# Shape from Caustics: Reconstruction of 3D-Printed Glass from Simulated Caustic Images

Marc Kassubeck[1] Florian Bürgel[2] Susana Castillo[1] Sebastian Stiller[2] Marcus Magnor[1]
[1] Institut für Computergraphik, TU Braunschweig graphics.tu-bs.de
[2] Institute for Mathematical Optimization, TU Braunschweig mo.tu-bs.de

## Abstract

*We present an efficient and effective computational framework for the inverse rendering problem of reconstructing the 3D shape of a piece of glass from its caustic image. Our approach is motivated by the needs of 3D glass printing, a nascent additive manufacturing technique that promises to revolutionize the production of optics elements, from lightweight mirrors to waveguides and lenses. One important problem is the reliable control of the manufacturing process by inferring the printed 3D glass shape from its caustic image. Towards this goal, we propose a novel general-purpose reconstruction algorithm based on differentiable light propagation simulation followed by a regularization scheme that takes the deposited glass volume into account. This enables incorporating arbitrary measurements of caustics into an efficient reconstruction framework. We demonstrate the effectiveness of our method and establish the influence of our hyperparameters using several sample shapes and parameter configurations.*

## 1. Introduction

Few examples in manufacturing have garnered as much attention as additive manufacturing techniques. Even mainstream media reports about 3D printed cars and houses [12, 4]. These techniques promise near limitless control over geometry, little to no manual post-processing, and fast adaptability to changing requirements to the extent that even personalized manufacturing seems possible. Recent research is very much concerned with expanding the capabilities of additive manufacturing techniques w.r.t. the utilized materials [20], in particular glass [38], due to its many interesting mechanical and optical properties. This work is inspired by the challenges associated with a recent approach to produce optical grade glass structures with laser-based fused glass deposition (LGD) [38, 32], which melts quartz fiber onto glass substrate and can subsequently build more complex structures by layering of small fibers. The resulting

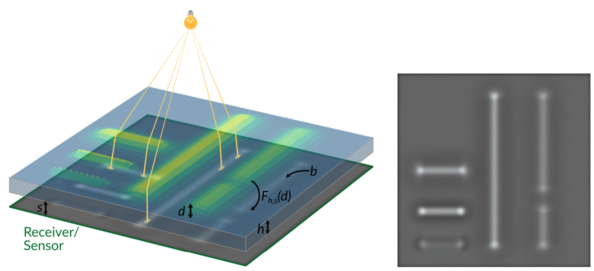


(a) Simulation through light transport (b) Simulation result

Figure 1: Depiction of the experimental setup.

surface is dependent on underlying parameters which influence is highly nonlinear and, thus, still a current subject of research. Furthermore, measurements of achieved surface geometries are not done *in-situ*, but with dedicated measurement equipment like confocal microscopy or tactile profilometry. Therefore, only indirect shape feedback is feasible at present. This circumstance prevents smarter approaches that correct for small errors or discard the workpiece as soon as irrecoverable shape errors have been encountered. Our goal is to incorporate caustic cues into shape estimation tasks arising from glass manufacturing processes. In particular, we provide an efficient and accurate multi-spectral caustic renderer, which can be used to predict real-world measurements inside the manufacturing environment when considering simple transmitter-workpiece-receiver geometries. The renderer is further equipped to provide gradients for efficient local search of the workpiece geometry. To this end, we incorporate knowledge of the production process into the reconstruction. The main difference between this work and related work on caustic design is our focus on real-world production feedback loops. Whereas caustic design usually asks to find a producible, i.e., low-curvature shape, we measure the caustic to reconstruct the real shape. To show the effectiveness of our solution to this problem, we simulate complex application setups with varying amount of optical elements. In addition, we demonstrate successful application to real-world data.

Our main contributions are: first, providing code[1] for

[1] https://bit.ly/3eqSzxR

a powerful, easy to use out of the box rendering framework; second, employment of a sparsity-based reconstruction scheme and incorporation of known process conditions, i.e., the volume of deposited material and the axis-accuracy; and third, demonstration of the approximation capabilities of our reconstruction for providing feedback in glass manufacturing processes. Furthermore, to the best of our knowledge, our method is the first to demonstrate reconstruction in complex scenes with arbitrary amount of refraction events and consistent handling of wavelength-dependent effects.

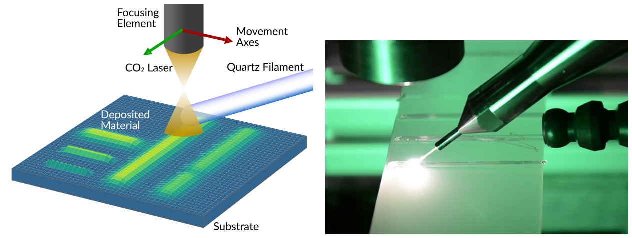


(a) Schematic glass printing process (b) Photo of LGD process [38]

Figure 2: LGD process. Process variables are all computer controlled. Color of the mesh (a) depicts total height.

## 2. Related Work

Some works to solve the ill-posed inverse refractor problem use extra information such as structured light [39, 23], or multi-view cues [26, 34, 23]. In contrast, we use only a single uniform light source and camera to facilitate integration into production feedback loops with existing machine tools. This relatively unconstrained setup relates our work to the field of caustic design [19, 5, 35] but is different as it is tailored to the needs of 3D glass printing.

As we consider the features of the produced shape to be significantly larger than the wavelength of light employed to measure it, our work is based upon simulation methods considering light transport using geometric optics [40, 18]. In Ray Tracing and unidirectional Path Tracing [18], rays are usually traced from the camera through the scene until arriving at a light source, making them unsuited for caustics. To tackle this problem, Photon Mapping approaches [16, 17] have been proposed. Similar to Path Tracing, these methods emit small energy packets from the light sources and store resulting irradiances in a spatial data structure at intersection points with the scene. These methods model caustic effects well, but usually require a significant number of photons to produce noise-free results. Furthermore, unlike Path Tracing or Bidirectional Path Tracing [37], these methods are not progressive, which has been addressed in later approaches [15, 14]. Photon Mapping and Bidirectional Path Tracing methods have been combined into a unified framework called Vertex Connection and Merging [13], which allows to achieve fast global illumination computation including caustics in general complex scenes. For our proposed imaging scenario, however, generalization to very complex scene geometry is not the most pertinent concern. Instead, we need to achieve low variance estimates of the imaging process and their respective gradients w.r.t. the known caustic distribution. At the same time, our goal is to still be able to model real-world scenarios (like partial occlusion of the light or stray caustics caused by inter-reflections) inside the manufacturing machine. Thus, we adopt the method of Frisvad et al. [11] for rendering our caustics, as it allows to simulate sharp and detailed caustics for short light paths with a reduced amount of emitted photons.

