

Indoor Airflow Imaging Using Physics-Informed Schlieren Tomography

Teh, Arjun; Ali, Wael H.; Rapp, Joshua; Mansour, Hassan

TR2025-035 March 08, 2025

Abstract

Remote temperature sensing of volumetric flows has a variety of applications, such as promoting thermal comfort, heat dissipation, or data center cooling. The emergence of background-oriented schlieren (BOS) imaging in recent years has enabled transparent flow visualization at minor costs. In this paper, we develop a framework for non-invasive volumetric indoor airflow estimation from a single viewpoint using BOS measurements and physics-informed reconstruction. Our framework utilizes a light projector that projects a pattern onto a target back wall and a camera that observes small distortions in the light pattern due to the change in the refractive index of the air as a result of the temperature variation. While the single-view BOS tomography problem is severely ill-posed, we regularize the reconstruction using a physics-informed neural network (PINN) that ensures that the reconstructed airflow is consistent with the coupled Boussinesq approximation of the incompressible Navier–Stokes and the heat transfer equations.

IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP) 2025

Indoor Airflow Imaging Using Physics-Informed Schlieren Tomography

Arjun Teh*
Carnegie Mellon University
Pittsburgh, PA, USA
ateh@cs.cmu.edu

Wael H. Ali, Joshua Rapp, Hassan Mansour
Mitsubishi Electric Research Laboratories
Cambridge, MA, USA
{wali, rapp, mansour}@merl.com

Abstract—Remote temperature sensing of volumetric flows has a variety of applications, such as promoting thermal comfort, heat dissipation, or data center cooling. The emergence of background-oriented schlieren (BOS) imaging in recent years has enabled transparent flow visualization at minor costs. In this paper, we develop a framework for non-invasive volumetric indoor airflow estimation from a single viewpoint using BOS measurements and physics-informed reconstruction. Our framework utilizes a light projector that projects a pattern onto a target back wall and a camera that observes small distortions in the light pattern due to the change in the refractive index of the air as a result of the temperature variation. While the single-view BOS tomography problem is severely ill-posed, we regularize the reconstruction using a physics-informed neural network (PINN) that ensures that the reconstructed airflow is consistent with the coupled Boussinesq approximation of the incompressible Navier–Stokes and the heat transfer equations.

Index Terms—background-oriented schlieren, physics-informed neural network, refractive index, ray tracing, airflow sensing

I. INTRODUCTION

Understanding the airflow in indoor spaces is crucial for improving the comfort and efficiency of heating, ventilation, and air conditioning (HVAC) systems [1]–[4]. However, three-dimensional (3D) airflow sensing is challenging since hardware sensors only measure localized spatial regions around the sensors [5], and computationally-aided systems rely on expensive computational fluid dynamics (CFD) simulations to predict airflow regimes [6]. Alternatively, existing imaging techniques require expensive, precise optics for schlieren imaging [7] or laser-induced fluorescence [8], or the injection of particles for particle image velocimetry (PIV) [9].

One promising development is background-oriented schlieren (BOS) tomography, which uses images of a patterned background to observe distortions due to changes in the refractive index of a transparent medium [10]–[14]. While BOS has been effective for quantitative measurement of gas flows with high spatial resolution, the tomographic inverse problem is only well-posed with a sufficient number of view angles to resolve spatial ambiguities. For instance, the experiments by Atcheson et al. [15] use 16 cameras and a large background on the order of several meters for a small $30 \times 15 \times 15$ cm measurement volume. Other experiments likewise use multiple cameras spread across large view angles to capture small volumes of interest [16]–[18] or assume axisymmetric flows [19], [20]. These configurations are impractical for room-scale scenes, as the placement of multiple cameras and textured backgrounds would be highly invasive. Recent work by Zhao et al. [21] uses an implicit neural representation (INR) of a refractive index field and demonstrates that the field can be unambiguously recovered from single-view measurements, so long as the pattern is embedded within the 3D volume. However, if the pattern is constrained to a plane behind the refractive volume, as in standard BOS measurements, the implicit regularization of the network is not sufficient to correct for the inherent ambiguity of the single-view.

*Arjun Teh conducted this work during an internship at MERL.




Fig. 1: Illustration of the imaging setup where a projector illuminates the back-wall with a textured pattern and a camera observes the distorted pattern as a result of density changes within the airflow.

In this paper, we combine several recent advances in refractive field reconstruction to achieve 3D airflow reconstructions for room-scale scenes from a single camera view. First, we integrate the approach of Zhao et al. [21] into a physics-informed neural network (PINN) framework [22], which adds the partial differential equations governing airflow as regularization. PINNs have previously been used for 3D BOS tomography, but only in multi-view [18] or axisymmetric scenarios [23]–[25]. The governing equations plus boundary information alone are not enough to predict the airflow, but combined with BOS measurements they sufficiently constrain the ill-posed optimization. Additionally, we employ improved ray tracing according to the refractive radiative transfer equation (RRTE) [26]. Finally, we explore the use of a light source projecting a pattern onto the wall [27], which may be more practical than a fixed background. The projector allows for longer path lengths through the refractive index field, and a spatial offset between the camera and projector encodes more angular information into the single-view setup.

