

# Transparent Object Reconstruction via Coded Transport of Intensity

Chenguang Ma<sup>1</sup> Xing Lin<sup>1</sup> Jinli Suo<sup>1</sup> Qionghai Dai<sup>1</sup> Gordon Wetzstein<sup>2</sup>  
<sup>1</sup>Department of Automation, Tsinghua University <sup>2</sup>MIT Media Lab

## Abstract

*Capturing and understanding visual signals is one of the core interests of computer vision. Much progress has been made w.r.t. many aspects of imaging, but the reconstruction of refractive phenomena, such as turbulence, gas and heat flows, liquids, or transparent solids, has remained a challenging problem. In this paper, we derive an intuitive formulation of light transport in refractive media using light fields and the transport of intensity equation. We show how coded illumination in combination with pairs of recorded images allow for robust computational reconstruction of dynamic two and three-dimensional refractive phenomena.*

## 1. Introduction

Imaging and recovering refractive surfaces and volumes has been an active area of computer vision research for decades. Whereas some applications seek reconstructions of geometry or reflectance of natural phenomena, such as gases, fluids, solids, or flames [12], others evaluate aerodynamic properties of vehicles [26], or synthesize composites of captured phenomena and renderings [35]. Reconstructing phenomena with complex refractive properties, however, is a very challenging problem.

Existing solutions can be roughly classified as intrusive or non-intrusive methods. For example, intrusive methods cover objects of interest with diffuse coatings [7] or immerse them in fluorescent liquids [9] whereas non-intrusive methods measure the distortion of a reference pattern by the refractive phenomenon [22, 34, 24, 3, 1].

In this paper, we propose a non-intrusive technique that builds on the transport of intensity equation (TIE) [28, 27]. Derived using wave optics models with coherent illumination, the TIE is known in the optics community and has applications in phase microscopy. We present a new derivation of the TIE using incoherent light fields, which makes it intuitive and practical for computer vision applications. Further, we modify widely-used TIE acquisition setups by introducing coded illumination that facilitates improved robustness of tomographic refractive volume reconstructions. Finally, we evaluate the proposed computational imaging system in simulation and with a range of physical experiments.

## 2. Related Work

**Fluid and Phase Imaging** Imaging and reconstructing fluids is a diverse area of research. An overview of techniques such as particle image velocimetry and other flow estimation techniques can be found in [19]. Phase contrast microscopy and differential interference contrast are standard techniques in microscopy [23]. Schlieren photography is a non-intrusive way of encoding refractive phenomena as changes in intensity or color of a captured image [26]. Traditionally, this is done with intricate optical setups; modern approaches combine optical flow with coded backgrounds [6]. More recently, coded 4D light field illumination has been shown to replace both complicated optical setups in traditional setups and expensive computation in background-oriented Schlieren methods [32, 33].

A popular technique for phase imaging with coherent illumination is the transport of intensity equation (TIE) [28, 27]. Using a first or higher-order [31] approximation of the light transport operator, 2D phase information of an object can be reconstructed from a focal stack containing two or more images. An interpretation of these techniques in phase-space was presented in [25]. The methods discussed in this paper build on existing coherent TIE formulations, but derive forward and inverse methods using incoherent light fields along with optical coding schemes to make these techniques practical for computer vision applications.

**Refractive and Specular Surface Reconstruction** In the computer vision community, refractive surface recovery has been an extremely active area. A recent survey [12] presents a comprehensive overview of the field, whereas [17] classifies the problem space based on required acquisition setup and number of refractive events. A body of work has tackled the problem using a shape from distortion approach [22, 34, 24, 3, 1]. Active, coded illumination can be used instead of observing the distortion of static calibration patterns [21]. A simple approach to acquiring surfaces with complex reflectance properties is to either apply a diffuse coating and scan them with standard techniques [7] or to immerse [9] or dye [11] objects of interest with fluorescent substances. However, these are highly intrusive approaches. Finally, polarization properties can also be observed to infer information about refraction [20, 10]. Our approach is

related to other shape from distortion techniques that use active illumination, but we introduce the TIE to computer vision, enhance it with coded illumination, and combine it with tomographic volume reconstruction.

**Refractive Tomography** When imaging three-dimensional, inhomogeneous refractive volumes, conventional surface reconstruction techniques usually fail. Computed tomography methods [16], however, have been adopted to recover volumes of flames [13], gases [2, 15], and solids [30, 29]. The main difference between these methods is whether the tomographic reconstruction is performed from observed light attenuation [13, 30] or refraction [2, 29, 15].

Our method is closely related to [29] by employing the TIE, but we follow [15] by applying a regularizer based on Fermat’s principle to the tomographic reconstruction. To the best of our knowledge, this is the first time that the TIE is derived in the geometrical optics domain using a light field notation. We also augment the TIE with coded illumination that allows us to estimate ray-ray correspondences for more robust tomographic reconstructions.

### 3. Image Formation and Reconstruction

#### 3.1. Light Propagation in Refractive Media

Using a geometrical optics model, the propagation of light in materials with a heterogeneous refractive index field  $n$  is governed by the Eikonal equation or ray equation of geometrical optics [4]:

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

As illustrated in Figure 1,  $\mathbf{x}$  is the lateral position of a light ray on some plane that is defined to be zero along the optical axis and  $s$  is its trajectory through the medium. We can write this as two ordinary differential equations [14]:

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

The local direction of propagation is  $\nu$ . Equation 2 states that the change in direction a light ray undergoes in the medium is equal to the lateral gradient of the refractive index field.

