

Reconstructing Gas Flows Using Light-Path Approximation

Yu Ji Jinwei Ye Jingyi Yu

University of Delaware, Newark, DE 19716, USA

{yuji, jye, yu}@cis.udel.edu

Abstract

Transparent gas flows are difficult to reconstruct: the refractive index field (RIF) within the gas volume is uneven and rapidly evolving, and correspondence matching under distortions is challenging. We present a novel computational imaging solution by exploiting the light field probe (LF-Probe). A LF-probe resembles a view-dependent pattern where each pixel on the pattern maps to a unique ray. By observing the LF-probe through the gas flow, we acquire a dense set of ray-ray correspondences and then reconstruct their light paths. To recover the RIF, we use Fermat’s Principle to correlate each light path with the RIF via a Partial Differential Equation (PDE). We then develop an iterative optimization scheme to solve for all light-path PDEs in conjunction. Specifically, we initialize the light paths by fitting Hermite splines to ray-ray correspondences, discretize their PDEs onto voxels, and solve a large, over-determined PDE system for the RIF. The RIF can then be used to refine the light paths. Finally, we alternate the RIF and light-path estimations to improve the reconstruction. Experiments on synthetic and real data show that our approach can reliably reconstruct small to medium scale gas flows. In particular, when the flow is acquired by a small number of cameras, the use of ray-ray correspondences can greatly improve the reconstruction.

1. Introduction

Accurately reconstructing transparent phenomena such as fluid wavefronts and gas flows remains as one of the most challenging problems in computer vision. The challenges are multi-folded. First, unlike opaque objects, transparent objects do not have their own images: they borrow appearance from nearby objects. Therefore traditional shading or texture based reconstruction solutions are not directly applicable. Second, compared with static transparent objects, transparent gas flows are even more difficult to reconstruct since the refractive index field (RIF) within the gas volume is uneven and rapidly evolving, introducing large image distortions. Finally, the acquisition system needs to be 1) non-intrusive to avoid affecting the dynamics and 2) portable to

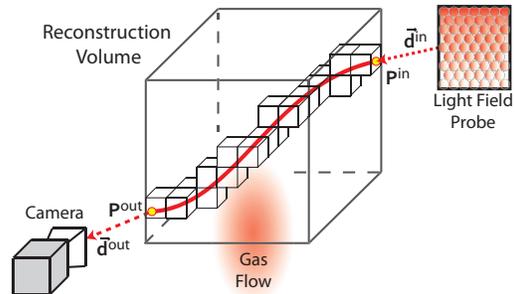


Figure 1. We acquire ray-ray correspondences to first estimate their light paths and then use them to recover the refractive index field (RIF) of the gas volume.

support on-site acquisition.

Most previous approaches are based on establishing point-pixel correspondences. Usually, a known reference pattern is placed near the transparent surface and robust tracking is applied to establish correspondences between a feature point on the pattern and its image in the camera. One then sets out to find the optimal surface that best matches the acquired correspondences. It is well-known that point-pixel correspondences are under-constrained even for single reflection or refraction. To resolve this ambiguity, additional constraints such as the planarity assumption [11, 15, 21], surface smoothness prior [25], surface integrability constraints [27], and most recently multi-view constraints [4, 20] need to be imposed. The seminal work of Atcheson *et al.* [2] uses 16 synchronized camcorders to capture the distorted wavelet noise patterns placed several meters away from the target flow. It then uses Background Oriented Schlieren (BOS) to measure deflections and applies tomographic reconstruction for recovering the gas flow. Their solution aims to reconstruct relatively large scale flows and their system tends to be bulky. We, in contrast, present a portable solution for acquiring small to medium scale gas flows.

Our solution exploits the light field probe (LF-Probe) [28, 29] which serves as a view-dependent reference pattern. A LF-probe, in essence, is an “inverted” light field camera [22] where each pixel on the pattern maps to a

unique ray. By acquiring the LF-probe through reflection-/refractions, one can establish ray-ray correspondences, *i.e.*, a pattern ray from the LF-probe will be mapped to a pixel ray in the camera, as shown in Fig. 1. Recent studies have shown that ray-ray correspondences greatly benefit specular surface reconstruction. For example, assuming single refraction through uniform media, Wetzstein *et al.* [29] and Ye *et al.* [30] use ray-ray correspondences to infer surface heights and normal fields of 2D dynamic wavefronts. We demonstrate how to use ray-ray correspondences for inferring light paths [16] through the RIF within the gas volume.

Under Fermat’s Principle, each light path corresponds to the shortest Optical Path Length (OPL). By using variational method, we show that each light path and the RIF is related via a Partial Differential Equation (PDE). We then develop an iterative optimization scheme to solve for all light-path PDEs in conjunction. Specifically, we initialize the light paths by fitting Hermite splines [8] to ray-ray correspondences, discretize their PDEs onto voxels, and solve a large, over-determined PDE system for the RIF. The RIF can then be used to refine the light paths. Finally, we alternate the RIF and light-path estimations to improve the reconstruction. Experiments on synthetic and real data show that our approach can reliably reconstruct small to medium scale gas flows. In particular, when the flow is acquired by a small number of cameras, the use of ray-ray correspondences can greatly improve reconstruction quality.

