



# MIT Open Access Articles

## *Computational Schlieren Photography with Light Field Probes*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

<b>Citation</b>	Wetzstein, Gordon, Wolfgang Heidrich, and Ramesh Raskar. "Computational Schlieren Photography with Light Field Probes." <i>International Journal of Computer Vision</i> 110, no. 2 (August 20, 2013): 113–127.
<b>As Published</b>	<a href="http://dx.doi.org/10.1007/s11263-013-0652-x">http://dx.doi.org/10.1007/s11263-013-0652-x</a>
<b>Publisher</b>	Springer US
<b>Version</b>	Author's final manuscript
<b>Citable link</b>	<a href="http://hdl.handle.net/1721.1/106832">http://hdl.handle.net/1721.1/106832</a>
<b>Terms of Use</b>	Creative Commons Attribution-Noncommercial-Share Alike
<b>Detailed Terms</b>	<a href="http://creativecommons.org/licenses/by-nc-sa/4.0/">http://creativecommons.org/licenses/by-nc-sa/4.0/</a>

# Computational Schlieren Photography with Light Field Probes

Gordon Wetzstein · Wolfgang Heidrich · Ramesh Raskar

Received: 27 February 2013 / Accepted: 1 August 2013 / Published online: 20 August 2013  
© Springer Science+Business Media New York 2013

**Abstract** We introduce a new approach to capturing refraction in transparent media, which we call light field background oriented Schlieren photography. By optically coding the locations and directions of light rays emerging from a light field probe, we can capture changes of the refractive index field between the probe and a camera or an observer. Our prototype capture setup consists of inexpensive off-the-shelf hardware, including inkjet-printed transparencies, lenslet arrays, and a conventional camera. By carefully encoding the color and intensity variations of 4D light field probes, we show how to code both spatial and angular information of refractive phenomena. Such coding schemes are demonstrated to allow for a new, single image approach to reconstructing transparent surfaces, such as thin solids or surfaces of fluids. The captured visual information is used to reconstruct refractive surface normals and a sparse set of control points independently from a single photograph.

**Keywords** Computational photography · Light transport · Fluid imaging · Shape from x

## 1 Introduction

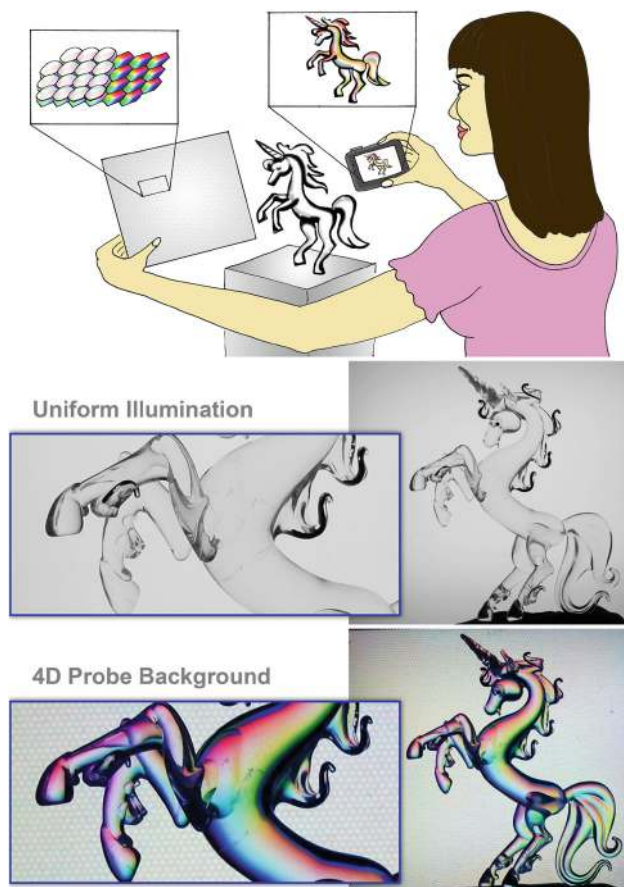
The acquisition of refractive phenomena caused by natural objects has been of great interest to the computer vision and graphics community. Co-designing optical acquisition setups and reconstruction algorithms can be used to acquire refractive solids, fluids, and gas flows (Ihrke et al. 2010), ren-

der complex objects with synthetic backgrounds (Zongker et al. 1999), or validate flow simulations with measured data. Unfortunately, standard optical systems are not capable of recording the nonlinear trajectories that photons travel along in inhomogeneous media. In this paper, we present a new approach to revealing refractive phenomena by coding the colors and intensities of a light field probe. As illustrated in Fig. 1, the probe is positioned behind an object of interest and the object and probe are photographed by a camera. Due to refractions caused by the medium, apparent colors and intensities of the probe change with the physical properties of the medium, thereby revealing them to the camera or a human observer.

The idea of optically transforming otherwise invisible physical quantities into observed colors and changes in intensity is not new. In fact it occurs in nature in the form of caustics. These types of phenomena are generally referred to as Shadowgraphs but reveal only limited information of the underlying physical processes (Settles 2001). More sophisticated techniques to visualizing and photographing gas and fluid flows, refractive solids, and shock waves were developed in the 1940s (Schardin 1942). Some of the phenomena that were depicted for the first time include the shock waves created by jets breaking the sound barrier and bullets flying through the air, or the heat emerging from our bodies. As illustrated in Fig. 2 (left, red lines), traditional Schlieren setups require collimated illumination, which is then optically disturbed by changes in the refractive index of a medium. A lens deflects all light rays so that the “regular” or unrefracted rays are focused on one specific point (usually the center) of some plane. The “irregular” or refracted rays intersect that plane at different points, which are determined by the angle and magnitude of the refraction. Optical filters such as knife edges or color wheels can be mounted in that plane to encode these properties in color or intensity changes. Further

G. Wetzstein (✉) · R. Raskar  
MIT Media Lab, Cambridge, MA, USA  
e-mail: gordonw@media.mit.edu

W. Heidrich  
University of British Columbia, Vancouver, BC, Canada



**Fig. 1** Light field probes—when included into the background of a scene—allow otherwise invisible optical properties to be photographed. In this example, the probe contains a classic Rainbow Schlieren filter that codes the angles and magnitudes of refractive events in hue and saturation

light propagation optically transforms the rays back to their “normal” distribution and an image can be recorded with a camera. Each pixel on the sensor is focused on a specific point in the refractive medium (Fig. 2, left, blue lines), which allows an image to be formed along with intensity or color changes caused by the refraction. Although recent improvements have made traditional Schlieren setups more practical (Settles 2010), fundamentally, these approaches require precise calibration and high-quality optical elements that are at least as big as the observed objects. Therefore, these systems are often bulky, expensive, and mostly constrained to laboratory environments.

With the increase of computational power, background oriented Schlieren imaging (BOS) (Dalziel et al. 2000) was invented to overcome the difficulties of traditional Schlieren photography. In BOS, a digital camera observes a planar high-frequency background through a refractive medium. Optical flow algorithms are used to compute a per-pixel deflection vector with respect to an undistorted reference

background. This type of optical flow estimation requires the background to be diffuse or photo-consistent.