II. BOS IMAGING FORMULATION

A. Airflow Imaging Setup

We consider a BOS imaging scenario comprising an air-filled room, a camera, and either a patterned background wall or a light source that can project a pattern on the back-wall, as shown in Fig. 1. An inlet on the side wall injects an airflow into the room with a temperature that differs from the ambient temperature. When no air is flowing, the camera captures a reference image I_{ref} of the back-wall pattern. When the inlet blows the airflow into the room, the change in density of the air induces a gradient in its refractive index η , causing light rays passing through the air to bend. A classical BOS measurement computes a displacement in the pixels between the reference image I_{ref} and the image I_{flow} obtained of the back-wall in the presence of the airflow. In contrast, we adopt a ray tracing framework, similar to [21], where luminance from the back-wall is traced through an estimate of the refractive field and compared to the flow image I_{flow} .

Let Φ denote the 3D airflow volume that is parameterized by the temperature, pressure, and velocity fields, i.e., $\Phi := (T, p, \mathbf{u})$, where T and p are scalar fields, and \mathbf{u} is a vector field in \mathbb{R}^3 . For a gas, the refractive index depends linearly on the density ρ of the gas via the Gladstone–Dale equation $\eta = 1 + G\rho$, where G is the Gladstone–Dale coefficient [28]. Assuming the pressure variation of air in a room is small, then by the ideal gas law, the refractive index is related to temperature using

$$\eta(T) = 1 + \rho_0 G \frac{T_0}{T}, \quad (1)$$

where ρ_0 is the ambient density and T_0 is the ambient temperature [18]. It can be seen from (1) that changes in the air temperature cause changes in the density and thus the refractive index.

B. Ray Tracing Formalism

We rely on ray tracing to model the propagation of light through a field of changing refractive index. We therefore use the image formation model described by refractive radiative transfer equation (RRTE) [26], which we describe in this section.

1) *Nonlinear Ray Tracing*: A light ray parameterized by a position, \mathbf{x} , and direction \mathbf{v} , will propagate through an inhomogeneous medium according to

$$\frac{d\mathbf{x}}{dt} = \mathbf{v}, \quad \frac{d\mathbf{v}}{dt} = \eta \nabla \eta. \quad (2)$$

In the absence of absorption and emission, basic radiance is conserved along a ray.

Define a light ray $r_{s \leftrightarrow w}$ as the set of points traversed by light between its endpoints $(\mathbf{x}_s, \mathbf{v}_s)$ at the camera sensor, and $(\mathbf{x}_w, \mathbf{v}_w)$ at the back-wall. We can formulate the bijective ray tracing mapping $\mathcal{T}_{s \rightarrow w}$ from the camera sensor to the back-wall as

$$(\mathbf{x}_w, \mathbf{v}_w) = \mathcal{T}_{s \rightarrow w}(\mathbf{x}_s, \mathbf{v}_s, \eta) := \begin{cases} \mathbf{v} = \int_{r_{s \leftrightarrow w}} \eta(\mathbf{x}(\tau)) \nabla \eta(\mathbf{x}(\tau)) d\tau \\ \mathbf{x} = \int_{r_{s \leftrightarrow w}} \mathbf{v}(\tau) d\tau \end{cases} \quad (3)$$

where τ denotes an infinitesimal step along the ray path. Note that for a given initial endpoint $(\mathbf{x}_s, \mathbf{v}_s)$ and refractive field map η , the ray path $r_{s \leftrightarrow w}$ and endpoint $(\mathbf{x}_w, \mathbf{v}_w)$ are fully determined using (3).

Ray tracing is performed using Monte Carlo estimation by shooting rays from the camera towards random points on the back-wall and evaluating their trajectories using (3). However, directly integrating the ray tracing equations can be an expensive operation. Instead, we use a quasi-linear approximation that is faster than direct integration with little deviation from the original, true path.

2) *Quasi-linear Ray Tracing Approximation*: Consider an initial endpoint $(\mathbf{x}_0, \mathbf{v}_0)$ located at the camera sensor and pointing toward an arbitrary location on the back-wall. Since the change in refractive index is small, we can approximate the trajectory of the light ray as a straight line, $\mathbf{x}(t) \approx \bar{\mathbf{x}}(t) = \mathbf{x}_0 + t\mathbf{v}_0$, for $t \in (0, D]$, where D is the total distance from the camera to the back-wall. Then the direction of a ray along the path $\bar{\mathbf{x}}$ is approximated by

$$\tilde{\mathbf{v}}(t) \approx \int_0^t \eta(\bar{\mathbf{x}}(\tau)) \nabla \eta(\bar{\mathbf{x}}(\tau)) d\tau + \mathbf{v}_0. \quad (4)$$

Finally, given the computed $\tilde{\mathbf{v}}(t)$ values, the final position $\mathbf{x}(t)$ is approximated by

$$\tilde{\mathbf{x}}(t) \approx \int_0^t \tilde{\mathbf{v}}(\tau) d\tau + \mathbf{x}_0. \quad (5)$$

Importantly, the linear approximation to the path is only used for querying the refractive index field. The resulting position and direction at the back-wall are used in the pixel intensity calculations, i.e., $(\mathbf{x}_w, \mathbf{v}_w) = \tilde{\mathcal{T}}_{s \rightarrow w}(\mathbf{x}_s, \mathbf{v}_s, \eta) := (\tilde{\mathbf{x}}(D), \tilde{\mathbf{v}}(D))$.