2. Related Work

Reconstructing transparent objects/phenomena such as fluids and gas flows is an important problem to oceanography and fluid mechanics and has recently attracted much attention from computer vision.

Static Reflective and Refractive Surfaces. Earlier approaches [3] have focused on modeling variations of reflection highlights for recovering surface geometry and reflective properties. Sankaranarayanan *et al.* [24] use point-pixel correspondences to first estimate the specular flow and then apply quadric approximations to recover mirror-type surfaces. A common issue in point-pixel based solutions is ambiguity: a pixel corresponds to a ray from the camera while the specular surface can lie at any position along the ray. Tremendous efforts have been focused on adding additional constraints [15, 21, 25, 27] for resolving this ambiguity. Bonfort and Sturm [4] use images captured by multiple calibrated cameras to reconstruct specular surface via space carving. Kutulakos and Steger [16] discover that by analyzing the piecewise linear light paths in homogeneous refractive medium, one can view surface reconstruction as a generalized triangulation problem. Their work showcases the usefulness of light paths. In this paper, we demonstrate using non-linear light paths for recovering inhomogeneous

refractive media.

Transparent Wavefronts. The problem of acquiring dynamic wavefronts is relatively new to computer vision. Morris and Kutulakos [20] track the corners of a checkerboard pattern over time in a stereo camera setting and then impose the refractive disparity constraint to iteratively solve for surface heights and normals. In a similar vein, Ding *et al.* [10] construct a camera array system to obtain multi-view point-pixel correspondences. Robustly tracking feature points, however, can be challenging as the observed image can exhibit severe distortions and motion blurs. Ye *et al.* [30] point out that to robustly recover light paths, it is important to establish ray-ray correspondences. For example, they propose using Bokode [19], a special pinhole projector, to acquire ray-ray correspondences for directly recovering fluid surface normals. Closest to our approach, Wetzstein *et al.* [29] replace the conventional checkerboard pattern with a LF-probe with 4D spatial and angular coding to obtain more accurate ray-ray correspondences. We study non-linear light paths using ray-ray correspondences.

Gas Flows. One of the most challenging transparent objects is 3D gas flows. In mechanical engineering, the widely adopted solution is Schlieren photography. Schardin [26] uses a knife edge to partially block rays proportional to their refracted directions to visualize dynamic gas flows, refractive solids, and shock waves. The original Schlieren photography, however, aims to visualize rather than reconstruct flows. Howes [12] modifies the traditional Schlieren to conduct quantitative evaluation of refractive index distribution by encoding the hue. Dalziel *et al.* [7] propose a much simpler and inexpensive solution called Background Oriented Schlieren (BOS) that images the refractive media via the distortion of a high-frequency background. It then calculates per-pixel deflection vectors using the optical flow. This class of methods [23, 18] can acquire the shifts of pattern positions due to refractions but not directions. The 3D tomography technique by Atcheson *et al.* [2] captures distorted wavelet noise patterns through the gas volume from multiple viewpoints. It then measures deflections caused by gas refraction to correlate the incident ray (*i.e.*, from the camera) to the exit point (*i.e.*, the feature point on the background pattern). The ray-point correspondences work well for distant patterns (*e.g.*, around 3 meters in their experiments) and is suitable for acquiring large scale gas flow. We develop a low-cost, portable solution for acquiring gas flows of small to medium scales.

3. Acquisition System

3.1. Light Field Probe (LF-Probe)

The core of our approach is to acquire ray-ray correspondences using the LF-probe [28, 29]. A LF-probe, as

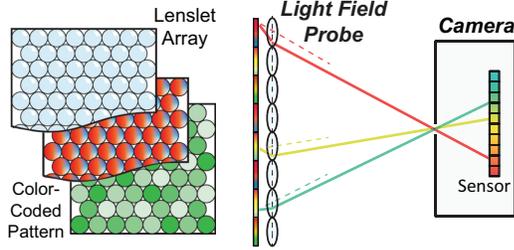


Figure 2. The LF-probe. Left: The assembly of a LF-probe. Right: Each pattern pixel maps to a unique ray.

shown in Fig. 2, can be viewed as an “inverted” light field camera Lytro (www.lytro.com). In Lytro, a microlenslet array is placed in front of the camera sensor to acquire the 4D light field, where the sensor-lens distance is identical to the microlens’ focal length. Each microlens serves as a virtual pinhole camera and the lenslet array serves as a camera array. The LF-probe keeps the same design except replacing the sensor with a specially designed color pattern. Uniform backlight is then used to illuminate the pattern and each pixel on the pattern maps to a unique ray.

To obtain a dense set of correspondences, similar to [28, 29], we use color-coded pattern to encrypt the 4D ray positions and directions emitted by the probe. In particular, we use a combination of horizontal red gradient and vertical blue gradient behind each microlens to discriminate rays of different directions. The red/blue gradients are identical for each microlens unit. To determine 2D positions, we use the variation of green channel. We first discretize the position to blocks based on microlenses by assuming the block behind each microlens has the same position as the center of the lens. We then randomly choose the green intensity for each block to form a random noise pattern as a whole. To find out the position shifted due to refraction, we perform optical flow between the refracted pattern and the original one on the green channel to estimate the deflection vectors. In this way, by matching color in the captured image, we can directly acquire the positions and directions of rays emitting from the LF-probe.

