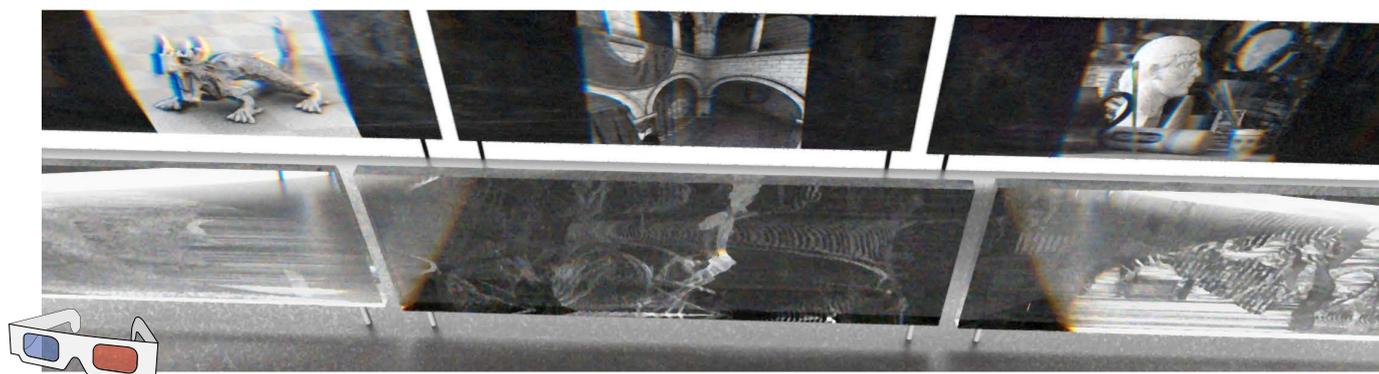


# Anaglyph Caustics with Motion Parallax

M. Lancelle<sup>1</sup>, T. Martin<sup>1</sup>, B. Solenthaler<sup>1</sup> and M. Gross<sup>1,2</sup>

<sup>1</sup>Department of Computer Science, ETH Zurich, Switzerland

<sup>2</sup>Disney Research Zurich, Switzerland



**Figure 1:** Simulated results of computed lenses. When viewed with anaglyph glasses, their caustic images produce a stereoscopic depth impression. Also without glasses, the caustic exhibits a 3D parallax effect when moving the lens towards or away from the screen. The anaglyph glasses symbol indicates that the figure contains stereoscopic image content and thus can be viewed with anaglyph glasses.

## Abstract

In this paper we present a method to model and simulate a lens such that its caustic reveals a stereoscopic 3D image when viewed through anaglyph glasses. By interpreting lens dispersion as stereoscopic disparity, our method optimizes the shape and arrangement of prisms constituting the lens, such that the resulting anaglyph caustic corresponds to a given input image defined by intensities and disparities. In addition, a slight change of the lens' distance to the screen causes a 3D parallax effect that can also be perceived without glasses. Our proposed relaxation method carefully balances the resulting pixel intensity and disparity error, while taking the subsequent physical fabrication process into account. We demonstrate our method on a representative set of input images and evaluate the anaglyph caustics using multi-spectral photon tracing. We further show the fabrication of prototype lenses with a laser cutter as a proof of concept.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Three-Dimensional Graphics and Realism—Display Algorithms

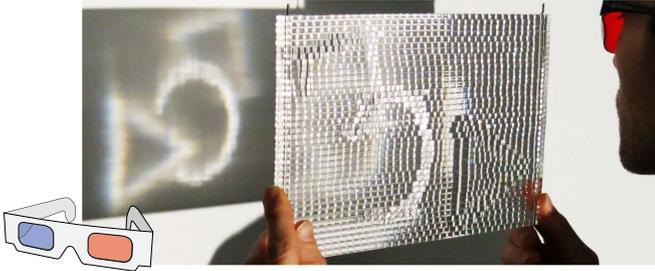
## 1. Introduction

A fascinating light phenomena, caustics, emerges when light is reflected or refracted by a material, projecting patterns with locally varying intensities onto another surface. They are especially interesting as the resulting light distribution cannot be easily deduced from the medium that produces the caustic. Perhaps the most spectacular optical effect, chromatic dispersion, is observed when light breaks into its spectral components resulting in a rainbow of colors. This rather subtle effect manifests itself as a fringe of colors along the boundaries of the caustic.

In computer graphics, caustics and chromatic dispersion effects

have been well studied in rendering, for example, to visualize materials such as gemstone [GS04] or to simulate rainbow formation [SML\*12]. While these works are characterized by taking an input surface and computing the resulting dispersion, the inverse problem, i.e., computing the surface geometry which diverts the light such that a desired caustic is produced, has recently gained increased attention. This finds particular application in architectural design, e.g., to generate precise caustic images [KEN\*12], or in entertainment, e.g., to develop a toy that decodes hidden images [PHN\*12].

Similar to previous work, we explore the computational design



**Figure 2:** The produced lens for the circle example casts an anaglyph caustic onto the wall, producing a 3D impression when viewed with anaglyph glasses.

of lenses to create the desired caustic images. In contrast, we explore the generation of stereoscopic 3D images. For this, we add chromatic dispersion to the design process and leverage the relative dispersion of the cyan and red wavelengths to reveal a 3D image when viewed through anaglyph glasses (see Fig. 2). As we show in this paper, a 3D parallax effect can be observed without wearing anaglyph glasses by moving the lens with respect to the screen. The additional value of the 3D effects may be applied in artistic or educational settings. The setup is simple and designed such that a user can explore the effects by interactively experimenting and playing with the lens.

The input of our method is an image with intensities and disparities. By relating input disparity and refracted light rays, we optimize the lens surface, i.e., the size and arrangement of the lens' prisms, such that the resulting anaglyph caustic corresponds to the input image. Our relaxation method allows to maximize light throughput while balancing local intensity and disparity errors. We further consider curved lens surfaces to better focus light and thus to counteract horizontal blur. We demonstrate the results from our physically-based simulation with various input images. Further, we present prototype lenses for real-world anaglyph caustics.

## 2. Related Work

**Stereo Methods.** Anaglyph stereoscopy introduced by Rollmann [Rol53] is a method to achieve a 3D effect: Two images, one for each eye, are overlaid and red/cyan glasses reveal a stereoscopic image. As inexpensive method it plays a role in entertainment and some engineering applications, especially those targeting the mass market. Possible issues are ghosting artifacts and loss of some or all color perception which are addressed by Ideses et al. [IY04] and Woods et al. [WR04]. These problems do not occur in our method as we work with a continuous spectrum display and monochrome input images. An overview of different display and stereo technologies is given in the survey of Masia et al. [MWDG13].

**Rendering Caustic Effects.** The motivation to render physically-based dispersion effects is often driven by simulating optical effects of materials such as prisms [Mus89], gemstones [GS04], water scenes [JC98] or of other common materials [CT82]. Optical effects such as rainbows or halos in participating media are studied in [Gla96, SML\*12]. Hostettler et al. [HHGJ15] use two prisms to split white light into its spectral components. By inserting a mask, color images can be projected. In this context, Snell's Law and the Fresnel equations play an

important role [HTSG91, WNO98] and are relevant for this work. While physically correct rendering of caustics is time intensive, a commonly used approximation to speed up the computation of the Fresnel terms is described in [Sch94]. To further reduce computation time, a GPU method was presented [SKP07] that renders caustics based on projective texture mapping. Günther et al. [GRR\*16] produce a caustic pattern of choice by ignoring the laws of physics and redirecting the final light rays. In our work, we use a physically-based simulation to evaluate the caustics produced by the lens generated with our method.

