

SINGLE VIEW CAMERA-BASED DYNAMIC AIRFLOW SENSING

Tandi, Kevin; Ali, Wael H.; Rapp, Joshua; Mansour, Hassan

TR2026-038 March 28, 2026

Abstract

Background-oriented schlieren (BOS) tomography has emerged as an effective tool for quantitatively visualizing and reconstructing spatio-temporal volumetric thermal flows. Existing solutions rely on capturing multiview snapshots of a volumetric flow and solving a time-resolved inverse problem to reconstruct the flow. In this work, we propose a single view BOS imaging system that, when constrained by partial differential equations (PDEs) that characterize the airflow temporal dynamics, allows us to accurately reconstruct the time-resolved flow field. Our framework leverages an image rendering schlieren loss metric coupled with a physics-informed neural network (PINN) representation of the target flow fields that minimize the residuals of the coupled Boussinesq approximation of the time-dependent incompressible Navier–Stokes and the heat transfer equations. We further investigate a data-driven closure strategy in which effective thermal transport coefficients are learned directly from BOS data, thereby compensating for model mismatch between prescribed molecular properties and the unresolved turbulent transport present in the true flow.

*IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)
2026*

SINGLE VIEW CAMERA-BASED DYNAMIC AIRFLOW SENSING

Kevin Tandi^a, Wael H. Ali^b, Joshua Rapp^b, Hassan Mansour^b

^a Department of Electrical and Computer Engineering, University of California San Diego, CA, USA

^b Mitsubishi Electric Research Laboratories, Cambridge, MA, USA
ktandi@ucsd.edu, {wali, rapp, mansour}@merl.com

ABSTRACT

Background-oriented schlieren (BOS) tomography has emerged as an effective tool for quantitatively visualizing and reconstructing spatio-temporal volumetric thermal flows. Existing solutions rely on capturing multiview snapshots of a volumetric flow and solving a time-resolved inverse problem to reconstruct the flow. In this work, we propose a single view BOS imaging system that, when constrained by partial differential equations (PDEs) that characterize the airflow temporal dynamics, allows us to accurately reconstruct the time-resolved flow field. Our framework leverages an image rendering schlieren loss metric coupled with a physics-informed neural network (PINN) representation of the target flow fields that minimize the residuals of the coupled Boussinesq approximation of the time-dependent incompressible Navier–Stokes and the heat transfer equations. We further investigate a data-driven closure strategy in which effective thermal transport coefficients are learned directly from BOS data, thereby compensating for model mismatch between prescribed molecular properties and the unresolved turbulent transport present in the true flow.

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

1. INTRODUCTION

Non-intrusive sensing of transient three-dimensional (3D) thermal flows has uses in the control and efficiency of heating, ventilation, and air conditioning (HVAC) systems [1–4]. Background-oriented schlieren (BOS) imaging has emerged as a promising technique for sensing transparent flows that exhibit thermal gradients. BOS offers quantitative reconstruction, a wide field of view, and high sensitivity for a low cost [5–8]. In BOS tomography, images of a patterned background are observed using a camera to detect distortions resulting from changes in the refractive field of the flow [9–13]. Although BOS enables high-resolution, quantitative flow diagnostics, the tomographic reconstruction problem becomes well-posed only when a sufficiently large number of viewing directions are available to mitigate spatial ambiguities [14–17]. Such requirements limit the practicality of BOS tomography for room-scale measurements, where deploying large backgrounds and numerous cameras becomes onerous.

