

High-fidelity modeling of shadowgraphy and schlieren imaging

By **D. Rossinelli, G. Li, A. Voci, L. Marshall, Jr., J. Hu, D. Brouzet, T. Fan, C. Williams, O. Martin, M. Khanwale, G. Vignat, T. Zahtila, M. Cutforth**
AND **G. Iaccarino**

We propose a novel algorithm that bridges the existing representation gap between simulations and experiments. From refractive index fields obtained by CFD, the algorithm produces shadowgraphy and schlieren images. Modeling fidelity is improved by a comprehensive image formation pathway, from light propagation through the medium, to optics and sensor response. With appropriate scale-resolving simulation capabilities, our results may be visually indistinguishable from their experimental counterparts. This allows the creation of a unified space where CFD simulations and experimental observations can be compared one-to-one. Of equal importance, the space facilitates CFD-driven augmentation of experimental data for the study of underlying dynamics, in turn offering a principled avenue for model inversion of imaged experiments.

1. Introduction and background

Unparalleled pace in semiconductor technology progress is driving a paradigm shift toward synergistically fusing CFD data with signals acquired in flow experiments. In experiments, flow visualization methods not only elucidate flow physics but also produce large volumes of data. Traditionally, the use of schlieren imaging and shadowgraphy is limited to qualitative diagnostics of flow features, and is not directly relatable to CFD quantities. At the same time, these imaging modalities are non-intrusive, ubiquitous, and able to produce information-rich multidimensional signals (Settles & Hargather 2017; Traldi *et al.* 2018).

In the past three decades, there have been a plethora of studies seeking to reproduce the experimental observations using computational modeling. Such a task is generally referred to as the computational flow imaging (CFI) problem (Havener 1994). CFI also enables an additional degree of control in visualization, which otherwise would be difficult or impossible to achieve from actual measurements (Brownlee *et al.* 2010). Separation of concerns (e.g., third-party ray tracers as black boxes, off-the-shelf screen-space filtering tools), and dedicated consumer workstations, seem to be fundamental traits common to the computational models proposed so far (Brownlee *et al.* 2010; Gori & Guardone 2018; Luthman *et al.* 2019; Chen *et al.* 2024).

Radically departing from conventional practice, we fuse knowledge from CFD, signal processing, and optics, to propose a novel vertical approach for *in silico* shadowgraphy and schlieren imaging. Moreover, we target the same HPC infrastructure used for CFD (e.g., many-core CPUs, GPUs, and interconnects) to produce images as seamless post-processing workloads.

1.1. Contributions and novelty.

This brief makes the following contributions:




Figure 1. Conceptual stages of numerical schlieren imaging: (1) collimation, (2) propagation through the index of refraction, (3) focusing and filtering of light by optical elements, and (4) sensing at the detector.

- A novel algorithm that takes CFD flow state as input and produces numerical schlieren and shadowgraph images as output,
- A qualitative assessment of the algorithm output on several flow problems, and
- A quantitative demonstration that such an algorithm creates a unified space where CFD and experimental images can consistently coexist.

Furthermore, it is underpinned by and describes the following three novel ideas:

- Meshing of beams to model illumination,
- Using a numerical scheme for solving the ray equation without inconsistencies, and
- Relying on a sampling scheme to compute sensed irradiance at the detector.

To the best of our knowledge, these three ideas have not been proposed or discussed previously, and they have proven to be central to high-fidelity optical imaging modeling. The algorithm has been built around contemporary computing microarchitectures, and targets both CPU-only and GPU-accelerated systems. Details about these aspects, however, will be elaborated elsewhere. This document is organized as follows: Section 2 reviews prior work. Section 3 presents the modeling and numerical schemes. In Section 4, the numerical solution of the ray equation is verified. Section 5 reports the experimental comparison and latent space analysis. Section 6 offers conclusions and an outlook.