C. Image Formation Model

For a given pixel j on the sensor, the intensity I_j can be described as

$$I_j = \int_A \int_{\Omega} W_j(\mathbf{x}_s) L_{\text{wall}}(\mathbf{x}_w, \mathbf{v}_w) \frac{\langle \hat{\mathbf{n}}_w, \mathbf{v}_w \rangle}{\|r_{s \leftrightarrow w}\|} d\mathbf{v}_s d\mathbf{x}_s. \quad (6)$$

where $(\mathbf{x}_w, \mathbf{v}_w)$ are obtained using $\tilde{\mathcal{T}}$, W_j is the camera filter function, $L_{\text{wall}}(\mathbf{x}, \mathbf{v})$ is the luminance of the back plane, $\langle \hat{\mathbf{n}}_w, \mathbf{v}_w \rangle$ is the cosine of the angle between the back-wall normal $\hat{\mathbf{n}}_w$ and the incident ray direction at the back-wall \mathbf{v}_w , $\|r_{s \leftrightarrow w}\|$ is the length of the ray path. The total intensity for the pixel integrates over all incoming directions, Ω , as well as over the area of the sensor pixel, A . Here, we approximate the camera filter with the triangular function

$$W_j(\mathbf{x}) = \max\left(\frac{p_w}{2} - |x - x_j|, 0\right) \cdot \max\left(\frac{p_h}{2} - |y - y_j|, 0\right) \quad (7)$$

where p_w and p_h are the pixel width and height, respectively, and (x_j, y_j) is the pixel center. This filter has the benefit of having non-zero gradient within the pixel's extents.

D. Light Source Model

We consider two cases, one in which the back-wall is a textured light source, and another where the back-wall is fully diffuse and is illuminated by an emitter, e.g., a light projector. In the case where the back-wall is a source, L_{wall} is known and can be queried directly. In practice, this would be the same as looking up a texture.

If we model the wall as being illuminated by a pinhole projector with a finite focal length, then a point, \mathbf{x}_w , on the back plane will be illuminated by a single point from the projector. Supposing that the projector is located at position \mathbf{x}_e , then

$$L_{\text{wall}}(\mathbf{x}_w, \mathbf{v}_w) = \mathcal{P}(\mathbf{v}_e(\mathbf{x}_w, \mathbf{v}_w^*, \mathbf{x}_e)) \frac{\langle \hat{\mathbf{n}}_w, \mathbf{v}_w^* \rangle}{\|r_{w \leftrightarrow e}\|}, \quad (8)$$

where \mathbf{v}_e is the ray's direction on the emitter, \mathbf{v}_w^* is the incoming ray direction from the emitter at the back-wall, and \mathcal{P} is the image displayed by the projector.

Above, the known quantities are \mathbf{x}_w and \mathbf{x}_e . We use ray tracing to find the incoming ray direction \mathbf{v}_w^* at the wall and the outgoing ray direction \mathbf{v}_e at the emitter as a function of $\mathbf{x}_w, \mathbf{x}_e$ and η . In order to find the path between the emitter and the wall, we use the shooting method as in [29] to solve for the initial direction from the wall that would reach the emitter position:

$$\mathbf{v}_w^* = \arg \min_{\mathbf{v}_w^*} \|\mathbf{x}_e - \tilde{\mathcal{T}}_{w \rightarrow e}(\mathbf{x}_w, \mathbf{v}_w^*, \eta)\|^2. \quad (9)$$

Once we find a valid path $r_{w \rightarrow e}$, we can then evaluate the outgoing direction \mathbf{v}_e at the emitter and consequently sample the image \mathcal{P} . Note that since the tracing procedure, $\tilde{\mathcal{T}}$, is linear and differentiable, we can solve (9) using a single linear solve and differentiate with respect to η using implicit differentiation.

Finally, we can express the overall ray tracing procedure that generates the flow image $I_{\text{flow}} = \mathcal{B}(\mathbf{x}_s, \mathbf{v}_s, \mathbf{x}_e, \mathcal{P}, \eta)$ using a combined BOS operator \mathcal{B} , such that,

$$\begin{aligned} \mathcal{B}(\mathbf{x}_s, \mathbf{v}_s, \mathbf{x}_e, \mathcal{P}, \eta) \\ := \int_A \int_{\Omega} W_j(\mathbf{x}_s) \mathcal{P}(\mathbf{v}_e(\mathbf{x}_w, \mathbf{v}_w^*, \mathbf{x}_e)) \frac{\langle \hat{\mathbf{n}}_w, \mathbf{v}_w^* \rangle}{\|r_{w \leftrightarrow e}\|} \frac{\langle \hat{\mathbf{n}}_w, \mathbf{v}_w \rangle}{\|r_{s \leftrightarrow w}\|} d\mathbf{v}_s d\mathbf{x}_s, \end{aligned} \quad (10)$$

where $(\mathbf{x}_w, \mathbf{v}_w) = \tilde{\mathcal{T}}_{s \rightarrow w}(\mathbf{x}_s, \mathbf{v}_s, \eta)$, and \mathbf{v}_w^* is given by (9).