Single view BOS tomography becomes feasible when prior information can be imposed on the reconstructed flow in the form of regularization using, for instance, a low-dimensional spatio-temporal representation, or by satisfying governing equations of the flow dynamics. Both prior structures may be realized using neural

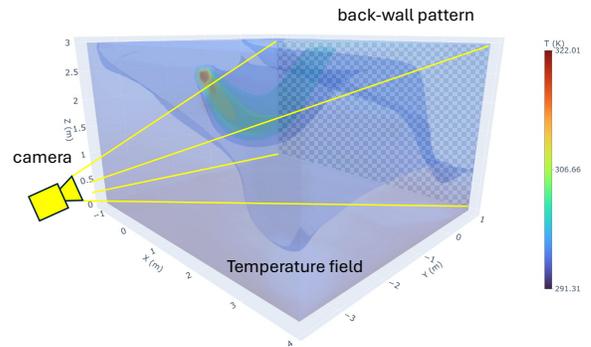


Fig. 1. Illustration of the imaging setup where a camera observes a patterned back-wall through a time-varying heterogeneous temperature field.

implicit representations [18–22]. In neural implicit refractive tomography (NIRT), a coordinate-based neural network takes spatial inputs and predicts the target field, such as the refractive index or density distribution. Zhao et al. [18] and Molnar et al. [19] examined NIRT for single-view reconstructions, while He et al. [20] and Li et al. [21] expanded the method to the multi-view setting. Molnar et al. [22] introduced a spatiotemporal formulation that jointly reconstructs 3D fields over an entire time series of projections, thereby embedding prior knowledge of flow dynamics into the recovery process.

Maintaining consistency with the BOS measurements is critical for ensuring accurate flow reconstruction [23]. The prevailing techniques include Tomographic BOS (TomoBOS) and Unified BOS (UBOS). TomoBOS consists of deflection sensing, spatial reconstruction, and a Poisson solve, whereas UBOS merges these into a single optical flow–deflection inversion leading to potentially improved density field estimation [5]. More recently, a rendering-based BOS metric developed by Teh et al. [24] demonstrated improved sensitivity to small gradients in the refractive field. They also incorporated a PINN to parameterize a steady state flow field and regularize the reconstruction.

In this paper, we adopt the rendering-based BOS technique and improve upon Teh et al.’s formulation by utilizing a time-dependent PINN architecture for modeling dynamic flow fields. We regularize the flow reconstruction using the residuals of the coupled Boussinesq approximation of the time-dependent incompressible Navier–Stokes and heat transfer equations. Moreover, we explore a simple data-driven closure strategy in which the effective thermal transport coefficients are learned over time from the BOS measurements, leading to more robust reconstructions when the coefficients prescribed in the regularization are not consistent with those of the true flow.

Kevin Tandi conducted this work during an internship at MERL.

2. RENDERING-BASED BOS IMAGE FORMATION

We consider a room containing a camera, an air inlet, and a patterned back-wall as illustrated in Fig. 1. The inlet introduces airflow at a temperature that differs from the ambient temperature, creating density variations in the flow field. When no flow is present, the camera captures a reference image I_{ref} of the background. When the flow is on, the refractive index η of the air changes, bending light rays and producing a distorted image I_{flow} . The refractive index η is linked to the gas density ρ through the Gladstone–Dale equation $\eta = 1 + G\rho$, where G is the Gladstone–Dale constant [25]. For small pressure variations indoors, the ideal gas law allows an approximation that expresses η directly in terms of temperature T :

$$\eta(T) = 1 + \rho_0 G T_0 / T, \quad (1)$$

where ρ_0 and T_0 are the ambient density and temperature, respectively [17]. Equation (1) highlights that changes in air temperature induce variations in density and thus refractive index, which are the source of optical distortions in BOS.

The rendering-based BOS formulation [24] differs from the classical BOS displacement fields by adopting a ray tracing perspective based on the refractive radiative transfer equation (RRTE) [26]. The trajectory of a ray with position \mathbf{x} and direction \mathbf{v} in an inhomogeneous medium satisfies

