

Imaging and 3D Tomographic Reconstruction of Time-varying, Inhomogeneous Refractive Index Fields

Bradley Atcheson¹, Ivo Ihrke², Derek Bradley¹, Wolfgang Heidrich¹, Marcus Magnor³, Hans-Peter Seidel²
1) The University of British Columbia, 2) Max-Planck-Institut für Informatik, 3) Braunschweig University

Technical Report
UBC CS TR-2007-06



Figure 1: Measurements of the refractive index field caused by rising hot air from a gas burner. Far left: magnitude of refraction for each pixel in one view. Center left: vector field of 2D displacements. Center: distortion measurements from one view are used as an environment matte to distort the background, data from another view is used to cast a caustic (shadowgraph). Center right: maximum intensity projection volume rendering of a 3D refractive index field recovered with tomography from 8 views. Far right: volume rendering of the 3D reconstruction.

Abstract

We present a technique for 2D imaging and 3D tomographic reconstruction of time-varying, inhomogeneous refractive index fields. Our method can be used to perform three-dimensional reconstruction of phenomena such as gas plumes or liquid mixing. We can also use the 2D imaging results of such time-varying phenomena to render environment mattes and caustics.

To achieve these results, we improve a recent fluid imaging technique called Background Oriented Schlieren imaging, and develop a novel theory for tomographic reconstructions from Schlieren images based on first principles of optics. We demonstrate our approach with two different measurement setups, and discuss example applications such as measuring the heat and density distribution in gas flows.

1 Introduction

Time varying, inhomogeneous refractive index fields are everywhere, from hot air rising above a fire (Figure 1) and mirages above hot roads in the summer, to gases flowing from pressurized containers, and mixtures of different liquids, such as tea and sweetener.

From a computer graphics point of view, refractive index fields can distort camera rays, locally change the focus of cameras, and cast caustics on surfaces (Figure 1, center). Under certain conditions, refractive index fields directly relate to the density field, or to the spatial temperature distribution. It is therefore interesting not only to capture real-world refractive index fields for use in rendering algorithms, but also to provide comparison data for fluid simulations.

In this paper, we present a novel method for imaging and volumetric reconstruction of time-varying refractive index fields. Our method is based on the Background Oriented Schlieren technique developed in the fluid imaging community [Meier 2002; Richard and Raffel 2001]. However, we extend it in several significant

ways, including a new approach for tracking larger refractions, and a novel theory for tomographic reconstruction of refractive index fields that does not require paraxial approximations. We describe measurement setups, and discuss several example applications of the method.

In particular, this paper makes the following contributions:

- a robust method for computing optical flow under strong isotropic and anisotropic scaling,
- a dynamic environment matting method that makes use of the optical flow algorithm to compute mattes of animated objects,
- a novel theory for the tomographic reconstruction of 3D refractive index fields from a number of 2D optical flow images, and finally
- applications of these methods to capturing 3D gas and fluid flows, 3D density and heat distributions, as well as rendering of environment mattes and caustics (shadowgraphs).

In the following, we first review work related to our approach, and then provide a more detailed overview of our method and contributions in Section 3. Section 4 deals with the 2D imaging of refractive index fields, and Section 5 with their tomographic reconstruction. In Section 6 we describe the specific measurement setups used in our experiments, and we conclude the paper with results and applications in Section 7.

2 Related Work

The measurement of transparent objects and phenomena has a long history in computer graphics, computer vision, fluid imaging, and other disciplines. For the purpose of conciseness, we discuss here only those methods that capture refractive effects. This includes 2D imaging methods such as environment matting, as well as 3D volumetric reconstruction methods, but excludes matting techniques that do not deal with refraction.

2D environment matting is one of the key 2D imaging methods for capturing transparent effects with refraction. It proceeds by observing a sequence of structured background patterns through the object under investigation. In the simplest case, the background is a binary encoding of regions that can be used to determine the 4-dimensional matrix of contributions of background pixels to foreground pixels [Zongker et al. 1999; Chuang et al. 2000]. Due to the large number of images required for this approach, later work has focused on reducing this number by using a hierarchy of backgrounds [Peers and Dutré 2003], or by using unstructured backgrounds in combination with a learning method [Wexler et al. 2002]. Environment matting techniques can also be extended to recover an approximate 3D shape, which can be combined with image-based deflection information to form a description of the 3D object [Matusik et al. 2002].

Despite a significant reduction of the number of images required for environment matting in recent years, it is still not possible to produce an environment matte with a single image and no prior information. Therefore, these methods do not apply to dynamic effects such as gas or liquid flows.

Optical flow algorithms provide a means to measure local image distortions due to refractive index variations. These algorithms can be used to measure the deflection of camera rays as they pass through a refractive medium. Since the pioneering work on local and global optical flow reconstruction by Lucas and Kanade [1981] and Horn and Schunck [1981], respectively, a multitude of computational approaches have been devised and various different fields of application discovered [Barron et al. 1994; Baker and Matthews 2004].

Related to our goal of recovering inhomogeneous 3D refractive index fields, Agarwal et al. [2002] extend the optical flow equation to reconstruct 2D refractive index variations from a video sequence with moving background, much like the work on environment matting. Still, one general challenge for any optical flow algorithm is the case where image regions undergo some noticeable change in scale, which frequently occurs in the presence of refractive index variations.