2. Prior work

2.1. Background

Schlieren images and shadowgraphy rely on collimated light propagating through a test section (stage 2 in Figure 1), where the flow experiment is taking place, through the optical elements (stage 3) and onto a detector (stage 4), also called the film plane (Yates 1993; Brownlee *et al.* 2010), where the image is formed. Variations in the index of refraction (IoR, also denoted as n) in the test section causes light to deflect and result in an interplay of bright and dark shades at the film plane. If the test section is long enough, or ∇n is strong enough, incident light at the detector creates the so-called hidden phase in the wavefront, where light beams originated at different source positions hit the film plane at the same location. Moreover, in schlieren imaging, after the test section, beams are focused and then filtered with a blocker (stage 3 in Figure 1), such as a knife edge, before sensing. Depending on its shape and position, the blocker can behave like a low-pass, band-pass, or high-pass filter.

2.2. Simplifications

Rudimentary attempts at reproducing schlieren and shadowgraph images involve plotting functions of the magnitude of density gradients at spanwise flow locations or applying edge detection filters to the density field (Samtaney & Zabusky 2012). However, these approaches are insufficient for quantitative CFD-to-experiment validation, as comparisons between experimental and simulated data can only be carried out using visual perception metrics (Luthman *et al.* 2019).

For quantitative reproductions, light propagation through the flow field must be modeled explicitly. Aggressive analytical simplifications are possible when modeling the net light deflection as a function of density variations along the light path (Havener 1994). Alternatively, applying the paraxial approximation (Yates 1993) allows for very efficient computation. However, the underlying assumption of these approaches, that overall deflection is small, rarely holds in compressible or multiphase flow problems.

2.3. Photorealism

A conventional approach for representing light is to consider it as a pointwise photon beam source, with the goal of tracing a large number of beams from the source to the film plane. The intensity at each pixel is obtained by resampling the trajectories nearby (Gori & Guardone 2018; Brownlee *et al.* 2010; Luthman *et al.* 2019). However, even with careful sampling procedures, the underlying pointwise representation is inherently prone to spurious artifacts. A convergence analysis on the number of rays sampled is typically needed (Dias Paulino da Costa *et al.* 2021).

Some older work considered the light source in a more homogeneous manner, such that the light source is split into elements with non-zero surface area. The element deforms as it traverses the flow field, and the change in cross-sectional area is proportional to the variation of energy and intensity of the light element (Yates 1993). This method is conceptually similar to transporting a streamtube through the flow field (Ihrke *et al.* 2007). Finally, the effect of the knife edge is typically modeled by attenuating the radiance carried by the trajectory.

3. Method

A schematic of the image formation pathway proposed herein is shown in Figure 1, which involves modeling of

- Light beams (Section 3.1),
- Light propagation (Section 3.2),
- Refractive index (Section 3.3),
- Optical elements (Section 3.4), and
- Sensing (Section 3.5).

This section details the computational procedure required for each component of the framework.

3.1. Light beams

Conventionally, illumination is modeled with discrete point light sources (Figure 2, left). Here, we instead discretize the collimated light by employing a conforming mesh of (non-overlapping) triangular beams, as depicted in Figure 2 (right). Each beam is bounded by three boundary rays, its vertices. Adjacent beams share vertices, creating a watertight partition, as in beam tracing (Heckbert & Hanrahan 1984). Vertices, because of their




Figure 2. Light modeled as point sources by prior work (Luthman *et al.* 2019; Brownlee *et al.* 2010), conception of the beam mesh (center), and subsequent improvement to make illumination more isotropic (right).

infinitesimal size, no longer represent light sources. Light trajectories originate from vertices and obey the governing ray equation.