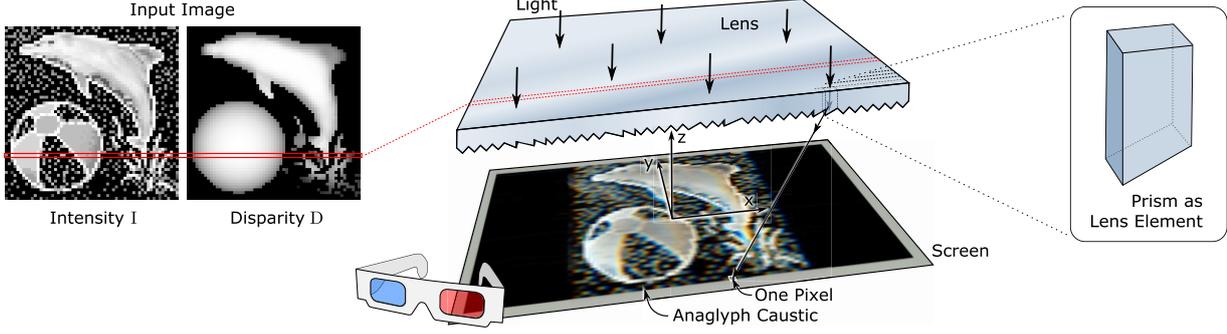
**Lens Modeling and Fabrication.** The challenge of constructing a lens to model and design a desired caustic has recently received attention in computer graphics. With architectural applications in mind, Kiser et al. [KEN\*12] present a method which optimizes a lens surface to control caustic patterns for a desired input shape. It optimizes the reflective and refractive surface generating the caustics, and enables color caustics by incorporating additional color filters. This method has been extended by Schwartzburg et al. [STTP14] and Yue et al. [YIC\*14] to allow high contrast caustics. While the rotation of such lens on a turntable produces interesting caustics, the light patterns do not move related to the image depth. In contrast, with our method a correct parallax effect can be obtained.

Weyrich et al. [WPMR09] propose a system to manufacture a physical surface which exhibits a desired surface appearance. Yue et al. [YIC\*12] show how to achieve caustics by arranging a set of prefabricated lens elements. They can be rearranged to show another caustic image. The simple geometric shape of the lens elements is similar to our approach but does not take dispersion into account. Related to their work is a system proposed by Papas et al. [PJJ\*11] to design and manufacture lens surfaces consisting of micro lenses that produce desired caustics. The method is based on non-negative image decomposition and uses a set of possibly overlapping anisotropic Gaussian kernels for each micro lens. A key property of the method is that caustic images can be achieved with a small number of patches. Finckh et al. [FDL10] propose a related solution to construct a lens surface for a given light source and a desired caustic. The method is based on simultaneous perturbation stochastic approximation which is used as a global optimization to construct the lens.

As the methods cited above are designed to solely model image intensities, it is unclear how to extend them to obtain motion parallax or to produce anaglyph caustics via disparity. These phenomena introduce additional constraints, affecting the lens design and particularly the optimization of the lens surface. Therefore, we introduce a new relaxation method which does not only output the lens geometry, but also enables control of resulting intensity and disparity or parallax errors.

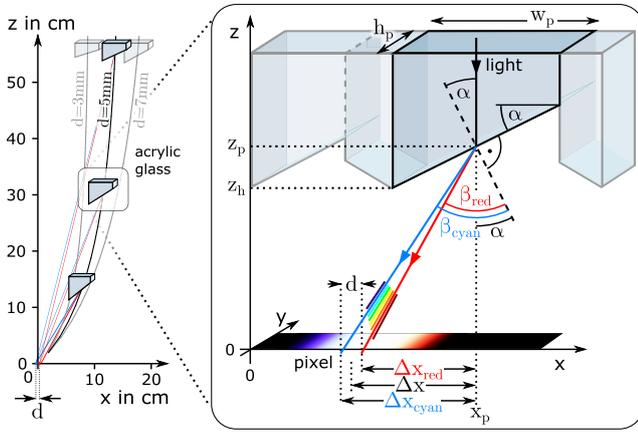
## 3. Anaglyph Caustics

In this section, we introduce our setup and outline how the prism properties control the resulting image. We first describe the anaglyph caustics and then examine how motion parallax is handled. Our proposed method outputs a lens geometry consisting of a collection of prisms, which generates a desired image on the screen.



**Figure 3:** With parallel light from above, the lens produces an anaglyph caustic image on the screen below. Our method computes the prism geometry of each row of the lens separately aiming at obtaining a horizontal disparity ( $x$ -direction). Each prism controls the intensity and disparity of one pixel in the caustic image.

An overview of our setup is illustrated in Figure 3. A close-up view of a single prism and its parameters is shown in Figure 4, right. Each prism creates one pixel of the caustic image on the screen.



**Figure 4:** Computation of the prism locations  $x_p$  and angles  $\alpha$  for a given disparity  $d$ . Left: All prisms along a curve map to the same disparity, respectively, with  $x$  and  $\alpha$  varying. The central curve shows possible prism locations  $(x_p, z_p)$  for  $d = 5\text{mm}$ . Right: Schematics.

Its parameters control the intensity and disparity as discussed in the course of this section. Prisms are perpendicular to the light and parallel to the screen, thus light is not refracted when it enters the prism. We set the prism height  $h_p$  to constant values with subsequent fabrication in mind. The prism width  $w_p$  can vary and depends on the targeted pixel intensity, as discussed in more detail below.

Since we are only interested in horizontal parallax, the prism's  $y$  position is the same as the  $y$  position of the caustic pixel. To achieve a desired disparity, the prism can be placed along the curve shown in Figure 4, left. For the values used in this example and a reasonable screen distance ( $z_p \approx 20\text{cm}$ ) the  $z$  position has hardly any influence on the resulting disparity. Thus, we set the distance from the bottom part of the prisms to the screen,  $z_h$ , to a constant value. Our algorithm changes the prism arrangements only in  $x$ -direction (discussed in Section 4).

When light leaves a prism it is refracted into a rainbow of colors, each corresponding to a specific wavelength as illustrated in Figure 4, right. Given our application to stereoscopic 3D, we are interested in the cyan and red wavelengths,  $\lambda_{\text{cyan}}$  and  $\lambda_{\text{red}}$ , respectively. The distance between the intersection points of these rays on the screen defines the resulting disparity  $d$  of a pixel. The disparity of a pixel depends on several parameters, which are the prism location  $x_p$ , the orientation  $\alpha$  of its bottom surface, and its position along  $z$ , referred to as  $z_p$ .