3.2. System Setup

Fig. 3 shows our gas flow acquisition system. We place 3 LF-probes to surround the target gas flow and 3 synchronized cameras to capture the corresponding LF-probe through the gas volume. Our goal is to first acquire a dense set of correspondences between rays entering the gas volume and the ones exiting the volume and then use these incident-exit ray pairs for estimating the light paths and the RIF.

Calibration. In order to reliably correlate the incident-exit ray pairs, we conduct two calibration procedures, one between each LF-probe and its viewing camera and the sec-

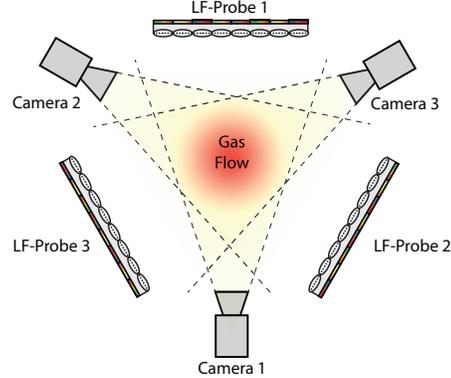


Figure 3. An illustration of our gas flow acquisition system.

ond between cameras. As discussed in Sec. 3.1, we determine ray-ray correspondences via color matching. Therefore we first calibrate colors between the printed LF-probe pattern and the observed image. To do so, we mount the LF-probe on a rotation table and then capture images of the probe at different viewing angles, as shown in Fig. 4(a). An additional checkerboard pattern is placed next to the probe for measuring the orientation of the probe w.r.t. the camera. Once we determine the direction β of a ray collected by the camera and the angle α between the probe’s normal and the camera’s principal axis, we can then compute the ray’s direction as $\gamma = \alpha + \beta$, as shown in Fig. 4(b). For each lenslet within the probe, we obtain a curve between each RGB channel and the directions (decomposed into vertical and horizontal components). A sample curve from a specific lenslet is shown in Fig. 4(c). By matching colors of red/blue channels for each lenslet, we can map each observed pixel to an incident ray direction. To determine the incident ray’s origin, we use a random noise pattern on the green channel so that we can associate a patch of pixels with a location. We obtain the ground truth by acquiring the LF-probe without any gas flows. When capturing the gas flows, we then use the optical flow for tracking the pattern.

Calibrating the exit rays is equivalent to calibrating the camera intrinsics and we directly apply Zhang’s algorithm [31]. Calibration between cameras is more challenging. Notice each viewing camera has a narrow Field-of-View (FoV) in order to capture a high resolution image of the LF-Probe. Therefore, the view frustums of the three cameras barely overlap. To resolve this issue, we place three additional cameras between the viewing cameras and conduct pair-wise camera calibrations. Once we finish the calibration process, each camera is able to acquire a dense set of ray-ray correspondences w.r.t. the LF-probe. For the rest of paper, we use $(\mathbf{P}^{in}, \vec{d}^{in})$ to represent rays emitting from the LF-Probes (the incident rays) and $(\mathbf{P}^{out}, \vec{d}^{out})$ for the rays entering the camera (the exit rays) as shown in Fig. 1, where \mathbf{P} and \vec{d} are the origin and direction of a ray respectively.

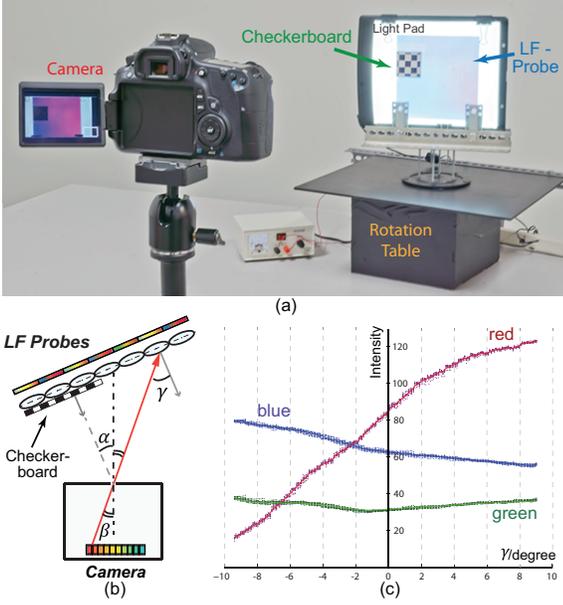


Figure 4. We use a rotation table (a) to calibrate ray directions (b) by matching observed colors to the color calibration curves (c).

4. Volumetric Gas Reconstruction

Given a dense set of ray-ray correspondences across the gas volume, our goal is to recover the RIF that best matches these correspondences. Previous approaches conduct light-path analysis via the Eikonal Equation [13, 14]. We instead derive how RIF is correlated with the light path using *Fermat's principle*: the light always travels along the path with the shortest *Optical Path length* (OPL) [5]. Assuming an arbitrary path c , the OPL S of c is computed as the weighted path length w.r.t. the refractive index n at every point $p(x, y, z)$ (or voxel in the discrete case) on the path c :

$$S = \int_c n(p) ds \quad (1)$$

We can further parameterize $p(x, y, z)$ as function of the time t that light reaches (x, y, z) as $p(x(t), y(t), z(t))$, then we have:

$$S = \int_c L dt, \quad L = n(p) \sqrt{x_t^2 + y_t^2 + z_t^2} \quad (2)$$

where $x_t = \frac{\partial x}{\partial t}$, $y_t = \frac{\partial y}{\partial t}$, $z_t = \frac{\partial z}{\partial t}$ and $n(p)$ is the refractive index at p . L is often referred to as the optical Lagrangian [17].