Using the notion of light fields [18], we can describe the effects of a refractive medium to an arbitrary distribution of light rays in an intuitive manner. A light field  $l(x, \nu)$  fully defines a radiance distribution using rays as its primitives, that are parameterized by a position  $x$  and the propagation direction  $\nu = \tan(\theta)$ , where  $\theta$  is the angle of propagation. We use an intuitive 2D light field notation throughout the paper; full 4D formulations and extended illustration can be found in the supplement.

The refraction of a light field  $l^{(-)}(x, \nu)$  incident on a thin

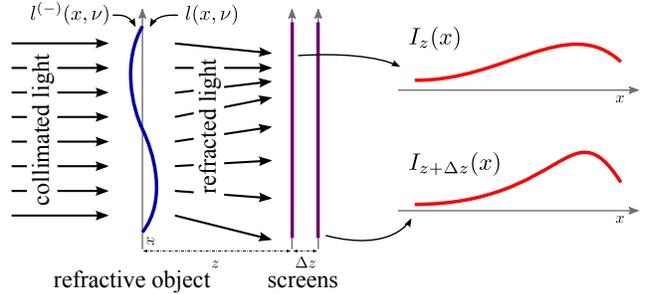


Figure 1. Ray schematic. Collimated illumination is refracted by a thin surface. The resulting intensity is measured at two distances,  $z$  and  $z + \Delta z$ . The transport of intensity equation relates the refractive surface gradients with the difference of those measurements, allowing the surface to be recovered.

refractive object is then

$$l(x, \nu) = l^{(-)}(x, \nu - \nabla n(x)), \quad (3)$$

where  $l(x, \nu)$  is the light field on the same location but after refraction. As illustrated in Figure 1, the refractive event basically shifts the angles of the light field. For volumetric index fields, one must trace a light ray through the volume along its curved trajectory  $c$ . From the Eikonal equations, we know that the outgoing angle of a ray is

$$\nu_{out} = \nu_{in} + \int_c \nabla n(x, z) ds, \quad (4)$$

for any ray with incident angle  $\nu_{in}$ . The volumetric refraction a light field undergoes is expressed as

$$l_0(x, \nu) = l_0^{(-)} \left( g(x, \nu), \nu - \int_c \nabla n(x, z) ds \right), \quad (5)$$

where  $g(x, \nu)$  is a general function describing the lateral displacement of a ray by the refractive volume.

We note that this formulation neglects effects such as reflection, absorption, scattering, diffraction, and dispersion. The refractive volume is assumed to be differentiable. Abrupt changes in a refractive index field, such as at the intersection of air and water, will be smoothed out in this formulation unless an additional sparse gradient prior is employed, as suggested in Section 3.4.

#### 3.2. Transport of Intensity

The transport of intensity equation (TIE) [28, 27] is often used in the optics community to model light propagation through refractive index fields using a first-order approximation of the transport operator. Whereas formulations exist using wave models of light, this section derives what we believe to be the first formulation of the TIE using light fields as a geometrical optics model.

Let us consider transport of a light field in free space from  $z$  by a distance  $\Delta z$ . This is a shear along the spatial coordinates:

$$l_{z+\Delta z}(x, \nu) = l_z(x - \Delta z \nu, \nu). \quad (6)$$

Following Semichaevsky and Testorf [25], the propagation operator can be approximated to first order using a Taylor expansion:

$$l_z(x - \Delta z\nu, \nu) \approx l_z(x, \nu) - \Delta z\nu \frac{\partial}{\partial x} l_z(x, \nu). \quad (7)$$

We note that this is a good approximation under paraxial assumptions, i.e. when the light field is propagating with small deviations along the optical axis  $z$ . Higher-order approximations can be employed if this is not the case [31].

The transport of intensity equation requires two intensity measurements to be taken at a slight axial distance  $\Delta z$  between each other. These can be recorded with a bare camera sensor or by photographing a white, diffuse screen with a conventional camera. The intensity on some axial position  $z$  is defined as the projection of a light field over angular domain  $\Omega_\nu$ :

$$I_z(x) = \int_{\Omega_\nu} l_z(x, \nu) d\nu. \quad (8)$$

Combining Equations 6, 7 and 8 yields an expression for the image intensity measured at  $z + \Delta z$ :

$$I_{z+\Delta z}(x) \approx \int_{\Omega_\nu} l_z(x, \nu) d\nu - \Delta z \int_{\Omega_\nu} \nu \frac{\partial}{\partial x} l_z(x, \nu) d\nu. \quad (9)$$

A forward finite differences approximation of intensity along the optical axis  $z$  is then

$$\frac{\partial I_z(x)}{\partial z} \approx \frac{I_{z+\Delta z}(x) - I_z(x)}{\Delta z} \quad (10)$$

Hence, combining Equations 8, 9, and 10 yields:

$$\frac{\partial I_z(x)}{\partial z} \approx - \int_{\Omega_\nu} \nu \frac{\partial}{\partial x} l_z(x, \nu) d\nu, \quad (11)$$

which is the foundation of the transport of intensity equation. Semichaevsky and Testorf [25] refer to the quantity  $\int_{\Omega_\nu} \nu l_z(x, \nu) d\nu$  as the first-order angular moment of the light field.