Since the light source is collimated, it can be represented as a Kronecker delta aligned to the collimating direction, i.e. the source radiance, and the latter is assumed to be uniform. The beam radiant intensity is proportional to its surface area A_{src} , formed by the planar triangle, and propagates in the direction normal to its surface. When a beam hits the detector, a triangle will be projected in the film plane. There will be a mismatch between the triangle area in its initial configuration, and the sensed one, A_{det} . The contribution to the incident radiation E_{det} from the beam at the detector is determined by an area ratio such that

$$E_{\text{det}} \propto \frac{A_{\text{src}}}{A_{\text{det}}}. \quad (3.1)$$

This novelty enables transformative improvements in image quality: we capture arbitrarily smooth or sharp intensities in a continuous representation, avoiding any need to deal with point sources. We are unsure why, to the best of our knowledge, this approach has not been proposed before. One possible explanation is that off-the-shelf software does not support it. One would have to extend the ray-tracing algorithms to enforce beam connectivity and consistent ray ordering, as it involves software modifications across multiple abstraction layers.

3.1.1. *Slit*

Numerical shadowgraphy works well even with idealized collimation, as beams will form smooth variations in the intensity gradient, and create an interplay of patterns with hidden phases, like in caustics. The light-blocking nature of schlieren, however, requires a slit of finite size at the light source, i.e. slightly uncollimated light, as explained in Settles (2001). To model a finite-size source aperture, we consider hundreds of slightly tilted directions. These directions are replicated in the triangle mesh with small angular perturbations that sample the source’s solid angle, and their detector contributions are summed. In the sampling limit, this reproduces the blur and penumbra of an extended emitter.

Outside the flow region, the light paths are linear due to the IoR being completely uniform, such as in the ambient environment, and experience a sudden deflection at the optical elements.

3.1.2. *Adaptive mesh refinement*

To ensure that the quality of the image is not significantly influenced by the discretization of the light source, a-posteriori adaptive mesh refinement of the light element




Figure 3. Light cone from the slit is modeled with hundreds of slightly different directions (left). A dedicated beam mesh is assigned to each tilted direction, since it represents a collimated light source (right).

is performed before generating the final image. This procedure first generates a schlieren image with an initial beam mesh, then identifies and refines the beams that contribute most to the error in the film plane. The procedure is repeated until the discrete nature of the triangular mesh has a negligible influence on the resulting image.

3.2. Light propagation

A light ray is modeled as a wave in the complex plane, and its propagation is governed by the Helmholtz equation, which ensues from Maxwell's system (Magnanini 2022; Gonzalez-Acuna & Chaparro-Romo 2024). Focusing on the real component of the Helmholtz equation, in the case of small wavelength, this simplifies to the eikonal equation (Luthman *et al.* 2019; Magnanini 2022; Lehujeur & Chevrot 2020; Gonzalez-Acuna & Chaparro-Romo 2024):

$$|\nabla\Psi| = n. \quad (3.2)$$

The solution, Ψ , is related to the optical path length and is the position of the wavefronts (Ihrke *et al.* 2007; Gonzalez-Acuna & Chaparro-Romo 2024) in the local refractive index n . Thus, the eikonal equation is a partial differential equation with nonlinearity that represents the behavior of the Helmholtz equation in the large wavenumber asymptote. This problem formulation assumes that for the wavelength under consideration, there is neither absorption nor emission, and is aligned with scenarios involving weakly participating media.

Light rays described by the eikonal equation travel at constant speed and form a wavefront. A reparameterization of a ray path by its arc length leads to the ray equation (Ihrke *et al.* 2007; Gonzalez-Acuna & Chaparro-Romo 2024):

$$\frac{d}{ds} \left(n \frac{d\vec{r}}{ds} \right) = \nabla n. \quad (3.3)$$

Here, s is the arc length parameter and \vec{r} represents the position of the light ray. Hence, $d\vec{r}/ds$ corresponds to the direction in which light propagates. Under arc length parameterization and wavefront tracking, the properties of $\frac{d\vec{r}}{ds}$ are captured by

$$\frac{d\vec{r}}{ds} = \frac{\nabla\Psi}{|\nabla\Psi|} = (\cos\alpha, \cos\beta, \cos\theta), \quad (3.4)$$

where α , β , and θ are the angles between the ray direction and the Cartesian axes. The above equation indicates that $d\vec{r}/ds$ is a unit length vector and its components are the optical direction cosines. Solving for Eq. (3.3) is the fundamental procedure for describing the trajectory of light in the flow field. To this end, different methods have been proposed