To compute  $\alpha$ , we consider Snell's law, which defines the relationship between the angles of incidence  $\alpha$  and refraction  $\beta$ , given the indices of refraction of the prism material and air,  $n_{\lambda}^p$  and  $n_{\lambda}^a$ . For our two wavelengths of interest,  $\lambda_{\text{cyan}}, \lambda_{\text{red}}$ , we then get

$$\sin(\beta_{\text{cyan}}) = \sin(\alpha) n_{\lambda_{\text{cyan}}}^p / n_{\lambda_{\text{cyan}}}^a, \quad (1)$$

$$\sin(\beta_{\text{red}}) = \sin(\alpha) n_{\lambda_{\text{red}}}^p / n_{\lambda_{\text{red}}}^a. \quad (2)$$

The wavelengths  $\lambda_{\text{cyan}}$  and  $\lambda_{\text{red}}$ , and their respective two indices of refraction,  $n_{\lambda_{\text{cyan}}}^p$  and  $n_{\lambda_{\text{red}}}^p$ , are determined by an actual measurement of the used anaglyph glass filter foils. By using trigonometric laws, we can define the disparity  $d$  of an image pixel  $(u, v)$ , with  $u \in [1..W]$  and  $v \in [1..H]$ , where  $W$  and  $H$  correspond to the image width and height, respectively, as

$$d_{u,v} = \Delta x_{\text{cyan}} - \Delta x_{\text{red}} \quad (3)$$

$$= z_p \left( \tan(\beta_{\text{cyan}} - \alpha) - \tan(\beta_{\text{red}} - \alpha) \right). \quad (4)$$

By substituting Equations 1 and 2 into Equation 4, we can determine  $\alpha$  numerically. With  $\alpha$ , we can then compute the prism location relative to the projected pixel as

$$\Delta x = \frac{\Delta x_{\text{cyan}} + \Delta x_{\text{red}}}{2}. \quad (5)$$

As  $\beta_{\text{cyan}}$  cannot exceed  $90^\circ$  degrees with  $\alpha$  at the total internal reflection limit, the maximum theoretical disparity can be approximated by

$$d_{\text{lim}} \approx z_p \left( \tan \left( 90^\circ - \arcsin \left( \frac{1}{n_{\lambda_{\text{cyan}}}^p} \right) \right) - \tan \left( \arcsin \left( \frac{n_{\lambda_{\text{red}}}^p}{n_{\lambda_{\text{cyan}}}^p} \right) - \arcsin \left( \frac{1}{n_{\lambda_{\text{cyan}}}^p} \right) \right) \right) \quad (6)$$

when assuming  $n_\lambda^a \approx 1$ . In case of acrylic glass (PMMA), this is around  $d_{im}^{PMMA} \approx 0.22z_p$ .

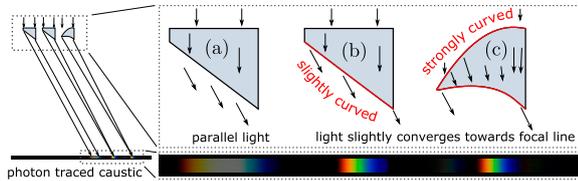
The prism area  $w_p \times h_p$  controls the amount of light of a caustic pixel. While we use a constant value for  $h_p$  to ease fabrication,  $w_p$  can vary. The resulting intensity  $i_{u,v}$  of an image pixel  $(u, v)$  is given as

$$i_{u,v} = \frac{w_p c_t}{S}, \quad (7)$$

where  $c_t$  denotes the transmission coefficient

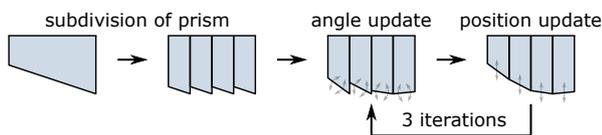
$$c_t = 1 - \frac{1}{2} \left( \left( \frac{\alpha_\emptyset n_\emptyset^a - \beta_\emptyset n_\emptyset^p}{\alpha_\emptyset n_\emptyset^a + \beta_\emptyset n_\emptyset^p} \right)^2 + \left( \frac{\alpha_\emptyset n_\emptyset^p - \beta_\emptyset n_\emptyset^a}{\alpha_\emptyset n_\emptyset^p + \beta_\emptyset n_\emptyset^a} \right)^2 \right) \quad (8)$$

according to [WNO98], where we use the average wavelength  $\lambda_\emptyset = 0.5(\lambda_{red} + \lambda_{cyan})$ .  $S$  is the global intensity parameter. For the maximum input intensity and  $S = 100\%$ , the caustic pixel receives the same amount of light as without a lens. A larger  $S$  will result in a brighter caustic pixel and requires a larger  $w_p$ .



**Figure 5:** Different types of computed lens elements drawn at scale. (a) Planar prisms. (b) Single curved prisms (slightly convex bottom). (c) Double curved prism, creating an ideal rainbow.

Until now we consider straight prism segments. Note that for each wavelength the light exits parallel over the whole prism width. For  $S > 100\%$  this leads to a horizontally blurred image (see Figure 5, (a)). Further causes for horizontal blur are fabrication errors and the wide band pass filters of the red/cyan glasses. To counteract horizontal blurring, a curved shape with a focused rainbow can be computed numerically by piece-wise planar extensions of a start prism segment (c). This means that the rainbows from rays across the whole lens have the same width and focus on the same position. However, the curvature is very strong and the achievable prism width is quite limited. When tolerating a small disparity error, just a convex bottom surface can be used instead (b). This leads to focused light which in practice produces results almost as good as in (c). For (b) and (c), the curvature is more pronounced for prisms that are wider or closer to the screen. To simplify subsequent fabrication, we use single curved prisms (b) and approximate each of them numerically by a number of small planar prisms. In our examples, this causes a maximum disparity error of below 1.1% of the disparity range. This subdivision process is illustrated in Figure 6.



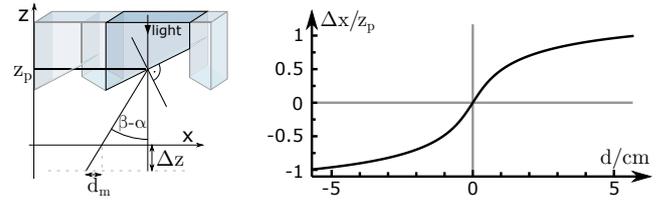
**Figure 6:** Illustration of the subdivision process to compute the geometry of a curved lens element as shown in Figure 5 (b). First, a prism is split into several small prisms. Then, their angles and positions are updated iteratively to remove concave areas.

### 3.1. Using Anaglyph Caustics for Motion Parallax

For a second application we ignore the disparity and are just interested in a parallax effect which is also visible without glasses. The input image has an intensity component  $i$  and a parallax motion component  $d_m$  which specifies the amount of desired parallax motion for each pixel. When changing the distance of the screen to the lens by  $\Delta z$ , the projected light moves along the  $x$ -axis by  $d_m = \Delta z \tan(\beta - \alpha)$  (see Fig. 7, left). Substituting  $\beta$  this yields

$$d_m = \tan \left( \arcsin \left( \frac{n_\lambda^p}{n_\lambda^a} \right) - \alpha \right). \quad (9)$$

Similar as above, to obtain the initial prism parameters,  $\alpha$  and  $\Delta x$  are computed numerically.