### 3.3. TIE-based Refractive Object Reconstruction

Using Equation 11, we can derive the transport of intensity equation for both thin and volumetric refractive phenomena. However, collimated incident illumination is required, i.e.  $l_0^{(-)}(x) = \delta(\nu) I(x)$ . Following Equation 3, the light field after refraction by a thin object is given as

$$l_0(x, \nu) = \delta(\nu - \nabla n(x)) I(x). \quad (12)$$

The right-hand side of Equation 11 now becomes

$$\begin{aligned} & - \int_{\Omega_\nu} \nu \frac{\partial}{\partial x} l_0(x, \nu) d\nu \\ &= - \frac{\partial}{\partial x} \left( I(x) \int_{\Omega_\nu} \nu \delta(\nu - \nabla n(x)) d\nu \right) \\ &= - \frac{\partial}{\partial x} (I(x) \nabla n(x)), \end{aligned} \quad (13)$$

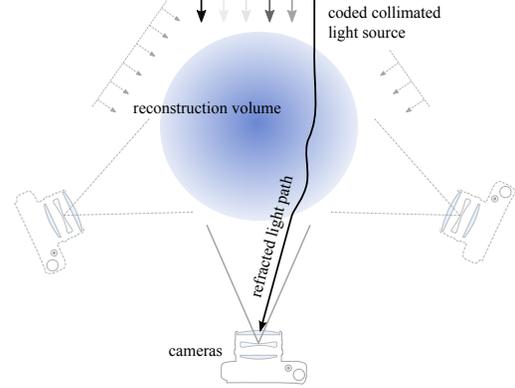


Figure 2. Tomographic volume reconstruction. The volumetric refractive index field is observed from multiple camera perspectives. A transport of intensity equation is solved for each independently. The resulting projections of the volume are then processed using computed tomography to recover the volume.

resulting in the transport of intensity equation for thin refractive phenomena:

$$\frac{\partial I(x)}{\partial z} = - \frac{\partial}{\partial x} (I(x) \nabla n(x)). \quad (14)$$

This formulation is closely related to the wave models derived in [28, 27] and to interpretations in phase-space [25]. Figure 1 illustrates the described setup using a ray diagram.

Imaging and reconstructing volumetric refractive index fields is also possible. Again, for collimated incident light  $l_0^{(-)}(x, \nu) = \delta(\nu) I(x)$ , this results in a variant of the TIE:

$$\frac{\partial I(x)}{\partial z} = - \frac{\partial}{\partial x} \left( I(x) \int_c \nabla n(x, z) ds \right). \quad (15)$$

### 3.4. Tomographic Volume Reconstruction

For the specific case of imaging volumetric refractive index fields, the TIE (Eq. 15) can only be solved for the accumulated amount of refraction that each ray underwent in the volume. This is a projection of all refractive index gradients along the optical path  $\int_c \nabla n(x, z) ds$ . As illustrated in Figure 2 and proposed in previous work (e.g., [2, 15, 29]), computed tomography methods [16] (e.g., the simultaneous algebraic reconstruction technique (SART)) can be used to solve for the volumetric index field from projections captured from different camera perspectives. For this purpose, the refractive index field is represented as a linear combination of  $i = 1 \dots k$  basis functions  $\phi_i$  with associated weights  $\alpha_i$ , leading to the following expression:

$$\int_c \nabla n ds = \int_c \left( \sum_i \alpha_i \phi_i \right) ds = \sum_i \alpha_i \left( \int_c \phi_i ds \right). \quad (16)$$

Common choices for  $\phi_i$  are cubic voxels or radially symmetric basis functions [2]; in either case, the number of basis functions is equal to the number of unknown voxels.

Using this formulation in hand, Equation 15 can be solved independently for multiple different camera perspectives to give one projection of the 3D index field each. Stacking all resulting projections in a vector  $\mathbf{b} \in \mathbb{R}^l$ ,  $l$  being the total number of all camera pixels, allows us to formulate a tomographic reconstruction problem that solves for the unknown coefficients  $\alpha \in \mathbb{R}^k$

$$\mathbf{b} = \Phi\alpha + \Gamma(\alpha). \quad (17)$$

Here, column  $i$  of matrix  $\Phi \in \mathbb{R}^{l \times k}$  contains the precomputed light path integrals over the basis functions  $\int_c \phi_i ds$  for some measurement. An additional visual hull restriction and total variance regularizer  $\Gamma(\cdot)$  can be employed if the equation system is ill-posed. The visual hull can be used to determine the basis functions with possibly nonzero coefficients [2] and the regularizer is a direct 3D extension of the commonly-used 2D TV regularizer. Detailed explanation can be found in the supplement. Equation 17 can be computed with any standard linear solver.

### 3.5. Incorporating Light Path Approximation

Unfortunately, the actual ray paths through the volume are usually unknown. Common assumptions include straight ray approximation [29] or approximating the path by a single refractive event close to the center of the volume [2]. Ji et al. [15] recently proposed a method that optically acquires the correspondences between light rays incident on the volume and the same, but refracted rays emerging on the other side using light field probes [32]. These ray-ray correspondences have been demonstrated to achieve higher-quality reconstructions by imposing Fermat’s principle as a regularizer to the tomography problem. We believe this to be an effective way of constraining the possible set of solutions for volumetric refractive index fields.