(Sharma *et al.* 1982; Sakamoto 1987; Ihrke *et al.* 2007; Brownlee *et al.* 2010). In Euler methods, the equation is reparameterized through a differential transformation for

$$dt = n^l ds, \quad (3.5)$$

where l is an integer constant. Hence, the optical ray vector can be defined as:

$$\vec{v} = n \frac{d\vec{r}}{ds} = n^{l+1} \frac{d\vec{r}}{dt}. \quad (3.6)$$

Thus (3.3) can be recast into:

$$\begin{aligned} \frac{d\vec{r}}{dt} &= \frac{\vec{v}}{n^{l+1}}, \\ \frac{d\vec{v}}{dt} &= \frac{n\nabla n}{n^{l+1}}. \end{aligned} \quad (3.7)$$

This system of ordinary differential equations (ODEs) can be solved with numerical recipes, such as the fourth-order Runge-Kutta (RK4) scheme. Depending on the choice of l , the reparameterization leads to various inherent properties, and the independent parameter t is not necessarily the arc length. We focus on the case of $l = 0$, which is one of the most popular choices since, here $t = s$ is the arc length.

This widely used approach, however, suffers from a consistency problem in the presence of strong gradients in the refractive index: the optical vector \vec{v} may completely reverse direction, in turn forming images with sudden and unrealistic decreases in intensity. We eliminate this fundamental inconsistency by proposing a Taylor series-based method, which enforces the unity of the optical ray vector and introduces the concept of a curvature vector. By defining $\vec{w} = \frac{d\vec{r}}{ds}$ as the directional cosine vector, Eq. (3.3) can be decomposed into another system of equations:

$$\begin{aligned} \frac{d\vec{r}}{ds} &= \vec{w}, \\ \frac{d\vec{w}}{ds} &= \frac{1}{n} [\nabla n - (\nabla n \cdot \vec{w}) \vec{w}]. \end{aligned} \quad (3.8)$$

This system can also be solved with various numerical schemes. This novelty offers the consistency advantage that even in the presence of strong IoR gradients, crucially important in the case of flame proximity, the optical direction remains consistent and does not experience any reversal. To the best of our knowledge, this approach has not been proposed before.

In addition to the RK4 integrator for the trajectories, we exploit the reversible nature of Eq. (3.3) by leveraging a local, three-step back-and-forth error compensation and correction (BFEC) scheme (Dupont & Liu 2003) to considerably improve accuracy.

3.3. Refractive Index

The refractive index field is used to compute trajectories of light propagating through the flow field, and is numerically reconstructed from the CFD simulation output. The local refractive-index field n is computed using

$$n = 1 + \sum_k \rho Y_k G_k, \quad (3.9)$$

where Y_k denotes the species mass fractions and G_k the Gladstone-Dale coefficient of the k -th species modeled in the flow field. Since the path of the rays will most likely not




Figure 4. Impulse sampling of irradiance field contributed by a single projected beam (left), box sampling for the same configuration (center), and a conceptual illustration of how to carry out the intersection quadrature efficiently (right).

intersect with grid nodes, interpolative schemes compute the IoR value and its gradients at arbitrary spatial locations. In prior work, computing the IoR gradient at nodes typically requires approximation procedures that directly operate on the discrete nodal samples (Samtaney & Zabusky 2012; Dias Paulino da Costa *et al.* 2021; Gao *et al.* 2022; Kouchi *et al.* 2025). Different interpolating kernels have been considered, including the box kernel (Brownlee *et al.* 2010; Gori & Guardone 2018), linear tent (Tamura & Fujii 1990; Jiang & Takayama 1999; Huang *et al.* 2016), radial basis function (Dias Paulino da Costa *et al.* 2021), and distance-averaged interpolation function (Liu *et al.* 2020). Efficient search algorithms may assist locating the position of the light beam with respect to vertex positions of the computational grid, which tracks the intersection of the beam with cells as it traverses through the flow medium (Tamura & Fujii 1990; Yates 1993; Jiang & Takayama 1999), by involving data structures such as kd-tree or max-min Octree (Gori & Guardone 2018; Brownlee *et al.* 2010).