3D reconstruction of transparent objects in computer graphics and computer vision has, with a few notable exceptions, focused on solids, largely due to the large number of views required to obtain the necessary information. A thorough analysis of ray paths can provide information about the object shape [Ben-Ezra and Nayar 2003; Miyazaki and Ikeuchi 2005; Morris and Kutulakos 2005], but these methods are fundamentally limited to ray geometries that interact with the object surface at no more than two points [Kutulakos and Steger 2005]. Trifonov et al. [2006] avoid this fundamental limitation by suspending the object in a fluid of the same refractive index, thereby eliminating ray bending. They use tomographic reconstruction based on volumetric absorption, rather than the changes in ray geometry, to determine the 3D shape. However, like the other methods, their work only applies to solids with a homogeneous refractive index.

Only recently, Ihrke and coworkers have considered certain classes of dynamic transparent objects, including liquids [Ihrke et al. 2005] and flames [Ihrke and Magnor 2004]. In the former case, the refractive index of the object is required to be constant, while in the latter refraction is ignored altogether.

Schlieren imaging and tomography is an imaging technique for dynamically changing refractive index fields that has been de-

veloped in the fluid imaging community over several decades.¹ The original setups developed in the 1940s require high-quality optical mirrors and lenses of a size comparable to the volume under investigation [Schardin 1942]. Although these kinds of systems are still in use even today [Settles 2001], they are expensive and difficult to set up. It is also difficult to perform *quantitative* rather than qualitative measurements with them, although it is possible with careful calibration [Howes 1984].

These shortcomings have prompted the development of the much simpler Background Oriented Schlieren (BOS) technique in recent years [Dalziel et al. 2000; Richard and Raffel 2001; Meier 2002; Elsinga et al. 2004]. In the BOS setup, a digital camera observes a high-frequency background through the volume under investigation. Optical flow algorithms are used to compute a per-pixel deflection vector with respect to an undistorted reference background. Since only a single image is required for each geometry, the method is suitable for dynamic refractive index fields.

If 2D deflection data for *multiple* viewpoints is available from either BOS, the original Schlieren setup, or other sources, then it is possible to reconstruct a volumetric estimate of the refractive index field [Venkatakrishnan and Meier 2004; Schwarz 1996; McMackin et al. 1999]. These methods make use of an approximation that is only valid for small ray deflections, as described further in Section 3. Due to the difficulty of acquiring deflection images of dynamic volumes from multiple viewpoints, some researchers use tomographic reconstruction from a single image assuming rotational symmetry of the volume [Agrawal et al. 1999; Faris and Byer 1988].

In this paper, we develop a new theory for tomographic reconstruction of refractive index fields from 2D deflection images. We use this theory to design a new reconstruction algorithm that does not require the approximations of previous methods, and we demonstrate it on BOS images for rotationally symmetric flows as well as flows captured from multiple viewpoints.

3 Overview

Our method for capturing dynamic, spatially-varying refractive index fields consists of two primary components: the 2D imaging of ray deflections due to a 3D refractive index field, and the tomographic reconstruction of that 3D field from a number of deflection images captured from different positions.

2D deflection sensing involves a BOS imaging setup, using digital video cameras observing a high-frequency background through the refractive index field under investigation. The per-pixel deflections caused by the refraction (Figure 2) are captured with an optical flow algorithm, which can be directly used to visualize these distortions.

Our contribution to BOS imaging is the development of a robust optical flow method that finds matches in the background pattern even if the refraction is strong enough to locally cause large isotropic and anisotropic scaling of the background pattern. It should be noted, however, that there are limits to what is possible with optical flow methods. If the refractive index differences are strong enough to cause total internal reflection or drastic changes in the focus of the optical setup, then optical flow may not be the best choice. Our method is therefore not well suited for recovering solids made of

¹The German word “Schlieren” means “streaks” and refers to the optical distortions caused by inhomogeneous refractive index fields.

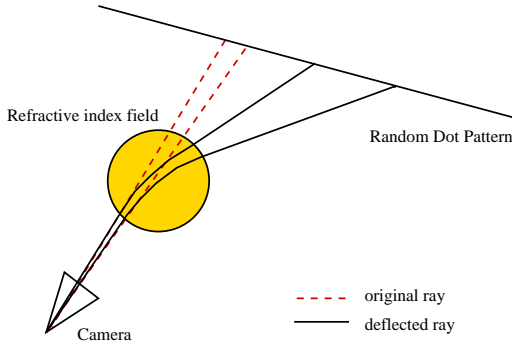


Figure 2: Principle of the deflection sensor: A plane with a high-frequency dot pattern is placed behind the scene of interest and an image is recorded without the object (dashed red lines). Then the inhomogeneous refractive index field is inserted between the camera and the background plane. Another image is taken and the deflection of the light rays in the image plane is computed using optical flow.

glass or other transparent materials. It may, however, be possible to use our method if such solids are immersed in fluids of comparable refractive index, similar to the work of Trifonov et al. [Trifonov et al. 2006], but without the need for a very precise match of refractive indices. We leave this application for future work.