Our goal is to solve for both the light path c and the RIF n . Base on the Fermat's Principle, each light path c corresponds to the shortest OPL. In the context of calculus of variations, this can be written as $\delta S = 0$ or:

$$\delta \int_c L dt = 0 \quad (3)$$

4.1. RIF Estimation

If we have the light paths, we can then estimate the RIF. By Eqn. 3, L should satisfy the Euler-Lagrange equation:

$$\left(\frac{\partial L}{\partial x}, \frac{\partial L}{\partial y}, \frac{\partial L}{\partial z} \right) = \frac{d}{dt} \left(\frac{\partial L}{\partial x_t}, \frac{\partial L}{\partial y_t}, \frac{\partial L}{\partial z_t} \right) \quad (4)$$

Let us only consider the x component. By substituting L of Eqn. 2 into Eqn. 4 and changing the order of derivatives, we have:

$$\frac{\partial n(p)}{\partial x} \sqrt{x_t^2 + y_t^2 + z_t^2} = \frac{d}{dt} \left(n(p) \frac{x_t}{\sqrt{x_t^2 + y_t^2 + z_t^2}} \right) \quad (5)$$

Recall that $x_t / \sqrt{x_t^2 + y_t^2 + z_t^2}$ is the x component normalized direction at position p .

Assuming the light ray reaches the gas volume at $t = t_0$ and leaves at $t = t_1$. Since at each time instance, the shortest OPL constraint should be satisfied, we can integrate Eqn. 5 from t_0 to t_1 :

$$\int_{t_0}^{t_1} \frac{\partial n(p)}{\partial x} \sqrt{x_t^2 + y_t^2 + z_t^2} dt = n(p) \left(\frac{x_t}{\sqrt{x_t^2 + y_t^2 + z_t^2}} \right) \Bigg|_{t_0}^{t_1} \quad (6)$$

Substituting the incident-exit ray pair: $(\mathbf{P}^{in}, \vec{d}^{in}; \mathbf{P}^{out}, \vec{d}^{out})$ into Eqn. 6, we have:

$$\int_{\mathbf{P}^{in}}^{\mathbf{P}^{out}} \frac{\partial n(p)}{\partial x} ds = n(\mathbf{P}^{out}) \vec{d}_x^{out} - n(\mathbf{P}^{in}) \vec{d}_x^{in} \quad (7)$$

Similar derivations hold for the y and z components of p . To recover the RIF, we discretize 3-D space into voxels and estimate the refractive index at each voxel. Specifically, we can predefine the gas volume and discretize Eqn. 7 for each light path c as:

$$\sum_c l(p) \begin{pmatrix} n(p_{x+1,y,z}) - n(p_{x,y,z}) \\ n(p_{x,y+1,z}) - n(p_{x,y,z}) \\ n(p_{x,y,z+1}) - n(p_{x,y,z}) \end{pmatrix} = n_{air} \begin{pmatrix} \vec{d}_x^{out} - \vec{d}_x^{in} \\ \vec{d}_y^{out} - \vec{d}_y^{in} \\ \vec{d}_z^{out} - \vec{d}_z^{in} \end{pmatrix} \quad (8)$$

where $l(p)$ is the Euclidean distance of a ray inside the voxel p . For voxels at the bounding faces of the gas volume, we assume their refractive indices are equal to $n_{air} = 1.000293$, *i.e.*, the boundary condition for the PDEs. Using all light paths, we form a PDE system of the RIF. We further impose an additional constraint that $\forall p, n(p) > 1$, and solve this over-determined PDE system using the Pre-conditioned Conjugate Gradient (PCG) method.

The RIF estimation method presented above requires knowing the light paths. However, we only have ray-ray correspondences at the first iteration. Therefore, we initialize light paths by fitting a Hermite spline [8] to each ray-ray correspondence $(\mathbf{P}^{in}, \vec{d}^{in}; \mathbf{P}^{out}, \vec{d}^{out})$. We choose the Hermite spline because it directly accounts for the locations and the tangent directions at two control points that are directly provided by the ray-ray correspondence.

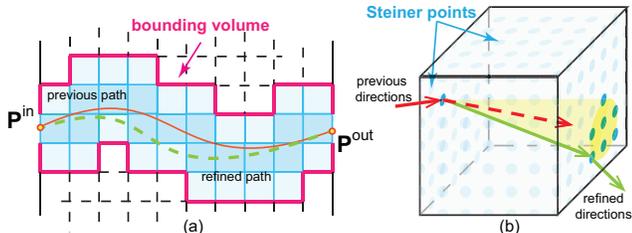


Figure 5. Accelerating light-path refinement via volume pruning. (a) We prune voxels outside the bounding volume; (b) We restrict the search direction within a cone.

