

Indoor Airflow Imaging Using Physics-Informed Background-Oriented Schlieren Tomography

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

TR2025-071 June 06, 2025

Abstract

We develop a framework for non-invasive volumetric indoor airflow estimation from a single viewpoint using background-oriented schlieren (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. While the single-view BOS tomography problem is severely ill-posed, our proposed framework addresses this using: (1) improved ray tracing, (2) a physics-based light rendering approach and loss formulation, and (3) a physics-based regularization using a physics-informed neural network (PINN) to ensure that the reconstructed airflow is consistent with the governing equations for buoyancy-driven flows.

International Symposium on Computational Sensing (ISCS) 2025

© 2025 MERL. This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Indoor Airflow Imaging Using Physics-Informed Background-Oriented Schlieren Tomography

Arjun Teh^{*1}, Wael H. Ali[†], Joshua Rapp[†], and Hassan Mansour[†]

^{*} Carnegie Mellon University, Pittsburgh, PA, USA

[†] Mitsubishi Electric Research Laboratories, Cambridge, MA, USA

Abstract—We develop a framework for non-invasive volumetric indoor airflow estimation from a single viewpoint using background-oriented schlieren (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. While the single-view BOS tomography problem is severely ill-posed, our proposed framework addresses this using: (1) improved ray tracing, (2) a physics-based light rendering approach and loss formulation, and (3) a physics-based regularization using a physics-informed neural network (PINN) to ensure that the reconstructed airflow is consistent with the governing equations for buoyancy-driven flows.

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, 2]. However, three-dimensional (3D) airflow sensing is challenging since hardware sensors only measure localized spatial regions around the sensors [3]. One promising imaging technique 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 [4–8]. 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 [9–12].

In this work, we combine several recent advances in refractive field tomography to achieve 3D airflow reconstructions for room-scale scenes from a single camera view. First, we use a physics-informed neural network (PINN) framework [13] as a regularizer to ensure that the reconstructed field adheres to the partial differential equations (PDEs) governing airflow. This approach is similar to those suggested in [12, 14–17] but applied for the first time in a single view room-scale setting. Additionally, we employ improved ray tracing and use a physics-based rendering approach according to the refractive radiative transfer equation (RRTE) [18, 19]. We also explore the use of a light source projecting a pattern onto the wall [20], which may be more practical than a fixed background. This extended abstract summarizes our approach, with a comprehensive version of this work [21] set to appear in the proceedings of ICASSP 2025.

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(a). When no air is flowing, the camera captures a reference image I_{ref} of the back-wall pattern. When the inlet on the side wall blows the airflow into the room, the captured image I_{flow} appears distorted due to the change in density of the air which induces a gradient in its refractive index η , as described by the Gladstone–Dale equation $\eta = 1 + G\rho$, where G is the Gladstone–Dale coefficient [22]. Assuming the pressure variation of air in the room is small, then by the ideal gas law, the refractive

index is related to temperature T as $\eta(T) = 1 + \rho_0 G \frac{T_0}{T}$, where ρ_0 is the ambient density and T_0 is the ambient temperature [12]. It can be seen that changes in the air temperature cause changes in the refractive index causing light rays passing through the air to bend.

B. Ray Tracing Formalism

The propagation of light inside a medium with continuously varying refractive index can be described using the ray tracing ordinary differential equations (ODEs) [23]

$$\frac{d\mathbf{x}(t)}{dt} = \mathbf{v}(\mathbf{x}(t)), \quad \frac{d\mathbf{v}(t)}{dt} = \eta(\mathbf{x}(t)) \nabla \eta(\mathbf{x}(t)). \quad (1)$$

Due to the implicit dependence of the refractive index η on the ray position $\mathbf{x}(t)$, the ODE system is fully coupled and its solution is referred to as *nonlinear ray tracing*. The fully coupled nature of this ODE system introduces significant computational complexity. To address this, we employ a *quasi-linear ray tracing* approach to decouple the velocity ODE from the position ODE, simplifying the computation with little deviation from the original, true path. Specifically, since the change in refractive index is small, we can approximate the trajectory of the light ray as a straight line starting from \mathbf{x}_0 along \mathbf{v}_0 . The refractive index field is queried along this linear path to compute an approximate velocity $\tilde{\mathbf{v}}(t)$ that is then used in the position ODE to obtain the updated ray path $\tilde{\mathbf{x}}(t)$.