In summary, the tracing procedure proceeds as follows. Sample a ray on the camera plane and trace it to the back wall through the volume. Then from the wall point, solve the minimization problem as described by (9). Once we have the connection to the projector, we use (10) to calculate the intensity on the sensor pixel.

III. PHYSICS-INFORMED BOS TOMOGRAPHY

Reconstructing the airflow Φ can be formulated as a tomographic inverse problem given BOS measurements and boundary conditions specifying the flow parameters of the room at a boundary region Γ . Since the single-camera BOS tomography problem is highly ill-posed, we propose to regularize the inversion using a physics-informed loss that imposes a Boussinesq approximation to the incompressible Navier–Stokes equation coupled with the heat-transfer equation.

A. Inverse problem

Let the true airflow Φ be parameterized by the fields (T^*, p^*, \mathbf{u}^*) and recall that the refractive index field η^* is related to T^* according to (1). We propose to reconstruct Φ by solving the following constrained optimization problem:

$$\begin{aligned} \min_{T, p, \mathbf{u}} \quad & \lambda_1 \mathcal{L}_{\text{BOS}}(\eta) + \lambda_2 \mathcal{L}_{\Gamma}(T, p, \mathbf{u}) + \lambda_3 \mathcal{L}_{\text{PDE}}(T, p, \mathbf{u}) \\ \text{subject to} \quad & \eta = 1 + \rho_0 G \frac{T_0}{T}, \end{aligned} \quad (11)$$

where the schlieren loss is defined as

$$\mathcal{L}_{\text{BOS}}(\eta) := \sum_j \left\| I_{\text{flow}}(j) - \sum_{\mathbf{x}_s \in \mathcal{N}_j} \mathcal{B}(\mathbf{x}_s, \mathbf{v}_s, \mathbf{x}_e, \mathcal{P}, \eta) \right\|_2^2, \quad (12)$$

where j denotes the pixel index, and \mathcal{N}_j is spatial region belonging to the pixel j . Differentiating the schlieren loss with respect to the refractive field η can be performed efficiently using automatic differentiation with adjoint-gradient [21], [30] for the $\tilde{\mathcal{T}}_{s \rightarrow w}$ operator and exploiting the implicit function theorem [31] in evaluating the gradient of \mathbf{v}_w^* with respect to η through equation (9).

The boundary loss is a Euclidean distance between the computed fields and the true fields at collocation points on the boundary

$$\mathcal{L}_{\Gamma}(T, p, \mathbf{u}) = \left\| (T_n^*, p_n^*, \mathbf{u}_n^*) \Big|_{\Gamma} - (T_n, p_n, \mathbf{u}_n) \Big|_{\Gamma} \right\|_2^2, \quad (13)$$

where subscript n denotes the field divided by its maximum value.

B. Physics-informed Loss

The airflow is assumed to be a steady, incompressible, Newtonian fluid that is governed by the Boussinesq approximation for buoyancy-driven flows. The underlying physics of the airflow is then imposed by combining loss functions obtained from the nondimensional steady-state Navier–Stokes equations in the Boussinesq approximation. These are defined in terms of the mass conservation, momentum conservation, and heat transfer equation residuals as [18], [32]–[34]

$$r_{\text{mass}}(\mathbf{x}) = \nabla \cdot \mathbf{u}, \quad (14a)$$

$$\mathbf{r}_{\text{mom}}(\mathbf{x}) = (\mathbf{u} \cdot \nabla) \mathbf{u} + \nabla p - \frac{1}{Re} \nabla^2 \mathbf{u} + Ri T_{\text{nd}} \mathbf{e}_g, \quad (14b)$$

$$r_{\text{heat}}(\mathbf{x}) = (\mathbf{u} \cdot \nabla) T_{\text{nd}} - \frac{1}{Pe} \nabla^2 T_{\text{nd}}. \quad (14c)$$

Here, ∇ and ∇^2 are the spatial gradient and Laplacian operators in 3D, respectively, \mathbf{x} is the nondimensional coordinate scaled by a characteristic length scale L , $\mathbf{u} = (u, v, w)$ is the nondimensional velocity scaled by a characteristic velocity scale U , and p is the nondimensional pressure deviation from hydrostatic equilibrium scaled by $p_0 = \rho_0 U^2$, where ρ_0 is a reference density. The nondimensional temperature fluctuation T_{nd} is obtained from the dimensional temperature T as $T_{\text{nd}} = \frac{T - T_0}{T_{\text{in}} - T_0}$, where T_{in} is the inlet temperature and T_0 is a reference temperature. Additional parameters include the acceleration due to gravity g and its unit vector \mathbf{e}_g , kinematic viscosity ν , coefficient of thermal diffusivity α , and coefficient of