3D tomographic reconstruction is based on a set of deflection images taken from different viewpoints. The existing methods in literature are, without exception, based on the *paraxial approximation* [Faris and Byer 1988; Venkatakrishnan and Meier 2004], i.e. the assumption that the deflections do not cause significant changes in the ray path through the reconstruction volume, even though the deflection angle is large enough to be measured. Under this approximation, consider a camera ray along the z -axis. We can write the angular deflection in the horizontal (i.e. x) direction as a *line integral* of differential horizontal changes in index of refraction along the ray

$$\phi_x = \frac{1}{n_0} \int \frac{\partial n}{\partial x} dz, \quad (1)$$

where n_0 is the refractive index of the surrounding environment. A similar equation holds for the vertical deflection angle ϕ_y .

Under this paraxial view, the pixel measurements correspond directly to line integrals of refractive index gradients. Consequently, a volume of refractive index gradients can be reconstructed with standard tomographic methods, such as Fourier slice reconstruction or algebraic reconstruction (ART) [Kak and Slaney 2001]. Finally, this gradient volume can be integrated into a refractive index field by solving a Poisson equation. For more detail on this approach, please refer to the work by Faris and Byer [1988].

Unfortunately, this method breaks down if inhomogeneities in the refractive index field are strong enough to cause ray deflections comparable to or larger than the voxel spacing used in the reconstruction. For this reason, we derive a novel theory for tomographic reconstruction that does not neglect changes in the geometry of the ray path. This theory is directly derived from the Eikonal equation and the ray equation, and gives rise to an efficient reconstruction algorithm.

In the following, we first discuss our contributions to Background Oriented Schlieren imaging and the new tomographic reconstruction

method, before we describe example measurement setups, as well as results and applications.

4 Background Oriented Schlieren Imaging

As mentioned before, the recently developed Background Oriented Schlieren technique [Meier 2002; Richard and Raffel 2001; Elsinga et al. 2004; Venkatakrishnan and Meier 2004] measures the per-pixel ray refraction caused by a volume by observing a high-frequency background through that volume, and computing the optical flow relative to an undistorted image (Figure 3). This method represents a significant reduction in cost and complexity compared to traditional Schlieren imaging.




Figure 3: Left: reference image, middle: distorted image, right: absolute difference (contrast enhanced). The candle plume is being blown from the side with a can of compressed air.

However, some challenges remain, especially if the object to be measured exhibits comparatively large differences in refractive index. In this case, the background distortions become too severe for simple optical flow algorithms to find matching regions. One challenge is that the frequency characteristics of the random dot patterns used in most BOS implementations [Richard and Raffel 2001] degenerate quickly with even uniform scaling: magnification of such patterns results in larger regions of uniform black or white color, while shrinking quickly results in a uniform medium grey color. In both cases, it is difficult for the optical flow algorithm to pick up matching regions. For this reason, we use Wavelet Noise [Cook and DeRose 2005] as our background pattern, since it contains details in a wide range of frequency bands, and therefore degrades much more gracefully.

Our optical flow algorithm is designed to use this high-frequency information to find matches under distortions. We use window-based normalized cross-correlation to compute the flow, which can be adapted to handle these distortions by searching for matches in isotropically or anisotropically scaled versions of the images. The basic optical flow is computed using spatial convolution on a per-window basis using the standard normalized cross-correlation formula

$$C(m, n) = \frac{\sum_{i,j=1}^N (f_{i,j} - \bar{f})(g_{i-m,j-n} - \bar{g})}{\sqrt{\sum_{i,j=1}^N (f_{i,j} - \bar{f})^2 \sum_{i,j=1}^N (g_{i,j} - \bar{g})^2}}, \quad (2)$$

where $f_{i,j}$ and $g_{i,j}$ are pixel values from the distorted and the background image, respectively, and \bar{f} and \bar{g} are the mean intensities over the comparison window.

The cross-correlation results in a matrix $C(m, n)$ of correlation scores for each translation $m, n \in [-N + 1 \dots N - 1]$. In the neighborhood of the maximum of these correlation scores we fit 3-point Gaussians in the horizontal and vertical dimensions to the scores, which helps us locate the match with sub-pixel accuracy. We also compute a signal-to-noise ratio for each window as the ratio between the peak correlation value and the average correlation score

of all other pixels. We use this ratio as a reliability metric for filtering in a post-processing stage.

This basic algorithm is adapted to handle significant distortions by iterative refinement. The optical flow field estimated in one iteration is used to determine an affine transformation for each neighborhood. In the next iteration, each extracted window is warped with the corresponding transformation before computing the correlation scores. In addition, we adjust the window size from iteration to iteration, starting with a large window size for robustness, and ending with a small window size that better captures detail in the data.

The final vector field is improved by filtering. Spurious vectors that differ by more than a fixed threshold from the global mean or their local median vectors are removed and filled in by linear interpolation. Bilateral filtering is used to smooth the resultant field while preserving the reasonably sharp boundaries that do occur in practice. A Gaussian spatial smoothing kernel is modulated by another Gaussian in the vector magnitude range. The width of this second Gaussian is set on a per-vector basis according to the signal-to-noise ratio, so that highly reliable vectors are not smoothed over, while poor quality ones have their neighbors weighted heavily.