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

For a given sensor position and direction $(\mathbf{x}_s, \mathbf{v}_s)$, the nonlinear mapping to the back-wall is defined as $(\mathbf{x}_w, \mathbf{v}_w) = \mathcal{F}_{s \rightarrow w}(\mathbf{x}_s, \mathbf{v}_s, \eta)$, where the integrals along the path account for changes in η . Since exact integration is expensive, Teh et al. [24] introduce a quasi-linear approximation, in which the initial spatial trajectory is approximated by a straight line path $\bar{\mathbf{x}}(d) \approx \mathbf{x}_0 + d\mathbf{v}_0$, $d \in (0, D]$, in order to sample the refractive field along the line, where D is the distance to the back-wall. The ray directions $\tilde{\mathbf{v}}(d)$ are then computed for the η values sampled along the straight path $\bar{\mathbf{x}}$ and the final position of the ray $(\mathbf{x}_w, \mathbf{v}_w)$ is evaluated by integrating the path along the directions $\tilde{\mathbf{v}}$:

$$\tilde{\mathbf{x}}(d) \approx \int_0^d \tilde{\mathbf{v}}(\sigma) d\sigma + \mathbf{x}_0. \quad (3)$$

The resulting position and direction at the back-wall are used in the pixel intensity calculations, i.e., $(\mathbf{x}_w, \mathbf{v}_w) = \tilde{\mathcal{F}}_{s \rightarrow w}(\mathbf{x}_s, \mathbf{v}_s, \eta) := (\tilde{\mathbf{x}}(D), \tilde{\mathbf{v}}(D))$. This enables efficient evaluation of refractive field queries while retaining accuracy in estimating ray deflections.

The image formation model expresses the intensity of pixel j as

$$I_j = \int_{A_j} \int_{\Omega_j} W_j(\mathbf{x}_s) L_{\text{wall}}(\mathbf{x}_w, \mathbf{v}_w) \frac{\langle \hat{\mathbf{n}}_w, \mathbf{v}_w \rangle}{\|r_{s \leftrightarrow w}\|} d\mathbf{v}_s d\mathbf{x}_s, \quad (4)$$

where W_j denotes the pixel filter, L_{wall} the wall luminance, $\langle \hat{\mathbf{n}}_w, \mathbf{v}_w \rangle$ is a cosine factor related to the back-wall normal $\hat{\mathbf{n}}_w$, and $\|r_{s \leftrightarrow w}\|$ is the ray path length. The total intensity for the pixel integrates over all incoming directions, Ω_j , and the pixel area, A_j .

The full BOS forward model is encapsulated by the operator \mathcal{B} :

$$I_{\text{flow}} := \mathcal{B}(\eta), \quad (5)$$

which integrates camera geometry, refractive field, and source characteristics to generate the measured flow image. The operator \mathcal{B} evaluates the integral in (4) for all pixels j , thus providing a differentiable, rendering-based function of the airflow-induced temperature fields to synthesize the BOS measurements.

3. TIME-RESOLVED BOS TOMOGRAPHY

We define the dynamic airflow reconstruction problem as the task of determining the airflow fields denoted by the time varying temperature, pressure, and velocity fields $(T^*(t), p^*(t), \mathbf{u}^*(t))$ over some time interval $t \in [0, \tau]$. The airflow fields are observed by a set of BOS images $I_{\text{flow}}(k)$, $k \in \{1, \dots, K\}$ that capture temporal BOS snapshots at discrete time instances t_k . The BOS measurements are complemented with boundary conditions $(T^*(t), p^*(t), \mathbf{u}^*(t))|_{\Gamma}$ sampled at a boundary region Γ that includes the air inlet and a zero-velocity condition on all walls. Finally, the inversion is regularized by minimizing the PDE residuals of the coupled Boussinesq approximation of the time-dependent incompressible Navier–Stokes and heat transfer equations.

3.1. Inverse problem losses

To overcome the ill-posedness of single-view BOS tomography due to the inherent ambiguities along the view direction [18], we define an inverse problem with three loss terms: a schlieren loss \mathcal{L}_{BOS} , a boundary loss \mathcal{L}_{Γ} , and a physics-informed loss \mathcal{L}_{PDE} .