Since our problem—the shape reconstruction of printed glass—is a nonlinear, ill-posed parameter identification problem, it is notoriously difficult to solve. To cope with the ill-posedness, we add penalty terms that are suited to a priori information about the reconstruction quantity [9]. The considered inverse rendering problem is related to the inverse scattering problem, which deals with the reconstruction of the refractive index, and thus also the shape, of a penetrable medium from the measurements of the waves scattered from that medium (see, e.g., [6]). We employ terms in the underlying regularization scheme that enforce sparsity and are based on physical bounds; in particular, area-based physical bounds as we take into account the axis accuracy of the 3D printer. In addition, we employ a term for volume conservation as we know the correct volume of the printed glass piece. The positive effect of combining terms that enforce sparsity in the pixel basis and are based on total variation as well as on physical bounds is shown in [2, 3]. The corresponding regularization scheme relies on a primal-dual algorithm. However, as we do not include the total variation, we achieve this by employing the thresholded, nonlinear Landweber scheme as a base [8, 21], together with its extension to area-based physical bounds [1], and a newly developed heuristic to approach the conservation of the volume. The Landweber scheme can be used with respect to the pixel basis or a wavelet basis, see [8]. We will present both variants and, in addition, use the combined effect of the sparsity in a wavelet basis and in the pixel basis.

Methods arising from computational caustic design share similarities to our work. However, the problem formulation is different to the one considered here: for a given caustic, these methods find a sensible and manufacturable shape of the refractor or reflector. This is in contrast to our case, where the shape (and thus the caustic) is the result of a specific manufacturing process, which constrains the possible set of shapes. As one of the earliest works on caustic design, Finckh et al. [10] use physically based rendering to find globally optimal shapes from caustics by means of gradient-free optimization [36]. In this setting the problem is to find a surface that optimally transports irradiance at the receiving surface from a source distribution to a target distribution. This problem has been approximated by decomposing the image into gaussians and the surface into

(a) Oblique view (b) Side view (c) Oblique view (d) Side view

Figure 3: *Fusion 360* renderings of complex scenes. (a)/(b): Two added lenses. (c)/(d): Prism and lens array.

micropatches [28], or casting it as a Poisson reconstruction problem [41]. Schwartzburg et al. [33] directly solve the *optimal transport problem* on power diagrams before computing a conforming surface, which allows for sharper caustic reproduction. Finally, Meyron et al. [24] extend the optimal transport formulation to four problems arising in refractor and reflector design and directly optimize for the target surface. None of these works consider multiple refractions or non-monochromatic light sources. While recent work on general differentiable rendering systems [27] also tackled the problem of caustic design in the presence of colored light sources we handle multiple wavelengths and refractions consistently.

## 3. Methods

In the following, we will discuss the technical details underlying our simulation and reconstruction. We mainly focus on a simple setup (Fig. 1(a)) and optimization strategy that is directly informed by the production process (Fig. 2). Despite this, our method is general enough to conduct reconstructions in the presence of additional optical elements (Figs. 3, 4).

### 3.1. Simulation

As we wish to obtain shape information from resulting caustics, we set up our scene as depicted in Fig. 1(a). A point light source illuminates the scene at position $\ell = (\ell_x, \ell_y, \ell_z)^T$. The glass substrate has a fixed thickness $h \in \mathbb{R}$ and added material is modeled as a height field $d \in \mathbb{R}^{m\times m}$. The receiving plane is placed at a distance $s$ and the frequency-dependent irradiance (with $c$ channels) is denoted as $b \in \mathbb{R}_+^{c\times n\times n}$. We denote the forward simulation as $F_{h,s}\colon \mathbb{R}^{m\times m} \to \mathbb{R}_+^{c\times n\times n}$. We envision that this setup can be directly integrated into the production process, as there is enough space under the substrate to fit an imaging system (see Fig. 2(b)). To achieve this, we adapt a Photon Mapping algorithm [11] in PyTorch [30, 31] to gain the benefit of GPU acceleration and compute gradients of a loss term based on caustics discrepancy (see M1 in Sec. 3.2). The representation of our reconstruction variable, i.e. the height, as a height field limits the expressiveness of achievable structures, especially considering that the height field representation has to be converted to a surface mesh before passing it to the Nvidia OptiX libary [29] (see next paragraph). We still adopt this representation as a concession to the production process (see Fig. 2), which can only layer additional material on top of existing structures.

#### 3.1.1 Forward Simulation Operator

Our scene setup follows the same general layout: we place a point light source at a fixed distance above the refractive specimen and the receiver or sensor plane below the substrate at distance $s$. The light source emits a user-defined number of photons $N_r$ for reference generation and $N_o$ for each reconstruction iteration. These photons are then path-traced through the setup and refracted or reflected at medium transitions. This continues until they reach the receiver screen at which point we record their energy into a texture map. To speed up this process we implemented the scene intersection calls using the Nvidia OptiX library [29]. We further assign each photon a wavelength from a user-defined set of wavelengths. Additionally, we allow the definition of material-specific frequency-dependent refractive index curves to be able to simulate real-world transmissive materials. Note that the formulation of this operator as full recursive light tracing is a distinction from most other works [19, 28, 33] and allows us to simulate caustics in less restrictive setups (see Figs. 3, 4), facilitating the deployment for *in-situ* measurements in production environments.

To generate the caustic images on the receiver plane, we adopt the photon differential splatting method of [11]. This allows us to achieve sharp caustics from our print line structures with fewer photons. Following their work (and notation) we compute the irradiance at the receiver plane as

$$E(x) \approx \sum_{P=1}^{k} \pi K(\|M_P(x - x_P)\|_2) E_P, \tag{1}$$

where $x \in \mathbb{R}^3$ and $x_P \in \mathbb{R}^3$ are a position on the receiver plane and the position of intersection of photon $P$ with the receiver plane respectively. $M_P$ is the change of basis matrix which transforms the elliptical into a spherical footprint and is computed from the photon differential vectors. Finally, $K$ is the Silverman's second-order kernel as in [11] and $E_P$ is the irradiance of photon $P$.