All stages of the optical flow algorithm are designed such that they can be implemented on a GPU. Processing times are on the order of a few minutes per frame for multiple iterations on 512×512 images with window sizes up to 64×64 . Figure 4 shows the optical flow for a plume of hot air and a 75 mm lens computed in this fashion.




Figure 4: Results from the BOS algorithm. Left: deflection field caused by a plume of hot air. Right: extreme ray deflections created by a (chipped) 75 mm lens, viewed from slightly off-axis.

The optical flow fields recovered with BOS can be used directly in rendering, for example to distort camera rays in environment mapping applications, or for deflecting light rays to compute caustics (see Figure 1, center and Section 7).

4.1 Computing Deflection Angles

The tomographic reconstruction algorithm described in the following section requires deflection angles rather than optical flow vectors as its input. We therefore have to convert the 2D vectors obtained with BOS into estimates of world-space angles. In a calibrated setup with known spacing between object and background pattern, this angle could be computed from the optical flow and the exact location of the point at which the deflected ray exits the volume under consideration. Unfortunately, the latter information is not readily available. A simple estimate used in the BOS literature is to approximate the precise location with the point where the original (undistorted) camera ray would exit the volume (Figure 5). This approximation is valid if the object diameter is significantly

smaller than the distance between object and background, which is the case for all our measurements.




Figure 5: The deflection direction can only be measured with two planes behind the object. An approximation is the direction from the intersection of the undeflected ray with the visual hull of the object to the position on the background plane that was measured. This approximation is valid as long as the extent of the object is considerably smaller than the distance to the background plane.

5 Tomographic Reconstruction

In the following, we discuss the tomographic reconstruction of a 3D refractive index field from the deflection angles measured with the BOS algorithm. We first derive the theory for a tomographic reconstruction of the gradient field that does not require the use of the paraxial approximation explained in Section 3. We then describe a practical implementation of this theory (Section 5.2), and finally describe how to integrate the gradient field to obtain the refractive index field (Section 5.3).

5.1 Gradient Field Tomography

The derivation of our reconstruction method starts from the Eikonal equation

$$(\nabla S)^2 = n^2, \quad (3)$$

which can itself be derived from Fermat's principle of least time or alternatively from the Maxwell equations. In this equation, S describes the time it takes for light to arrive at a particular point in space, and n is the refractive index as before. Unfortunately S is not a function since it is multi-valued: light can reach a point in space via multiple paths of varying length, causing the solution to branch at various loci in space. These places are known as caustics. To solve Equation 3 uniquely, the so-called viscosity solution is usually considered. This solution of the Eikonal equation describes the time of *first arrival* via any path at a particular point in space. The solution of Equation 3 depends on the initial and boundary conditions that are used to set the positions of light sources, or, via Helmholtz reciprocity, the location of cameras.

Since iso-surfaces of S describe regions of constant time of first arrival, and light rays are described by the path of least time, light 'particles' travel normal to the iso-surfaces of S , along ∇S . We therefore have

$$n \frac{d\mathbf{x}}{ds} = \nabla S, \quad (4)$$

where the factor n enters the equation because ∇S is not normalized. From Equation (3) it is clear that the magnitude of ∇S is n , the reciprocal of the relative speed of light in the medium.

We do not have to solve Equation (3) explicitly for S and then integrate using Equation (4). Instead, Equations (3) and (4) can be combined to obtain a differential equation for the particle position in terms of the refractive index n . For this we take the gradient of Equation (3):

$$\nabla(\nabla S \cdot \nabla S) = \nabla(n \cdot n) \quad (5)$$

end thus

$$2(\nabla \cdot \nabla S) \nabla S = 2n \nabla n. \quad (6)$$

Inserting Equation (4) together with $\nabla \cdot \frac{d\mathbf{x}}{ds} = \frac{d}{ds}$ yields

$$\frac{d}{ds} \left(n \frac{d\mathbf{x}}{ds} \right) = \nabla n. \quad (7)$$

Equation (7) is known as the *ray equation of geometric optics*. It is a second order ordinary differential equation that can be written as a set of first order ordinary differential equations.

Starting from Equation (7) and setting

$$n \frac{d\mathbf{x}}{ds} = \mathbf{d}, \quad (8)$$

we obtain

$$\frac{d\mathbf{d}}{ds} = \nabla n. \quad (9)$$

Observing that

$$\frac{dn}{ds} = \frac{dn}{d\mathbf{x}} \cdot \frac{d\mathbf{x}}{ds} = \nabla n \cdot \frac{d\mathbf{x}}{ds}, \quad (10)$$

we find that Equations (8) - (10) define a coupled system of ordinary differential equations for the path of a bent ray in a medium of non-uniform refractive index in terms of ∇n .

We integrate Equation (9) to obtain an equation that relates the unknown gradient of the refractive index field to our measurements, i.e. the outgoing ray directions

$$\mathbf{d}^{out} = \int_c \nabla n ds + \mathbf{d}^{in}. \quad (11)$$

Equation (11) forms the basis of our reconstruction scheme. We discretize the unknown vector function ∇n using a set of basis functions ϕ_i with unknown coefficient vectors \mathbf{n}_i ,

$$\int_c \sum_i \mathbf{n}_i \phi_i ds = \sum_i \mathbf{n}_i \int_c \phi_i ds = \mathbf{S} \mathbf{n}_i = \mathbf{d}^{out} - \mathbf{d}^{in}. \quad (12)$$

Solving this linear system of equations for the coefficient vectors \mathbf{n}_i allows us to recover the gradient of the refractive index field ∇n . Observe that Equation (12) is vector valued and describes three linear systems with three sets of coefficients, one for each coordinate of the gradient field.