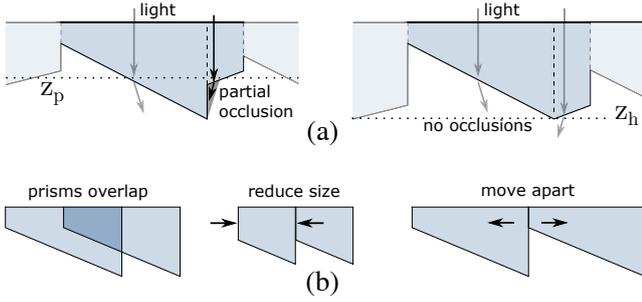
**Figure 7:** Left: Changing the distance of the lens to the screen by  $\Delta z$  causes the projected light to move by  $d_m$ , creating a parallax effect. Right: The amount of 3D parallax motion for anaglyph caustics, depending on input disparity. For a linear relation between depth and disparity, a straight line would be ideal.

Since both applications are very related, we now look at the parallax motion when using a lens computed for anaglyph caustics. Figure 7 shows that the relation between disparity and parallax motion is non-linear and therefore not ideal. However, it increases monotonically and is approximately linear for small values of  $|d|$ . We found that in our experiments, the lens computed for anaglyph caustics can equally be used for both anaglyph and motion parallax caustics. To emphasize or balance use cases, the angle can also be computed with a weighted average between both approaches.

For a slight rotation of the lens, a similar effect can be achieved with again a different computation of  $\alpha$  and  $\Delta x$ . However, the amount of parallax caused by rotation has only a monotonic relation to the disparity or depth of the lenses optimized for above use cases, when  $d$  is restricted to only positive or only negative values. Also, the range of rotation is limited as the total internal reflection may be reached after just a few degrees of rotation.

## 4. Lens Computation

Our algorithm takes as input an intensity and disparity image  $I(u, v)$  and  $D(u, v)$  and computes the optimal arrangements of the lens prisms such that a corresponding anaglyph caustic is produced. The prisms are initially computed as described in Section 3. As each prism is placed independently, multiple prisms can shadow or overlap each other as illustrated in Figure 8. In this section we show how such conflicts can be resolved. As we are only aiming at horizontal disparity ( $x$ -direction), the problem is reduced to finding



**Figure 8:** (a) To avoid self-occlusion (left), conflicted prisms are aligned to their minimal  $z$  value (right). (b) Overlapping prisms (left) are resolved by either reducing their sizes (center) or by moving them apart and recomputing  $\alpha$  (right).

---

**Algorithm A1** Lens Computation

---

- 1: **Input:** intensity and disparity images  $I(u, v), D(u, v)$ ;  $S$  (controls initial overlaps);  $\omega^l$  (balances  $e_i^l$  and  $e_d^l$ );  $z_h$
  - 2: **Output:** position, size and  $\alpha$  of all prisms
  - 3: **for** each row of the lens **do**
  - 4:   initialize prisms using Equations 4,5,7 // overlaps occur
  - 5:   **while** overlaps exist **do**
  - 6:     **for** each overlapping pairs of prisms  $p_i, p_j$  **do**
  - 7:       scale  $p_i, p_j$  //  $e_i^l$  increases locally
  - 8:       move  $p_i, p_j$ , update  $\alpha$  //  $e_d^l$  increases locally
  - 9:     **end for**
  - 10:   **end while**
  - 11:   fabrication specific enhancements // gaps and alignment
  - 12:   subdivide prisms // approximate curved prisms
  - 13:   fill gaps // add light blocking prisms
  - 14: **end for**
- 

an optimal prism configuration for each row of the lens separately. We first discuss the initialization and describe the resulting prism conflicts and their influence on image errors. Then, we describe the iterative relaxation algorithm and discuss post-processing steps. The algorithm is summarized in Algorithm A1. The parameter  $\omega^l$  is balances two types of errors and is explained later.

**Initialization** Prisms are arranged such that they optimally match the given input intensity and disparity (line 4 of Algorithm A1), likely causing some prism overlaps. By choosing a small value for the global intensity parameter  $S$ , the number of initial overlaps can be decreased. This has the advantage that fewer overlaps have to be resolved in the relaxation, and thus local image errors introduced in the relaxation are decreased, too.

**Conflicting Prisms.** There are two types of conflicts: overlaps and self-occlusions. The latter can be avoided by aligning prisms to a common distance  $z_h$  as shown in Figure 8 (a). As discussed before, changing the distance to the screen  $z_p$  has almost no influence on the disparity. The overlap issue is illustrated in Figure 8 (b). Our method to resolve these overlaps is to accept small errors in disparity by horizontally moving conflicting prisms apart ((b), right), and to reduce the scale of the conflicting prisms ((b), center), thereby reducing the intensity of the affected caustic pixels. Both strategies

are used in the relaxation, which allows to balance intensity versus disparity errors. Thus, we iteratively resolve prism overlaps by locally scaling and translating prisms until all conflicts are resolved. The degrees of freedom which are affected by these changes are the position along the  $x$ -direction  $x_p$  and the width  $w_p$ . Consequently,  $\Delta x$  and  $\alpha$  need to be updated.

**Image Errors.** Here, we quantify the image errors that are introduced by the method. During the relaxation, a local intensity error  $e_i^l$  is caused by scaling a prism, changing the intensity of the corresponding pixel  $I(u, v)$  to the optimized intensity  $I_o(u, v)$ . We define the error as

$$e_i^l(u, v) = 1 - \frac{I_o(u, v)}{I(u, v)}. \quad (10)$$

Additionally, a local disparity error  $e_d^l$  is introduced by moving a prism. The disparity of the corresponding pixel  $D(u, v)$  changes to  $D_o(u, v)$ . The error is defined relative to the maximum absolute disparity and is given as

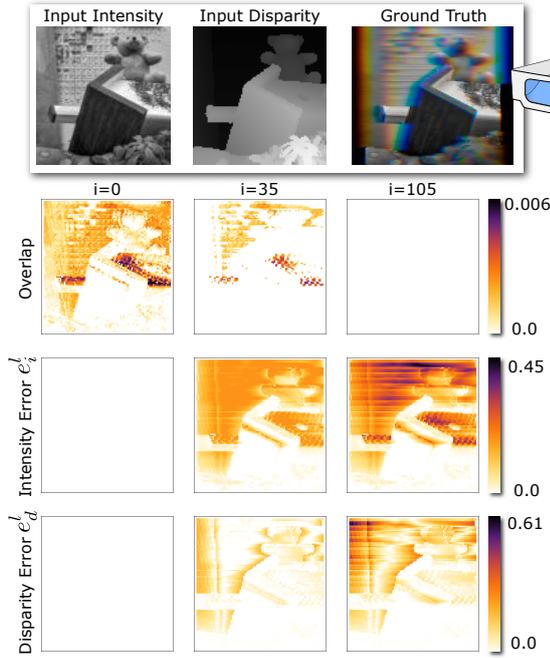
$$e_d^l = \frac{D_o(u, v)}{D(u, v)} \frac{1}{\max|D(u, v)|}. \quad (11)$$

**Relaxation.** Overlaps are then resolved in the iterative, local relaxation by scaling and moving prisms (lines 7 and 8 of Algorithm A1). Reducing the size of the two prisms also reduces image intensity locally and thus increases the local intensity error  $e_i^l$ . We balance between the local intensity and disparity errors,  $e_i^l$  and  $e_d^l$  with the weight  $\omega^l$ .