4.2. Light-Path Refinement

Once we obtain the initial estimation of the RIF, we refine the light paths within the gas volume using Fermat’s Principle. Specifically, we set out to find the shortest OPL from \mathbf{P}^{in} to \mathbf{P}^{out} . This is a classical shortest path problem on polyhedra and NP-hard [6]. Our approach is to approximate the solution by mapping the problem to a planar graph so that existing efficient algorithms such as the Dijkstra’s [9] can be directly used. It is important to note that constructing the graph is non-trivial. A brute-force construction is to directly treat the centroid of each voxel as a node and impose six-direction connectivity. However, such graph approximation does not consider how much the distance light travels *inside* each voxel. We therefore include two types of additional nodes: 1) the corners of voxels and 2) the Steiner points [1] on the faces and edges of voxels. In our experiments, we generate 48×48 Steiner points on each face and 48 on each edge. The weight between two nodes is computed as: $w = n \cdot l$, where n is the estimated refractive index of the voxel and l is the Euclidean distance between the nodes.

The resulting graph, however, contains too many links and directly applying the Dijkstra’s algorithm is computational expensive. For example, if we discretize the volume to $32 \times 32 \times 32$ voxels, we obtain a graph of 10^6 nodes. To speed up shortest path computation, we aggressively prune the nodes based on the observation that light paths only slightly deviate from a linear path. Specifically, we impose a “bounding volume” along the previously estimated light path as shown in Fig. 5(a) and only conduct searching within the bounding volume. Furthermore, inside each voxel, we prune a large amount of corners/Steiner points from the search by assuming that the light direction will not change drastically after the refinement. Specifically, we fit a cone of 1° with the incident direction into the voxel as its central axis and the start point as its apex and only consider those nodes falling within the cone, as shown in Fig. 5(b).

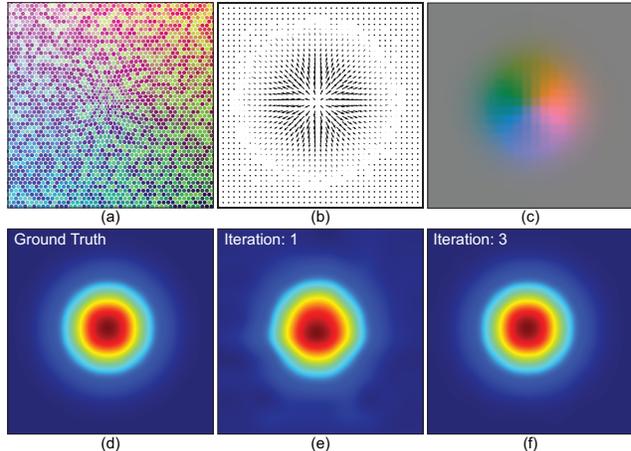


Figure 6. A simple Gaussian gas flow. (a) A rendered LF-probe image through the flow; (b) Optical flow results on the green channel; (c) Directional variations of rays; (d) The ground truth RIF; (e) Our estimated RIF after 1 iteration; (f) Our final result.

5. Experiments

We have validated our approach on both synthetic and real data. All experiments are conducted on a PC with an Intel i7-3930K CPU (3.20GHz 6-core) and 16G memory.

5.1. Synthetic Scene Simulations

We first test our solution on a static gas volume whose RIF follows Gaussian distribution: $n(x, y, z) = n_{air} - (n_{air} - 1)e^{-((x-x_0)^2 + (y-y_0)^2 + (z-z_0)^2)/2}$, where (x_0, y_0, z_0) is the center of flow. We discretize the gas volume to $23 \times 23 \times 23$ voxels. The resulting RIF has 12,167 unknowns. We emulate a LF-probe with a lenslet array of 60×60 microlenses and color pattern as described in Sec. 3.1. To capture the LF-probe image, we have implemented a voxel-based Ray-tracer that can trace along non-linear light paths within the volume. Fig. 6(a) shows a sample of rendered LF-probe image. If we assume that each microlens is seen once by the viewing camera, each LF-probe provides at least 3,600 ray-ray correspondences. In this synthetic scene, we use three LF-probes surrounding the gas volume to mimic the real setup. This leads to a total of 10,800 ray-ray correspondences. By Eqn. 8, we form three equations (for the x, y, and z components respectively) for each correspondence and we have 32,400 equations in total for 12,167 unknowns.

Since our synthesized images are noise-free, we directly map the red/blue channels to incident ray directions and apply optical flow on the green channel to determine incident ray origins. The measured optical flows and angular variations are shown in Fig. 6(b) and (c). Once we acquire the ray-ray correspondences, we apply the reconstruction algorithm described in Sec. 4 to iteratively estimate the light paths and the RIF. We stop the iteration when the change