Since the curved rays c are initially unknown, we develop an iterative solution that updates the ray geometry along with the volume estimate as described next.




Figure 6: The matrix entries are computed while performing Euler integration to determine the ray path with a fixed spatial step size. The values at the ray sampling positions (black dots) approximate the ray integral over the radially symmetric basis functions (blue).

5.2 Reconstruction Algorithm

Our reconstruction algorithm for the gradient of the refractive index field takes the following form:

1. compute curved rays c , Equation (7),
2. set up linear system, Equation (12),
3. solve unconstrained linear system for ∇n ,
4. until convergence go to step 1.

The initial guess for the gradient field is $\nabla n = 0$, i.e. we initially start from the straight ray geometry also assumed in the paraxial approximation. The rays are computed by a discretization of Equations (8) - (10). The ray integrals can be performed using any standard ODE integration scheme, such as the Runge-Kutta family. After recomputing the curved rays for the current volume estimate, we can also recompute the estimate of deflection angles from the optical flow data (Section 4.1).

The linear systems, Equation (12), are of the following form:

$$\begin{pmatrix} \mathbf{d}_1^{out} - \mathbf{d}_1^{in} \\ \mathbf{d}_2^{out} - \mathbf{d}_2^{in} \\ \vdots \\ \mathbf{d}_{n_m}^{out} - \mathbf{d}_{n_m}^{in} \end{pmatrix}^j = \begin{pmatrix} \int_{c_1} \phi_1 ds & \dots & \int_{c_1} \phi_{n_b} ds \\ \int_{c_2} \phi_1 ds & \dots & \int_{c_2} \phi_{n_b} ds \\ \vdots & \vdots & \vdots \\ \int_{c_{n_m}} \phi_1 ds & \dots & \int_{c_{n_m}} \phi_{n_b} ds \end{pmatrix} \mathbf{n}^j. \quad (13)$$

The index j refers to the coordinates x, y and z . We have one equation for every measurement \mathbf{d}^{out} ; \mathbf{d}^{in} is obtained from the camera calibration. The values n_m and n_b denote the number of measurements and the number of basis functions respectively. To set up the linear system we have to determine the line integrals for the line of sight of each pixel over the basis functions ϕ_i . We use radially symmetric Kaiser-Bessel basis functions for a high-quality reconstruction. The matrix entries are computed while performing the ray integration, see Figure 6. We are using a visual hull restricted tomography algorithm similar to the work by Ihrke and Magnor [2004] to solve the linear system in Equation (12) independently for each coordinate. Therefore, those columns of the matrix containing integrals over basis functions outside the visual hull [Laurentini 1994] are removed from the computations. However, in our case this cannot be done as a post-processing step as suggested by Ihrke and Magnor, because of the non-linear nature of the rays c . For this reason we have to compute the visual hull of the object before setting up the linear systems.

Similar to Ihrke and Magnor [2004], the linear systems can be solved using CGLS [Hansen 1998] if the number of cameras recording the scene is moderate. In the case of more than 8 cameras

it is necessary to resort to out-of-core algorithms because of the memory requirements for storing matrix \mathbf{S} . For larger numbers of viewpoints we use a re-formulation of the simultaneous algebraic reconstruction technique (SART) [Kak and Slaney 2001; Trifonov et al. 2006] in terms of matrix-vector products. The SART iteration can be performed on partial matrices \mathbf{S}_k , e.g. corresponding to equations generated by the measurements of one view. It takes the following form:

$$\mathbf{n}_{k+1}^j = \mathbf{n}_k^j + \mathbf{S}_k^T \frac{(\mathbf{d}^{out} - \mathbf{d}^{in})^j - \mathbf{S}_k \mathbf{n}_k^j}{m \text{diag}(\mathbf{S}_k \mathbf{S}_k^T)}, \quad (14)$$

where m is the number of equations in matrix \mathbf{S} , j is the coordinate index, and k is the iteration count for the tomographic reconstruction. The SART iteration has an intuitive explanation as the back-projected, rescaled residual error in the image plane, that is, the solution is updated with a rescaled version of the back-projected residual error. The matrix-vector formulation of SART readily accounts for the visual hull restriction that enables reconstructions of good quality from a small number of projections.

5.3 Integration of the Gradient Field

After the computation of ∇n , the final step is to find the refractive index field n from the gradient information. Computing n from ∇n is similar to computing a surface from its normals. We use the definition of the Laplacian operator

$$\Delta n = \nabla \cdot \nabla n \quad (15)$$

to compute n . The left hand side of Equation (15) is discretized and the right hand side of it is computed using our recovered ∇n . The resulting Poisson equation is solved for n . The Poisson equation can be solved most efficiently with multi-grid solvers. Since we have to perform the integration only once per frame we use a less efficient but easier to implement Jacobi-preconditioned Conjugate Gradient method [Barrett et al. 1994] to solve the linear system, Equation (15).