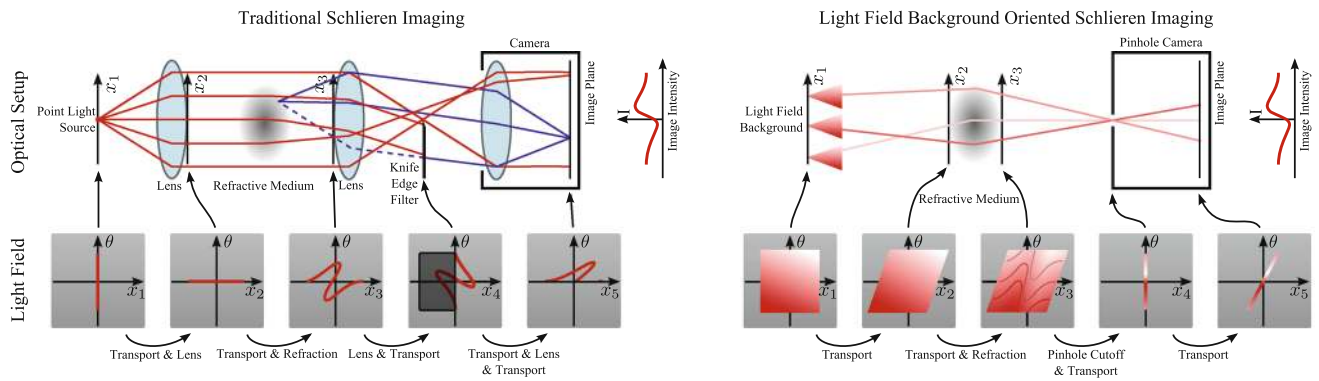
In this paper, we introduce light field background oriented Schlieren photography (LFBOS), which also employs a background probe, but rather than coding only two dimensions, we can encode up to four dimensions of the light field: spatial and angular variation (Wetzstein et al. 2011). Figure 3 shows a comparison of the most important characteristics of LFBOS compared with BOS and traditional Schlieren imaging. The asterisk indicates that although BOS backgrounds have a high spatial resolution and no angular variation, these patterns are usually placed at a large distance to the object so that they effectively become angular-only probes. This increases the size of the setup and often results in focus discrepancies between background and object. Our probes have a small form factor and do not suffer from focus mismatches.

We also demonstrate that LFBOS allows for quantitative measurements. In particular, we aim for a single camera, single image method more similar in spirit to photometric stereo (Woodham 1980), and especially to single-image variants using colored light sources (Hernandez et al. 2007). For this purpose, we propose to reconstruct transparent surfaces from the observed distortion of high-dimensional reference light field probes (Wetzstein et al. 2011). These probes can encode the 2D spatial and the 2D angular domain on their surface; possible implementations include lenslet arrays, parallax-barriers, or holograms. The distortion of a light field emitted by such a probe allows us to simultaneously reconstruct the normals and a sparse set of absolute 3D points representing either a single refractive boundary surface or a thin refractive solid.

We demonstrate successful acquisitions of refractions in solids, fluids, and gases. Specifically, we make the following contributions:

- We introduce the concept of computational light field probes for recording and processing new kinds of visual information with off-the-shelf cameras.
- We present a new type of Background Oriented Schlieren imaging that is portable, alleviates the problem of focus discrepancies, and allows spatial and angular information to be coded.
- Facilitated by coded high-dimensional light field probes, we demonstrate a single image surface reconstruction technique for refractive objects.
- We build a prototype light field background oriented Schlieren system and demonstrate its feasibility with a variety of qualitative and quantitative experiments.

Furthermore, we discuss application-specific optimality criteria for designing light field probes, compare our work with previous approaches, discuss limitations of our technique, and present a variety of application scenarios.



**Fig. 2** Illustration of optical setups and light field propagation for traditional Schlieren imaging (left) and light field background oriented Schlieren photography (right). Optical paths for forward light propagation to the camera are red, whereas backward propagation paths from

a camera pixel are blue. Please note that the 4D light field probe illustrated on the right only codes the two angular dimensions in this case (Color figure online)

	Trad. Schlieren	BOS	LFBOS
Coded Dimensions	2D	2D	4D
Setup Complexity	moderate-difficult	easy	easy
Angular Resolution	very high	N/A*	high
Spatial Resolution	N/A	very high*	moderate
Processing	low	very high	low
Cost	high	low	low
Form Factor	usually bulky	usually bulky*	small

**Fig. 3** Overview of light field background oriented Schlieren photography compared to traditional and background oriented Schlieren photography (\*see text)

**2 Related Work**

*Fluid imaging* is a wide and active area of research. Generally, approaches to measure fluid flows can be categorized into optical and non-optical methods. Non-optical methods make velocity fields observable by inducing particles, dye, or smoke; alternatively, the surface shear, pressure forces, or thermal reactions can be measured on contact surfaces that are coated with special chemicals. An extensive overview of fluid imaging techniques can be found in the book by Merzkirch (1987).

*Schlieren and phase imaging* are non-intrusive, optical approaches to visualize and quantify refraction in transparent media, such as fluids. These techniques have been developed in the fluid mechanics community over centuries (Settles 2001). Approaches to phase-contrast microscopy (Murphy 2001), such as Zernike phase contrast and differential interference contrast (DIC), also encode refractions caused by transparent specimen in changes of intensity and color. Traditional and background-oriented Schlieren (Dalziel et al. 2000) techniques are successful in coding two-dimensional light ray deflections with a high precision (Howes 1984). Light field probes, as employed in this paper, can encode up to four dimensions. We show that

such high-dimensional coding techniques allow for visualizing refractive media and reconstructing both normals and some positions of thin refractive surfaces independently. As opposed to phase-contrast methods, light field probes do not require coherent illumination and are fabricated from off-the-shelf parts.

*Light field displays* have been used for a variety of different applications, predominantly 3D image display (Okano et al. 1999; Wetzstein et al. 2012). The work that is probably closest to ours is the light field microscope (Levoy et al. 2009). This device has an integrated light field illumination mechanism that can produce exotic lighting effects on the specimen. The light field microscope uses 4D illumination for reflective, microscopic objects; LFBOS on the other hand optically codes the background behind a transparent medium to visualize and quantify refractions caused by fluids in macroscopic environments.

*Transparent and specular object reconstruction* has recently gained a lot of traction (Ihrke et al. 2010). Kutulakos and Steger (2005) analyze the space of these reconstructions based on acquisition setup and number of refractive events in the optical path of light rays. Generally, refractive object capture and reconstruction can be performed using a single camera but multiple images or, alternatively, using multiple cameras. Ben-Ezra and Nayar (2003) reconstruct smooth, parameterized refractive objects from the distortions of a diffuse background in an image sequence from a fixed viewpoint. Optical flow can be formulated to account for refraction (Agarwal et al. 2004; Shimizu and Okutomi 2006) or reflection (Gao and Ahuja 2006). Miyazaki and Ikeuchi (2005) and Huynh et al. (2010) exploit the polarization of refracted light to estimate transparent surfaces. A tomographic reconstruction of transparent solids from multiple images was proposed by Trifonov et al. (2006). Ihrke et al. (2005) compute the shape of flowing water by dyeing it with fluorescent chemicals. Range scanning can be used for the acquisition of refractive solids, if they are

immersed in a fluorescent liquid (Hullin et al. 2008). Morris and Kutulakos (2007) show that the surface of complex refractive objects can be reconstructed from multiple photographs with changing illumination. Furthermore, specular objects can be acquired using shape from distortion (Tarini et al. 2005; Bonfort et al. 2006). Multiple cameras have been used for dynamic refractive stereo (Morris and Kutulakos 2005) and for the reconstruction of smooth gas flows (Atcheson et al. 2008). As opposed to all of these approaches, the reconstruction method presented in Sect. 6 of this paper only requires a single image.