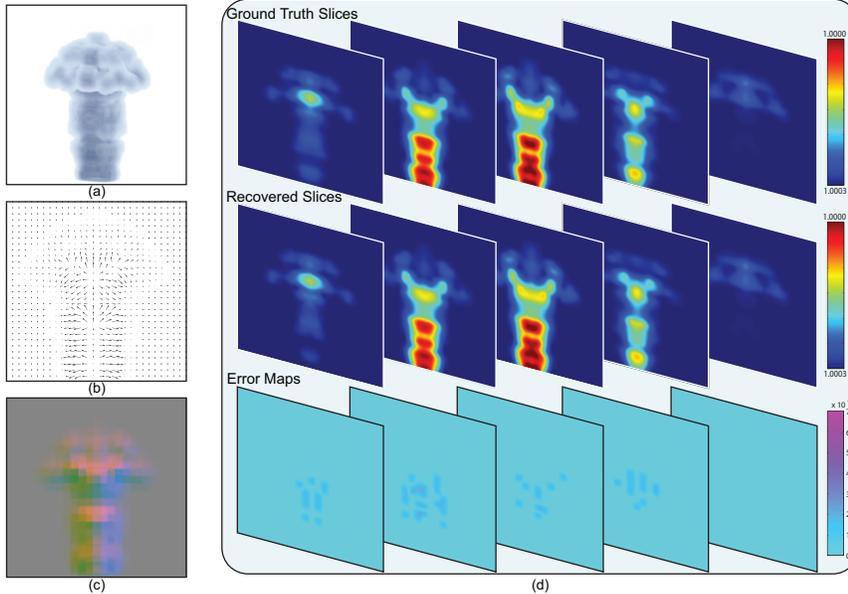


Figure 7. The fuel injection dataset. (a) A volume rendering of the ground truth data; (b) Optical flow results on the green channel; (c) Directional variations of rays; (d) Slices of our reconstructed RIF vs. the ground truth.

from the previous iteration is below a threshold. In this example, our estimated RIF converges after 3 iterations and each iteration takes around 5 minutes. To visualize our results, we show the central slice of the reconstructed RIF in Fig. 6(f). To evaluate accuracy, we use two error measures: the Averaged Percentage Error (APE) and the Maximum Percentage Error (MPE) over all voxels. Our reconstruction achieves high accuracy: APE = 1.19%, MPE = 1.74%.

Next, we test on a more realistic fuel injection dataset from DFG SFB 382 (www.volvis.org). Since this volume data has more complex geometric structures, we discretize the volume into $32 \times 32 \times 32$ voxels. We also increase the number of microlenses in the lenslet array to 160×160 in order to acquire dense enough ray-ray correspondences. We apply the same approach to iteratively recover the RIF and the light paths. In this example, the RIF converges after 5 iterations. Fig. 7(d) compares our reconstruction results with the ground truth data. Our reconstruction produces low error throughout the volume (APE = 1.84%, MPE = 5.29%).

Finally, we compare our ray-ray vs. traditional point-pixel correspondence based solutions. Specifically, we have implemented the multi-camera BOS gas flow reconstruction algorithm [2] and conducted comparisons under various configurations: w.r.t the number of camera-pattern pairs and gas-pattern distances. When testing different numbers of cameras, we fix the gas-pattern distance to $5 \times$ the gas volume size; when studying the impact of gas-pattern distance, we use 4 camera-pattern pairs. The results are reported in Table 1. With even a small number of camera-LF-Probe pairs, our approach produces accurate reconstruction. In contrast, BOS works well when using a large number of

| Errors w.r.t the Number of Camera-Pattern Pairs | | | | | |
|---|-----|--------|--------|--------|--------|
| # of cameras | | 3 | 4 | 8 | 16 |
| Atcheson <i>et al.</i> [2] | APE | 5.09% | 4.61% | 3.67% | 3.56% |
| | MPE | 17.52% | 13.78% | 12.78% | 12.96% |
| Ours | APE | 1.84% | 1.13% | 0.93% | 0.92% |
| | MPE | 5.29% | 3.20% | 2.57% | 2.55% |

| Errors w.r.t the Pattern-Gas Distance | | | | | |
|---------------------------------------|-----|------------|------------|------------|-------------|
| Distance [†] | | 2 \times | 4 \times | 8 \times | 16 \times |
| Atcheson <i>et al.</i> [2] | APE | 10.34% | 4.82% | 2.08% | 1.19% |
| | MPE | 25.73% | 14.51% | 10.49% | 3.33% |
| Ours | APE | 1.28% | 1.09% | 1.12% | 1.08% |
| | MPE | 2.93% | 2.77% | 3.04% | 2.85% |

[†] refers to the pattern to gas distance in units of the gas volume size.

Table 1. Reconstruction errors using our approach vs. [2] under different configurations. (APE: Averaged Percentage Error. MPE: Maximum Percentage Error.)

camera-pattern pairs and distant patterns. Its performance drops with either fewer camera-pattern pairs or closer patterns. This is because point-pixel correspondences are insufficient for determining the light paths unless the pattern is placed far away. However, distant patterns will cover smaller FoV and therefore more cameras would be needed.

5.2. Real Scene Experiment

Fig. 8 shows our acquisition system. We use 3 Pointgrey Flea2 cameras to capture the LF-probes through the gas. To assemble a LF-probe, we use a FresnelTech

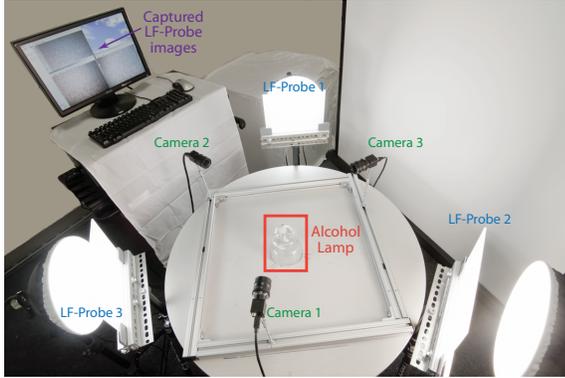


Figure 8. Our LF-probe based gas flow acquisition system.