Establishing ray-ray correspondences in the proposed framework is straightforward: instead of illuminating the sample object with uniform collimated light, i.e.  $I(x) = const$ , we can code the intensity distribution  $I(x)$  such that ray-ray correspondences can be estimated using optical flow. We employ a simple color gradient, although more sophisticated illumination patterns, such as gray codes, are possible. We note that ray-ray correspondences require both position and direction of an incident ray to be matched with one of the refracted rays. Whereas [15] achieve this using 4D light field probes, our method relies on collimated illumination thereby unambiguously identifying incident rays using purely spatial illumination codes.

## 4. Analysis and Evaluation

In this section, we evaluate the proposed methods. We proceed by simulating and evaluating system parameters and noise resilience for coded acquisition and reconstruction of a thin, two-dimensional refractive object. We also

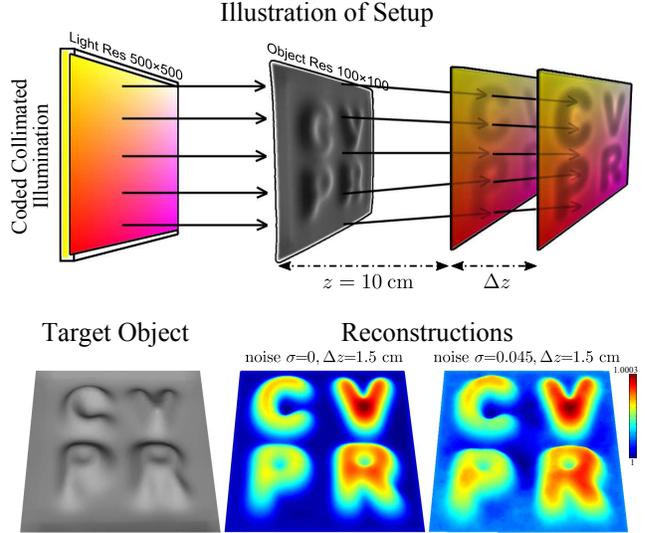


Figure 3. Reconstruction of thin refractive object. Two images, separated by distance  $\Delta z$ , are captured (top) and processed to recover the refractive index (bottom, right).

show tomographic reconstructions of three-dimensional volumetric scenes, both for homogeneous and inhomogeneous refractive index fields. Finally, we compare our method with previous work.

**Thin Refractive Object** Figure 3 illustrates the setup of a synthetic experiment. A thin refractive object distorts the incident collimated illumination, which is then measured at two different distances from the sample object. By solving the transport of intensity equation (Eq. 14), we can recover the refractive index field. Figure 3 shows two reconstructions with different levels of iid Gaussian sensor noise. For this experiment, the object is simulated to be  $4 \times 4$  cm, the incident illumination is coded with a simple color gradient, and the refractive index varies smoothly between 1 and 1.0003. Equation 14 is solved by converting the measured color images to grayscale.

Reconstruction performance w.r.t to sensor noise and a varying distance  $\Delta z$  is further evaluated in Table 1. Whereas the peak signal-to-noise ratio (PSNR), measured between 3D ground truth index field and reconstruction, improves with a decreasing distance in the absence of noise, the subtle changes in images recorded at such small distances are easily perturbed by noise. An optimal tradeoff has to be found depending on the employed camera’s noise characteristics.

**Homogeneous Volumetric Refractive Index Field** To recover a full, three-dimensional refractive index field, multiple projections of the volume need to be captured from different perspectives. We simulate a voxelized bunny with a refractive index of 1.4 surrounded by air in Figure 4. Image pairs are recorded from eight different camera perspectives and individually processed with Equation 14 to yield

$\Delta z$ in cm	Sensor Noise Standard Deviation $\sigma$					
	0	0.02	0.03	0.035	0.040	0.045
0.25	28.1	22.5	19.8	20.8	17.9	16.3
0.75	28.3	26.4	24.9	22.0	19.3	17.9
1.5	27.2	26.2	25.5	23.8	21.1	21.0
3	24.5	23.8	22.8	21.9	19.8	18.8

Table 1. Evaluation of resilience to sensor noise and variations in distance between the two captured images. Reconstruction errors are given in peak signal-to-noise ratio (PSNR). Based on these simulations, we see that a smaller distance is favorable for lower sensor noise, but larger distances are more resilient to noise.

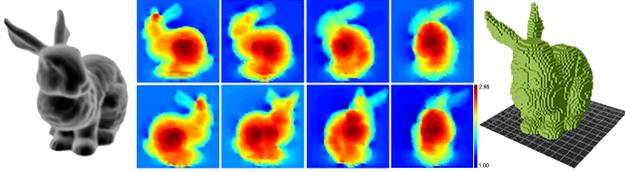


Figure 4. Reconstruction of a homogeneous refractive volume. From left: target object, projections reconstructed via Eq. 14 from eight different perspectives, tomographically-reconstructed 3D volume.

projections of the refractive index volume. Using the tomographic reconstruction method outlined in Equation 17, we can recover the refractive volume.