Fig. 2: Proposed MLP framework for evaluating the BOS, PDE, and boundary losses.

thermal expansion β , leading to the nondimensional Reynolds, Péclet, and Richardson numbers, defined as follows:

$$Re = \frac{UL}{\nu}, Pe = \frac{UL}{\alpha}, Ri = \frac{g\beta(T_{\text{in}} - T_0)L}{U^2}. \quad (15)$$

Following the PINNs framework [22], [35], the residual equations (14) are combined to form the physics-informed loss

$$\mathcal{L}_{\text{PDE}} = \sum_{i=1}^{N_c} \gamma_1 r_{\text{mass}}^2(\mathbf{x}_i) + \gamma_2 \|\mathbf{r}_{\text{mom}}(\mathbf{x}_i)\|_2^2 + \gamma_3 r_{\text{heat}}^2(\mathbf{x}_i), \quad (16)$$

where $\gamma_{1,2,3}$ are scalar multipliers that balance the weight of each residual, and $\mathbf{x}_{i=1, \dots, N_c}$ are coordinates of the collocation points uniformly sampled in the computational domain at each iteration.

C. Neural Representation of the Airflow Fields

Single-view 3D BOS has inherent ambiguities along the view direction [21], therefore, we propose to use a PINN framework so the physics of airflow can regularize the reconstruction [18], [23]–[25]. The PINN consists of a multilayer perceptron (MLP) whose outputs are the T , p , and \mathbf{u} fields. The input coordinates are first transformed via random Fourier feature embeddings [36], followed by three fully connected layers of width 64 and SIREN activations [37]. The last layer is a tanh activation that maps the outputs to their respective ranges. Fig. 2 illustrates the computational framework of the MLP architecture and how the outputs are used in the different optimization losses. Since ray tracing through the MLP is slow, we use an intermediate step that first samples the MLP on a voxel grid that is then used to perform ray tracing and compute the schlieren loss \mathcal{L}_{BOS} . On the other hand, the MLP is directly sampled at collocation points throughout the computational domain to compute the physics-informed loss \mathcal{L}_{PDE} , and the MLP is sampled in the boundary region Γ to determine \mathcal{L}_{Γ} .

IV. EXPERIMENTAL SETUP

A. Airflow simulation

We use the experimental setup shown in Fig. 1 where we simulate a room of size 5.1m×4.6m×3m, with the bottom left corner located at $(-1.2, -3.6, 0)$ m, using the computational fluid dynamics (CFD) solver OpenFOAM based on the `comfortHotRoom` tutorial [38]. The room is ventilated using an inlet pointing in the positive x direction of size 30cm×20cm located in the center of the $y-z$ plane at a height of 2.2m. An outlet of size 30cm×10cm is located above the inlet at a height of 2.6m. In order to image the flow in the entire room, we simulate a hypothetical camera and projector pointing in the positive y direction with a large focal length that are located at $(0.85, -29, 1.5)$ m and $(1.85, -29, 1.5)$ m, respectively.

We perform a Reynolds-Averaged Navier–Stokes (RANS) CFD simulation with OpenFOAM of a turbulent flow with $k-\epsilon$ coefficients $C_\mu = 0.09$, $C_1 = 1.44$, $C_2 = 1.92$, and $\sigma_\epsilon = 1.11$. The ambient temperature and density are set to $T_0 = 300$ K and $\rho_0 = 1$ kg/m³,