6 Physical Measurement Setup

We use two different camera setups for our experiments with real-world measurements (Figure 7):

- a single video camera setup that can be used to acquire 2D data for use as environment mattes, or for the reconstruction of rotationally symmetric flows. In this setup, we use a Prosilica black-and-white 1.5 megapixel C-MOUNT camera, and
- an array of 8 Imperx MDC-1004 video cameras with 1 megapixel resolution that can be used to capture non-symmetric flow for 3D reconstruction.

The Wavelet Noise patterns that we use as a background are printed either on paper, or on overhead transparencies with a laser printer. Since short camera exposure times are required to capture fast or very turbulent flows, it helps to use transparencies that are backlit by a lightbox or another bright and well-diffused light source. For slow or laminar flow, the patterns can be printed on paper and used in a reflective setting.



Figure 7: The single camera setup (left) and the camera array (right) that we use for acquisition.

We calibrate the external and internal camera parameters using the algorithms of Tsai’s [1987] or Zhang’s [1999] after removing radial lens distortion. For fluid imaging, we use a rectangular water tank with clear plastic walls of homogeneous thickness. The camera is located outside the tank, observing it through one of the walls. The ray bundle emerging from the perspective camera is refracted on a planar interface between water and tank, yielding a different perspective view with a virtual center of projection behind the location of the camera. This virtual center of projection and the associated intrinsic and extrinsic (virtual) camera parameters can again be found with Tsai’s algorithm.

7 Results and Applications

Our evaluation of the BOS algorithm and the tomographic reconstruction is based on both simulations with synthetic data, and real-world measurements.

BOS imaging on synthetic datasets. For ground truth experiments with the BOS algorithm, we applied known 2D flow fields to a Wavelet Noise pattern, and recovered the original flow from these distorted images. We found that we can obtain very good results (relative RMS error below 1%) for high local isotropic scaling factors and anisotropic scaling around 4 : 1. This is more than sufficient accuracy for the flow fields we are considering here, as well as for low-curvature solids such as the 75 mm lens from Figure 4. High curvature solids could result in larger distortions, which would require different methods.

BOS imaging on real measurements. We acquired BOS data for a large number of different flows. Figure 8 shows the displacement magnitude plots of a small selection from that set, including both laminar flows and more turbulent ones.

These kinds of datasets can be used directly in rendering, for example as environment mattes to distort the camera rays. Similarly, the data can be used to distort light rays, for example in a photon mapping algorithm. This will cast caustics on the receiving surface, which are also known as *shadowgraphs*. Real-world shadowgraphs are related to Schlieren imaging, and can also be used to image flows [Settles 2001]. Both the environment matting and the shadowgraph rendering are depicted in Figure 9.

Tomographic reconstruction of simulated data. To obtain quantitative results for the robustness of our tomographic reconstruction method, we ran the algorithm on synthetic data. We tested both smooth datasets and ones with high spatial frequencies, and analyzed the relative RMS error of the reconstruction depending on




Figure 8: Displacement magnitude images for various flows. Top left: low turbulence gas jet from a spray can. Top right: turbulent hot air rising from a candle lantern. Bottom left: turbulent interaction of a hot air plume above a candle with a jet of compressed air. Bottom right: turbulent mixture of water with corn syrup.




Figure 9: Environment matte and shadowgraph rendering of two BOS measurements. Left: hot air, Right: the chipped lens from Figure 4. Also see video.

the number of views and the differences in refractive index. For the smooth datasets, we used Gaussian blobs of different variance. Publicly available volume datasets with volume densities linearly mapped to refractive indices were used for the high-frequency analysis.

Table 1 shows the relative RMS error for the HIPIP dataset from the UNC CHVRTD volume collection for various configurations. We define relative RMS error as $\text{RMS}/\Delta n$, with $\Delta n = n_{\max} - n_{\min}$. In all tests we set $n_{\min} = 1$. As expected, the RMS error drops with the number of views used for reconstruction. However, even with only 8 views and a comparatively large Δn of 0.1, the RMS error across the volume is below 3% for refractive index fields, which indicates that our camera array is adequate for measurements of gases with reasonable precision. As expected, the reconstruction error increase with the difference in refractive index, but it does so slowly, which shows that our method can deal with a wide range of refractive indices. Values of Δn around 0.1 are already challenging for state-of-the-art methods requiring the paraxial approximation. For the kinds of datasets we explore, we expect values of Δn up to 0.15; higher values are common with fluid mixtures.

It is worth noting that most of the error is concentrated in regions of strong gradients, which appear blurred in the reconstruction. A larger number of views reduces this blur, and therefore the recon-

# Cameras ($\Delta n = 0.1$)	RMS Error	Δn (64 cams.)	RMS Error
8	2.89%	0.1	1.98%
16	2.69%	0.2	2.01%
32	2.38%	0.3	2.02%
64	1.97%	0.4	2.04%
128	1.54%	0.5	2.06%

Table 1: RMS error for the tomographic reconstruction of the synthetic HIPIP volume from optical flow images.

struction error.