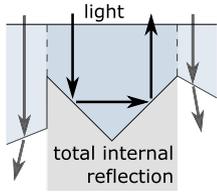
In the implementation, we resolve each colliding prism pair sequentially. Each iteration step consists of two passes. We first traverse all prism pairs of a row from left to right, only shifting the respective right prism to the right. In a second step, we traverse the prisms from right to left, only shifting the left prism to the left. The amount of overlap of a prism pair  $i$  and  $j$  with positions  $x_i < x_j$  is  $\delta = (x_i + w_i/2) - (x_j - w_j/2)$ . In case of an overlap, we move one prism by the amount of translation  $s = \max(0.48\text{pixel size}, \delta)$  and scale both prisms by  $(\omega^l)^{s/p_s}$ , where  $p_s$  is the size of the pixel of the caustic image. We continue with this process until all overlaps are resolved. Note that the order of the prisms is not changed by this process. Figure 9 visualizes how the values of the error terms change during the optimization. While overlaps are being decreased, the disparity and local intensity errors increase.

**Post-processing.** After all overlaps have been resolved, the subdivision of each prism (line 12 of Algorithm A1, Figure 6) is performed to better focus the light along the  $x$ -direction. Each prism is broken up into many smaller prisms with iteratively updated angles and positions.

Rows can contain many small prism segments (see Figure 10 (a)). When aligning lenses to their neighbors, concave areas can be reduced. This also results in a shorter tool path and can be beneficial for some fabrication processes. However, thick lenses result (see Figure 10 (c)) which increase material cost and lens weight. A compromise is to only use vertical segments when the thickness of a group of prisms exceeds a user defined threshold (see Figure 10 (b)). Note that vertical surfaces between rows along the  $y$  direction are necessary as otherwise significant amounts of undesired vertical disparity would result.



**Figure 9:** The amount of overlaps as well as disparity and intensity errors are visualized for the iteration steps 0, 35, and 105. With decreasing overlaps, the local intensity and disparity errors increase.

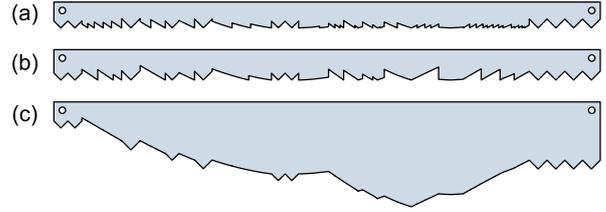


Finally, empty regions between prisms are filled by a  $45^\circ$  zigzag pattern (line 13 of Algorithm A1) using the total internal reflection to block light (see illustration on the left). This means that only materials with  $n_{red}^p > \frac{1}{\sin(45^\circ)} \approx 1.414$  can be used for the lens.

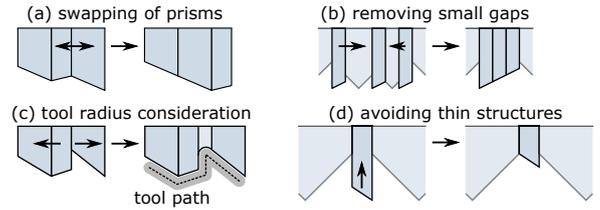
### 5. Simulation and Fabrication

We assess the quality of anaglyph images created by our optimized surfaces in simulations using the Mitsuba [Jak10] ray tracer. We use multiple passes to obtain a multi-spectral result, taking the non-linear relationship between the wavelength of the light and the corresponding index of refraction of the material into account. We follow [KSIN07] to approximate the index of refraction of a material given a discrete set of wavelengths. As we intend to use the lenses with sunlight, we approximate sunlight in the simulation with an area light source corresponding to an angular diameter of  $0.53^\circ$ . Photon tracing then reveals the luminance corresponding to a specific wavelength. The multi-spectral image channels are transformed to sRGB using the look up tables and transformation from [SG31], which relates wavelengths of light to calibrated display systems according to human perception.

For the purpose of validating our framework, we fabricate lens prototypes with laser cutting as a cheap fabrication process. Each row of a lens is cut individually. The lens is then assembled by stacking the rows on top of each other, mounting them with threaded bolts and screws. We apply this strategy to 4mm acrylic



**Figure 10:** Alignment approaches showing geometries automatically computed by our algorithm with a representative example. All approaches result in a similar caustic image. (a) Alignment to  $z_h$  as described above. (b) Partially avoiding concave areas to ease some fabrication processes (used for the cougar lens). (c) Avoiding all vertical segments leads to a thick lens.

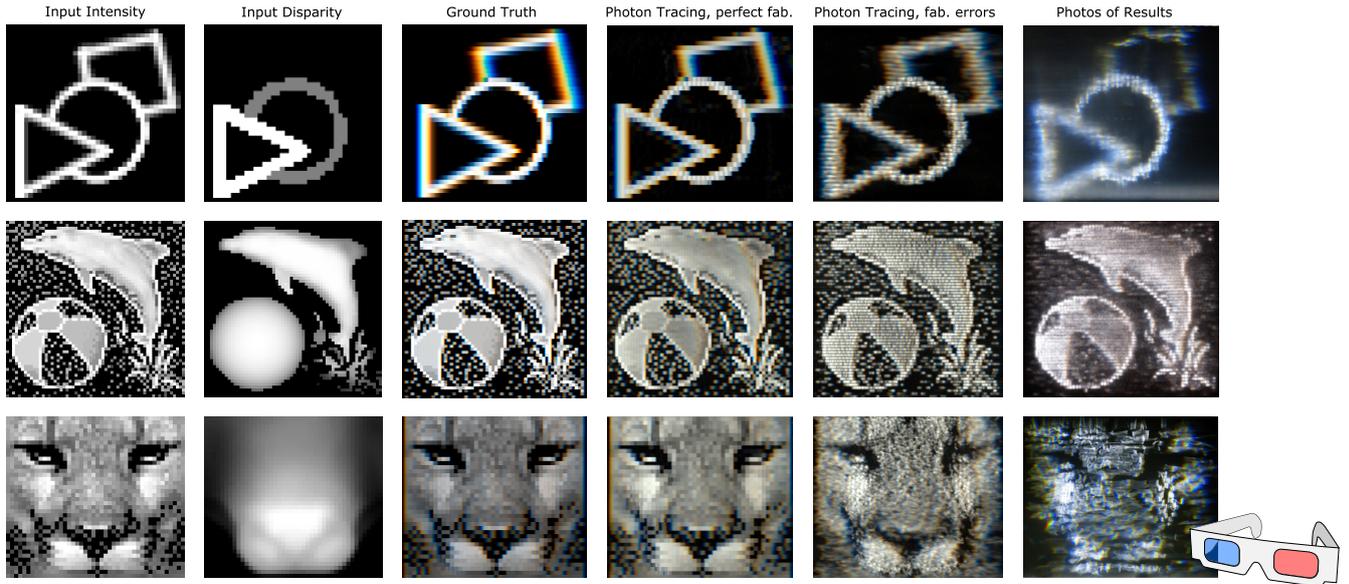


**Figure 11:** Post processing can improve fabrication. In some cases prisms can be swapped (a) or grouped (b) to obtain longer edges with similar angles. Long thin structures can be a mechanical issue and can be avoided (c). For a significant tool radius, prisms must be moved apart at concave corners (d).

glass (PMMA). The main advantage of this fabrication process is the fact that laser cutting is fast and that the prism surface along  $y$  is reproduced with high accuracy. The minimal feature size of the rows depends on the laser radius which is around  $0.15mm$  on the Epilog Helix laser cutter we use. The most significant errors occur due to the melting of the surface, causing rounded edges and a slightly wavy ripple structure.