A naive implementation of (1) using PyTorch's `scatter_add` method leads to prohibitive memory consumption and run time, when increasing the photon footprint. Thus, we chose to implement this method in a custom CUDA kernel. As in the original method, we evaluate this equation by looping over the photons and splatting their contribution onto the pixels in the support of the photon footprint (i.e., $x \in \operatorname{supp}(K, M_P, x_P)$) instead of looping over pixels and gathering from contributing photons. Our final forward simulation operator thus consists of chaining a path tracing operator $P_{h,s}(d)$ for photons from the light source $\ell$ and photon splatting: $b \approx F_{h,s}(d) = E(P_{h,s}(d))$.

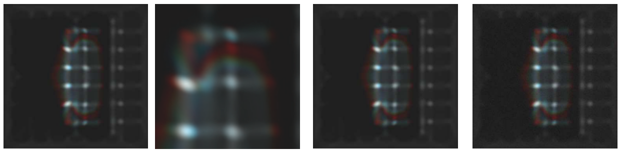

(a) Low noise reference (b) Crop of (a) (c) High noise reference (d) (a) with added noise

Figure 4: Reference caustic of prism and lens array scene as in Fig. 3(c)/(d) with three wavelengths mapped to RGB: (a) generated with $N_r \approx 1 \cdot 10^8$ photons and intrinsic noise level of $0.6\%$. (b) shows fine details and dispersion effects. (c) generated with only $N_r \approx 6 \cdot 10^6$ photons and intrinsic noise level of $2.6\%$. (d) adds Gaussian noise to (a) with total noise level of $\delta = 40\%$.

#### 3.1.2 Backward Simulation Operator

For the backward splatting operator we need to compute the partial derivatives of $E$ as follows:

$$\begin{aligned} \frac{\partial E(x)}{\partial E_P} &\approx \pi K(\|M_P(x - x_P)\|_2), \\ \frac{\partial E(x)}{\partial x_P}(\hat{x}) &\approx -E'_P M_P \hat{x}, \\ \frac{\partial E(x)}{\partial M_P}(\hat{M}) &\approx E'_P \hat{M}(x - x_P). \end{aligned} \tag{2}$$

Note that we have provided the partial derivatives of $E$ w.r.t. $x_P$ and $M_P$ as a linear mapping, i.e., $\frac{\partial E(x)}{\partial x_P}: \mathbb{R}^3 \to \mathbb{R}$ and $\frac{\partial E(x)}{\partial M_P}: \mathbb{R}^{2\times 3} \to \mathbb{R}$. In addition, the derivative of the kernel function is $K'(x) = -\frac{12}{\pi}\, x\,(1-x^2)$ if $x < 1$ and $K'(x) = 0$ otherwise. The term $E'_P$ follows from chain rule:

$$E'_P = \pi K'(\|M_P(x - x_P)\|_2) E_P \left( \frac{M_P(x - x_P)}{\|M_P(x - x_P)\|_2} \right)^T.$$

When evaluating the full backward splatting operator on a general differentiable loss function $L$, these partial derivatives have to be summed over all pixels that were affected by photon $P$:

$$\frac{\partial L}{\partial D_P} = \sum_{x \in \operatorname{supp}(K, M_P, x_P)} \frac{\partial L}{\partial E(x)} \frac{\partial E(x)}{\partial D_P} \tag{3}$$

for $D_P \in \{E_P, x_P, M_P\}$. Note that this can be expressed as gather operation over all affected pixels by each photon and is thus the inverse of the splat operation defined above (1).

This approach produces correct gradients for the whole simulation with the exception of areas that rely on differentiable visibility. Therefore, our derivatives are incorrect in the presence of geometry discontinuities, which could be remedied by an approach such as the one described in [22]. In practice, this error is mitigated by a light source located far away and by rather gentle slopes of the height fields, which is our main application scenario.

### 3.2. Reconstruction

A real-world dataset will always include errors in measurement to be dealt with during reconstruction. In our case, the data is the frequency-dependent irradiance (brightness) $b \in \mathbb{R}_+^{c\times n\times n}$ (Fig. 1(a)). We denote data with noise by $b^\delta$, where $\delta > 0$ indicates the relative noise level, i.e., $\|b^\delta - b\|_{\mathrm{F}} \le \delta \|b\|_{\mathrm{F}}$ with $\|\cdot\|_{\mathrm{F}}$ being the Frobenius norm.

As already mentioned, the task is to reconstruct the shape of the printed glass from one or several frequency-dependent caustic images. Note that we will limit our following discussion to one wavelength. Hence, we are interested in the height field $d$ of printed glass on top of a glass substrate with known height $h$ and known sensor distance $s$ such that $F_{h,s}(d)$ matches the data $b^\delta$ as well as possible. We assume that the refractive index of the material is uniform and given. We start all reconstructions with height field $d$ set to zero. When the desired shape is provided as part of the design process, even better results could be achieved, as the specified shape can serve as a better initial solution.

#### 3.2.1 Naive Approach (M1)

The naive reconstruction approach is to minimize the difference between $F_{h,s}(d)$ and $b^\delta$ by solving $\min_d \frac{1}{2}\|F_{h,s}(d) - b^\delta\|_{\mathrm{F}}^2$ for $d \in \mathbb{R}^{m\times m}$. It is straightforward to employ the gradient descent method with step size $\tau_p > 0$ to receive a result. However, the result contains oscillations and differs strongly from the original.

As already mentioned, the underlying problem is ill-posed and we have to add penalty terms to receive a sensible reconstruction. A well-known approach is to minimize the following Tikhonov functional with penalty terms in the convex functional $\mathcal{P}$, i.e., $d \mapsto \frac{1}{2}\|F_{h,s}(d) - b^\delta\|_{\mathrm{F}}^2 + \mathcal{P}(d)$. In the following, we successively add penalty terms taking into account a priori information about the shape of the printed glass. These terms enforce sparsity, keep the result within physical bounds and take into account the conservation of the volume of the printed glass. So, we start with a simple cost function and successively add more sophisticated terms. At the same time we introduce suitable reconstruction schemes and extend them to our needs.