Alternative *single image reconstruction* techniques include the seminal work by Murase (1990), where a wavy water surface is reconstructed by observing the distortions of a diffuse probe under water with an orthographic camera. Zhang and Cox (1994) also reconstruct a water surface with an orthographic camera by placing a big lens and a 2D screen at its focal length in the water. This allows the surface gradients to be measured, which can subsequently be integrated to compute the surface shape. For both approaches the mean water level needs to be known. Savarese and Perona (2002) present an analysis of single image reconstruction of smooth mirroring objects using shape from distortion. Compared to these techniques, the approach presented in Section 6 also assumes that there is only a single refractive or reflective event; however, no constraints are placed on the camera setup. Furthermore, we demonstrate how to reconstruct both surface points and normals simultaneously from a captured photograph.

### 3 Theory

#### 3.1 Image Formation

The propagation of light in inhomogeneous refractive media is governed by the ray equation of geometric optics (Born and Wolf 1999):

$$\frac{\partial}{\partial s} \left( n \frac{\partial \mathbf{x}}{\partial s} \right) = \nabla n, \quad (1)$$

where  $\mathbf{x}$  is the position on a trajectory in space,  $\partial s$  is the differential path length along the trajectory, and  $n$  is the (spatially varying) refractive index field. The wavelength dependency of  $n$  is disregarded in this model. Equation 1 can be formulated as a coupled system of first-order ODEs (Ihrke et al. 2007; Atcheson et al. 2008):

$$n \frac{\partial \mathbf{x}}{\partial s} = \mathbf{d}, \quad \frac{\partial \mathbf{d}}{\partial s} = \nabla n, \quad (2)$$

with  $\mathbf{d}$  being the local direction of propagation. Integrating Eq. 2 leads to an expression for the global directional deformation within a refractive object (Atcheson et al. 2008):

$$\mathbf{d}^{out} = \mathbf{d}^{in} + \int_c \nabla n ds. \quad (3)$$

In our setup, the refractive index field is observed against a known light field background, as shown in Fig. 2 (right). In order to understand the distortion of this known 4D probe, Eq. 2 can be solved numerically, for instance with forward Euler schemes (Ihrke et al. 2007). However, in our case it is more intuitive to trace the light field on the sensor plane back to the probe. Considering a pinhole camera, as shown in Fig. 2 (right), the light field probe  $l(x, y, \theta, \phi)$  is sampled in the following way:

$$i(\mathbf{x}_p) = l \left( \zeta(\mathbf{x}_p, \mathbf{d}_p), \varphi \left( \mathbf{d}_p + \int_c \nabla n ds \right) \right). \quad (4)$$


Here,  $i$  is the sensor image and  $\mathbf{d}_p$  the normalized direction from a pixel  $\mathbf{x}_p = (x_p^x, x_p^y)$  to the camera pinhole. The function  $\varphi(\mathbf{d})$  maps a direction  $\mathbf{d} = (d^x, d^y, d^z)$  to the angular parameterization of the light field, i.e.  $\varphi_{\theta, \phi}(\mathbf{d}) = (\tan^{-1}(d^x/d^z), \tan^{-1}(d^y/d^z))$ . The position and direction of a light ray incident on a pixel can be mapped to a position on the probe by the function  $\zeta$ ; this depends on the distances between camera and object, object and probe, and the ray displacement within the object. If the ray displacement is negligible, the refractive medium can be approximated as a thin element that causes a single refractive event, such as a thin lens. We explore reconstructions of such media in Sect. 6.

Although pinhole cameras sample discrete rays of the refracted light field, in practice cameras usually have a finite aperture implying that each pixel integrates over a small range of light field directions on the probe. We assume that the distance between the light field probe and the camera is large compared to the size of the aperture and that the distance between the probe and the refractive object is small, which yields a good approximation of a pinhole camera.

#### 3.2 Designing Light Field Probes

The goal of LFBOS is to encode the directions and locations of a 4D probe with color and intensity variations so that the former parameters can be inferred from the colors in a photograph. The immediate question that arises is how one can design a light field probe that encodes positions and directions in a meaningful way. To answer this question, we need to consider a number of application-specific parameters that are discussed in the following.

Although a pinhole camera is a reasonable approximation under the above mentioned assumptions, finite pixel sizes and small apertures along with strong distortions of the wavefront caused by refraction often amplify the integration area of each camera pixel in the space of the light field probe. In order to compensate for this effect, the distribution of *color tones*



**Fig. 4** Light field probe design. A lenslet array is mounted at its focal distance to a color-coded background. The same sub-pixels under each lenslet code the color and intensity of one emitted light direction (*left*). Assigning a unique color to each lenslet array allows for the location on the lenslet plane to be coded (*right*)

and intensities in the probe should be *smooth* in the 4D spatio-directional domain. This also implies graceful degradation of the captured colors in case the pinhole assumption breaks down.

If the absorption of light within the refractive medium is not negligible, the color distributions should ideally be independent of intensity changes. This can, for instance, be implemented by encoding the desired parameters only in hue and saturation, i.e. constant values in HSV color space. Doing so also ensures resilience to vignetting and other possible intensity changes caused by the lenslets of the probe.

Positions and directions of a background light field can either be encoded in the *absolute* 4D reference frame of the probe or *relative* to a fixed camera position. This is not necessary for orthographic cameras, but to *compensate for the perspective* of a non-orthographic camera.


Figure 4 illustrates two 2D light field probes implemented with lenslet arrays. While a periodic pattern of a red color gradient under each lenslet allows for the emitted light directions to be coded (left, red), additionally coding the area under each lenslet allows for the emitted light ray location on the probe to be coded (right, green).

### 4 Experimental Prototype

We have implemented prototypes of our light field probes using both lenticular sheets with cylindrical lenses and lenslet arrays with hexagonal grids of spherical lenses. These make a tradeoff between spatial and angular resolution. The specific lenslet arrays we have used in our experiments along

**Table 1** Technical specifications of the lenslet arrays used in our experiments

Lenslet type	f (in)	d (in)	fov (°)	Angular resolution [°]	
				600 dpi	2,038 dpi
MicroLens animation 10	0.11	0.1	48	0.80	0.24
MicroLens 3D 20	0.10	0.05	29	0.97	0.28
FresnelTech hexagonal 300	0.12	0.09	42	0.77	0.23
FresnelTech hexagonal 310	1.00	0.94	51	0.09	0.03



**Fig. 5** Prototype setup. The light field probe (*left*) comprises an inkjet-printed transparency that is taped on an LED light box behind a lenslet array. The probe is placed behind a refractive object and photographed by a camera (*right*)

with their spatial and angular resolution and fields-of-view are listed in Table 1. Alternative implementations of 4D probes include holograms or dynamic parallax barrier displays (e.g., Wetzstein et al. 2012); these allow for probes with a very high spatial and angular resolution.