We noticed that despite the alignment considerations mentioned above, in rare cases, small prisms still occur. They suffer most from fabrication inaccuracies. Inspired by dithering algorithms, small prisms can be avoided by merging several neighboring dark pixels with a similar depth value to a single brighter pixel. This is applied to the input image as a pre-processing step and consequently allows larger input resolutions for the same fabrication tolerance.

In case many small surfaces with large angle variation pose a problem for fabrication, further processing can be done to obtain more continuous rows. First, if a prism with a significantly different angle is between two prisms with a similar angle (within  $5^\circ$ ), we swap the middle prism with the one on the side where the swapping introduces the lower disparity error (see Figure 11 (a)). To avoid single small prisms, small groups of prisms, or small gaps between prisms, the concerned gaps can be iteratively reduced (see Figure 11 (b)). While these strategies increase the disparity error, they may facilitate fabrication or lead to a better result. However, in cases where the disparity error would exceed a threshold ( $0.2 \max|D(u, v)|$ ), these strategies are not applied. For a 2D fabrication process with a tool radius that is significant with respect to the smallest lens width, some prisms have to be moved apart by the



**Figure 12:** From left to right: Input intensity, input disparity, ground truth result, results from photon tracing without and with simulation of fabrication errors, photographs of real-world results by using our fabricated lenses. Original cougar photo by Art G., digitalART2, flickr.

required amount (see Figure 11 (c)). This triggers the application of the relaxation scheme for a second time. Finally, in rare cases, long thin structures are present that could break in some fabrication processes. We remove them as depicted in Figure 11 (d). After most of the described operations,  $\alpha$  and  $z_h$  need to be recomputed and alignment may have to be adjusted in few iterations.

The cost effective laser cutting that we use as a proof-of-concept, introduces small ripples in the individual lens surfaces, resulting in blurred caustics. Fabrication accuracy and thus image quality can be significantly improved by using high end milling and injection molding instead of our inexpensive process. Please refer to the supplemental material for an investigation on alternative fabrication processes.

## 6. Results

### 6.1. Implementation

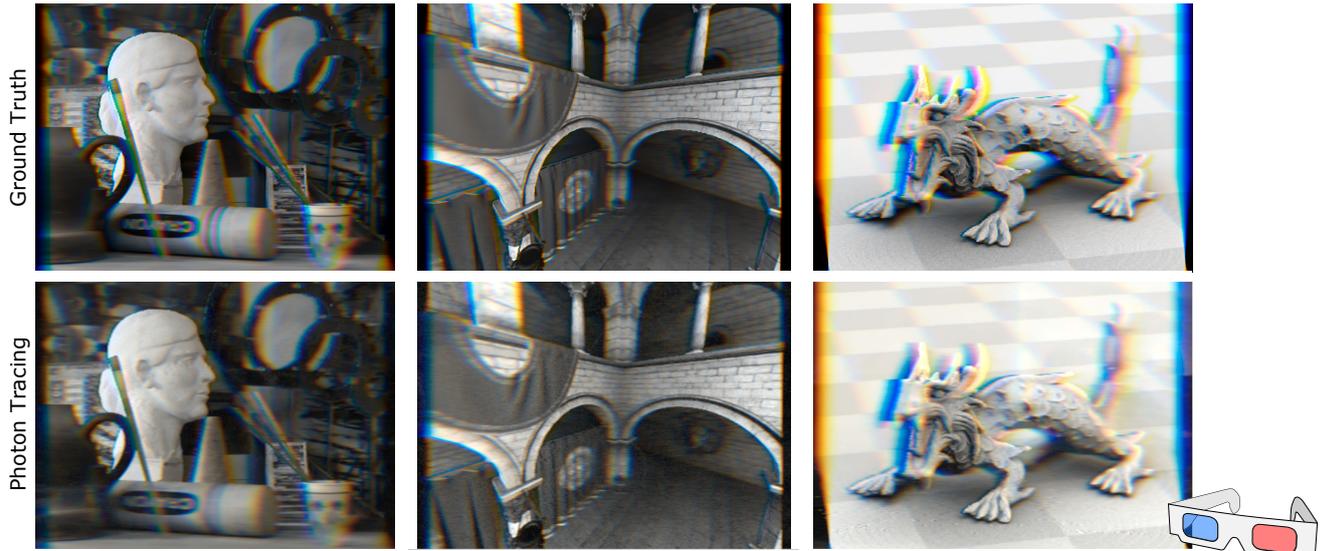
With larger distances to the screen  $z_h$ , larger disparities can be achieved. This also influences the total width of the lens which is in practice limited by the fabrication process. Prisms with large disparities ( $\max|D(u, v)| \gtrsim 0.3 d_{tm}$ ) have surface angles close to the total internal reflection where small surface variations result in larger errors. This amplification of errors results in a blurred image. All of our simulated images were obtained with  $\alpha < 40^\circ$  (where total internal reflection starts at around  $42^\circ$ ). A reasonable choice – used in all our examples – is when  $z_h$  is in the order of magnitude of the dimensions of the actual caustic image. The parameters for all examples shown in this paper are summarized in the supplemental material. With these parameters, the time to optimize the lens surface was less than one minute using a conventional desktop computer. With the low power laser cutter we use, a row of a lens typically takes around 6 minutes to cut, depending on its size.

### 6.2. Anaglyph Caustic Results

We have applied our method to various examples and present the anaglyph caustics from simulation and results from our lens prototypes. Except for the *circle* example, all lenses are computed with single curved prisms to control horizontal blur.

Figure 12 presents three examples, *circle*, *dolphin* and *cougar*, where the input intensity and disparity maps are shown on the left. The ground truth solutions are provided in the following column. The next two columns show the results of our physical simulation with photon tracing, once assuming perfect fabrication and once using a bump map texture to mimic surface deviations caused by fabrication errors. The horizontal stripes are caused by slightly rounded edges along  $x$ . Small ripple structures on the surface along  $y$  cause horizontally spread speckle structures. Photographs from the actual anaglyph caustics generated from the fabricated lenses are presented in the last column. It can be seen that our algorithm produces images close to the ground truth data and that we can well predict most fabrication errors in our simulated results. Especially for the caustic from *dolphin* lens, the stereoscopic 3D effect is reduced compared to the corresponding ground truth solution. This is due to errors introduced during the fabrication process which is visible as image blur. Further, ambient light is responsible for lower contrast. Both issues lower the perception of depth. Nevertheless, the disparity present in the anaglyph caustic images is sufficient to produce the stereoscopic 3D effect. The parallax effect is more robust to fabrication errors and is clearly visible when moving the lens with respect to the screen. In addition, even with a small misalignment of the lens the caustic shows the scene with a parallax.