The schlieren loss $\mathcal{L}_{\text{BOS}}(\eta_k)$ is defined as the sum of the squared Euclidean distance between the observed images and the rendered images obtained using (5), i.e.

$$\mathcal{L}_{\text{BOS}}(\{\eta_k\}_{k=1}^K) = \sum_{k=1}^K \left\| I_{\text{flow}}(k) - \mathcal{B}(\eta_k) \right\|_2^2, \quad (6)$$

with $\eta_k = 1 + \rho_0 G \frac{T_0}{T(t_k)}$. The schlieren loss can be efficiently differentiated with respect to the refractive field η using automatic differentiation with adjoint-based gradients [18, 24, 27].

The boundary loss is imposed by minimizing a Euclidean distance between the computed fields and the measured fields at collocation points on the boundary Γ and for times $t \in \mathcal{T} \subset [0, \tau]$ in some collocation set \mathcal{T} :

$$\mathcal{L}_{\Gamma}(T, p, \mathbf{u}) = \sum_{t \in \mathcal{T}} \left\| (T_n^*(t), p_n^*(t), \mathbf{u}_n^*(t)) \Big|_{\Gamma} - (T_n(t), p_n(t), \mathbf{u}_n(t)) \Big|_{\Gamma} \right\|_2^2, \quad (7)$$

where subscript n denotes the fields normalized by dividing by their corresponding maximum value.

Lastly, we describe the physics-informed loss \mathcal{L}_{PDE} . Since we are dealing with time-varying flows, we assume the airflow to be a time-dependent, incompressible, Newtonian fluid that is governed by the Boussinesq approximation for buoyancy-driven flows. Specifically, the loss functions correspond to the norms of the residuals of the mass conservation, momentum conservation, and heat transfer equations as [17, 28–30]

$$r_{\text{mass}}(\mathbf{x}) = \nabla \cdot \mathbf{u}, \quad (8a)$$

$$\mathbf{r}_{\text{mom}}(\mathbf{x}, t) = \frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} + \nabla p - \frac{1}{Re} \nabla^2 \mathbf{u} + Ri T_{\text{nd}} \mathbf{e}_g, \quad (8b)$$

$$r_{\text{heat}}(\mathbf{x}, t) = \frac{\partial T_{\text{nd}}}{\partial t} + (\mathbf{u} \cdot \nabla) T_{\text{nd}} - \frac{1}{Pe} \nabla^2 T_{\text{nd}}. \quad (8c)$$

Above, ∇ and ∇^2 denote the three-dimensional gradient and Laplacian operators, respectively. The variable \mathbf{x} represents the nondimensional spatial coordinate normalized by a characteristic length L , and t corresponds to the nondimensional time normalized by a characteristic time τ . The velocity field $\mathbf{u} = (u, v, w)$ is nondimensionalized using a characteristic velocity U , while p corresponds to

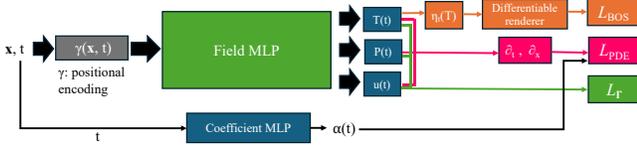


Fig. 2. Proposed architecture of the spatio-temporal neural representation for evaluating the BOS, PDE, and boundary losses.

the nondimensional pressure deviation from hydrostatic equilibrium, scaled by $p_0 = \rho_0 U^2$, with ρ_0 being a reference density. The nondimensional temperature perturbation T_{nd} is defined from the dimensional temperature T as $T_{\text{nd}} = \frac{T - T_0}{T_{\text{in}} - T_0}$, where T_{in} denotes the inlet temperature and T_0 the reference temperature. Other relevant quantities include the gravitational acceleration g and its unit vector e_g , the kinematic viscosity ν , the thermal diffusivity α , and the coefficient of thermal expansion β . These parameters lead to the standard nondimensional groups: the Reynolds, Péclet, and Richardson numbers, given by

$$Re = \frac{UL}{\nu}, \quad Pe = \frac{UL}{\alpha(t)}, \quad Ri = \frac{g\beta(T_{\text{in}} - T_0)L}{U^2}. \quad (9)$$

A key challenge in physics-informed BOS tomography is that the molecular thermal diffusivity α_{mol} does not adequately capture real indoor turbulent flows. To reduce mismatch between BOS data and PDE residuals, we learn a time-dependent scalar $\alpha(t)$, while keeping ν fixed since it cannot be reliably inferred from single-view BOS measurements.