As shown in Fig. 5, the lenslet arrays are mounted on a light box that provides a uniform background illumination. We print the probe codes at a resolution of 1,200 dpi on transparencies that are manually aligned with the lenslet arrays before mounting them on the light box. For an increased contrast, multiple transparencies can be stacked. Care must be taken with the non-linear color transformations between specified digital images, the printer gamut, and the color gamut of the camera. We define our patterns within the device-specific CMYK color gamut of the employed printer (RGB and CMYK ICC profiles of printers are usually available from the manufacturer), then print them with device-internal color mappings disabled, and record camera images in sRGB space. This process allows us to transform the captured photographs to any desired color space as a post-processing step.

### 5 Visualizing Refraction of Transparent Media

#### 5.1 Capturing Refractive Solids

Figures 1, 6, 7, 8, and 9 demonstrate a variety of different intensity and color filters encoded in the probes; these reveal angular variation caused by refractive solids. For the experi-

ments discussed below, the encoded light field is pre-distorted to account for the perspective of the camera. An undistorted view of the probe from the camera position therefore shows the center of a specific angular probe code.

*Intensity gradients* correspond to knife edge filters in traditional Schlieren imaging. These filters usually sacrifice half of the background illumination by coding undistorted rays in gray; the magnitudes of light deflection in a particular direction are coded in increasing or decreasing intensity. An example of one-dimensional gradients is shown in Figs. 6 and 7. Here, lenticular sheets are used as probes where the gradient under each cylindrical lens is equal except for a varying lens-to-lens pitch that corrects for the perspective of the camera. As an alternative to 1D gradients, a circular gradient under each circular lens of a hexagonal lenslet array can encode the magnitudes of deflected rays as seen in Fig. 8 (center right).

*Annular bright field and dark field filters* are inspired by microscopy (Murphy 2001; Levoy et al. 2009) where similar illumination patterns can be used to illuminate reflective or refractive specimen. The annular bright field probe code is a uniform circular pattern with a small black outline that does not effect undeflected rays or those that underwent only small amounts of refraction. Strong refractions, as seen on the object boundaries of the unicorn in Fig. 8 (second from left), are completely blocked. The inverse of the annular bright field is the annular dark field, which only allows rays with strong refractions to reach the camera (see Fig. 8, third from left).

*Directional cutoffs* can be used to completely block ray deflections into a certain direction. An example of this filter type can be seen in the vertical-only and horizontal-only cutoff shown in Fig. 8 (the two rightmost images).

*Color filters* are popular in traditional Schlieren setups. Usually, these approaches are called Rainbow Schlieren (Howes 1984) and allow the magnitude and the angle of ray deflection to be encoded in different color gradients. The HSV color wheel is a popular choice for the filter. Undeflected rays are coded in white, whereas the deflection angle is coded in hue and the magnitude in saturation of the color wheel. Examples of these filters are shown in Figs. 1 and 9.

## 5.2 Capturing Liquids

We also show that the proposed technique is able to capture refractive liquids. For the experiment in Fig. 10, for instance, we record the interactions of liquids with different refractive indices. For this experiment, we mix water and clear cornstarch. The latter has a higher refractive index resulting

in stronger light ray deflections than the water. Figure 11 shows air bubbles rising in clear corn starch. The difference in refractive index results in ray deflections, which are optically coded in varying color tone and saturation.

## 5.3 Gas Capture


Capturing transparent, refractive gases is particularly challenging. The subtle changes in refractive index caused by temperature changes in air, for instance, require a very high angular resolution of the employed probes. We successfully capture the dynamic refractive index field above the flame of a butane burner. Figure 12 shows several frames of a captured video (contrast enhanced by 50%). The complex shape of the hot air plume is clearly visible. For this experiment, we use an angular probe encoding a grayscale intensity gradient. As illustrated in Fig. 12 (right), the probe is constructed from a single lens with a long focal length; the intensity gradient is located at its focal distance and displayed by a conventional computer monitor. The size of this probe is significantly larger than that of probes used for all other experiments in this paper. Yet, it is necessary to provide a high-enough angular resolution to visualize the small changes in the refractive index of this particular experiment. High-resolution printing processes, such as offset printing, may facilitate such experiments with smaller lenslet arrays and shorter focal lengths in the future.

## 6 Refractive Shape From Light Field Distortion

### 6.1 Coding Light Field Illumination


For the purpose of single-shot transparent object reconstruction, the color and intensity codes emitted by a light field probe need to satisfy two important criteria. First, the patterns are required to uniquely encode position and angle on the probe surface, so that a camera measures this information in a single image. Second, in order to account for slight miscalibrations of the probe prototype, the color codes should be smooth in the 4D spatio-angular domain. We restrict our prototype to readily available hardware, as illustrated in Fig. 5, and limit the feasible colors and intensities to the combined printer and camera gamut and dynamic range.

The most intuitive coding scheme satisfying the above requirements are color gradients. In our implementation, we use red, blue, and green gradients to code the 2D directions and the 1D vertical position, respectively. As demonstrated in Sect. 6.3, the missing second spatial dimension can be recovered through geometric constraints in post-processing. This encoding is illustrated for a 1D case in Fig. 13. Here, the incident angle is coded in a shade of red and the position on the probe surface is coded in green.




**Fig. 6** A lenticular light field probe encoding 1D directional light variation of the refraction created by a lens with an intensity gradient. This type of filter resembles the knife edge filter of traditional Schlieren photography in the specific direction. *Left* probe encoding a horizontal

intensity gradient without any object; *center left* convex lens with uniform background illumination; *center right* lens in front of the probe; *right* lens in front of the rotated probe




**Fig. 7** A plate in front of a uniform background (*left*), and a light field probe that encodes directional variation caused by refraction with a horizontal (*center left*) and a vertical (*center right*) intensity gradient. The magnifications (*right*) and the structure on the plate show how other-

wise invisible information is revealed with our probes. The color bar, mapping colors to magnitudes of refraction, is computed from the field of view of the lenticulars and a calibration gradient that is cropped from the photographs



**Fig. 8** A variety of directional filters can be encoded in our probes. From *left* uniform background, annular bright field, annular dark field, circular intensity gradient, horizontal cutoff, and vertical cutoff




**Fig. 9** Additional results for refractive solids with an angular hue-saturation probe. The magnifications show some of the objects under uniform illumination






**Fig. 10** The classic Rainbow Schlieren filter encoded in our light field probe can visualize magnitudes and angles of refractions in complex media such as this mix of clear corn syrup and water




**Fig. 11** An angular hue-saturation probe behind a bottle filled with clear corn starch. Light rays are refracted by the bottle, its content, and rising air bubbles



**Fig. 12** Gas capture. An angular intensity probe is placed behind an open flame (right). The hot air plume refracts light, which is coded in intensity changes by the probe (left and center). Changes in the

refractive index field are very subtle; a single, large lens with a long focal length is necessary to provide sufficient angular resolution for this experiment



**Fig. 13** Schematic showing how both position and incident angle of a refracted ray are color coded by a light field probe

This simple, yet effective coding scheme allows both angle and position of light rays to be encoded in observed colors and intensities. Without refraction in the optical path, the measured colors at each pixel of a calibrated camera correspond to the information predicted by the calibration, but in the presence of refraction these differ. In the following subsections we show how to reconstruct refractive surfaces from such measurements. The employed color codes ignore

the wavelength-dependency of refraction as well as attenuation and scattering caused by the medium.