#### 3.2.2 Pixel-Based Landweber Approach (M2)

We assume that the height field $d$ is sparse, i.e., described by few non-zero coefficients, in the pixel basis. This is taken into account by a penalty term weighted by a regularization parameter $\alpha_p > 0$, $\alpha_p \|d\|_1$. Then, we have to solve the minimization problem

$$\min_{d \in \mathbb{R}^{m\times m}} \frac{1}{2}\|F_{h,s}(d) - b^\delta\|_{\mathrm{F}}^2 + \alpha_p \|d\|_1. \tag{4}$$

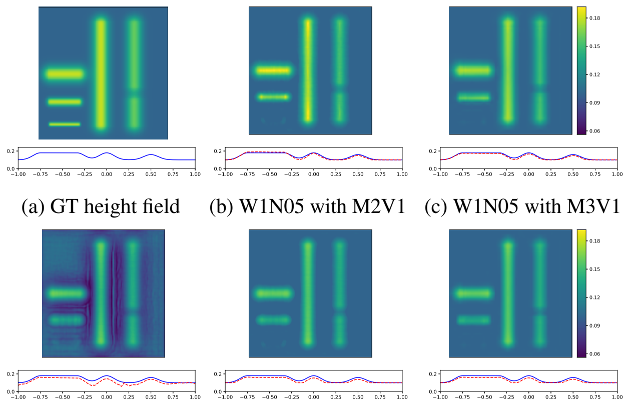

(a) GT height field (b) W1N05 with M2V1 (c) W1N05 with M3V1

(d) W1N05 with M1 (e) W1N05 with M2V0 (f) W1N05 with M3V0

Figure 5: Height field and vertical center slice comparison of reconstruction methods in the case of one (W1) wavelength for simulated data with 5% noise (N05). For better comparison, the center slices contain the predefined height field (blue) and the reconstruction (red).

This can be solved with the well-known thresholded, nonlinear Landweber scheme, see, e.g., [8, 21],

$$d^{[n+1]} = \mathbb{S}\left(d^{[n]} - \tau_p^{[n]} \nabla f_{\mathrm{dis,p}}(d^{[n]}), \tau_p^{[n]}\, \alpha_p\right), \qquad (5)$$

where $\tau_p^{[n]} > 0$ is the step size, $f_{\mathrm{dis,p}}(d) := \frac{1}{2}\|F_{h,s}(d) - b^\delta\|_{\mathrm{F}}^2$, i.e., $\nabla f_{\mathrm{dis,p}}(d) = [F'_{h,s}(d)]^*[F_{h,s}(d) - b^\delta]$, and $\mathbb{S}$ is the soft-shrinkage operator: $\mathbb{S}(d(x), \kappa)$ is $d(x) + \kappa$ if $d(x) \leq -\kappa$, it is 0 in the case of $d(x) \in (-\kappa, +\kappa)$, and $d(x) - \kappa$ for $d(x) \geq \kappa$. For the derivation of the already known Landweber scheme, we refer to the supplementary material.

*Extended Thresholded, Nonlinear Landweber Scheme* In addition to enforce sparsity, it is useful to take into account further a priori information about the height field $d$ of printed glass. We assume that we know physical bounds, i.e., lower and upper bound. Therefore, an indicator function $\delta_{[p_1,p_2]}$ is defined: $\delta_{[p_1,p_2]}(d)$ equals infinity if one or more entries of $d$ are outside a reasonable interval $[p_1, p_2]$ and it equals zero otherwise, see [2]. Then, the minimization problem is

$$\min_{d \in \mathbb{R}^{m\times m}} \frac{1}{2}\|F_{h,s}(d) - b^\delta\|_{\mathrm{F}}^2 + \alpha_p\|d\|_1 + \delta_{[p_1,p_2]}(d).$$

This problem can be solved by the extended thresholded, nonlinear Landweber scheme, see [2] for a more general case and [1] for this case. Similar to (5) it is

$$d^{[n+1]} = \mathcal{I}_{[p_1,p_2]}\left(\mathbb{S}\left(d^{[n]} - \tau_p^{[n]} \nabla f_{\mathrm{dis,p}}(d^{[n]}), \tau_p^{[n]}\, \alpha_p\right)\right),$$

where $\mathcal{I}_{[p_1,p_2]}$ is the interval projection operator defined for real-valued $x$ and $p_1 \leq p_2$ by $\mathcal{I}_{[p_1,p_2]}(x)$ resulting in $p_1$ if $x < p_1$, in $x$ in the case of $x \in [p_1, p_2]$, and in $p_2$ for $x > p_2$. Of course, the element-wise application is assumed if $x$ is a vector. For the derivation of the extended Landweber scheme, we refer to the supplementary material.

*Area-Based Physical Bounds* As additional a priori information we use the knowledge of the axis-accuracy, i.e., where the material was approximately deposited. Hence, we extend the physical bounds to area-based physical bounds: we separate the suspected print area from the suspected unprinted area; see [1] for a more general case. The physical bounds must be zero in the unprinted area. The printed part of the height field $d \in \mathbb{R}^{m\times m}$ of printed glass is denoted by $\check{d}$, the unprinted part by $\mathring{d}$. In the same sense we separate the indices $(j_1, j_2)$ of the height field $d$ into pairs $j := (j_1, j_2)$ of printed glass by $\check{N}$ and unprinted glass by $\mathring{N}$. Then, the minimization problem is

$$\min_{d \in \mathbb{R}^{m\times m}} \frac{1}{2}\|F_{h,s}(d) - b^\delta\|_{\mathrm{F}}^2 + \alpha_p\|d\|_1 + \delta_{[p_1,p_2]}(\check{d}) + \delta_{[0,0]}(\mathring{d}). \qquad (6)$$

We define $w^{[n]} := d^{[n]} - \tau_p^{[n]} \nabla f_{\mathrm{dis,p}}(d^{[n]})$. Then, the solution is

$$d_j^{[n+1]} = \begin{cases} \mathcal{I}_{[p_1,p_2]}\left(\mathbb{S}(w_j^{[n]}, \tau_p^{[n]}\, \alpha_p)\right) & \text{if } j \in \check{N}, \\ \mathcal{I}_{[0,0]}\left(\mathbb{S}(w_j^{[n]}, \tau_p^{[n]}\, \alpha_p)\right) & \text{if } j \in \mathring{N}. \end{cases} \qquad (7)$$

The implementation of $\mathcal{I}_{[p_1,p_2]}\,(\mathbb{S}(w, \kappa))$ for $\kappa > 0$ is simply $\max\{p_1, \min\{p_2, \mathrm{sign}(w)\max\{|w| - \kappa, 0\}\}\}$.