The physics-informed loss is finally defined as a weighted sum of the residual norms as follows

$$\mathcal{L}_{\text{PDE}}(T, p, \mathbf{u}, \alpha) = \sum_{(\mathbf{x}, t) \in \mathcal{X} \times \mathcal{T}} \gamma_1 r_{\text{mass}}^2(\mathbf{x}) + \gamma_2 \|\mathbf{r}_{\text{mom}}(\mathbf{x}, t)\|_2^2 + \gamma_3 r_{\text{heat}}^2(\mathbf{x}, t; \alpha(t)), \quad (10)$$

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

3.2. Parameterization using PINN

Several works in the literature have shown that parameterizing the (T, p, \mathbf{u}) fields using a physics informed neural network (PINN) composed of a multilayer perceptron (MLP) architecture allows for a seamless imposition of the PDE constraints while also providing additional spatial regularization to the problem [17, 18, 24, 31]. We adopt this perspective, defining a neural operator \mathcal{M}_θ that maps from a 3D spatial position \mathbf{x} and time t input to the five dimensional field values $(T(t), p(t), \mathbf{u}(t))$ and thermal diffusivity parameter $\alpha(t)$ at time t :

$$(T(t), p(t), \mathbf{u}(t), \alpha(t)) = \mathcal{M}_\theta(\mathbf{x}, t), \quad (11)$$

where θ denotes the set of all network weights included in the architecture. The position and time inputs are first mapped using random Fourier feature embeddings [32], and then passed through three fully connected layers of width 64 with SIREN activations [33]. The final layer applies a tanh activation to constrain the outputs within their respective ranges. The time input t is also mapped through a sub-module in the network, dubbed the coefficient branch, to predict suitable values of the thermal diffusivity parameter $\alpha(t)$. The coefficient branch is composed of three fully connected layers of width 16 and tanh activations. The final layer also applies a sigmoid activation to constrain α to physically reasonable values. Figure 2 depicts

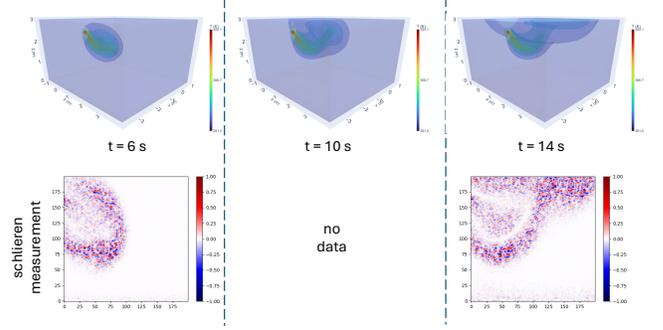


Fig. 3. Airflow visualization at $t = 6$ s, $t = 10$ s, and $t = 14$ s. The **top row** shows isosurfaces of the temperature field from the URANS simulation. The **bottom row** shows the corresponding synthetic BOS measurements at $t = 6$ s and $t = 14$ s. No BOS measurement is available at $t = 10$ s, and this intermediate state is used solely for evaluation.

the MLP architecture and illustrates how its outputs contribute to the different optimization losses. Because direct ray tracing through the MLP is computationally expensive, we adopt an intermediate procedure [24] in which the MLP is first sampled on a voxel grid; this grid representation is then used to perform ray tracing and evaluate the schlieren loss \mathcal{L}_{BOS} . In contrast, the MLP is sampled directly at spatio-temporal collocation points to compute the physics-informed loss \mathcal{L}_{PDE} and boundary loss \mathcal{L}_Γ .