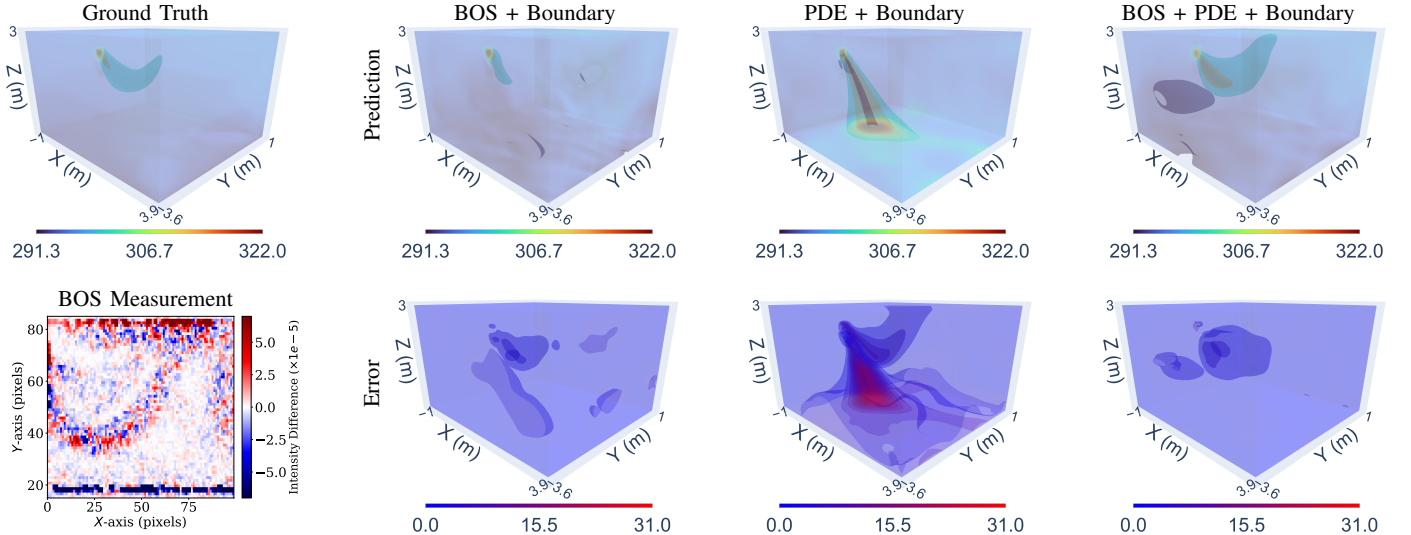


Fig. 3: Airflow reconstruction results using the projector and camera BOS acquisition setup. Top row from left to right shows the ground truth flow and temperature reconstructions (in K) using different combinations of the losses. Bottom left figure illustrates the active pixels of the acquired BOS image. The remaining bottom row plots illustrate the absolute error of each reconstruction regime.

while the coefficient of thermal expansion $\beta = 3 \times 10^{-3} \text{ K}^{-1}$. The transport properties are $\mu = 10^{-5} \text{ kg/m/s}$ and Prandtl number $\text{Pr} = 0.7$. The inlet blows air with a velocity $\mathbf{u}_{\text{in}} = (2, 0, -1.96) \text{ m/s}$ and temperature 321 K.

B. Airflow reconstruction

We evaluate the reconstruction performance on a $64 \times 64 \times 64$ voxelized grid. We implement our method including the forward rendering and optimization in JAX with the Equinox, Difffrax, and Optax libraries [39]–[42]. The back-wall is illuminated using a wavelet noise pattern [43] with a frequency determined by the resolution of the camera. We render a 100×100 pixel image with 2 samples per pixel during each iteration of the optimization. The boundary domain Γ is defined as the $y-z$ plane at $x = -1.2 \text{ m}$ containing the inlet/outlet. We use the non-dimensional physics-informed loss in (16) with the characteristic quantities $U = 3.26 \text{ m/s}$ and $L = 5.1 \text{ m}$. The dimensional parameters ρ_0 , ν , α , and β are identical to the ones used in the OpenFOAM simulation. Note that there is a model and resolution mismatch between the incompressible flow equations imposed by the physics-informed loss and those of the OpenFOAM RANS simulation. We set the weights of the loss functions as follows: $\lambda_1 = \frac{10^5}{\|I_{\text{ref}} - I_{\text{flow}}\|_{\infty}}$, $\lambda_2 = 10^4$, $\lambda_3 = 1$, $\gamma_1 = 100$, $\gamma_2 = 10$, and $\gamma_3 = 100$. Optimization is performed using the Optax adabelief solver with an initial learning rate of 10^{-3} and an exponential decay schedule. We run 80,000 iterations of the optimizer with mini-batch updates where the batch sizes are 8192 spatial points for the PDE loss, 5000 pixels for the BOS loss, and 4096 boundary points for the boundary loss.

V. RESULTS AND CONCLUSION

We compare the performance of reconstructing the airflow volume using our proposed BOS+PDE+boundary losses with two reconstruction regimes: BOS+boundary and PDE+boundary. The BOS+boundary regime is equivalent to the INR parameterization of [21], whereas the PDE+boundary regime evaluates the importance of using the BOS measurement. Table I summarizes the root mean squared error (RMSE) of the reconstructed volumes evaluated on

TABLE I: Reconstruction RMSE for the different regimes.

	\mathcal{L}_{BOS}	\mathcal{L}_{PDE}	T (K)	p (Pa)	\mathbf{u} (m/s)
camera	✓	✗	2.33	8.59	2.98×10^{-1}
only	✓	✓	1.07	3.81×10^{-2}	7.84×10^{-2}
projector	✓	✗	1.19	6.77	2.96×10^{-1}
+ camera	✓	✓	1.47	5.51×10^{-2}	7.71×10^{-2}
	✗	✓	4.11	5.65×10^{-2}	8.67×10^{-2}

the $64 \times 64 \times 64$ grid. The results in the table demonstrate that incorporating the BOS loss function is critical for accurately reconstructing the temperature field in the flow, especially since the BOS is primarily affected by the refractive index η which is coupled to the temperature field through (1). This can also be seen in the qualitative reconstruction results shown in Fig. 3 where the combination of the three losses results in fewer artifacts. However, reconstructing the pressure and velocity fields can only be performed by incorporating the PDE constraints. Furthermore, we observe small errors in the reconstructed temperature field of the BOS+PDE+boundary regime that we attribute to the mismatch in modeling the effects of turbulence and the compressibility of the flow between the PDE constraint and the OpenFOAM simulation.