### Inhomogeneous Volumetric Refractive Index Field

Using the fuel injection dataset DFG SFB 382 ([www.volvis.org](http://www.volvis.org)), we evaluate the proposed method for an inhomogeneous refractive index field. Figure 5 shows a 3D rendering of the volume (left) as well as individual slices (right). We simulate acquisition by seven equally-spaced perspectives in a half ring setup. From each perspective, two images are recorded at slightly different axial positions (one shown in Fig. 5, left); their difference is used to solve for the 2D projection of the refractive index field via Equation 14. The coded illumination allows for ray-ray correspondences to be estimated using optical flow (Fig. 5, left), which makes a tomographic reconstruction of the 3D volume more robust. For this experiment, we employ a 3D total variation (TV) regularizer for the reconstruction (Eq. 17) and assume that the surrounding medium has the refractive index of air.

The reconstruction quality of the proposed 3D tomographic method w.r.t. the number of camera perspectives is evaluated in Table 2. As expected, an increasing number of perspectives makes the volumetric reconstruction more robust and leads to lower root-mean-square errors (RMS) and peak signal-to-noise ratios (PSNR). The proposed reconstruction is similar to SART [16]. Whereas direct inversion methods, such as filtered backprojection using the Radon transform, could be significantly faster, these are only applicable when a dense set of projections of a full-ring setup are available. We also implemented the methods

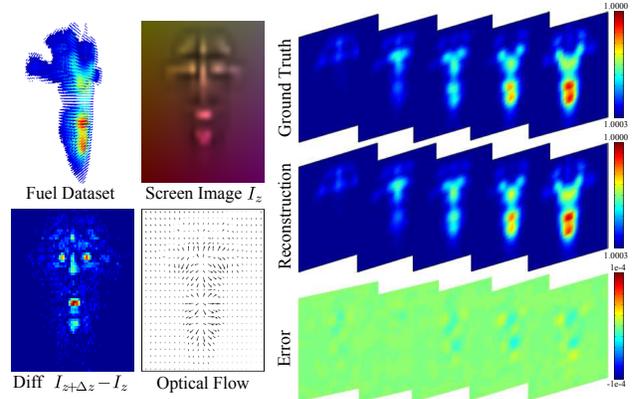


Figure 5. Volumetric reconstruction. Left: fuel dataset (ground truth), one of the two images captured at slightly different axial positions, their difference, and estimated optical flow from coded illumination. Right: slices of original volume, reconstructed volume, and error.

Simulation Errors w.r.t. Number of Cameras						
# of cameras		3	5	8	16	32
Proposed Method	RMS	0.051	0.049	0.041	0.036	0.030
	PSNR	25.8	26.25	27.77	28.85	30.57
Acheson et al. 2008	RMS	0.067	0.060	0.054	0.045	0.040
	PSNR	23.43	24.41	25.37	27.03	27.98
Ji et al. 2013	RMS	0.057	0.053	0.045	0.041	0.036
	PSNR	24.49	25.81	26.75	28.22	29.30

Table 2. Evaluation of tomographic reconstruction using a varying number of cameras in a half-ring setup. Compared to Acheson et al. [2] and Ji et. at [15], the proposed method always performs better under the simulated conditions.

proposed by Acheson et al. [2] and Ji et. at [15]. The results show that our method has a higher RMS/PSNR than theirs for all simulated configurations. This is intuitive, because we follow [15] by using a light path approximation prior but our setups allows us to capture higher-resolution data. Additional results and comparisons can be found in the supplement.

## 5. Implementation

**Hardware** Figure 6 shows our prototype capture device. We use a Texas Instruments LightCrafter as a programmable light source, which provides a resolution of  $608 \times 684$  pixels and three primary color channels. A lens is mounted at its focal distance to the projector and collimates the projected optical code – a simple color gradient – which then rear-illuminates the refractive sample. A beam splitter separates the optical path on the other side of the sample such that two intensity projections at distances spaced by 1cm can be imaged by two machine vision cameras (Point-Grey Flea2-08S2,  $1024 \times 768$  px) simultaneously. The refractive sample is mounted on a rotation stage; for volumetric reconstructions, multiple images are captured showing the sample from different perspectives.

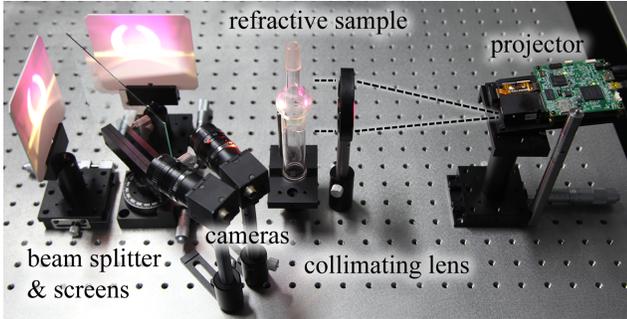


Figure 6. Prototype capture device. A collimated projector provides coded illumination that is refracted by the sample and imaged on the other side of the setup.

**Calibration** Both TIE and tomographic volume reconstruction require precise image registration. For this purpose, we use standard camera calibration techniques [5] to correct for intrinsic distortions. Further, we register the two captured images using features observed in a calibration image displayed by the projector.