*Conservation of the Volume* A further useful a priori information is the volume of the deposited material as it delivers information about the height field $d$ of the printed glass. Let $V$ be the volume of the glass used for printing and $V_\varepsilon$ the corresponding relative uncertainty. We define $v(d) := \sum_j d_j$ and denote the area of a pixel (on the glass substrate) by $a$. Then, we know that $v(d)$ must be between $q_1 = V \cdot (1 - V_\varepsilon)/a$ and $q_2 = V \cdot (1 + V_\varepsilon)/a$. To enforce this conservation of the volume, we add the penalty term $\delta_{[q_1,q_2]}(v(d))$. However, this penalty term remains elusive to establish in the Landweber scheme. Thus, we develop a heuristic to enforce the conservation of the volume. The effect will be an acceleration of the growing. The main idea is to enforce a slow growing of the volume at the beginning and the end, i.e., if $v(d)$ is close to zero and close to $q_1$. In between, a fast growth is accepted. Therefore, the sine function will be employed. Of course, a similar behaviour is useful if $v(d)$ is too high. Then, the heuristic has to decrease $d$ slowly if $v(d)$ is close to $q_2$. In addition, we ensure that the heuristic does not affect the unprinted area, i.e., we only modify $d_j$ with $j \in \check{N}$. We define the parameter $\gamma$ to influence the growing of the volume. Furthermore, as we are interested in a robust heuristic, we use the mean value of a pixel itself and the surrounding pixels with a radius of $r_V$ pixels and denote it by $\bar{d}$. In the following heuristic, we omit the number of iteration to avoid overloading notation.

| #wave-lengths | Noise level | Recon. method | Fig. 5 | Run time (min) | Rel. dis. | Rel. err. | #iter | $\tau_{dis}$ | $\tau_p$ | $\alpha_p$ | $\tau_w$ | $\alpha_w$ | $V_\varepsilon$ | $\gamma$ |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | 5% | (M2V1) | (b) | 2.95 | .0504 | .0648 | 72 | 1.01 | .1 | .002 | | | .25 | .1 |
| | | (M3V1) | (c) | 2.77 | .0505 | .0658 | 65 | 1.01 | .1 | .0005 | .1 | .01 | .25 | .1 |
| | | (M1) | (d) | 6.62 | .0549 | .1800 | 162 | 1.1 | .1 | | | | | |
| | | (M2V0) | (e) | 4.72 | .0600 | .1086 | 114 | 1.2 | .1 | .002 | | | | |
| | | (M3V0) | (f) | 4.38 | .0598 | .1065 | 102 | 1.2 | .1 | .0005 | .1 | .01 | | |

Table 1: Run times, relative discrepancies, and errors as well as chosen parameters for numerical experiments shown in Fig. 5. In M2 and M3 the physical bounds are limited between $p_1 = 0$ and $p_2 = 0.3$. In V1 we set the radius $r_V$ to 2.

This heuristic is intended to be used after each reconstruction iteration (7) and is defined for $j \in \check{N}$ by

$$d_j := \begin{cases} d_j + \bar{d}_j\,\gamma\,\sin\left(\pi\frac{v(d)}{q_1}\right) & \text{if } v(d) \leq q_1, \\ d_j - \bar{d}_j\,\gamma\,\sin\left(\pi\left[\frac{v(d)}{q_2} - 1\right]\right) & \text{if } v(d) \geq q_2. \end{cases} \quad (8)$$

**Stopping Criterion** Following the best practice for inverse problems, we stop the reconstruction iteration by Morozov's discrepancy principle, see [25], i.e., if $\|F_{h,s}(d) - b^\delta\|_{\mathrm{F}}/\|b^\delta\|_{\mathrm{F}} \leq \tau_{\mathrm{dis}}\,\delta$ with tolerance parameter $\tau_{\mathrm{dis}} > 1$.

**Solution of the Multi-Frequency Problem** The reconstruction of the shape of the printed glass from one caustic image is given by the minimization of (6) by (7) and (8). Of course, this single-frequency case can be extended to the multi-frequency case, i.e., we take several frequency-dependent caustic images into account. We provide a channel for each frequency. Hence, for $c$ channels, the forward operator $F_{h,s}$ has the codomain $\mathbb{R}_+^{c\times n\times n}$ instead of $\mathbb{R}_+^{n\times n}$ and the brightness $b$ is in $\mathbb{R}_+^{c\times n\times n}$ instead of $\mathbb{R}_+^{n\times n}$. To avoid overloading notation we will not consider the multi-frequency case in the following.

#### 3.2.3 Wavelet-Based Landweber Approach (M3)

The Landweber scheme can be used with respect to a wavelet basis instead of the pixel basis, see [8]. In addition, we use the combined effect of the sparsity in a wavelet basis with the sparsity in the pixel basis.

We introduce the discrete wavelet transform $W$ and its inverse $W^{-1}$. The corresponding wavelet coefficients are denoted by $w_d$, i.e., $w_d = Wd$ and $d = W^{-1}w_d$. Then, the minimization problem (4) with respect to a wavelet basis instead of the pixel basis is $\min_{w_d} \frac{1}{2}\|F_{h,s}(W^{-1}w_d) - b^\delta\|_{\mathrm{F}}^2 + \alpha_w\|w_d\|_1$, with regularization parameter $\alpha_w > 0$ to enforce the sparsity in a wavelet basis. The choice of the wavelet family depends on the application. In the case of printed glass we decided on Daubechies 3 wavelets.

Analogous to the definition of $f_{\mathrm{dis,p}}$, we define $f_{\mathrm{dis,w}}(w_d) := \frac{1}{2}\|F_{h,s}(W^{-1}w_d) - b^\delta\|_F^2$. Then, the corresponding Landweber scheme is similar to (5),

$$w_d^{[n+1]} = \mathbb{S}\left(w_d^{[n]} - \tau_w^{[n]}\nabla f_{\mathrm{dis,w}}(w_d^{[n]}), \tau_w^{[n]}\,\alpha_w\right), \quad (9)$$

where $\tau_w^{[n]} > 0$ is the step size. We implement this scheme using the PyTorch-wavelets library for $W$ and $W^{-1}$ [7].

To take into account sparsity in the pixel basis, area-based physical bounds and the conservation of the volume, we do the following steps after each reconstruction iteration (9). First, we compute the corresponding height field $d$ of printed glass for the wavelet coefficients $w_d$. Second, we use (7) with $d^{[n]}$ instead of $w^{[n]}$ as input. Third, the heuristic for conservation of volume as described in (8) is employed. Finally, we compute the corresponding wavelet coefficients from the result for the next reconstruction iteration.