### 6.2 Reconstructing Surface Normals

The normal of each surface point imaged by a camera pixel can be computed using Snell's law<sup>1</sup>:  $n_1 \sin \theta_{in} = n_2 \sin \theta_{out}$ . In our application, we seek the unknown normals given the incoming normalized rays  $v_{in}$ , which are known from camera calibration, and the refracted ray directions  $v_{out}$ , which are extracted from the imaged probe color (Fig. 13). The absolute angles  $\theta_{in}$  and  $\theta_{out}$  are unknown, but we can compute the difference  $\theta_d$  between the two as  $\cos \theta_d = v_{in} \cdot v_{out}$ . For known refractive indices of the two media  $n_1$  and  $n_2$ , the angle between incoming ray and surface normal is then given as

<sup>1</sup> The connection between Snell's law and the ray equation of geometric optics (Eq. 1) is well-known in physics, see e.g. [http://en.wikibooks.org/wiki/Introduction\\_to\\_Mathematical\\_Physics/Electromagnetism/Optics,\\_particular\\_case\\_of\\_electromagnetism](http://en.wikibooks.org/wiki/Introduction_to_Mathematical_Physics/Electromagnetism/Optics,_particular_case_of_electromagnetism).

$$\theta_{in} = \tan^{-1} \left( \frac{n_2 \sin \theta_d}{n_2 \cos \theta_d - n_1} \right). \quad (5)$$

Therefore, the surface normal  $\mathbf{n}$  can be computed independently for each camera pixel by rotating  $\mathbf{v}_{in}$  by the angle  $\theta_{in}$ . The rotation is performed on the plane spanned by  $\mathbf{v}_{in}$  and  $\mathbf{v}_{out}$ , so

$$\mathbf{n} = \mathbf{R}(\theta_{in}, \mathbf{v}_{in} \times \mathbf{v}_{out}) (-\mathbf{v}_{in}), \quad (6)$$

where  $\mathbf{R}(\theta, \mathbf{v})$  is a rotation matrix defined by angle  $\theta$  around an axis  $\mathbf{v}$ .

### 6.3 Point Cloud Estimation

In order to triangulate absolute 3D surface points for each camera pixel, we need to determine the intersection of the lines  $\mathbf{c} + t\mathbf{v}_{in}$  and  $\mathbf{p} + s\mathbf{v}_{out}$ . The camera position  $\mathbf{c}$  as well as the unrefracted ray directions  $\mathbf{v}_{in}$  are known from camera calibration and uniquely define a line in 3D space. The direction  $\mathbf{v}_{out}$  is estimated from the observed colors of the light field probe refracted by an object, however, only a single spatial coordinate is coded by the probe color, i.e.  $p^y$ . Nevertheless, the intersection problem for the two lines results in a linear system with three equations and three unknowns  $p^x$ ,  $s$ , and  $t$  because the origin of the coordinate system is defined on the plane of the probe, i.e.  $p^z = 0$ . Therefore, we can uniquely triangulate a 3D point per camera pixel as

$$t = \frac{1}{v_{in}^y - \frac{v_{in}^z v_{out}^y}{v_{out}^z}} \left( p^y + \frac{c^z v_{out}^y}{v_{out}^z} - c^y \right). \quad (7)$$

The triangulated positions are only numerically robust when significant refraction occurs along a ray; otherwise  $\mathbf{v}_{in}$  and  $\mathbf{v}_{out}$  are co-linear. At the same time, all measured ray directions  $\mathbf{v}_{out}$  will be noisy due to camera noise and possible color non-linearities of a fabricated probe. Therefore, we can only hope to robustly estimate a sparse set of 3D points from such measurements at camera pixels that observe a strong amount of refraction. The noise sensitivity of triangulated points is illustrated for a synthetic example in Fig. 14.

### 6.4 Surface Estimation From Normals and Points

While a normal field can be efficiently integrated to reconstruct surfaces (see e.g., Agrawal et al. 2006), including an additional set of sparse 3D control points can remove ambiguities in these integration schemes (Horowitz and Kiryati 2004; Nehab et al. 2005). For all of our reconstructions, we employ the integration method proposed by Ng et al. (2010), which uses an optimization with kernel basis functions.

We show synthetic results in Fig. 14. Here, a sinusoidal function acts as the original surface with a refractive index corresponding to water; 3D positions and normals of the original surface are shown in the left column. We simulated pho-

tographs of an orthogonal camera that show the surface in front of a light field probe with the color coding scheme discussed in Sect. 6.1 along with estimated normals, triangulated control points, and final reconstructions. While the extracted normals are relatively resilient to an increasing amount of camera noise, the triangulated positions quickly become less reliable. We mask out triangulated points that correspond to small angles between incoming and refracted rays for each pixel; the masks are shown in the insets of the second row.

### 6.5 Experimental Results



The prototype for the quantitative experiments discussed in this section (see Figs. 5, 13) is composed of a light box, two stacked transparencies, a lenslet array, and a camera. The light box is LED-based, as opposed to fluorescent, in order to maintain consistent lighting throughout the capture process even when using a short exposure time, such as in video. The lenslet sheet is a FresnelTech hexagonal lenslet array with a focal length of 0.12" and a lenslet diameter of 0.09". The transparencies are printed with an Epson stylus photo 2200 printer at 1,440 dpi, which, in combination with the lenslets, results in a theoretical angular resolution of 0.32°. This printer has six ink-based primaries; for improved contrast we stack two transparencies on top of each other. For still photographs we use a Canon D5 Mark II and for the videos a Prosilica EC1350C camera.

Intrinsic and extrinsic camera parameters are estimated in a pre-processing step (Atcheson et al. 2010). The gamma curves of the printer are also estimated as a pre-processing step and compensated in the measurements.