**Software** After this one-time calibration procedure, the system simultaneously records calibrated image pairs with a resolution of  $450 \times 450$  pixels. For thin refractive object, we further down-sample them to  $256 \times 256$  pixels and use a Fast Poisson Solver (<http://cs.nyu.edu/~harper/poisson.htm>) to solve Equation 14. This solver is implemented in Matlab and takes approx. 76 seconds on an Intel i7 PC with 8 GM RAM. For volumetric refractive index fields, we use the CVX toolkit [8] to solve Equation 17. Recovered volumes are discretized to  $64 \times 64 \times 64$  voxels. A reconstruction from eight different perspectives takes approx. 3 minutes. Both the Poisson solver and the tomographic reconstruction could be significantly accelerated using more advanced algorithms.

## 6. Experimental Results

In this section, we describe a number of experiments we have performed with the prototype setup described in the previous section.

### 2D Reconstruction of Static Refractive Index Fields

Figure 7 shows results of static objects recovered with the proposed technique. For this experiment, we drew characters on a thin glass<sup>1</sup> plate using clear optical adhesive with a refractive index of 1.48. Photographs of the target objects are shown in the top row, whereas the center row shows reconstructions of the refractive index field with a resolution of  $256 \times 256$ . Please note that the objects are not actually thin, so the reconstructions show integrals over the refractive index field as described by Equation 15. Height field renderings of the reconstructions are shown in the bottom

<sup>1</sup>The glass plate has a refractive index of 1.51.

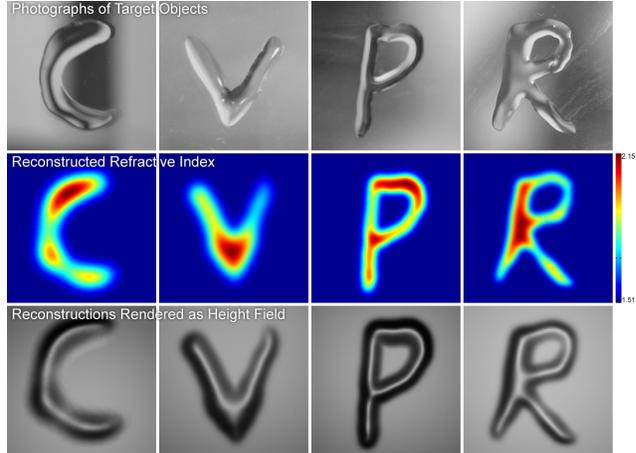


Figure 7. Reconstructions of refractive index fields. These characters are composed of clear optical adhesive with a refractive index that is different from the underlying glass plate. By capturing two differently-focused images from the same perspective, we can recover 2D projections of the refractive index fields, as illustrated in the center and bottom rows.

of Figure 7. The two focal planes are separated by a distance  $\Delta z = 1 \text{ cm}$ . Other system parameters are described in Section 5.

**2D Reconstruction of Dynamic Index Fields** Using the dual-camera setup described in Section 5, we can capture dynamic scenes from a single perspective but focused at two different axial planes. As shown in Figure 8, the hot air plume of a candle can be captured and recovered using the proposed method. Similar to the reconstructions of Figure 7, the transport of intensity equation actually computes a two-dimensional projection of the three-dimensional refractive index field. System parameters for this experiment are identical to the previous one.

### 3D Tomographic Reconstruction of Static Index Fields

Finally, we also show a tomographic reconstruction of a solid three-dimensional glass object in Figure 9. This is captured from multiple perspectives by mounting it on a rotation stage and recording photographs at seven settings that are equally spaced apart by  $30^\circ$ . The two focal planes for each perspective are separated by a distance of  $\Delta z = 0.5 \text{ cm}$ . For each of the perspectives, we recover the 2D projection of the volumetric refractive index field via Equation 15 and then solve the tomographic problem, i.e. Equation 17. The latter incorporates light path approximation constraints using the ray-ray correspondences that are estimated from the illumination coding scheme described in Section 3.5.

## 7. Discussion

In summary, we propose a technique to reconstruct two and three-dimensional refractive index fields using a new, geometrical optics formulation of the transport of intensity

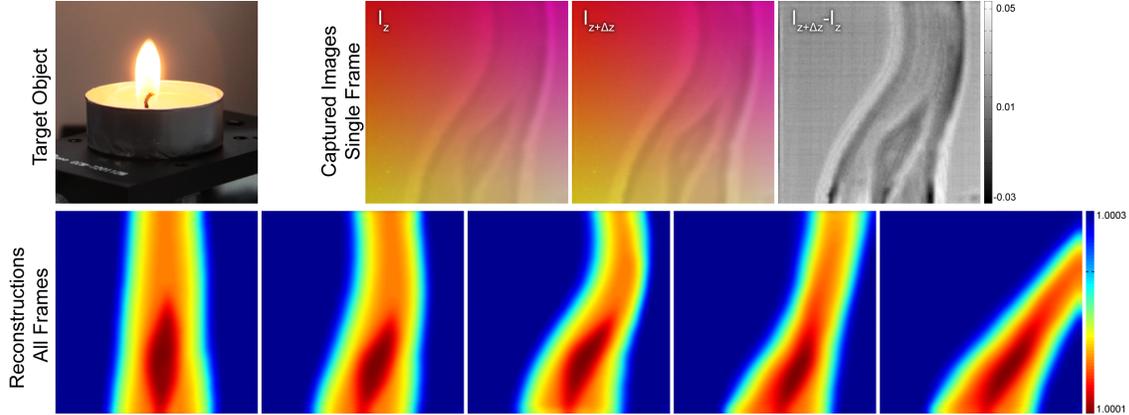


Figure 8. Reconstructions of a dynamic refractive index field. Even the small differences in refraction caused by the heat above a candle can be faithfully recovered via coded transport of intensity. In this experiment, we capture time-synchronized videos with two cameras focused at different focal planes (shown for a single frame of the animation, top) and reconstruct each frame separately (bottom).