To further evaluate our proposed framework, we apply our algorithm on a range of higher resolution input images, all with an image width of 400 pixels. Figure 13 shows the synthetic results once as ground truth (top row) and once simulated using multi-spectral photon tracing (bottom row). The leftmost image shows the *mid-*



**Figure 13:** Further examples with ground truth (top row) and the simulated results (bottom row) using multi-spectral photon tracing. From left to right: the middlebury, sponza and dragon data sets.

*middlebury* data set and the center image depicts the *sponza* scene. On the right we show the *dragon* from the Stanford 3D Scanning Repository. The results show that the simulated images correspond well to the ground truth solutions. Note that also without wearing anaglyph glasses a 3D parallax effect can be observed by moving the lens with respect to the wall. This is well visible in the simulation shown in the accompanying video.

The influence of the parameters  $S$  and  $\omega^l$  used in our method is illustrated in Figure 14. For this analysis, we use the *teddy* input image (Figure 9, top), which is a cropped image from the *middlebury* data set. The left part in Figure 14 displays the influence of the global intensity parameter  $S$  on the overall image brightness. It can be seen that a smaller value leads to darker images and thus to less local intensity and disparity errors. By increasing  $S$  and thus the image brightness, local errors increase. Larger  $S$  were chosen to better illustrate the influence of the errors. The influence of  $\omega^l$  is displayed in the right part in Figure 14. With a larger value for  $\omega^l$ , local disparity errors are preferred over intensity errors. A user can easily find a suitable balance between global intensity and introduced errors, as a preview of the resulting caustic image can be obtained within seconds. Note that the errors also depend on the distance to the screen  $z_h$ . With a larger distance, a wider lens could be used with less prism overlaps and thus with fewer errors. However, our examples and parameters are always chosen such that subsequent fabrication of prototype lenses is feasible.

With the *teddy* example, we show how the number of overlaps and local errors change during the relaxation in Figure 9. While the overlaps are maximal at the beginning, they are resolved with an increasing number of iterations, leading to local intensity and disparity errors.

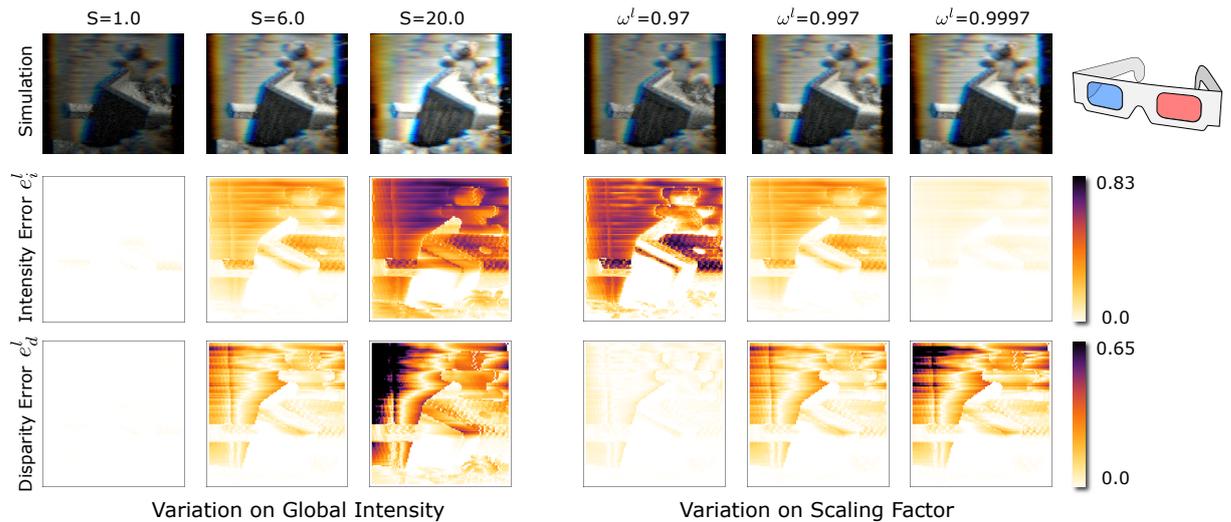
### 6.3. Discussion

**Horizontal Depth Discontinuities.** The intensity of each pixel is the same for both eyes and cannot be controlled independently. Thus, at a strong horizontal depth discontinuity, i.e., a vertical edge, the background that should be visible to one eye is missing. We partially addressed this problem in the dragon example by introducing a disparity dependent offset along  $x$ , thereby moving close parts further left for this example to minimize the gap. Another option is to use further pixels on a separate layer to add some background pixels.

For black intensity areas, no prisms are inserted. Therefore, no depth information can be perceived. Instead, the depth of neighboring pixels or of the screen is perceived. Hence, dark areas in the input image should be avoided, or the described dithering strategy should be applied.

Similarly, large areas with the same intensity should cause a uniformly lit area in the caustic. Thus, there may not be enough depth cues. Instead, on a structured screen the scene might be perceived at screen depth. This problem is resolved by adding high frequencies to the intensity input, e.g. by slightly changing the focusing. As our fabrication processes already introduce some high frequency intensity variations, this problem only appeared in our computed ground truth images.

**Anaglyph Glasses.** Due to the wide spectral transmission of the anaglyph glasses, a wide rainbow is perceived as wide stripes. Thus, large disparities, i.e., near and far image regions, lead to horizontally blurred objects. To reduce this effect we employ stronger focusing as described in Section 4. To avoid large disparities, we also use depth compression (see e.g. [DRE\*12] or [CHA\*14] for more details) on the input disparity images of the *dragon* and *sponza* scenes.



**Figure 14:** Left: Influence of the global intensity parameter  $S$ . A low value leads to very few errors but results in a dark image. At  $S=1.0$ , the brightest pixel intensity of the caustic matches the light intensity on the projection surface next to the lens. Right: Influence of the parameter  $\omega^l$  to trade off disparity errors with local intensity errors. A low value leads to more intensity errors, a higher value leads to more disparity errors.

## 7. Conclusion

In this paper we propose a computational framework and associated relaxation method to produce anaglyph caustics with motion parallax. The algorithm is designed to achieve a good balance between the total image intensity and local errors. Our method is the first to control and design dispersion for anaglyph caustics. We evaluate the proposed method on various photon traced simulations. With slight modifications of the algorithm to easy fabrication we have demonstrated that our computational method can be used to produce real world lenses. Future work could include to extend our method by considering adjacent rows of the lens in the optimization. Another direction of future investigations addresses fabrication aspects, in particular to further improve the lens surface. Lastly, given the rapid improvement of 3D printing technology, we plan to investigate more general lens surface geometries suitable for printing to produce anaglyph caustics.

## Acknowledgements

We thank Nicola Ranieri, Marios Papas, Emily Whiting, Christian Schüller, Cengiz Öztireli, Jakob Wenzel and Disney Research Zurich for their support.