## 4. Results

In this section we will discuss the results of our methods, including forward simulations by $F_{h,s}$, and reconstructions with multiple algorithmic variations. Our experiments were executed and run times taken on a workstation PC with an Intel i9-7980XE CPU with 2.6 GHz and 128 GB of RAM. The GPU was a Nvidia Titan RTX with 1.35 GHz and 24 GB of VRAM.

### 4.1. Simulation Results

Resulting caustic images from our simulations with the complex prism and lens-array scene from Fig. 3(c)/(d) can be seen in Fig. 4. This setup reveals light dispersion effects by our multi-spectral approach as well as inherent noise of the estimator in (1). We roughly estimate this intrinsic noise by computing the simulation twice as well as the corresponding relative discrepancies. We then add further Gaussian noise up to the user-defined noise level $\delta$ to generate a sensor readout that mimics the sensor response in a real-world production system. This is then used in all our reconstruction experiments as input data, excepting Sec. 4.5.

### 4.2. Ablation Study

As mentioned in Sec. 2, sparsity is well-known and area-based physical bounds have already been studied. Therefore, we focus our ablation study on the additional effect of the novel penalty terms (Fig. 5). We perform our ablation study on a predefined height field, that is depicted in Fig. 5 (cf. Fig. 1(a) for the full setup) as ground truth (GT) along with reconstruction results of our algorithm variations. This synthetic height field was generated to illustrate several problems in glass 3D printing like interruptions in the printing

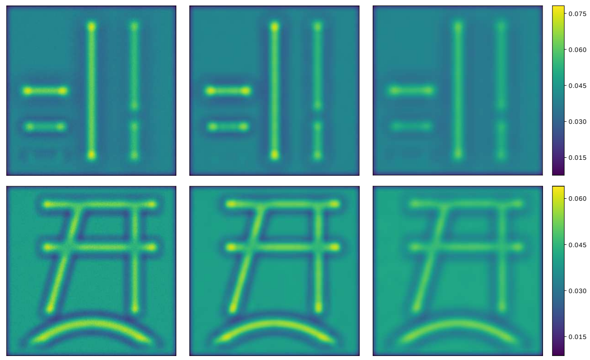

(a) GT caustic (b) M2V1 caustic (c) [33] caustic

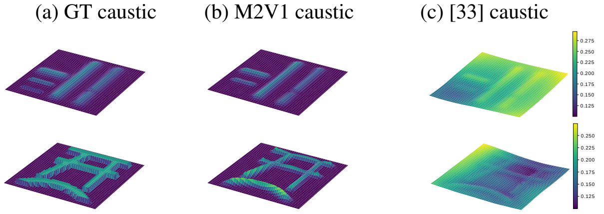

(d) GT surface (e) M2V1 reconstruction (f) [33] reconstruction

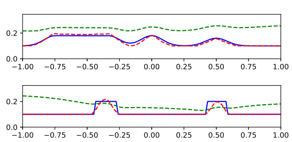

| | Discrepancies w.r.t. (a) | | Errors w.r.t. (d) | |
|---|---|---|---|---|
| | Lines | A | Lines | A |
| **[33]** | 0.0693 | 0.0862 | 1.0217 | 0.5596 |
| **M2V1** | 0.0544 | 0.0548 | 0.0702 | 0.1760 |

(g) Slice comparison (h) Measured deviations

Figure 6: Comparison of reconstructed caustics (first two rows) and surfaces (rows 3 and 4) from the *Lines* and *A* height fields. The last row shows: (g) the center slices of the reconstructions (blue: GT, green: [33], and red: ours); as well as a table (h) with the reconstruction errors from both algorithms.

line and lines of different thickness. Furthermore, reconstructing the thinnest print line is particularly challenging, as the receiver screen lies behind the focal point generated by the strong curvature. We compare three algorithmic variations: M1 naive gradient descent, M2 pixel-based Landweber and M3 wavelet-based Landweber with disabled (V0) and enabled (V1) volume heuristic, where applicable. We set the noise level $\delta = 5\%$ to simulate real-world sensor data. Corresponding numerical results and the chosen parameters are indicated in Tab. 1. Here discrepancy refers to the error in caustic images and error denotes height field errors. As expected, the reconstruction with M1 differs strongly from the ground truth and has the longest run time. Note that here the reconstructed form is partly below the substrate height of 0.1 and exhibits oscillating features outside the print regions. This clearly indicates that regularization is necessary for physically plausible reconstruction results. Both pixel-based (M2) and wavelet-based Landweber (M3) enforce physical bounds as well as sparsity, thus delivering close to optimal reconstructions. Our experiments indicate that both techniques perform similarly well. (Further comparisons of (M2) and (M3) are in the supplementary material.) A significant positive impact can be observed from the volume heuristic (V1), as it decreases the height field error from 11% to 7% and the run time from 5 min to 3 min. None of the method variations here are able to fully reconstruct the smallest printline. Nevertheless, experiments using three wavelengths shows hints of the thinnest printline as discussed in the supplementary material.

Of course, using three images (one image for each wavelength) provides more information than the case we usually consider (one image with one wavelength). Note that a setup with images from multiple viewing angles would increase the amount of data and improve the reconstruction quality but cannot be realized for the intended *in-situ* measurements in production. Therefore, the problem requires a specific numerical algorithm and the input of some known parameters.

### 4.3. Comparison with Caustic Design

We compare our reconstructed shapes against the caustic design approach from Schwartzburg et al. [33]. Even though caustic design has a different objective from our method, the approach has common inputs and outputs, facilitating quantitative and qualitative comparisons in Fig. 6. This figure shows two datasets, one denoted as *Lines* depicting different sizes and orientations of printed lines and the other denoted as *A*, which was chosen to emphasize the effect of line features that are not aligned to the grid as well as steeper slopes of the height field. For these results we use the reference simulation as in M2V1 with the same amount of noise $\delta = 5\%$ as input to both methods. Further information about the Schwartzburg parameters can be found in the supplementary material. Both approaches are able to construct shapes which produce caustics with low errors for the sample height fields. However, there are clear differences in the surface reconstruction as demonstrated by the lower surface error of our method. On the one hand, this further highlights the ill-posedness of the problem. On the other hand, it shows two possible solutions: Schwartzburg et. al. solve it by constructing a shape which is closer to the light source and smoother overall. Our regularization, however, prevents us from finding shapes, which are implausible in the 3D printing context.