C. Image Formation Model

We use a physics-based rendering approach – specifically, the path integral expression of the RRTE [18, 19] – to compute the intensity I_j for a given pixel j on the sensor plane as

$$I_j = \int_A \int_{\Omega} W_j(\mathbf{x}_s) L_{\text{wall}}(\mathbf{x}_w, \mathbf{v}_w) \frac{\langle \hat{\mathbf{n}}_w, \mathbf{v}_w \rangle}{\|r_{s \leftrightarrow w}\|} d\mathbf{v}_s d\mathbf{x}_s, \quad (2)$$

where W_j is a triangular camera filter function, $(\mathbf{x}_w, \mathbf{v}_w)$ are the wall position and velocity of a traced ray starting from $(\mathbf{x}_s, \mathbf{v}_s)$ on the sensor plane, L_{wall} is the luminance of the back-wall, $\langle \hat{\mathbf{n}}_w, \mathbf{v}_w \rangle$ is the cosine of the angle between the back-wall normal $\hat{\mathbf{n}}_w$ and \mathbf{v}_w , and $\|r_{s \leftrightarrow w}\|$ is the length of the ray path. The total intensity for the pixel integrates over all starting ray velocities, $\mathbf{v}_s \in \Omega$, and all starting ray positions within the area of the sensor pixel, $\mathbf{x}_s \in A$. In practice, we use Monte Carlo sampling to evaluate this integral.

To determine the luminance of the back-wall, we consider two cases. In the case where the back-wall is a textured light source, L_{wall} is known and can be queried directly. If we model the wall as being illuminated by a pinhole projector, then a point on the back-wall will be illuminated by a single point from the projector. Further details regarding the back-wall-projector connection and the overall image formation model for this case can be found in [21].

D. Physics-Informed BOS Tomography

Reconstructing the temperature T , pressure p , and velocity \mathbf{u} fields of the airflow can be formulated as a tomographic inverse problem given BOS image measurements I_{flow} and boundary conditions $(T^*, p^*, \mathbf{u}^*)|_{\Gamma}$ at a boundary region Γ . Single-view 3D BOS has inherent ambiguities along the view direction [14]. Following the

¹Arjun Teh conducted this work during an internship at MERL

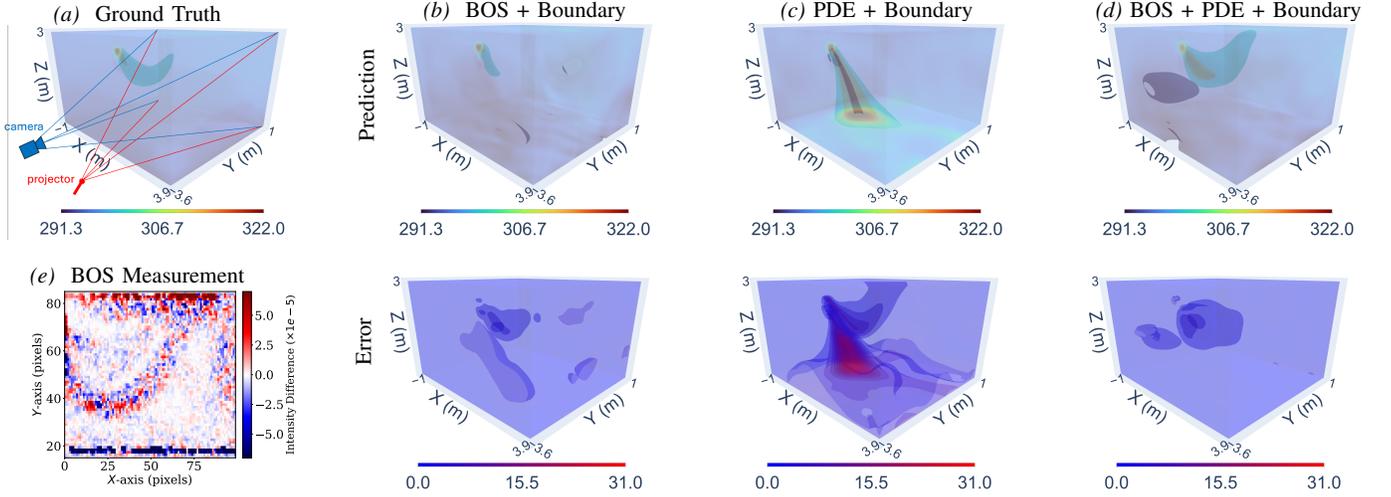


Fig. 1: Airflow reconstruction results using the projector and camera BOS acquisition setup. Top row from left to right shows the ground truth and reconstructions of the temperature field (in K) using different combinations of the losses. Bottom left figure illustrates the BOS image measurement. The remaining bottom row plots illustrate the absolute error of each reconstruction regime.