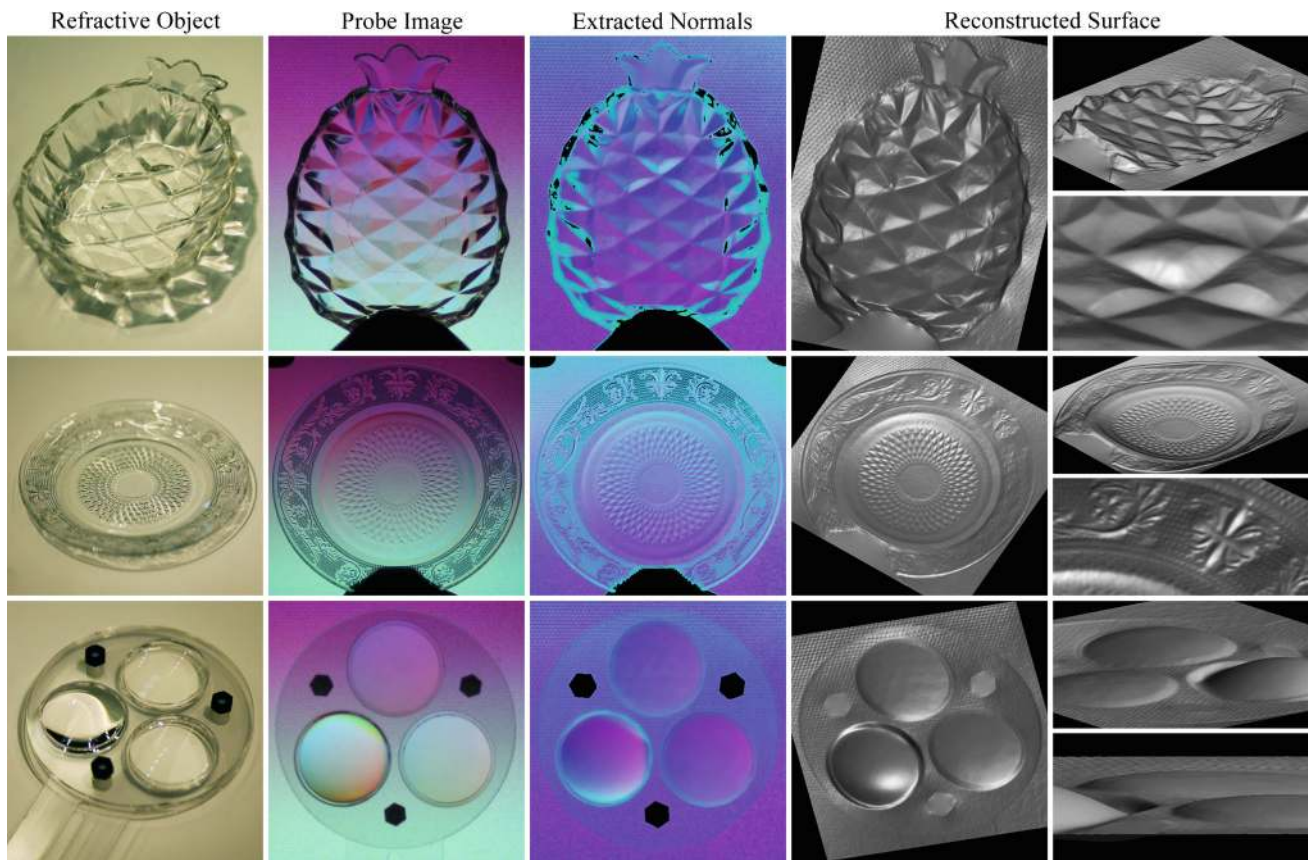
Reconstructions of water surfaces are shown in Fig. 15. Here, we positioned the probe underneath a rectangular water tank and filmed the scene from above (Fig. 15, rows 1 and 3). Secondary refractions from the glass tank bottom are negligible in this case. The results show a water drop falling into the tank in rows 1 and 2; rows 3 and 4 depict water being poured into the tank. Some high-frequency noise is visible in the reconstruction, which is due to the printer half-toning patterns on the transparencies that become visible as noise on the probe when the camera is focused on it. Alternative printing technologies, such as light valve technology ([www.bowhaus.com](http://www.bowhaus.com)), could alleviate this problem.

Figure 16 shows reconstructions of three thin solid objects from a single photograph each. Although two refractive events occur for each camera ray, one at the air-glass interface toward the camera and another one at the glass-air boundary on the other side, the objects are thin enough that ray displacements within the glass are negligible. This is a common assumption for thin lenses. The reconstructed normals (Fig. 16, column 3) for these examples therefore show the difference between front and back normal of the surface; for the plate and the pineapple, the front side is flat and parallel

**Fig. 14** Synthetic results for a refractive sinusoidal object. Normals and positions are shown for the original object (*left column*), and for reconstructions (*other columns*) from simulated camera images with an increasing amount of noise (*top row*)



**Fig. 15** Camera images and reconstructed geometry of dynamic water surfaces. The *upper rows* shows a drop falling into the water, whereas the *lower rows* depict water being poured into the tank



**Fig. 16** Three thin refractive objects under room illumination (*left column*) and in front of a light field probe (*center left column*). The distorted colors of the probe allow us to estimate refractive surface normals from

a single image (*center row*), which can be integrated to reconstruct thin shapes that approximate the geometry of transparent, refractive solids (*right*)

to the fine details on the rear side. The reconstructed surfaces (Fig. 16, right) only contain a flat triangle mesh and corresponding normals.

## 6.6 Evaluation


A quantitative evaluation of the proposed reconstruction algorithm with respect to camera noise and refractive index mismatches is shown in Fig. 17. In this experiment, we simulate the acquisition and reconstruction of a 1D parabolic surface. An orthographic camera observes the scene from above with a light field probe illuminating it from the bottom. The surface represents the boundary between two media, the upper one is air and the lower one has a refractive index of  $n = 1.5$ . We add zero-mean Gaussian noise to the simulated sensor measurements and evaluate reconstruction quality for different refractive index mismatches. Surface gradients (Fig. 17, center) are directly computed from the noisy sensor measurements and subsequently integrated to yield the actual surfaces (Fig. 17, left).

Based on these experiments, we can see that a mismatch in the refractive index results in a vertical shear of the gra-

dients (Fig. 17, center, purple line), which corresponds to low frequency distortions of the actual surface (Fig. 17, left, purple line). The mean squared error (MSE) between original surface and reconstruction is particularly high when the assumed refractive index is lower than that of the medium (Fig. 17, top right, purple line). Furthermore, there is an approximately linear relationship between sensor noise and the noise observed in both reconstructed gradients and surfaces (Fig. 17, right). The mean squared error plots on the right of Fig. 17 are averaged over 500 experiments, each exhibiting different random noise.

## 7 Comparison to BOS

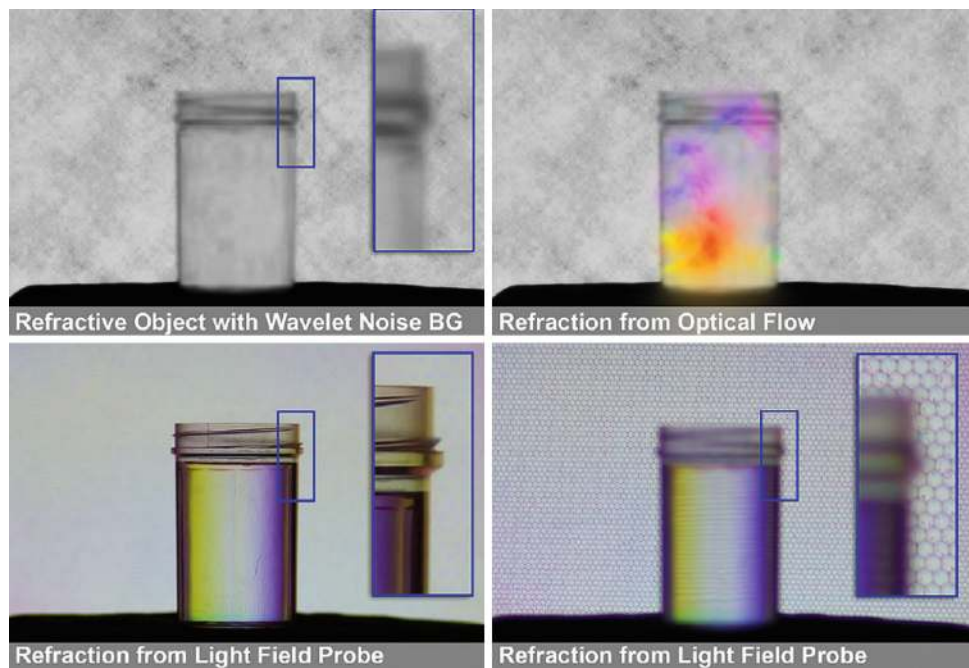
Background Oriented Schlieren setups usually require a high-frequency background to be located at a large distance to the object. In this way, the per-pixel displacement vectors estimated by optical flow are proportional to the angular ray deflections, which is related to the refractive index gradient (Eq. 3). The form factor of these setups is therefore usually large. Furthermore, the camera needs to be focused on the background pattern so that its distortion can be tracked by