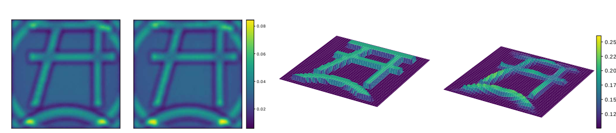

(a) GT caustic (b) Predicted caustic (c) GT surface (d) Reconstructed surface

Figure 7: Reconstruction for the two lenses setup

### 4.4. Unrestricted Setups

We provide additional reconstructions for more complex scenes with two lens elements in the light path as depicted in Fig. 3(a)/(b). In Fig. 7 we show the reference surface and caustic image as well as our results using M2V1. These lenses result in a zoom-in effect on the caustic image and produce a visible ring due to their outer border. Even in this challenging scenario we achieve reconstruction error of 17.1%, which is similar to the scene without the lenses. Due to the additional intersection events and decreased step size, the run time increases from 2.05 min to 8.55 min. In contrast to this example, the scene in Fig. 3(c)/(d) poses a greater challenge (see supplementary material). This is mainly due to light paths which deposit the energy far from their origin and paths which do not intersect the sensor plane (see Fig. 4(a)) and thus do not contribute to the loss. In these cases additional measurements will be needed to reach the same quality. However, to the best of our knowledge, we are the first to demonstrate shape reconstruction in the presence of arbitrary number of refraction events.

### 4.5. Real-World Sample

Finally we provide a result on a real-world sample as depicted in Fig. 8. While the sample suffers from production artifacts such as a broken edge and a weld on the backside of the substrate, we are still able to obtain a reconstruction result on a suitable crop of the piece. The photograph of the workpiece was obtained with a Sony α7R IV mirrorless camera and a Sony 90 mm F2.8 Macro lens focused on a semi-transparent paper screen. The light source was a single flash light with a diffuser and a green filter of 525 nm, with which we roughly approximate an equivalent point light source for reconstruction. As can be seen in the reconstructed height field in Fig. 8(b), the overall shape and position of the printed feature is accurate, but the small detail at the end of the line is not reconstructed and the overall width of the line is less consistent than in reality. Furthermore, the height of the printed structure is generally underestimated, with the reconstructed piece having a maximal height of 4.14 mm (approximately 0.67 units in Fig. 8(b)) and the real piece having a maximal height of 4.20 mm (without printing 4.08 mm). The run time for this reconstruction was 1.78 min and final relative discrepancy 0.1740. Further information can be found in the supplementary material. Considering the feature size of the printing material, which is only 0.4 mm in diameter, this may already enable some feedback loops in production, but for final deployment improvements such as a better screen material and optimized light sources are necessary to achieve the required accuracy for applications such as adaptive layer height adjustments.

## 5. Conclusion

In this paper we presented a reconstruction approach for feedback loops in glass printing processes, which is beyond the realm of usual caustic design methods. Although 3D glass printing is a forward-looking application of our reconstruction framework, it is not limited to this area. Its modular and flexible design, the robustness against highly perturbed data and the ability of the simulation to deal with arbitrary scattering obstacles facilitate its use in a wide range of applications. For example, by extending the method to deal with complex absorption coefficients, one could simulate and reconstruct a wider range of materials. Our experiments demonstrate the effectiveness of our approach on predefined shapes, perturbed data and 3D printable glass geometries. Moreover, we are able to successfully simulate caustics and reconstruct their corresponding shapes in scenarios containing additional elements such as various lenses. To the best of our knowledge, we are the first to successfully reconstruct shape from caustics in the presence of arbitrary number of refractions.

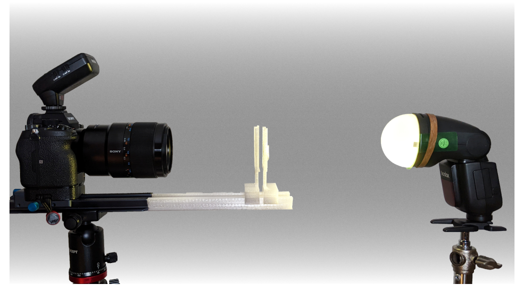

(a) Measurement setup

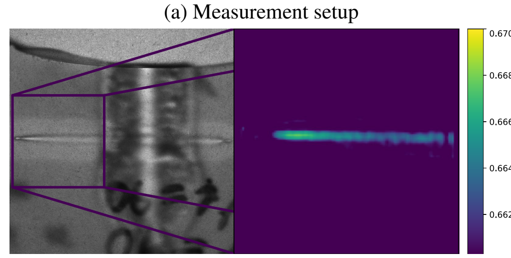


(b) Photo of caustic and reconstructed height field of specific crop

Figure 8: Measurement setup, input and reconstruction

**Acknowledgements** The authors would like to gratefully acknowledge funding from the German Science Foundation (DFG) under Germany's Excellence Strategy within the Cluster of Excellence PhoenixD (EXC 2122, Project ID 390833453), and from the German Federal Ministry of Education and Research (grant No. 05M18MBA-MOReNet).

## References

[1] F. Bürgel. *Effective and Efficient Reconstruction Schemes for the Inverse Medium Problem in Scattering*. Dissertation, Universität Bremen, 2019. Retrieved from `http://nbn-resolving.de/urn:nbn:de:gbv:46-00107828-15`.

[2] F. Bürgel, K.S. Kazimierski, and A. Lechleiter. A sparsity regularization and total variation based computational framework for the inverse medium problem in scattering. *Journal of Computational Physics*, 339:1–30, 2017.

[3] F. Bürgel, K.S. Kazimierski, and A. Lechleiter. Algorithm 1001: IPscatt—A MATLAB Toolbox for the Inverse Medium Problem in Scattering. *ACM Transactions on Mathematical Software*, 45(4):45:1–45:20, 2019.

[4] K. Canales. An austin startup can 3d-print tiny homes in 24 hours for a fraction of the cost of traditional homebuilding – here's how icon could revolutionize affordable housing, Oct. 2019. [Online; Last accessed 17-April-2020].

[5] D. Chen, D. Levin, P. Didyk, P. Sitthi-Amorn, and W. Matusik. Spec2fab: A reducer-tuner model for translating specifications to 3d prints. *ACM Trans. Graph.*, 32(4), July 2013.