Tomographic reconstruction of real measurements. On the right of Figure 10, we show a 3D reconstruction of a laminar candle plume from a single BOS image, assuming rotational symmetry. The BOS data (left) was projected from a total of 16 virtual camera positions. In the process, the slight asymmetries present in the BOS image are averaged out, so that the reconstructed volume is fully symmetric. The closing of the isosurfaces near the top is an artifact of the limited reconstruction volume since the top of the plume was outside the camera field of view.




Figure 10: Reconstruction of a rotationally symmetric flow from a single BOS image.

Figure 11 shows a few reconstructed frames from some of the real-world footage we captured with the camera array. The top two rows are different renderings of the same dataset: turbulent hot air rising from a gas burner. The first row shows a maximum intensity projection rendering of the refractive indices, and the second row an isosurface representation of the same information. Both sequences clearly show the advection of the hot air. We also encourage the reader to view the full sequence in the video.

The third row of Figure 11 shows a direct volume rendering of the *temperature distribution* in a plume of hot air above a candle. This flow is more laminar than the flow above the burner. These images are an example of the secondary information that can be extracted from the refractive index fields, depending on the dataset.

For gases, the refractive index is directly related to the volume density ρ through the Gladstone-Dale Equation:

$$\rho = \frac{n-1}{k(\lambda)}, \quad (16)$$

where k is the Gladstone-Dale constant, which depends on the material and weakly on the wavelength of light. For air and visible light, k has an approximate value of $0.23 \text{ cm}^3/\text{g}$.

Under certain assumptions, one can derive further quantities from the volumetric densities. For example, under constant pressure such as in the examples of hot air, the ideal gas law (see, e.g. [VanWylen




Figure 11: Results from the tomographic reconstruction process. Top: refractive index volume of turbulent hot air rising from a gas burner, rendered with maximum intensity volume rendering. Center: the same volumes rendered as isosurfaces. Bottom: direct volume rendering of the temperature distribution in a more laminar flow of hot air rising above a candle (for clarity, the volume rendering is restricted to the visual hull).

and Sonntag 1976]) can be used to infer the temperature distribution in the volume as

$$T = \frac{p \cdot M}{R \cdot \rho}, \quad (17)$$

where p is the pressure, M is the molecular mass of the gas ($\approx 29 \text{ g/mol}$ for air), and R is the gas constant ($\approx 8.31 \text{ J}/(\text{K} \cdot \text{mol})$).

In a similar fashion other secondary information can be inferred from the refractive index field. In the case of liquid mixtures, for example, the refractive index is an indicator of the local fluid concentrations. This kind of information could be useful for verifying fluid simulators, and furthering the understanding of certain types of flow in general.

8 Conclusions

In this paper we have presented a novel technique for capturing time-varying, inhomogeneous refractive index fields, such as the ones created by gas and liquid flows. Our major contributions are an improvement of the Background Oriented Schlieren method developed in fluid imaging, and a novel theory for tomographic reconstruction of 3D volumes from Schlieren images, which gives rise to a practical algorithm. With these methods, it is now possible to capture complex flows with very moderate hardware requirements. The data captured with this approach can be directly used in computer graphics for rendering camera distortions or caustics.

Maybe even more interesting in the long run is the fact that the recovered refractive index fields are indicators of other physical properties, such as the density or temperature distribution in gases, or concentrations of fluids in the scenario of liquid mixing. These and similar derived quantities could be excellent tools for furthering the understanding of fluids, and comparison of fluid simulations against reference data.

References

- AGARWAL, S., MALLICK, S., KRIEGMAN, D., AND BELONGIE, S. 2002. On Refractive Optical Flow. In *Proc. of ECCV*, 279–290.
- AGRAWAL, A., ALBERS, B., AND GRIFFIN, D. 1999. Abel inversion of deflectometric measurements in dynamic flows. *Applied Optics* 38, 15 (May), 3394–3398.
- BAKER, S., AND MATTHEWS, I. 2004. Lucas-Kanade 20 Years On: A unifying framework. *International Journal of Computer Vision* 56, 3 (Feb.), 221–255.
- BARRETT, R., BERRY, M., CHAN, T. F., DEMMEL, J., DONATO, J., DONGARRA, J., EIJKHOUT, V., POZO, R., ROMINE, C., AND DER VORST, H. V. 1994. *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods, 2nd Edition*. SIAM, Philadelphia, PA.
- BARRON, J., FLEET, D., AND BEAUCHEMIN, S. 1994. Performance of Optical Flow Techniques. *International Journal of Computer Vision* 12, 1 (Feb.), 43–77.
- BEN-EZRA, M., AND NAYAR, S. 2003. What does motion reveal about transparency? In *Proc. of ICCV*, vol. 2, 1025–1032.
- CHUANG, Y.-Y., ZONGKER, D. E., HINDORFF, J., CURLESS, B., SALESIN, D. H., AND SZELISKI, R. 2000. Environment matting extensions: Towards higher accuracy and real-time capture. In *Proc. of ACM SIGGRAPH*, 121–130.
- COOK, R. L., AND DEROSE, T. 2005. Wavelet Noise. *ACM Trans. Graphics (Proc. ACM SIGGRAPH)* 24, 3, 803–811.
- DALZIEL, S., HUGHES, G., AND SUTHERLAND, B. 2000. Whole-field density measurements by ‘synthetic schlieren’. *Experiments in Fluids* 28, 322–335.
- ELSINGA, G., OUDHEUSDEN, B., SCARANO, F., AND WATT, D. 2004. Assessment and application of quantitative schlieren