(www.fresneltech.com) hexagonal lenslet array sheet with 100×90 microlenses, each having a diameter 0.09" and focal length 0.12". Our color pattern is printed onto a transparency using a Canon PIXMA Pro9000 inkjet photo printer with 1200 dpi. Our probe has a FoV of 18.92° and an angular resolution of $18.92^\circ/48 = 0.39^\circ$ under 48 color tones suitable for acquisition by an 8-bit sensor. All cameras are synchronized and capture at a resolution of 1024×768 at 30fps. Each camera uses a lens with focal length 16mm to cover a FoV of 22° horizontally and 18° vertically. A small aperture is used to reduce defocusing. Since the gas flow is fast-evolving, we use a fast shutter of $1/320$ s. To reduce image noise due to fast shutter, we use ultra bright illuminations. Specifically, we place a lamp of 250W (6500 lumen) covered by a diffuser behind the pattern. To generate real 3D gas flows, we use an alcohol lamp whose flame temperature can reach around 600°C . We recover a cube right above the flame of size $10\text{cm} \times 10\text{cm} \times 10\text{cm}$ and discretize the volume into $32 \times 32 \times 32$ voxels. To reduce noise, we estimate each incident ray's direction and position by averaging over a patch of 5×5 pixels.

Fig. 9 shows our reconstructed gas flows at three different time instances. For each time instance, we also show one of the three LF-probe images and its corresponding optical flow and angular variations (Fig. 9(a) and (b)). The gas flow inside the volume is highly inhomogeneous. We then show 3 vertical slices for each time instance (Fig. 9(c)-(e)) to illustrate the RIF distribution inside the volume. Our reconstructed RIF indicates that the central portion of the volume has a lower refractive index, *i.e.*, higher temperature, which is consistent with the physical model. Additional results (the actual deflections, different volume slicing, *etc.*) can be found on the project web page www.eecis.udel.edu/~yuji/GasFlow. We refer the reader to our supplementary videos for the complete sequence of reconstructed RIF.

An important application of our solution is to synthesize

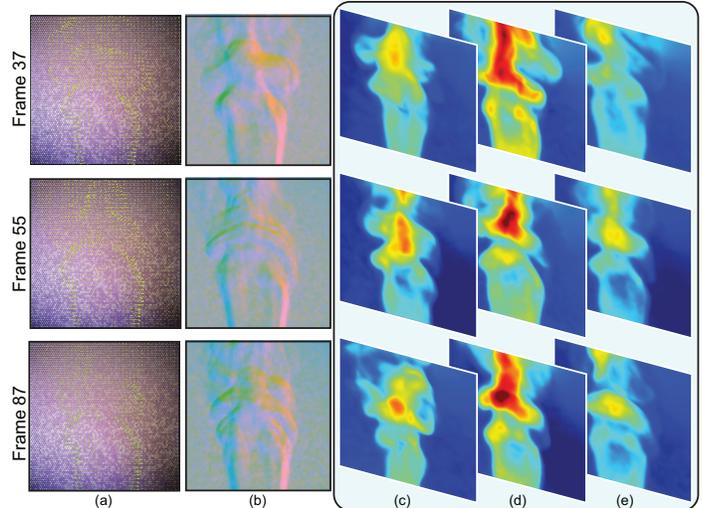


Figure 9. Reconstruction results on a real gas flow. Each row shows the reconstructed flow at a different time instance. (a) The captured LF-probe images; (b) Measured ray direction variations through the flow; (c)-(e) Three vertical slices of the reconstructed RIFs where (d) is the central slice.

gas motions and gas appearance for computer animation. Conventional physical-based fluid animation techniques are generally complex and computationally intensive. Our acquired results are real and can be directly composed into still images or film footage to achieve realistic visual effects. In Fig. 10, we compose the reconstructed flows from Fig. 9 onto a new background of a pool scene. The results are generated by ray-tracing through the RIFs. Our approach vividly synthesizes how heat distorts the background. The entire sequence can be found in our supplementary videos.



Figure 10. We compose the reconstructed flows onto a new background.

6. Discussions and Conclusions

We have presented a new computational imaging solution for reconstructing dynamic 3D gas flows. We use the LF-probe for generating view-dependent features: by coupling the LF-probe with a viewing camera, we establish reliable ray-ray correspondences. We have then used these

correspondences to iteratively estimate the light paths and the refractive index field of the gas volume. Experiments on synthetic and real data have shown that our approach provides a portable and reliable solution for reconstructing dynamic and inhomogeneous gas flows of small to medium scales.

One limitation of our approach is its sensitivity to noise. We use a fast shutter to capture fast evolving flows. To reduce image noise, we illuminate the LF-probes with ultra-bright light sources. In our experiments, we still observe color noise in the acquired LF-probe images and have to apply denoising filters. The LF-probe we use has a rather small FoV. As a result, we can only detect ray direction variations within a small range (e.g., $\pm 15^\circ$). This has limited our approach to small to medium scale flows. The issue can be alleviated by using microlens arrays with a wider FoV. On the reconstruction front, our light-path refinement step is still computationally expensive despite aggressive node pruning. A GPU-based solution may significantly accelerate our scheme for time-sensitive applications.