3.3. Optimization formulation

The time-varying airflow fields can be reconstructed by solving the following optimization problem in terms of the MLP weights θ :

$$\begin{aligned} \min_{\theta} \quad & \lambda_1 \mathcal{L}_{\text{BOS}}(\{\eta_k\}_{k=1}^K) + \lambda_2 \mathcal{L}_\Gamma(\mathcal{M}_\theta) + \lambda_3 \mathcal{L}_{\text{PDE}}(\mathcal{M}_\theta) \\ \text{where} \quad & \eta_k = 1 + \rho_0 G \frac{T_0}{T(t_k)}, \end{aligned} \quad (12)$$

where the MLP network \mathcal{M}_θ is evaluated at collocation points $(\mathbf{x}, t) \in \mathcal{X} \times \mathcal{T}$ for some spatial position set \mathcal{X} in the computational and the temporal instance set \mathcal{T} . The multipliers $\lambda_{1,2,3}$ are scalars that balance the weight of each loss term.

4. EXPERIMENTAL RESULTS AND CONCLUSION

4.1. Transient airflow simulation and BOS imaging setup

We simulate a time-dependent indoor airflow scenario using the setup illustrated in Fig. 1. This configuration closely follows our previous work [24], where a steady Reynolds-Averaged Navier–Stokes (RANS) solver was used to model the flow. In this paper, we extend that setup to an unsteady RANS (URANS) simulation in order to capture the transient dynamics required for time-resolved BOS reconstruction. The room measures $5.1 \times 4.6 \times 3$ m, with its bottom-left corner located at $(-1.2, -3.6, 0)$ m. Ventilation is provided through a rectangular inlet of size 0.30×0.20 m oriented along the positive y -axis and positioned at the center of the $x-z$ plane at a height of 2.2 m. An outlet of size 0.30×0.10 m is located directly above the inlet at a height of 2.6 m. For schlieren imaging, we model a single virtual pinhole camera whose optical axis is aligned with the positive y -axis and placed at $(1.35, -29, 1.5)$ m with a long focal length to ensure the full volume is visible.

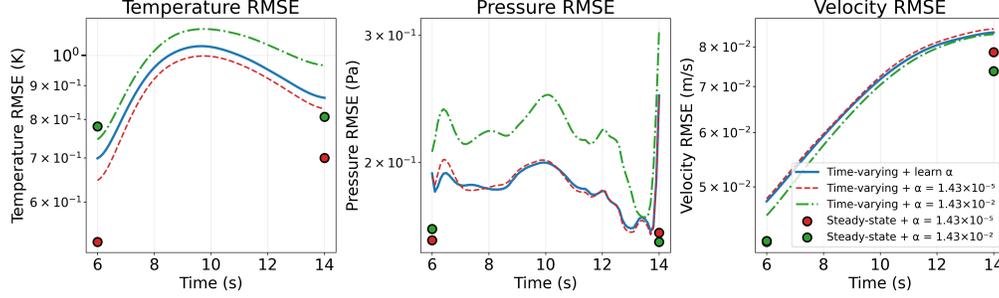


Fig. 4. RMSE of the reconstructed $(T(t), p(t), \mathbf{u}(t))$ fields over time for the adaptive and fixed thermal diffusivity α models. Also shown are the reconstruction errors from using a steady-state model evaluated at the two schlieren snapshots.

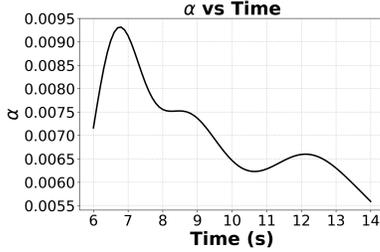


Fig. 5. Time profile of the learned thermal diffusivity α .

