\mathbf{p}})) dT \right) dA. \quad (10)$$

Here all projective cosines have been absorbed into the obliquity term $o_{\mathbf{p}} \doteq \cos^3 \theta_{Q,\mathbf{p}} \cos \theta_{A,\mathbf{p}}$, which echoes the \cos^4 law of the camera equation for Lambertian radiant surfaces (for isotropic emitters, $\ell_{\mathbf{p}}$ will contain a $\cos^{-1} \theta_{B,\mathbf{p}}$ term).

Differencing two instances of Eq. (10)—one for a corrected ($c \neq 0$) surface that provides the desired irradiance $I_{\mathbf{p}}$ and one for an uncorrected ($c = 0$) surface that provides the actual irradiance $\hat{I}_{\mathbf{p}}$ (as per the r.h.s. of Eq. (9))—yields the irradiance error:

$$I_{\mathbf{p}} - \hat{I}_{\mathbf{p}} = \int_{\Omega} \frac{n-1}{n} \ell_{\mathbf{p}} \frac{o_{\mathbf{p}}}{r_{\mathbf{p}}} \nabla^2 c dA. \quad (11)$$

Here we have assumed that the correction is small enough that the change to $\ell_{\mathbf{p}} o_{\mathbf{p}} / r_{\mathbf{p}}$ is negligible.

The r.h.s. of Eq. (11) reminds us that the irradiance error is a convolution of the light field restriction $\ell_{\mathbf{p}}$ and the curvatures $\nabla^2 c$ of the correction field. That convolution is nonstationary because the \mathbf{p} -subscripted terms are different for each test point. However, unlike the irradiance convolution, the curvature-correcting deconvolution is straightforward: We set up a least-squares problem for the Laplacian of the correction $\nabla^2 c$ by discretizing the lens surface into facets with area dA and choosing test points \mathbf{p}_i on the projection surface. For each test point \mathbf{p}_i , Eq. (11)




Fig. 4. Axonometric view of an “E”-projecting lens tailored via Eq. (11). The radiant surface (left) simulates two lens-embedded LEDs that subtend between 9° and 22° when viewed from points on the tailored exit surface (middle). The total distance from the radiant to the irradiant surface (right) is 2 lens-widths = $\sqrt{2}\phi$. The inset shows the irradiance obtained from a surface designed for a point source, comparably illuminated by the LEDs.

generates one linear equality constraint, in inner product form

$$\left\langle \text{vec} \left(\frac{n-1}{n} \ell_{\mathbf{p}} \frac{o_{\mathbf{p}}}{r_{\mathbf{p}}} dA \right), \text{vec}(\nabla^2 c) \right\rangle = I_{\mathbf{p}} - \hat{I}_{\mathbf{p}}, \quad (12)$$

where the first vectorized field is nonzero wherever the camera at \mathbf{p} sees light coming from the optical surface z . Iterating over all test points \mathbf{p}_i , this forms a large but sparse system of linear equations for the elements of $\nabla^2 c$. To tailor the surface z we solve the system of linear equalities for $\nabla^2 c$, solve the Poisson problem for the correction c , add c to surface z , and repeat until convergence.

Figure 4 shows a lens with two embedded LEDs that was tailored in this manner to project a white “E” on a black background. For $N = 128 \times 128$ height values the calculation took 25 seconds, and can be expected to grow at a rate between log-linearly (solution time is dominated by the Poisson problem) and cubically with N (solution time is dominated by the linear system), depending on how much of the optical surface is involved in irradiating each test point.

In general, the positive-étendue tailoring problem is a strongly non-convex optimization problem with no zero-error solution and many local minima. Good results depend on favorable initializations. We often find that a zero-étendue design is a useful starting point for high-étendue problems; such is the case in Fig. 4, where a high-fidelity design for a point light source produces a highly blurred irradiance but serves as a good initial surface for high-étendue tailoring.

4. Summary and open questions

Observing an essentially linear relationship between incident wavefront curvature, optical surface curvature, beam dilation, and irradiant intensity dilution, we formulated zero-étendue freeform irradiance tailoring as a Poisson problem. Algebraically eliminating the wavefront yielded a sequence of surface curvature corrections based on pullbacks of the irradiance; these rely on considerably more modest linearizations than the basic Poisson problem, and perform well in optical path geometries previously considered difficult. Reversing the roles of the radiant and irradiant surfaces in this formulation, we found a solution for general positive-étendue light-field tailoring, e.g., with multiple extended sources. In particular, the problem of discovering what parts of an optical surface are responsible for flaws in the irradiance—previously a poorly understood nonstationary deconvolution—is revealed to be a straightforward system of linear equations in the domain of curvature correction fields. This approach easily accommodates varied light transport phenomena; beyond the pedagogical examples in this paper, we have used it to tailor freeforms situated in optical paths with spatially separated extended light sources, additional optical elements, partial occlusions, internal reflections, and Fresnel losses.

Strictly speaking, all realistic tailoring problems involve positive-étendue sources, and almost all positive-étendue tailoring problems are infeasible—existence of exact solutions would imply entropy-reduction. Yet we find that many problems have good approximate solutions. Since approximation quality can be improved by discarding some radiance to reduce étendue, an open question of great interest is how to characterize and optimize this trade-off for specific irradiance targets and optical path geometries.