- methods: Calibrated color schlieren and background oriented schlieren. *Experiments in Fluids* 36, 309–325.
- FARIS, G., AND BYER, R. 1988. Three-dimensional beam-deflection optical tomography of a supersonic jet. *Applied Optics* 27, 24 (Dec.), 5202–5215.
- HANSEN, P. C. 1998. *Rank-Deficient and Discrete Ill-Posed Problems*. Society of Industrial and Applied Mathematics.
- HORN, B., AND SCHUNCK, B. 1981. Determining Optical Flow. *Artificial Intelligence* 17, 185–203.
- HOWES, W. L. 1984. Rainbow schlieren and its applications. *Applied Optics* 23, 4, 2449–2460.
- IHRKE, I., AND MAGNOR, M. 2004. Image-Based Tomographic Reconstruction of Flames. *Proc. of SIGGRAPH/Eurographics Symposium on Computer Animation* (June), 367–375.
- IHRKE, I., GOLDBLUECKE, B., AND MAGNOR, M. 2005. Reconstructing the Geometry of Flowing Water. In *Proc. of ICCV*, 1055–1060.
- KAK, A. C., AND SLANEY, M. 2001. *Principles of Computerized Tomographic Imaging*. Society of Industrial and Applied Mathematics.
- KUTULAKOS, K. N., AND STEGER, E. 2005. A Theory of Refractive and Specular 3D Shape by Light-Path Triangulation. In *Proc. of ICCV*, 1448–1455.
- LAURENTINI, A. 1994. The visual hull concept for silhouette-based image understanding. *IEEE Trans. on Pattern Analysis and Machine Recognition* 16, 2 (Feb.), 150–162.
- LUCAS, B., AND KANADE, T. 1981. An iterative image registration technique with an application to stereo vision. In *Proc. Seventh International Joint Conference on Artificial Intelligence*, 674–679.
- MATUSIK, W., PFISTER, H., ZIEGLER, R., NGAN, A., AND MCMILLAN, L. 2002. Acquisition and Rendering of Transparent and Refractive Objects. In *Proc. of the Eurographics Symposium on Rendering*, 267–278.
- MCMACKIN, L., HUGO, R., BISHOP, K., CHEN, E., PIERSON, R., AND TRUMAN, C. 1999. High speed optical tomography system for quantitative measurement and visualization of dynamic features in a round jet. *Experiments in Fluids* 26, 249–256.
- MEIER, G. 2002. Computerized Background-Oriented Schlieren. *Experiments in Fluids* 33, 181–187.
- MIYAZAKI, D., AND IKEUCHI, K. 2005. Inverse Polarization Raytracing: Estimating Surface Shapes of Transparent Objects. In *Proc. of CVPR '05*, vol. 2, 910–917.
- MORRIS, N. J. W., AND KUTULAKOS, K. N. 2005. Dynamic Refraction Stereo. In *Proc. ICCV*, 1573–1580.
- PEERS, P., AND DUTRÉ, P. 2003. Wavelet environment matting. In *Proc. of Eurographics Symposium on Rendering*, 157–166.
- RICHARD, H., AND RAFFEL, M. 2001. Principle and applications of the background oriented schlieren (BOS) method. *Measurement Science Technology* 12, 1576–1585.
- SCHARDIN, H. 1942. Die Schlierenverfahren und ihre Anwendungen. *Ergebnisse der Exakten Naturwissenschaften* 20, 303–439. English translation available as NASA TT F-12731, April 1970 (N70-25586).
- SCHWARZ, A. 1996. Multi-tomographic Flame Analysis with a Schlieren Apparatus. *Measurement Science Technology* 7, 406–413.
- SETTLES, G. S. 2001. *Schlieren and Shadowgraph Techniques*. Experimental Fluid Mechanics. Springer.
- TRIFONOV, B., BRADLEY, D., AND HEIDRICH, W. 2006. Tomographic Reconstruction of Transparent Objects. In *Proc. of Eurographics Symposium on Rendering 2006*, 51–60.
- TSAI, R. 1987. A versatile camera calibration technique for high accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses. *IEEE Journal of Robotics and Automation* 3, 4 (Aug).
- VANWYLEN, G., AND SONNTAG, R. 1976. *Fundamentals of Classical Thermodynamics*. John Wiley & Sons.
- VENKATAKRISHNAN, L., AND MEIER, E. 2004. Density measurements using the background oriented schlieren technique. *Experiments in Fluids* 37, 237–247.
- WEXLER, Y., FITZGIBBON, A., AND ZISSERMAN, A. 2002. Image-based environment matting. In *Proc. of Eurographics Symposium on Rendering*, 279–290.
- ZHANG, Z. 1999. Flexible camera calibration by viewing a plane from unknown orientations. international conference on computer vision. In *Proc. ICCV*, 666–673.
- ZONGKER, D., WERNER, D., CURLESS, B., AND SALESIN, D. 1999. Environment Matting and Compositing. In *Proc. of ACM SIGGRAPH*, 205–214.