The URANS simulation employs a $k-\epsilon$ turbulence closure with coefficients $C_\mu = 0.09$, $C_1 = 1.44$, $C_2 = 1.92$, and $\sigma_\epsilon = 1.11$. Ambient conditions are $T_0 = 300$ K and $\rho_0 = 1$ kg/m³, with a coefficient of thermal expansion $\beta = 3e-3$ K⁻¹. The transport properties correspond to the molecular values of air, with a kinematic viscosity $\nu = 1e-5$ m² s⁻¹ and a thermal diffusivity $\alpha = 1.43e-5$ m² s⁻¹. The inlet jet introduces air with velocity $\mathbf{u}_{in} = (0, 2, -1.96)$ m/s and temperature $T_{in} = 321$ K.

4.2. Reconstruction results and conclusion

As shown in Fig. 3, we extract BOS measurements from the URANS simulation at $t = 6$ s and $t = 14$ s, which serve as input frames for reconstruction. The intermediate state at $t = 10$ s is held out and used for quantitative evaluation of the reconstructed flow. We first compare in Table 1 the performance of temperature, pressure, and velocity reconstruction with and without the BOS measurements. As in [24], we confirm that combining BOS measurements with the PDE and boundary conditions leads to better recovery of each field than the PDE and boundary conditions alone.

We next compare our proposed time-varying flow model to the steady-state model of Teh et al. [24] and investigate the effect of having α be a learnable coefficient. Table 2 and Figure 4 summarize the RMSE of the temperature field at the start and end times (with BOS measurements), as well as the performance over the continuous time interval with no BOS measurement. First, we observe that the steady-state model achieves the best performance at $t = 6$ s and 14 s when a good α value (e.g., $1.43e-5$) has been selected. However, the steady-state RMSE increases if a poor choice of $\alpha = 1.43e-2$ is used; moreover, it cannot recover the flow field for times without BOS measurements. Our proposed time-varying approach shows similar behavior: better performance depends on knowing a good value of α in advance. Furthermore, we show that using an additional MLP to estimate α tends to approach the RMSE of the good α , without knowing or searching for that value in advance.

Table 1. RMSE for time-varying model with and without BOS

	\mathcal{L}_{BOS}	\mathcal{L}_{PDE}	T (K)	p (Pa)	\mathbf{u} (m/s)
learned α	✓	✓	0.932	0.185	0.069
$\alpha = 1.43e-5$	✗	✓	1.357	0.193	0.07

Table 2. Reconstruction RMSE for T (K) at different times

	$t = 6$ s	$t = 10$ s	$t = 14$ s
time-varying + $\alpha = 1.43e-5$	0.647	0.997	0.831
time-varying + $\alpha = 1.43e-2$	0.746	1.095	0.966
time-varying + learned α	0.699	1.030	0.863

Figure 4 further demonstrates the utility of having a learned α coefficient by displaying the RMSE for each field over the time window $t = 6$ to 14 s. Intuitively, the temperature field has lower error for the endpoints, when BOS measurements have been provided. The pressure and velocity fields, which are inferred only through the PDE constraints, are less directly affected by the presence of BOS measurements. Note that the fixed α with the best performance is different for temperature and pressure than it is for velocity. Learning the thermal diffusivity coefficient α enables the reconstruction to approach or match the performance of good coefficient choices and adapt to changes over time. Figure 5 shows the learned value of α as it varies with time. The coefficient rises sharply after $t = 6$ s, reaching a peak of approximately 9.3×10^{-3} before gradually decreasing toward 5.6×10^{-3} at $t = 14$ s. This trend is physically plausible: α is elevated during the early transient regime when turbulence and mixing are strongest, then declines as the flow diffuses. The network’s recovery of this temporal variation without prior specification demonstrates the utility of learning $\alpha(t)$ as an effective closure parameter in the BOS inversion problem.

While in this work we restricted α to be a scalar function of time, future extensions could explore more expressive parameterizations. For example, modeling $\alpha(\mathbf{x}, t)$ as a function of local flow features could capture spatial variations in turbulent mixing that are not reflected in a global scalar. We also plan to evaluate our reconstruction on experimental schlieren measurements of airflow. Finally, another benefit of our time-varying approach is its computational efficiency relative to the single snapshot solution. Solving for the parameters of the single snapshot PINN with 80K iterations at each time instance required 6 hours 39 minutes to complete. This is compared to the time-varying PINN which can be solved with 80K iterations in 22 hours 7 minutes to recover the complete continuous time dynamics.