[6] D. Colton and R. Kress. *Inverse Acoustic and Electromagnetic Scattering Theory*. Springer, New York, 2013.

[7] F. Cotter and S. McLaughlin. fbcotter/pytorch_wavelets: Zenodo release, Oct. 2019. doi 10.5281/zenodo.3471950.

[8] I. Daubechies, M. Defrise, and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied Mathematics*, 57(11):1413–1457, 2004.

[9] H.W. Engl, M. Hanke, and A. Neubauer. *Regularization of Inverse Problems*. Kluwer Academic Publishers, Dordrecht, Netherlands, 1996.

[10] M. Finckh, H. Dammertz, and H.P.A. Lensch. Geometry construction from caustic images. In *Computer Vision – ECCV 2010*, pages 464–477. Springer Berlin Heidelberg, 2010.

[11] J.R. Frisvad, L. Schjøth, K. Erleben, and J. Sporring. Photon differential splatting for rendering caustics. *Computer Graphics Forum*, 33(6):252–263, Sept. 2014.

[12] G. Gastelu. Local motors 3D-printed car could lead an american manufacturing revolution, July 2014. [Fox News; Last accessed 17-April-2020].

[13] I. Georgiev, J. Křivánek, T. Davidovic, and P. Slusallek. Light transport simulation with vertex connection and merging. *ACM Trans. Graph.*, 2012.

[14] T. Hachisuka and H.W. Jensen. Stochastic progressive photon mapping. *ACM Trans. Graph.*, Dec. 2009.

[15] T. Hachisuka, S. Ogaki, and H.W. Jensen. Progressive photon mapping. *ACM SIGGRAPH Asia 2008 papers*, 2008.

[16] H.W. Jensen. Global illumination using photon maps. *Rendering Techniques' 96*, 1996.

[17] H.W. Jensen. *Realistic image synthesis using photon mapping*. AK Peters/CRC Press, 2001.

[18] J.T. Kajiya. The rendering equation. In *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '86, pages 143–150, New York, NY, USA, 1986. ACM.

[19] T. Kiser, M. Eigensatz, and M. Pauly. Architectural caustics — controlling light with geometry. *Advances in Architectural Geometry 2012*, pages 91–106, 2013.

[20] P. Kürnsteiner, M. B. Wilms, A. Weisheit, B. Gault, E. A. Jägle, and D. Raabe. High-strength damascus steel by additive manufacturing. *Nature*, 582(7813):515–519, June 2020.

[21] A. Lechleiter, K.S. Kazimierski, and Mirza Karamehmedović. Tikhonov regularization in $L^p$ applied to inverse medium scattering. *Inverse Problems*, 29(7):075003, 2013.

[22] G. Loubet, N. Holzschuch, and W. Jakob. Reparameterizing discontinuous integrands for differentiable rendering. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 38(6):2019, 2019.

[23] Jiahui Lyu, Bojian Wu, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. Differentiable refraction-tracing for mesh reconstruction of transparent objects. *ACM Transactions on Graphics (Proceedings of SIGGRAPH ASIA 2020)*, 39(6):195:1–195:13, 2020.

[24] J. Meyron, Q. Mérigot, and B. Thibert. Light in power: A general and parameter-free algorithm for caustic design. *ACM Trans. Graph.*, 37(6), Dec. 2018.

[25] V.A. Morozov. *Methods for Solving Incorrectly Posed Problems*. Springer, New York, 1984. Translated from the Russian.

[26] N. J. W. Morris and K. N. Kutulakos. Dynamic refraction stereo. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 2, pages 1573–1580 Vol. 2, 2005.

[27] M. Nimier-David, D. Vicini, T. Zeltner, and W. Jakob. Mitsuba 2: A retargetable forward and inverse renderer. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 38(6), 2019.

[28] M. Papas, W. Jarosz, W. Jakob, S. Rusinkiewicz, W. Matusik, and T. Weyrich. Goal-based caustics. In *Computer Graphics Forum*, volume 30, pages 503–511, 2011.

[29] S.G. Parker, A. Robison, M. Stich, J. Bigler, A. Dietrich, H. Friedrich, J. Hoberock, D. Luebke, D. McAllister, M. McGuire, and K. Morley. OptiX: a general purpose ray tracing engine. *ACM Trans. Graph.*, 29(4):1, July 2010.

[30] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. In *NIPS 2017 Workshop on Autodiff*, 2017.

[31] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.

[32] K. Rettschlag and S. Kramprich. Printing complex glass structures in 3d using the laser, Apr. 2020. [Online; Last accessed 17-April-2020].

[33] Y. Schwartzburg, R. Testuz, A. Tagliasacchi, and others. High-contrast computational caustic design. *ACM Transactions on Graphics*, 2014.

[34] Q. Shan, S. Agarwal, and B. Curless. Refractive height fields from single and multiple images. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 286–293, 2012.

[35] V. Sitzmann, S. Diamond, Y. Peng, X. Dun, S. Boyd, W. Heidrich, F. Heide, and G. Wetzstein. End-to-end optimization of optics and image processing for achromatic extended depth of field and super-resolution imaging. *ACM Trans. Graph.*, 37(4):1–13, July 2018.

[36] J.C. Spall. A stochastic approximation algorithm for large-dimensional systems in the Kiefer-Wolfowitz setting. In *Pro-*

*ceedings of the 27th IEEE Conference on Decision and Control*, pages 1544–1548 vol.2, Dec. 1988.

[37] E. Veach and L.J. Guibas. Optimally combining sampling techniques for monte carlo rendering. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 419–428, 1995.

[38] P. von Witzendorff, L. Pohl, O. Suttmann, P. Heinrich, A. Heinrich, J. Zander, H. Bragard, and S. Kaierle. Additive manufacturing of glass: CO2-Laser glass deposition printing. *Procedia CIRP*, 74:272–275, Jan. 2018.

[39] G. Wetzstein, D. Roodnick, W. Heidrich, and R. Raskar. Refractive shape from light field distortion. In *2011 International Conference on Computer Vision*, pages 1180–1186, 2011.

[40] T. Whitted. An improved illumination model for shaded display. In *Proceedings of the 6th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '79, page 14, New York, NY, USA, Aug. 1979. Association for Computing Machinery.

[41] Y. Yue, K. Iwasaki, B. Chen, Y. Dobashi, and T. Nishita. Poisson-based continuous surface generation for goal-based caustics. *ACM Trans. Graph.*, 33(3), June 2014.