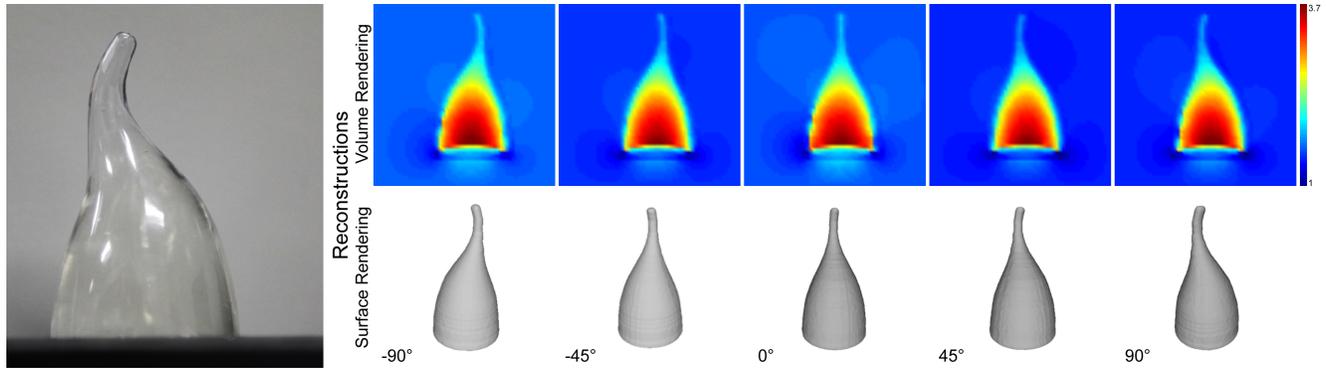


Figure 9. Three-dimensional tomographic reconstruction of a volumetric object (left). We capture pairs of photographs from seven perspectives (not shown) and recover the volumetric refractive index field as described in the text. Volume and surface renderings of the reconstructions are shown from perspectives that are different from those captured.

equation. In addition to deriving it using light field notation, we enhance the models used in the optics community by coded background illumination that allows us to incorporate more sophisticated regularizers, as recently proposed in the computer vision community [15]. We evaluate the proposed technique in simulation and demonstrate that it achieves higher-quality three-dimensional, tomographic reconstructions than alternative techniques. We also build a prototype device and show reconstructions of static and dynamic refractive index fields.

### 7.1. Limitations

The proposed method makes a number of assumptions that may not be met for arbitrary refractive objects. These include the absence of reflection, scattering, diffraction, and dispersion. Further, the projection of the refractive index field is assumed to be differential, hence smooth. While the transport of intensity equation inherently separates absorption from refraction, the proposed optical coding strategy requires that the medium is either non-absorptive or that absorption is constant over the imaged area. Otherwise, the

optical flow may fail to recover precise ray-ray correspondences.

Our prototype system has limitations as well. The finite projector aperture provides rear-illumination that is not perfectly collimated and the optical codes are limited by the minimal observable changes in pixel intensity and hue by the camera. Dynamic scenes captured with a finite camera exposure time may contain motion blur.

### 7.2. Future Work

In the future, we would like to experiment with an array of cameras that captures dynamic, three-dimensional index fields from multiple perspectives simultaneously. We could reduce the number of required devices by using a single light field camera for each perspective and synthesize multiple intensity projections at different focal distances. This approach, however, would come at the cost of reduced image resolution. Finally, capturing more than two intensity projections from each perspectives can achieve better reconstructions [31].

## 8. Acknowledgements

We thank the reviewers for their insightful feedback. Chenguang Ma, Xing Lin, Jinli Suo and Qionghai Dai were supported by the Project of NSFC (No. 61327902, 61120106003 and 61035002). Gordon Wetzstein was supported by an NSERC Postdoctoral Fellowship.