Lastly, the results in Table I demonstrate that using a projector provides similar reconstruction performance and is a feasible replacement to a camera with a fixed back-wall pattern. This is partly due to the long focal length of the camera and projector that we used in our simulations where the entire volume is visible by the camera. For future work, we plan on exploring the use of the projector to probe regions of the volume that are not observed by the camera when a shorter focal length is used.

ACKNOWLEDGMENTS

We thank Marcus Greiff for his support with the OpenFOAM simulation, Saviz Mowlavi for helpful discussions on PINNs, and Samuel Grauer for insightful conversations on BOS.

REFERENCES

- [1] S. V. Patankar, "Airflow and cooling in a data center," *J. Heat Transfer*, vol. 132, Art. no. 073001, Apr. 2010.
- [2] M. Abedi, F. Jazizadeh, B. Huang, and F. Battaglia, "Smart HVAC systems — adjustable airflow direction," in *Adv. Comput. Strat. Eng.*, I. F. C. Smith and B. Domer, Eds., 2018, vol. 10864 of *Lect. Notes Comput. Sci.*, pp. 193–209.
- [3] J. Wang, X. Liu, S. Chen, H. Jiang, G. Fang, W. Chen, and S. Deng, "Reduced-scale model study on cable heat dissipation and airflow distribution of power cabins," *Appl. Therm. Eng.*, vol. 160, Art. no. 114068, Sept. 2019.
- [4] H. Fang, K. Li, G. Wu, R. Cheng, Y. Zhang, and Q. Yang, "A CFD analysis on improving lettuce canopy airflow distribution in a plant factory considering the crop resistance and LEDs heat dissipation," *Biosyst. Eng.*, vol. 200, pp. 1–12, Dec. 2020.
- [5] K. Okamoto, T. Ohhashi, M. Asakura, and K. Watanabe, "A digital anemometer," *IEEE Trans. Instrum. Meas.*, vol. 43, no. 2, pp. 116–120, Apr. 1994.
- [6] C. Jiang, Y. C. Soh, and H. Li, "Sensor and CFD data fusion for airflow field estimation," *Appl. Therm. Eng.*, vol. 92, pp. 149–161, Jan. 2016.
- [7] A. Toepler, *Beobachtungen nach einer neuen optischen Methode: Ein Beitrag zur Experimentalphysik*, Bonn, 1864.
- [8] J. P. Crimaldi, "Planar laser induced fluorescence in aqueous flows," *Exp. Fluids*, vol. 44, no. 6, pp. 851–863, June 2008.
- [9] R. J. Adrian, "Scattering particle characteristics and their effect on pulsed laser measurements of fluid flow: speckle velocimetry vs particle image velocimetry," *Appl. Optics*, vol. 23, no. 11, pp. 1690–1691, June 1984.
- [10] M. Raffel, H. Richard, and G. E. A. Meier, "On the applicability of background oriented optical tomography for large scale aerodynamic investigations," *Exp. Fluids*, vol. 28, no. 5, pp. 477–481, May 2000.
- [11] S. B. Dalziel, G. O. Hughes, and B. R. Sutherland, "Whole-field density measurements by 'synthetic schlieren'," *Exp. Fluids*, vol. 28, no. 4, pp. 322–335, Apr. 2000.
- [12] G. E. A. Meier, "Computerized background-oriented schlieren," *Exp. Fluids*, vol. 33, pp. 181–187, July 2002.
- [13] G. S. Settles and M. J. Hargather, "A review of recent developments in schlieren and shadowgraph techniques," *Meas. Sci. Technol.*, vol. 28, 2017.
- [14] G. S. Settles, "On background-oriented schlieren (BOS) velocimetry," in *Proc. Int. Symp. Flow Vis.*, Zurich, Switzerland, June 2018.
- [15] 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. Graphics*, vol. 27, no. 5, Art. no. 132, Dec. 2008.
- [16] F. Nicolas, V. Todoroff, A. Plyer, G. Le Besnerais, D. Donjat, F. Micheli, F. Champagnat, P. Cornic, and Y. Le Sant, "A direct approach for instantaneous 3D density field reconstruction from background-oriented schlieren (BOS) measurements," *Exp. Fluids*, vol. 57, Art. no. 13, Dec. 2016.
- [17] S. J. Grauer, A. Unterberger, A. Rittler, K. J. Daun, A. M. Kempf, and K. Mohri, "Instantaneous 3D flame imaging by background-oriented schlieren tomography," *Combust. Flame*, vol. 196, pp. 284–299, Oct. 2018.
- [18] S. Cai, Z. Wang, F. Fuest, Y. J. Jeon, C. Gray, and G. E. Karniadakis, "Flow over an espresso cup: inferring 3-D velocity and pressure fields from tomographic background oriented Schlieren via physics-informed neural networks," *J. Fluid Mech.*, vol. 915, Art. no. 102, 2021.
- [19] L. Venkatakrishnan and G. E. A. Meier, "Density measurements using the Background Oriented Schlieren technique," *Exp. Fluids*, vol. 37, pp. 237–247, Aug. 2004.