**Fig. 17** Evaluation of reconstruction with respect to noise and refractive index mismatch. A 1D parabolic surface (*left, dotted red*) is simulated to be captured with a light field probe and reconstructed with different amounts of camera noise and mismatches in the refractive index of the medium (*left*). While noise results in high frequency artifacts,

a mismatch in the refractive index causes low frequency distortions. We show the mean squared error of surfaces (*top right*) and gradients (*bottom right*) for an increasing amount of sensor noise (Color figure online)

**Fig. 18** Light field background oriented Schlieren photography compared to a failure case of background oriented Schlieren imaging. Optical flow algorithms in BOS require the background to be focused, which places the object out-of-focus (*upper left*). In this case, the refractions are so strong that they blur out the background pattern and therefore prevent a reliable optical flow estimation (*upper right*). LFBOS works well for in-focus (*lower left*) and out-of-focus settings (*lower right*)



the optical flow algorithm. The total amount of light in BOS setups is often limited, which is why cameras typically need to use a large aperture for capture. Unfortunately, this places the object of interest out-of-focus as seen in Fig. 18 (upper left).


Strong refractions, for instance caused by fluids or solids, often lead to extreme distortions of the background pattern. These distortions may prevent a reliable optical flow estimation as shown in Fig. 18 (upper right). Although an optical flow algorithm for refractive objects has been proposed (Agarwal et al. 2004), this requires many frames of a video sequence to be analyzed and is not practical for dynamic media such as fluids.

In comparison, our approach requires only a single image and the light field background can be placed at close proxim-

ity to the object, which alleviates the focus mismatch problem (Fig. 18, lower left). Furthermore, if the light field probe encodes smooth gradients, as discussed in Sect. 3.2, even a defocused refractive object will reveal the mean of colors and intensities in the integration manifold of the 4D probe space to a camera pixel (Fig. 18, lower right).

## 8 Discussion

In summary, we have presented a new approach to capturing refractive phenomena using light field probes. Our method presents a portable and inexpensive alternative to traditional and BOS imaging; it works well with strong refractions, which is often not the case for BOS, and also alleviates the





**Fig. 19** Failure case: the field-of-view of the lenslet array is too narrow to properly encode the strong refractions near the sides of the glass. To overcome this, the lenslets should be chosen according to the amount of refraction in the scene

focus mismatch between background and objects of interest. Inspired by Schlieren imaging and microscopy, we have shown how a variety of different filters for visualizing refractive events can be encoded in our light field probes and recorded with off-the-shelf cameras. We also show how the observed light field distortions facilitate reconstruction of surface normals and triangulation of a sparse set of control points from a single photograph. While the normals are relatively resilient to sensor noise and allow for high-quality reconstructions, the triangulated control points are very sensitive to noise, but allow low-frequency ambiguities of the surface normals to be corrected.

### 8.1 Limitations

As in most Schlieren approaches, the size of the refractive volume is limited by the size of the employed optical elements, in our case the light field probe. An exception are traditional Schlieren setups that use the sun or other distant point lights as the source of the required collimated illumination and natural-background oriented Schlieren techniques (Hargather and Settles 2009), which also work outdoors.

**Fig. 20** Light field probes for coded Schlieren illumination in microscopy. Microscopic glass shards, each being about  $100\ \mu\text{m}$  in diameter, are imaged with a cellphone through a magnifying glass (left). The Schlieren illumination (right) is provided by a light field probe consisting of a single lenslet and a 2D intensity gradient mounted at its focal length. The contrast of the glass shards is significantly improved as compared to a uniform illumination (center)



Currently, the proposed approach is mostly limited by the employed off-the-shelf hardware. All transparencies in the presented experiments were printed using standard desktop printers, usually with 1,200 dpi. Color non-linearities and color cross-talk introduced by the printing process need to be calibrated. Alternative printing processes that would allow for higher resolution, improve contrast, and a larger color gamuts, include light valve technology ([www.bowhaus.com](http://www.bowhaus.com)) and professional offset printing.

All of our prototypes are implemented with lenslet arrays or lenticulars, which trade spatial and angular resolution. A very high spatial and angular resolution can be achieved with alternative technologies, such as holograms or compressive light field displays (Wetzstein et al. 2012). When the camera is focused on a lenslet-based probe, the space between individual lenses usually appears darker—this problem could also be overcome with alternative probe implementations. As the fields-of-view of the lenslets in our current prototypes are defined by the manufacturing process, refractions that exceed the field-of-view cannot be coded reliably with a single shot. An example of such a failure case is shown in Fig. 19. The same problem often occurs in parallax-barrier or lenslet-based auto-stereoscopic displays. For our application, the lenslets for a specific experiment should be chosen in accordance with the expected amount of refraction.

Disregarding the prototype, our approach is fundamentally limited by the light field coding scheme and the reconstruction algorithm. Although the employed color codes are optimized for single image capture and reconstruction, attenuation, emission, and scattering within the medium as well as wavelength-dependency of refraction are assumed to be negligible. Alternative, dynamic codes could overcome these limitations at the cost of requiring multiple photographs. The proposed reconstruction algorithm requires the refractive index of the medium to be known and restricts light rays to refract only once in the scene. In combination with

advanced coding schemes, novel algorithms could overcome these limitations as well.

## 8.2 Future Work

In the future, we would like to experiment with alternative technologies for fabricating light field probes, such as holograms, and test more sophisticated light field coding schemes. Applying temporal multiplexing with dynamic probes could lift current limitations; multi-spectral displays and cameras could improve the amount of coded information as well. We would like to explicitly separate attenuation and refraction caused by the medium and test our approach with multi-camera, multi-probe configurations. Schlieren illumination for microscopy is a promising avenue of future research. In Fig. 20, we show preliminary results of a cellphone-based microscope (Arpa et al. 2012) that uses a simple angular gradient as the background illumination—the contrast of refractive glass shards is significantly improved. Finally, 4D light field probes could be useful for capturing a variety of other phenomena, including BTDFs and BRDFs, and for de-scattering or separating local and global illumination.

**Acknowledgments** Gordon Wetzstein was supported by an NSERC Postdoctoral Fellowship. Ramesh Raskar was supported by an Alfred P. Sloan Research Fellowship and a DARPA Young Faculty Award.