## References

- [1] S. Agarwal, S. P. Mallick, D. Kriegman, and S. Belongie. On Refractive Optical Flow. In *Proc. ECCV*, 2004. **1**
- [2] B. Atcheson, I. Ihrke, W. Heidrich, A. Tevs, D. Bradley, M. Magnor, and H.-P. Seidel. Time-resolved 3d capture of non-stationary gas flows. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, 27(5):132, 2008. **2, 3, 4, 5**
- [3] M. Ben-Ezra and S. K. Nayar. What Does Motion Reveal About Transparency? In *Proc. ICCV*, pages 1025–1032, 2003. **1**
- [4] M. Born and E. Wolf. *Principles of Optics*. Cambridge University Press, 7 edition, 1999. **2**
- [5] J.-Y. Bouguet. Camera Calibration Toolbox for Matlab, 2010. **6**
- [6] S. Dalziel, G. Hughes, and B. Sutherland. Whole-Field Density Measurements by Synthetic Schlieren. *Experiments in Fluids*, 28(4):322–335, 2000. **1**
- [7] M. Goesele, H. P. A. Lensch, J. Lang, C. Fuchs, and H.-P. Seidel. DISCO: Acquisition of Translucent Objects. *Proc. SIGGRAPH*, 23:835–844, 2004. **1**
- [8] M. Grant, S. Boyd, and Y. Ye. CVX: Matlab software for disciplined convex programming, 2008. **6**
- [9] M. B. Hullin, M. Fuchs, I. Ihrke, H.-P. Seidel, and H. P. A. Lensch. Fluorescent Immersion Range Scanning. In *ACM Trans. Graph. (Siggraph)*, pages 87:1–87:10, 2008. **1**
- [10] C. P. Huynh, A. Robles-Kelly, and E. Hancock. Shape and Refractive Index Recovery from Single-View Polarisation Images. In *Proc. CVPR*, 2010. **1**
- [11] I. Ihrke, B. Goldluecke, and M. Magnor. Reconstructing the Geometry of Flowing Water. In *Proc. ICCV*, pages 1055–1060, 2005. **1**
- [12] I. Ihrke, K. N. Kutulakos, H. P. A. Lensch, M. Magnor, and W. Heidrich. Transparent and Specular Object Reconstruction. *Computer Graphics Forum*, 29(8), 2010. **1**
- [13] I. Ihrke and M. Magnor. Image-Based Tomographic Reconstruction of Flames. In *Proc. SCA*, pages 367–375, 2004. **2**
- [14] I. Ihrke, G. Ziegler, A. Tevs, C. Theobalt, M. Magnor, and H.-P. Seidel. Eikonal Rendering: Efficient Light Transport in Refractive Objects. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 26(3):59, 2007. **2**
- [15] Y. Ji, J. Ye, and J. Yu. Reconstructing Gas Flows Using Light-Path Approximation. In *Proc. CVPR*, 2013. **2, 3, 4, 5, 7**
- [16] A. C. Kak and M. Slaney. *Principles of Computerized Tomographic Imaging*. Society for Industrial Mathematics, 2001. **2, 3, 5**
- [17] K. N. Kutulakos and E. Steger. A Theory of Refractive and Specular 3D Shape by Light-Path Triangulation. In *Proc. ICCV*, page 1448–1455, 2005. **1**
- [18] M. Levoy and P. Hanrahan. Light Field Rendering. In *Proc. Siggraph*, pages 31–42, 1996. **2**
- [19] W. Merzkirch. *Flow Visualization*. Academic Press, 1987. **1**
- [20] D. Miyazaki and K. Ikeuchi. Inverse Polarization Raytracing: Estimating Surface Shapes of Transparent Objects. In *Proc. CVPR*, page 910–917, 2005. **1**
- [21] N. J. W. Morris and K. N. Kutulakos. Reconstructing the Surface of Inhomogeneous Transparent Scenes by Scatter Trace Photography. In *Proc. ICCV*, 2007. **1**
- [22] H. Murase. Surface Shape Reconstruction of an Undulating Transparent Object. In *Proc. ICCV*, pages 313–317, 1990. **1**
- [23] D. B. Murphy. *Fundamentals of Light Microscopy and Electronic Imaging*. Wiley-Liss, 2001. **1**
- [24] S. Savarese and P. Perona. Local Analysis for 3D Reconstruction of Specular Surfaces - Part II. In *Proc. ECCV*, pages 759–774, 2002. **1**
- [25] A. Semichaevsky and M. Testorf. Phase-space interpretation of deterministic phase retrieval. *J. Opt. Soc. Am. A*, 21(11):2173–2179, 2004. **1, 3**
- [26] G. S. Settles. *Schlieren and Shadowgraph Techniques*. Cambridge University Press, 2001. **1**
- [27] N. Streibl. Phase imaging by the transport equation of intensity. *Optics Communications*, 49(1):6 – 10, 1984. **1, 2, 3**
- [28] M. R. Teague. Deterministic phase retrieval: a green’s function solution. *J. Opt. Soc. Am.*, 73(11):1434–1441, 1983. **1, 2, 3**
- [29] L. Tian, J. Petrucci, Q. Miao, and G. Barbastathis. Compressive XRay Phase Tomography based Intensity Transport. In *OSA Imaging and Applied Optics*, 2013. **2, 3, 4**
- [30] B. Trifonov, D. Bradley, and W. Heidrich. Tomographic Reconstruction of Transparent Objects. In *Proc. EGSR*, pages 51–60, 2006. **2**
- [31] L. Waller, L. Tian, and G. Barbastathis. Transport of intensity phase-amplitude imaging with higher order intensity derivatives. *Opt. Express*, 18(12):12552–12561, 2010. **1, 3, 7**
- [32] G. Wetzstein, R. Raskar, and W. Heidrich. Hand-Held Schlieren Photography with Light Field Probes. In *IEEE International Conference on Computational Photography (ICCP)*, 2011. **1, 4**
- [33] G. Wetzstein, D. Roodnick, R. Raskar, and W. Heidrich. Refractive Shape from Light Field Distortion. In *International Conference on Computer Vision (ICCV)*, 2011. **1**
- [34] X. Zhang and C. S. Cox. Measuring the Two-dimensional Structure of a Wavy Water Surface Optically: A Surface Gradient Detector. *Exp. in Fluids*, 17:225–237, 1994. **1**
- [35] D. E. Zongker, D. M. Werner, B. Curless, and D. H. Salesin. Environment Matting and Compositing. In *ACM Trans. Graph. (Siggraph)*, pages 205–214, 1999. **1**