- [20] T. Shimazaki, S. Ichihara, and Y. Tagawa, "Background oriented schlieren technique with fast Fourier demodulation for measuring large density-gradient fields of fluids," *Exp. Therm Fluid Sci.*, vol. 134, pp. 110598, June 2022.
- [21] B. Zhao, A. Levis, L. Connor, P. P. Srinivasan, and K. L. Bouman, "Single view refractive index tomography with neural fields," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2024, pp. 25358–25367.
- [22] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *J. Comput. Phys.*, vol. 378, pp. 686–707, Feb. 2019.
- [23] J. P. Molnar, L. Venkatakrishnan, B. E. Schmidt, T. A. Siphens, and S. J. Grauer, "Estimating density, velocity, and pressure fields in supersonic flows using physics-informed BOS," *Exp. Fluids*, vol. 64, Art. no. 14, Jan. 2023.
- [24] J. P. Molnar, S. J. Grauer, O. Léon, D. Donjat, and F. Nicolas, "Physics-informed background-oriented schlieren of turbulent underexpanded jets," in *Proc. AIAA SciTech*, National Harbor, MD, Jan. 2023.
- [25] J. P. Molnar, E. J. LaLonde, C. S. Combs, O. Léon, D. Donjat, and S. J. Grauer, "Forward and inverse modeling of depth-of-field effects in background-oriented schlieren," arXiv:2402.15954 [physics], Feb. 2024.
- [26] M. Ament, C. Bergmann, and D. Weiskopf, "Refractive radiative transfer equation," *ACM Trans. Graphics*, vol. 33, no. 2, Art. no. 17, pp. 1–22, Apr. 2014.
- [27] J. M. Weisberger and B. F. Bathel, "Projection background-oriented schlieren," *Appl. Optics*, vol. 61, no. 20, pp. 6006–6015, July 2022.
- [28] J. H. Gladstone and T. P. Dale, "Researches on the refraction, dispersion, and sensitiveness of liquids," *Philos. Trans. R. Soc. London*, vol. 153, pp. 317–343, 1863.
- [29] A. Pediredla, Y. K. Chalmiani, M. G. Scopelliti, M. Chamanzar, S. Narasimhan, and I. Gkioulekas, "Path tracing estimators for refractive radiative transfer," *ACM Trans. Graphics*, vol. 39, no. 6, Art. no. 241, Dec. 2020.
- [30] A. Teh, M. O'Toole, and I. Gkioulekas, "Adjoint nonlinear ray tracing," *ACM Trans. Graphics*, vol. 41, no. 4, Art. no. 126, July 2022.
- [31] S. G. Krantz and H. R. Parks, *The implicit function theorem: history, theory, and applications*, Springer Science & Business Media, 2002.
- [32] A. Oberbeck, "Über die Wärmeleitung der Flüssigkeiten bei Berücksichtigung der Strömungen infolge von Temperaturdifferenzen," *Ann. Phys.*, vol. 243, no. 6, pp. 271–292, 1879.
- [33] J. Boussinesq, *Théorie analytique de la chaleur: mise en harmonie avec la thermodynamique et avec la théorie mécanique de la lumière*, vol. 2, Gauthier-Villars, 1903.
- [34] M. Lappa, "Incompressible flows and the Boussinesq approximation: 50 years of CFD," *C.R. Mec.*, vol. 350, no. S1, pp. 1–22, 2022.
- [35] S. Wang, S. Sankaran, H. Wang, and P. Perdikaris, "An expert's guide to training physics-informed neural networks," arXiv:2308.08468 [physics], Aug. 2023.
- [36] M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. Barron, and R. Ng, "Fourier features let networks learn high frequency functions in low dimensional domains," in *Adv. Neural Inf. Process. Syst.*, 2020, vol. 33, pp. 7537–7547.
- [37] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein, "Implicit neural representations with periodic activation functions," in *Adv. Neural Inf. Process. Syst.*, 2020, vol. 33, pp. 7462–7473.
- [38] "OpenFOAM tutorials / heatTransfer / buoyantSimpleFoam / comfortHotRoom," <https://develop.openfoam.com/Development/openfoam/-/tree/master/tutorials/heatTransfer/buoyantSimpleFoam/comfortHotRoom>.
- [39] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, "JAX: composable transformations of Python+NumPy programs," 0.3.13, <http://github.com/google/jax>, 2018.
- [40] P. Kidger and C. Garcia, "Equinox: neural networks in JAX via callable PyTrees and filtered transformations," arXiv:2111.00254 [cs], Oct. 2021.
- [41] P. Kidger, *On Neural Differential Equations*, Ph.D. thesis, University of Oxford, 2021.
- [42] M. Hessel, D. Budden, F. Viola, M. Rosca, E. Sezener, and T. Hennigan, "Optax: composable gradient transformation and optimisation, in JAX!," <http://github.com/deepmind/optax>, 2020.
- [43] R. L. Cook and T. DeRose, "Wavelet noise," *ACM Trans. Graphics*, vol. 24, no. 3, pp. 803–811, July 2005.