PINNs framework [13, 24], we propose to regularize the inversion using a physics-informed loss. This is done by solving the following optimization problem:

$$\begin{aligned} \min_{T, p, \mathbf{u}} \quad & \lambda_1 \mathcal{L}_{\text{BOS}}(\eta) + \lambda_2 \mathcal{L}_{\Gamma}(T, p, \mathbf{u}) + \lambda_3 \mathcal{L}_{\text{PDE}}(T, p, \mathbf{u}) \\ \text{subject to} \quad & \eta = 1 + \rho_0 G \frac{T_0}{T}, \end{aligned} \quad (3)$$

This objective function includes the following loss components:

- The BOS image loss $\mathcal{L}_{\text{BOS}} = \sum_j \left\| I_{\text{flow}}^j - I^j(\eta) \right\|_2^2$, where j denotes the pixel index, I_{flow} is the BOS image measurement intensity, and $I(\eta)$ is the predicted intensity from (2).
- The boundary loss $\mathcal{L}_{\Gamma} = \left\| (T_n^*, p_n^*, \mathbf{u}_n^*)|_{\Gamma} - (T_n, p_n, \mathbf{u}_n)|_{\Gamma} \right\|_2^2$, where subscript n denotes the field divided by its maximum value.
- The physics-informed loss $\mathcal{L}_{\text{PDE}} = \sum_{i=1}^{N_c} \gamma_1 r_{\text{mass}}^2(\mathbf{x}_i) + \gamma_2 \|\mathbf{r}_{\text{mom}}(\mathbf{x}_i)\|_2^2 + \gamma_3 r_{\text{heat}}^2(\mathbf{x}_i)$, where $\gamma_{1,2,3}$ are scalar multipliers that balance the weight of each residual, and $\mathbf{x}_{i=1, \dots, N_c}$ are collocation points uniformly sampled in the computational domain. r_{mass} , \mathbf{r}_{mom} , and r_{heat} are the nondimensional mass conservation, momentum conservation, and heat transfer equation residuals defined by the Boussinesq approximation for buoyancy-driven flows:

$$\begin{aligned} r_{\text{mass}}(\mathbf{x}) &= \nabla \cdot \mathbf{u}, \\ \mathbf{r}_{\text{mom}}(\mathbf{x}) &= (\mathbf{u} \cdot \nabla) \mathbf{u} + \nabla p - \frac{1}{Re} \nabla^2 \mathbf{u} + Ri T_{\text{nd}} \mathbf{e}_g, \\ r_{\text{heat}}(\mathbf{x}) &= (\mathbf{u} \cdot \nabla) T_{\text{nd}} - \frac{1}{Pe} \nabla^2 T_{\text{nd}}. \end{aligned}$$

Here, the nondimensional temperature fluctuation T_{nd} is obtained from the inlet T_{in} and reference T_0 temperatures as $T_{\text{nd}} = \frac{T - T_0}{T_{\text{in}} - T_0}$. Additional parameters include the acceleration due to gravity g and its unit vector \mathbf{e}_g , kinematic viscosity ν , coefficient of thermal diffusivity α , coefficient of thermal expansion β , and characteristic length L and velocity U scales leading to the nondimensional Reynolds, Péclet, and Richardson numbers: $Re = \frac{UL}{\nu}$, $Pe = \frac{UL}{\alpha}$, $Ri = \frac{g\beta(T_{\text{in}} - T_0)L}{U^2}$.

We use an implicit neural representation to parameterize the T , p , and \mathbf{u} fields using a multilayer perceptron (MLP), i.e. $(T, p, \mathbf{u}) = \text{MLP}(\mathbf{x}; \theta)$ where \mathbf{x} denotes the spatial location, and θ are the MLP parameters to be determined by solving the optimization problem (3).

Our forward rendering and optimization framework is implemented in JAX and uses the Equinox, Difffrax, and Optax libraries [25–28]. It is end-to-end differentiable where efficient gradient computation

is achieved using automatic differentiation and the adjoint state method [14, 29].

III. EXPERIMENTAL SETUP AND RESULTS

A. Airflow simulation

We obtain the ground truth airflow by performing a Reynolds-Averaged Navier-Stokes (RANS) simulation using OpenFOAM. The room setup and ground truth temperature field are shown in Fig.1(a) with an inlet pointing in the positive x direction. 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 $(0.85, -29, 1.5)\text{m}$ and $(1.85, -29, 1.5)\text{m}$, respectively. The projector illuminates the back-wall with a wavelet noise pattern [30]. The BOS image measurement shown in Fig. 1(e) is obtained using our physically-based renderer for a 100×100 sensor resolution with 2 samples per pixel.