## References

- [CHA\*14] CHAPIRO A., HEINZLE S., AYDIN T., POULAKOS S., ZWICKER M., SMOLIC A., GROSS M.: Optimizing stereo-to-multiview conversion for autostereoscopic displays. *Computer Graphics Forum, proc. of EG 33*, 2 (2014). 8
- [CT82] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. *ACM Trans. Graph.* 1, 1 (1982), 7–24. doi:10.1145/357290.357293. 2
- [DRE\*12] DIDYK P., RITSCHER T., EISEMANN E., MYSZKOWSKI K., SEIDEL H.-P.: Apparent stereo: the cornsweet illusion can enhance perceived depth. In *Human Vision and Electronic Imaging XVII, SPIE* (Burlingame, CA, 2012). 8

- [FDL10] FINCKH M., DAMMERTZ H., LENSCH H. P. A.: Geometry construction from caustic images. In *Proc. ECCV* (2010). 2
- [Gla96] GLASSNER A.: Computer-generated solar halos and sun dogs. *Computer Graphics and Applications* 16, 2 (1996). doi:10.1109/38.486687. 2
- [GRR\*16] GÜNTHER T., ROHMER K., RÖSSL C., GROSCH T., THEISEL H.: Stylized caustics: Progressive rendering of animated caustics. *Computer Graphics Forum (Proc. Eurographics)* (2016). 2
- [GS04] GUY S., SOLER C.: Graphics gems revisited: Fast and physically-based rendering of gemstones. *ACM Trans. Graph.* 23, 3 (2004), 231–238. doi:10.1145/1015706.1015708. 1, 2
- [HHGJ15] HOSTETTLER R., HABEL R., GROSS M., JAROSZ W.: Dispersion-based color projection using masked prisms. *Computer Graphics Forum (Proceedings of Pacific Graphics)* (2015). doi:10.1111/cgfm.12771. 2
- [HTSG91] HE X. D., TORRANCE K. E., SILLION F. X., GREENBERG D. P.: A comprehensive physical model for light reflection. In *Proc. of Computer Graphics and Interactive Techniques* (1991), SIGGRAPH '91, pp. 175–186. doi:10.1145/122718.122738. 2
- [IY04] IDESES I., YAROSLAVSKY L.: New methods to produce high quality color anaglyphs for 3-d visualization. In *Image Analysis and Recognition*, vol. 3212 of *Lecture Notes in Computer Science*. 2004, pp. 273–280. doi:10.1007/978-3-540-30126-4\_34. 2
- [Jak10] JAKOB W.: Mitsuba renderer, 2010. <http://www.mitsuba-renderer.org>. 6
- [JC98] JENSEN H. W., CHRISTENSEN P. H.: Efficient simulation of light transport in scenes with participating media using photon maps. In *Proc. Computer Graphics and Interactive Techniques* (1998), SIGGRAPH '98, pp. 311–320. doi:10.1145/280814.280925. 2
- [KEN\*12] KISER T., EIGENSATZ M., NGUYEN M., BOMPAS P., PAULY M.: Architectural caustics - controlling light with geometry. In *Advances in Architectural Geometry*. 2012, pp. 91–106. doi:10.1007/978-3-7091-1251-9\_7. 1, 2
- [KIN07] KASAROVA S. N., SULTANOVA N. G., IVANOV C. D., NIKOLOV I. D.: Analysis of the dispersion of optical plastic materials. *Optical Materials* 29, 11 (2007), 1481–1490. doi:10.1016/j.optmat.2006.07.010. 6

- [Mus89] MUSGRAVE F. K.: Prisms and rainbows: a dispersion model for computer graphics. In *Proc. of Graphics Interface* (1989). 2
- [MWDG13] MASIA B., WETZSTEIN G., DIDYK P., GUTIERREZ D.: A survey on computational displays: Pushing the boundaries of optics, computation, and perception. *Computers & Graphics* 37, 8 (2013). doi:10.1016/j.cag.2013.10.003. 2
- [PHN\*12] PAPAS M., HOUIT T., NOWROUZEZAHRAI D., GROSS M., JAROSZ W.: The magic lens: Refractive steganography. *ACM Trans. Graph. (SIGGRAPH Asia Proc.)* 31, 6 (2012), 186:1–186:10. doi:10.1145/2366145.2366205. 1
- [PJJ\*11] PAPAS M., JAROSZ W., JAKOB W., RUSINKIEWICZ S., MATUSIK W., WEYRICH T.: Goal-based caustics. *Computer Graphics Forum (Proc. Eurographics)* 30, 2 (2011). 2
- [Rol53] ROLLMANN W.: Zwei neue stereoskopische Methoden. *Annalen der Physik* 166, 9 (1853), 186–187. doi:10.1002/andp.18531660914. 2
- [Sch94] SCHLICK C.: An inexpensive BRDF model for physically-based rendering. *Computer Graphics Forum* 13, 3 (1994), 233–246. doi:10.1111/1467-8659.1330233. 2
- [SG31] SMITH T., GUILD J.: The c.i.e. colorimetric standards and their use. *Trans. of the Optical Society* 33, 3 (1931), 73. 6
- [SKP07] SHAH M., KONTTINEN J., PATTANAIK S.: Caustics mapping: An image-space technique for real-time caustics. *Visualization and Computer Graphics* 13, 2 (2007). doi:10.1109/TVCG.2007.32. 2
- [SML\*12] SADEGHI I., MUNOZ A., LAVEN P., JAROSZ W., SERON F., GUTIERREZ D., JENSEN H. W.: Physically-based simulation of rainbows. *ACM Trans. Graph. (SIGGRAPH Proc.)* 31, 1 (2012), 3:1â3:12. doi:10.1145/2077341.2077344. 1, 2
- [STTP14] SCHWARTZBURG Y., TESTUZ R., TAGLIASACCHI A., PAULY M.: High-contrast computational caustic design. *ACM Trans. Graph.* 33, 4 (July 2014). doi:10.1145/2601097.2601200. 2
- [WNO98] WOLFF L. B., NAYAR S. K., OREN M.: Improved diffuse reflection models for computer vision. *Int. J. Comput. Vision* 30, 1 (1998), 55–71. doi:10.1023/A:1008017513536. 2, 4
- [WPMR09] WEYRICH T., PEERS P., MATUSIK W., RUSINKIEWICZ S.: Fabricating microgeometry for custom surface reflectance. *ACM Trans. Graph. (SIGGRAPH Proc.)* 28, 3 (2009), 32:1–32:6. doi:10.1145/1531326.1531338. 2
- [WR04] WOODS A. J., ROURKE T.: Ghosting in anaglyphic stereoscopic images. In *Proc. SPIE 5291* (2004), pp. 354–365. 2
- [YIC\*12] YUE Y., IWASAKI K., CHEN B.-Y., DOBASHI Y., NISHITA T.: Pixel art with refracted light by rearrangeable sticks. *Computer Graphics Forum* 31 (2012). doi:10.1111/j.1467-8659.2012.03036.x. 2
- [YIC\*14] YUE Y., IWASAKI K., CHEN B.-Y., DOBASHI Y., NISHITA T.: Poisson-based continuous surface generation for goal-based caustics. *ACM Trans. Graph.* 33, 3 (June 2014). doi:10.1145/2580946. 2