## References

- Agarwal, S., Mallick, S.P., Kriegman, D., & Belongie, S. (2004). On refractive optical flow. In *Proceeding of ECCV* (pp. 483–494).
- Agrawal, A., Raskar, R., & Chellappa, R. (2006). What is the range of surface reconstructions from a gradient field? In *Proceedings of ECCV* (pp. 483–494).
- Arpa, A., Wetzstein, G., Lanman, D., & Raskar, R. (2012). Single lens off-chip cellphone microscopy. In *Proceedings of IEEE PROCAMS*.
- Atcheson, B., Heide, F., & Heidrich, W. (2010). CALTag: High precision fiducial markers for camera calibration. In *Proceedings of VMV*.
- Atcheson, B., Ihrke, I., Heidrich, W., Tevs, A., Bradley, D., Magnor, M., et al. (2008). Time-resolved 3D capture of non-stationary gas flows. *ACM Transactions on Graphics (Siggraph Asia)*, 27, 132.
- Ben-Ezra, M., & Nayar, S.K. (2003). What does motion reveal about transparency? In *Proceeding of ICCV* (pp. 1025–1032).
- Bonfort, T., Sturm, P., & Gargallo, P. (2006). General specular surface triangulation. In *Proceedings of ACCV* (pp. 872–881).
- Born, M., & Wolf, E. (1999). *Principles of Optics* (7th ed.). Cambridge: Cambridge University Press.
- Dalziel, S., Hughes, G., & Sutherland, B. (2000). Whole-field density measurements by synthetic schlieren. *Experiments in Fluids*, 28, 322–335.
- Gao, C., & Ahuja, N. (2006). A refractive camera for acquiring stereo and super-resolution images. In *Proceeding of CVPR* (pp. 2316–2323).
- Hargather, M. J., & Settles, G. S. (2009). Natural-background-oriented Schlieren imaging. *Experiments in Fluids*, 48, 59–68.
- Hernandez, C., Vogiatzis, G., Brostow, G., Stenger, B., & Cipolla, R. (2007). Non-rigid photometric stereo with colored lights. In *Proceeding of ICCV*.
- Horowitz, I., & Kiryati, N. (2004). Depth from gradient fields and control points: Bias correction in photometric stereo. In *Proceedings of ICIVC* (pp. 681–694).
- Howes, W. L. (1984). Rainbow schlieren and its applications. *Applied Optics*, 23, 2449–2460.
- Hullin, M. B., Fuchs, M., Ihrke, I., Seidel, H. P., & Lensch, H. P. A. (2008). Fluorescent Immersion Range Scanning. *ACM Transactions on Graphics (Siggraph)*, 87(1–87), 10.
- Huynh, C. P., Robles-Kelly, A., & Hancock, E. (2010). Shape and refractive index recovery from single-view polarisation images. In *Proceeding of CVPR*.
- Ihrke, I., Goldluecke, B., & Magnor, M. (2005). Reconstructing the geometry of flowing water. In *Proceedings of ICCV* (pp. 1055–1060).
- Ihrke, I., Kutulakos, K. N., Lensch, H. P. A., Magnor, M., & Heidrich, W. (2010). Transparent and specular object reconstruction. *Computer Graphics Forum*, 29, 2400–2426.
- Ihrke, I., Ziegler, G., Tevs, A., Theobalt, C., Magnor, M., & Seidel, H.P. (2007). Eikonal rendering: Efficient light transport in refractive objects. In *Proceeding of SIGGRAPH. ACM Transactions on Graphics*, 26:59.
- Kutulakos, K.N., & Steger, E. (2005) A theory of refractive and specular 3D shape by light-path triangulation. In *Proceeding of ICCV* (1448–1455).
- Levoy, M., Zhang, Z., & McDowall, I. (2009). Recording and controlling the 4D light field in a microscope. *Journal of Microscopy*, 235, 144–162.
- Merzkirch, W. (1987). *Flow visualization*. New York: Academic Press.
- Miyazaki, D., & Ikeuchi, K. (2005). Inverse polarization raytracing: Estimating surface shapes of transparent objects. In *Proceeding of CVPR* (pp. 910–917).
- Morris, N. J. W., & Kutulakos, K. N. (2005). Dynamic refraction stereo. In *Proceedings of ICCV*.
- Morris, N. J. W., & Kutulakos, K. N. (2007). Reconstructing the surface of inhomogeneous transparent scenes by scatter trace photography. In *Proceedings of ICCV*.
- Murase, H. (1990) Surface shape reconstruction of an undulating transparent object. In *Proceedings of ICCV* (pp. 313–317).
- Murphy, D. B. (2001). *Fundamentals of light microscopy and electronic imaging*. New York: Wiley.
- Nehab, D., Rusinkiewicz, S., Davis, J., & Ramamoorthi, R. (2005). Efficiently combining positions and normals for precise 3D geometry. *ACM Transactions on Graphics (Siggraph)* 24.
- Ng, H. S., Wu, T. P., & Tang, C. K. (2010). Surface-from-gradients without discrete integrability enforcement: A Gaussian Kernel approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32, 2085–2099.
- Okano, F., Arai, J., Hoshino, H., & Yuyama, I. (1999). Three-dimensional video system based on integral photography. *Optical Engineering*, 38, 1072–1077.
- Savarese, S., & Perona, P. (2002). Local analysis for 3D reconstruction of specular surfaces—part II. In *Proceedings of ECCV* (pp. 759–774).
- Schardin, H. (1942). Die Schlierenverfahren und ihre Anwendungen. Ergebnisse der Exakten Naturwissenschaften. (English translation available as NASA TT F-12731), April 1970 (N70–25586) 20, 303–439.
- Settles, G. S. (2001). *Schlieren and shadowgraph techniques*. Cambridge: Cambridge University Press.
- Settles, G. S. (2010). Important Developments in Schlieren and Shadowgraph Visualization during the Last Decade. In *Proceedings of International Symposium on Flow Visualization* (p. 267).
- Shimizu, M., & Okutomi, M. (2006) Reflection stereo—Novel monocular stereo using a transparent plate. In *Proceeding of CRV* (pp. 1–14).

- Tarini, M., Lensch, H. P., Goesele, M., & Seidel, H. P. (2005). 3D acquisition of mirroring objects using striped patterns. *Graphical Models*, 67, 233–259.
- Trifonov, B., Bradley, D., & Heidrich, W. (2006). Tomographic reconstruction of transparent objects. In *Proceedings of EGSR* (pp. 51–60).
- Wetzstein, G., Lanman, D., Hirsch, M., & Raskar, R. (2012). Tensor displays: Compressive light field synthesis using multilayer displays with directional backlighting. *ACM Transactions on Graphics (SIGGRAPH)*, 31, 1–11.
- Wetzstein, G., Raskar, R., & Heidrich, W. (2011) Hand-held schlieren photography with light field probes. In *Proceeding of ICCP* (pp. 1–8).
- Wetzstein, G., Roodnick, D., Raskar, R., & Heidrich, W. (2011). Refractive Shape from light field distortion. In *International Conference on Computer Vision (ICCV)*.
- Woodham, R. (1980). Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1), 139–144.
- Zhang, X., & Cox, C. S. (1994). Measuring the two-dimensional structure of a wavy water surface optically: A surface gradient detector. *Experiments in Fluids*, 17, 225–237.
- Zongker, D.E., Werner, D.M., Curless, B., & Salesin, D.H. (1999). Environment matting and compositing. In *Proceeding of SIGGRAPH*. Computer Graphics (pp. 205–214).