Appendix A: Numerical considerations

The Poisson problem $\nabla^2 c = f$, can be solved in $O(N \log N)$ time via FFT. On a rectangular domain, this takes a particularly simple form

$$K(c) = \mathcal{F}^{-1}(\mathcal{F}(K(f)) \cdot \mathcal{F}(\mathcal{L})^{-1}).$$

Here $\mathcal{F}(\cdot)$ and $\mathcal{F}^{-1}(\cdot)$ are the 2-dimensional DFT and inverse DFT with periodic boundary conditions; $K(f)$ reflects a field as

$$K : f \mapsto \begin{bmatrix} f & \mp \lambda \\ \mp \lambda & f \end{bmatrix};$$

\mathcal{L} is the discrete Laplacian operator; $\mathcal{F}(\mathcal{L})^{-1}$ is the pseudo-inverse of the Laplacian in the Fourier domain (an elementwise inverse of $\mathcal{F}(\mathcal{L})$ with a resulting infinity zeroed out); and $\mathcal{F}^{-1}(\cdot)$ is the inverse DFT. All field arithmetic is elementwise. The signs in $K(\cdot)$ are optionally positive to impose Neuman boundary conditions on the corrections (zero derivative in the direction normal to the boundary); this fixates the boundary of the projection in the case of a collimated source. Tailoring can also be done in polar or spherical coordinates using mixed Fourier/Chebyshev methods for the Poisson problem.

Pulled-back zeroes and near-zeroes in the target irradiance field can imply infinite local dilation and thus infinite curvature. To prevent numerical instabilities, we recommend the substitution $1/u \rightarrow 1/(u + \epsilon_t)$ with small $\epsilon_t > 0$ decreasing in each iteration of the fixpoint.

The linear system in light field tailoring Eq. (11) can have poor numerical condition or be underconstrained, e.g., if the irradiance is undersampled or if some optical surface area element is not on any path from a light source to a projection test point. In such cases, instead of $\mathbf{A}\mathbf{x} = \mathbf{b}$, the Tikhanov-regularized problem $(\mathbf{A}^\top \mathbf{A} + \epsilon^2 \mathbf{I})\mathbf{x} = \mathbf{A}^\top \mathbf{b}$ selects for flatter, smoother corrections.

Appendix B: Variates

Table 1. Variables and their meanings.

Symbol	Meaning	Symbol	Meaning
c	Optical surface correction	$r, r_{\mathbf{p}}$	z -component of transport
dA	Lens element	s	Incident flux density at z
dB	Projection plane element, ($G = 0$)	\bar{s}	Flux density at z along ∇W
$dB_{\mathbf{p}}$	Radiant area element, ($G > 0$)	u	Pullback of target irradiance
dT	Projection plane element, ($G > 0$)	\hat{u}	Pullback of actual irradiance
G	Étendue	W	Wavefront function
$I, I_{\mathbf{p}}$	Target irradiance	w	Local wavefront isosurface
$\hat{I}_{\mathbf{p}}$	Actual irradiance at \mathbf{p} , ($G > 0$)	z	Optical surface height field
$\ell, \ell_{\mathbf{p}}$	Radiance	$\theta_A, \theta_{A,\mathbf{p}}$	Angle to optical axis at z
n	Refractive index ratio	θ_B	Angle of incidence at dB
$o, o_{\mathbf{p}}$	Obliquity	$\theta_{B,\mathbf{p}}$	Angle of exitance at $dB_{\mathbf{p}}$
\mathbf{p}	Test point location	θ_Q	Ray angle to the z -axis

References

1. H. Ries and J. Muschaweck, "Tailored freeform optical surfaces," *J. Opt. Soc. Am. A* **19**, 590–595 (2002).
2. V. Oliker, "Mathematical aspects of design of beam shaping surfaces in geometrical optics," *Trends in Nonlinear Analysis*. Markus Kirkilionis and Susanne Kromker and Rolf Rannacher and Friedrich Tomi, eds. (Springer, 2003), chap. 3; 193–224.
3. D. Michaelis, P. Schreiber, and A. Bräuer, "Cartesian oval representation of freeform optics in illumination systems," *Opt. Lett.* **36**, 918–920 (2011).
4. C. Prins, J. ten Thije Boonkkamp, J. van Roosmalen, W. IJzerman, and T. Tukker, "A numerical method for the design of free-form reflectors for lighting applications," *Tech. Rep. CASA-report 1322* (2013).
5. Y. Schwartzburg, R. Testuz, A. Tagliasacchi, and M. Pauly, "High-contrast computational caustic design," *ACM Trans. Graph.* **33**, 74 (2014).
6. G. Damberg, J. Gregson, and W. Heidrich, "High brightness hdr projection using dynamic freeform lensing," *ACM Trans. Graph.* **35**, 24 (2016).
7. R. Wester, G. Müller, A. Völl, M. Berens, J. Stollenwerk, and P. Loosen, "Designing optical free-form surfaces for extended sources," *Opt. Express* **22**, A552–A560 (2014).
8. Z. Zhao, H. Zhang, H. Zheng, and S. Liu, "New reversing freeform lens design method for led uniform illumination with extended source and near field," *Opt. Commun.* **410**, 123–129 (2018).
9. R. Wu, Z. Feng, Z. Zheng, R. Liang, P. Benítez, J. C. Miñano, and F. Duerr, "Design of freeform illumination optics," *Laser Photonics Rev.* **12**, 1700310 (2018).
10. S. Sorgato, R. Mohedano, J. Chaves, M. Hernández, J. Blen, D. Grabovičkić, P. Benítez, J. C. Miñano, H. Thienpont, and F. Duerr, "Compact illumination optic with three freeform surfaces for improved beam control," *Opt. Express* **25**, 29627–29641 (2017).
11. D. Ma, Z. Feng, and R. Liang, "Deconvolution method in designing freeform lens array for structured light illumination," *Appl. Opt.* **54**, 1114–1117 (2015).
12. M. V. Berry, "Laplacian magic windows," *J. Opt.* **19** (2017).
13. W. L. Wolfe, *Introduction to radiometry* (SPIE, 1998).