5. 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] S. J. Grauer and A. M. Steinberg, "Fast and robust volumetric refractive index measurement by unified background-oriented schlieren tomography," *Exp. Fluids*, vol. 61, Art. no. 80, Mar. 2020.
- [6] A. Unterberger and K. Mohri, "Evolutionary background-oriented schlieren tomography with self-adaptive parameter heuristics," *Opt. Express*, vol. 30, no. 6, pp. 8592–8614, 2022.
- [7] H. M. Lang, K. Oberleithner, C. O. Paschereit, and M. Sieber, "Measurement of the fluctuating temperature field in a heated swirling jet with BOS tomography," *Exp. Fluids*, vol. 58, pp. 1–21, 2017.
- [8] H. Liu, C. Shui, and W. Cai, "Time-resolved three-dimensional imaging of flame refractive index via endoscopic background-oriented schlieren tomography using one single camera," *Aerosp. Sci. Technol.*, vol. 97, pp. 105621, 2020.
- [9] 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.
- [10] 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.
- [11] G. E. A. Meier, "Computerized background-oriented schlieren," *Exp. Fluids*, vol. 33, pp. 181–187, July 2002.
- [12] G. S. Settles and M. J. Hargather, "A review of recent developments in schlieren and shadowgraph techniques," *Meas. Sci. Technol.*, vol. 28, 2017.
- [13] G. S. Settles, "On background-oriented schlieren (BOS) velocimetry," in *Proc. Int. Symp. Flow Vis.*, Zurich, Switzerland, June 2018.
- [14] 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.
- [15] 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.
- [16] 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.
- [17] 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.
- [18] B. Zhao, A. Levis, L. Connor, P. P. Srinivasan, and K. L. Bouman, "Single view refractive index tomography with neural fields," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2024, pp. 25358–25367.
- [19] 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 Journal*, vol. 62, no. 11, pp. 4316–4329, 2024.
- [20] Y. He, Y. Zheng, S. Xu, C. Liu, D. Peng, Y. Liu, and W. Cai, "Neural refractive index field: Unlocking the potential of background-oriented schlieren tomography in volumetric flow visualization," 2024, <https://doi.org/10.48550/arXiv.2409.14722>.
- [21] J. Li, X. Meng, Y. Xiong, T. Jia, C. Pan, and J. Wang, "Neural deflection fields for sparse-view tomographic background oriented schlieren," *Physics of Fluids*, vol. 36, no. 12, pp. 121701, 2024.
- [22] J. P. Molnar and S. J. Grauer, "Algorithm for time-resolved background-oriented schlieren tomography applied to high-speed flows," in *Proc. AIAA SciTech Forum*, 2025.
- [23] B. E. Schmidt, B. F. Bathel, S. J. Grauer, M. J. Hargather, J. T. Heineck, and M. Raffel, "Twenty-five years of background-oriented schlieren: Advances and novel applications," in *Proc. AIAA SciTech Forum*, 2025.
- [24] 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, pp. 1–5.
- [25] 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.
- [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] A. Teh, M. O'Toole, and I. Gkioulekas, "Adjoint nonlinear ray tracing," *ACM Trans. Graphics*, vol. 41, no. 4, Art. no. 126, July 2022.
- [28] A. Oberbeck, "Über die Wärmeleitung der Flüssigkeiten bei Berücksichtigung der Strömungen infolge von Temperaturdifferenzen," *Ann. Phys.*, vol. 243, no. 6, pp. 271–292, 1879.
- [29] 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.
- [30] M. Lappa, "Incompressible flows and the Boussinesq approximation: 50 years of CFD," *C.R. Mec.*, vol. 350, no. S1, pp. 1–22, 2022.
- [31] 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, Art. no. 14, Jan. 2023.
- [32] 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.
- [33] 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.