There are a number of future directions we plan to explore. Our solution can benefit from denser ray-ray correspondences. We plan to use the light field camera such as Lytro (www.lytro.com) or Raytrix (www.raytrix.de) for acquiring the LF-probe. Compared with our current setting that each microlens on the LF-probe generates a single ray-ray correspondence, the new setup will map each pixel of the microlens pattern to a ray-ray correspondence. Our solution by far has not considered fluid dynamics: we reconstruct each frame individually. In the future, we will investigate using fluid dynamic models as the Navier-Stokes to enforce temporal coherence. At the same time, we can use the acquired gas flows to improve fluid simulations for producing more realistic animations.

Acknowledgements

This project was partially supported by the National Science Foundation under grants IIS-CAREER-0845268 and IIS-RI-1016395, and by the Air Force Office of Science Research under the YIP Award.

References

- [1] L. Aleksandrov, A. Maheshwari, and J.-R. Sack. Approximation algorithms for geometric shortest path problems. In *ACM symposium on Theory of Computing*, 2000.
- [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 SIGGRAPH Asia*, 2008.
- [3] A. Blake. Specular stereo. In *Proc. of international joint conference on Artificial intelligence*, pages 973–976, 1985.
- [4] T. Bonfort and P. Sturm. Voxel carving for specular surfaces. In *ICCV*, 2003.
- [5] M. Born and E. Wolf. *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. Pergamon Press, New York, 1959.
- [6] J. Canny and J. Reif. New lower bound techniques for robot motion planning problems. In *28th IEEE Symposium on Foundations of Computer Science*, pages 49–60, Oct. 1987.
- [7] S. B. Dalziel, G. O. Hughes, and B. R. Sutherland. Whole-field density measurements by synthetic schlieren. *Experiments in Fluids*, 28:322–335, 2000.
- [8] C. De Boor. *A practical guide to splines*. Springer, revised edition, 2001.
- [9] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959.
- [10] Y. Ding, F. Li, Y. Ji, and J. Yu. Dynamic fluid surface acquisition using a camera array. In *ICCV*, 2011.
- [11] Y. Ding and J. Yu. Recovering shape characteristics on near-flat specular surfaces. In *CVPR*, 2008.
- [12] W. L. Howes. Rainbow schlieren and its applications. *Appl. Opt.*, 23(14):2449–2460, Jul 1984.
- [13] I. Ihrke. *Reconstruction and Rendering of Time-Varying Natural Phenomena*. PhD thesis, Universität des Saarlandes, Department of Computer Science, 2007.
- [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. on Graphics (SIGGRAPH'07)*, pages 59–1 – 59–9, 2007.
- [15] K. Ikeuchi. Determining surface orientations of specular surfaces by using the photometric stereo method. *IEEE TPAMI*, 3(6):661–669, nov. 1981.
- [16] K. Kutulakos and E. Steger. A theory of refractive and specular 3D shape by light-path triangulation. In *ICCV*, 2005.
- [17] V. Lakshminarayanan, A. Ghatak, and K. Thyagarajan. *Lagrangian Optics*. Springer, Nov. 2001.
- [18] G. Meier. Computerized background-oriented schlieren. *Experiments in Fluids*, 33:181–187, 2002.
- [19] A. Mohan, G. Woo, S. Hiura, Q. Smithwick, and R. Raskar. Bokode: imperceptible visual tags for camera based interaction from a distance. In *ACM SIGGRAPH*, 2009.
- [20] N. Morris and K. Kutulakos. Dynamic refraction stereo. In *ICCV*, 2005.
- [21] H. Murase. Surface shape reconstruction of an undulating transparent object. In *ICCV*, 1990.
- [22] R. Ng, M. Levoy, M. Brédif, G. Duval, M. Horowitz, and P. Hanrahan. Light field photography with a hand-held plenoptic camera. Computer Science Tech Report CSTR 2005-02, Stanford University, Apr 2005.
- [23] M. Raffel, H. Richard, and G. E. A. Meier. On the applicability of background oriented optical tomography for large scale aerodynamic investigations. *Experiments in Fluids*, 28:477–481, 2000.
- [24] A. Sankaranarayanan, A. Veeraraghavan, O. Tuzel, and A. Agrawal. Specular surface reconstruction from sparse reflection correspondences. In *CVPR*, 2010.
- [25] S. Savarese and P. Perona. Local analysis for 3D reconstruction of specular surfaces. In *CVPR*, 2001.
- [26] H. Schardin. Die schlierenverfahren und ihre anwendungen. In *Ergebnisse der exakten naturwissenschaften*, volume 20, pages 303–439. 1942.
- [27] M. Tarini, H. P. A. Lensch, M. Goesele, and H.-P. Seidel. 3D acquisition of mirroring objects using striped patterns. *Graph. Models*, 67:233–259, July 2005.
- [28] G. Wetzstein, R. Raskar, and W. Heidrich. Hand-held schlieren photography with light field probes. In *ICCP*, 2011.
- [29] G. Wetzstein, D. Roodnick, R. Raskar, and W. Heidrich. Refractive Shape from Light Field Distortion. In *ICCV*, 2011.
- [30] J. Ye, Y. Ji, F. Li, and J. Yu. Angular domain reconstruction of dynamic 3d fluid surfaces. In *CVPR*, 2012.
- [31] Z. Zhang. A flexible new technique for camera calibration. *IEEE TPAMI*, 22(11):1330–1334, Nov. 2000.