B. Airflow reconstruction

We evaluate the reconstruction performance on a $64 \times 64 \times 64$ voxelized grid given the BOS image measurement and the boundary T, p and \mathbf{u} fields on the boundary $y - z$ plane at $x = -1.2\text{m}$ containing the inlet/outlet. 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 run 80,000 iterations of the Optax adabelief optimizer with mini-batch updates using 8192 spatial points for the PDE loss, 5000 pixels for the BOS loss, and 4096 points for the boundary loss.

We compare the performance of reconstructing the airflow volume using our proposed BOS+PDE+boundary losses with two alternative reconstruction regimes: BOS+boundary and PDE+boundary. The reconstruction results in Fig. 1 show that combining all three loss terms significantly reduces artifacts and is essential for accurately reconstructing the temperature field. We also show in [21] that the root mean squared error is 2 order of magnitudes lower for the p field and 1 order of magnitude lower for the \mathbf{u} field in the BOS+PDE+boundary reconstruction when compared with the BOS+boundary results. These findings highlight the advantages of our physics-informed and differentiable rendering framework in achieving high-accuracy BOS reconstructions of turbulent airflow.

REFERENCES

- [1] S. V. Patankar, "Airflow and cooling in a data center," *J. Heat Transfer*, vol. 132, Apr. 2010.
- [2] M. Abedi, F. Jazizadeh, B. Huang, and F. Battaglia, "Smart HVAC systems — adjustable airflow direction," in *Adv. Comput. Strat. Eng.*, ser. Lect. Notes Comput. Sci., I. F. C. Smith and B. Domer, Eds., vol. 10864, 2018, pp. 193–209.
- [3] 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.
- [4] 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.
- [5] 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.
- [6] G. E. A. Meier, "Computerized background-oriented schlieren," *Exp. Fluids*, vol. 33, pp. 181–187, Jul. 2002.
- [7] G. S. Settles and M. J. Hargather, "A review of recent developments in schlieren and shadowgraph techniques," *Meas. Sci. Technol.*, vol. 28, 2017.
- [8] G. S. Settles, "On background-oriented schlieren (BOS) velocimetry," in *Proc. Int. Symp. Flow Vis.*, Zurich, Switzerland, Jun. 2018.
- [9] 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, Dec. 2008.
- [10] 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, Dec. 2016.
- [11] 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.
- [12] 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, 2021.
- [13] 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.
- [14] 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. 25 358–25 367.
- [15] J. P. Molnar, L. Venkatakrishnan, B. E. Schmidt, T. A. Sipkens, and S. J. Grauer, "Estimating density, velocity, and pressure fields in supersonic flows using physics-informed BOS," *Exp. Fluids*, vol. 64, Jan. 2023.
- [16] 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.
- [17] 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," *AIAA J.*, vol. 62, no. 11, pp. 4316–4329, 2024.
- [18] M. Ament, C. Bergmann, and D. Weiskopf, "Refractive radiative transfer equation," *ACM Trans. Graphics*, vol. 33, no. 2, pp. 1–22, Apr. 2014.
- [19] 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, Dec. 2020.
- [20] J. M. Weisberger and B. F. Bathel, "Projection background-oriented schlieren," *Appl. Optics*, vol. 61, no. 20, pp. 6006–6015, Jul. 2022.
- [21] A. Teh, W. H. Ali, J. Rapp, and H. Mansour, "Indoor airflow imaging using physics-informed schlieren tomography," in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Process.*, 2025.
- [22] 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.
- [23] A. Sharma, D. V. Kumar, and A. K. Ghatak, "Tracing rays through graded-index media: a new method," *Appl. Optics*, vol. 21, no. 6, pp. 984–987, Mar. 1982.
- [24] 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.
- [25] 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.
- [26] P. Kidger and C. Garcia, "Equinox: neural networks in JAX via callable PyTrees and filtered transformations," arXiv:2111.00254 [cs], Oct. 2021.
- [27] P. Kidger, "On neural differential equations," Ph.D. dissertation, University of Oxford, 2021.
- [28] 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.
- [29] A. Teh, M. O'Toole, and I. Gkioulekas, "Adjoint nonlinear ray tracing," *ACM Trans. Graphics*, vol. 41, no. 4, Jul. 2022.
- [30] R. L. Cook and T. DeRose, "Wavelet noise," *ACM Trans. Graphics*, vol. 24, no. 3, pp. 803–811, Jul. 2005.