To achieve superior image quality, we leverage spectral upsampling schemes for n and spectral derivatives for ∇n (3D DCT types II and III for upsampling, and zero-padding + 3D DFT for derivatives). We note that high-accuracy spectral derivatives are limited to grids that have a Cartesian logical space, such as curvilinear coordinates on structured mesh.

To sample n and ∇n when solving the ray equation, we use the cardinal cubic spline as the interpolating function. The cardinal cubic spline does not have compact support, but it can be factorized as an infinite impulse response (IIR) filter cascaded with a fourth-order B-spline finite impulse response (FIR) filter (Unser *et al.* 2002), and the latter is based on a smooth piecewise tricubic volumetric reconstruction kernel with $4 \times 4 \times 4$ support.

IIR prefiltering as in Unser *et al.* (2002) computes the interpolating coefficients for n and ∇n (previously obtained through spectral differentiation). With a careful HPC implementation, this stage takes negligible time, as it leverages the regularity of the data access patterns. The implementation relies, in turn, on the separability of the tensor product used to extend B-splines to higher dimensions, data-level parallelism across sample runs, and fast transpositions.

We opt for cardinal cubic spline interpolation for two reasons. Firstly, the ∇ operator is a high-frequency amplifier, while the cardinal cubic function is a low-pass filter with a 0.88π /sample cutoff, which suppresses only near-Nyquist frequencies, which are presumably numerical noise from differentiation. Secondly, we accelerate online lookup by offloading workload once to the IIR prefilter. A single interpolation request can then be carried out very quickly ($4 \times 4 \times 4 = 64$ lookups), yet it is influenced by a cloud of about 4200 nearby grid points.

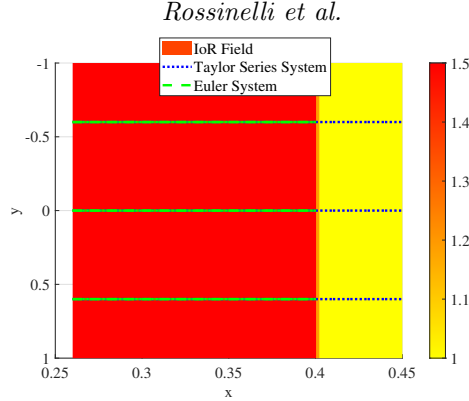


Figure 5. One-dimensional IoR with a sharp sigmoid function at $x = 0.4$. Trajectories of rays overlaid on top of the IoR are computed using the Taylor series system and the Euler system. These rays have different initial locations, but they all have initial directions that are perpendicular to the sharp change.

3.4. Optical elements

Lens focusing and blockers are a crucial stage of the algorithm, as illustrated in Figure 1 (stage 3). Once light propagates outside the flow problem under examination, we assume a uniform $n \approx 1$ surrounding medium. This stage takes very little computational time, yet it is as critical to outcome quality as the light propagation stage.

The lenses are assumed to be a homogeneous medium which enables light beams to traverse in a straight line, with $n = 1.5$. Hence, the interaction of individual rays with the lenses is completely governed by the sharp transition interfaces with undefined gradients, hence Eq. (3.3) is not applicable. We directly leverage Snell's law to compute the trajectories as light propagates across optical elements. Blockers, such as knife edges and orifices, create the necessary occlusion patterns to form schlieren images. This type of cutoff leads to a reduction in intensity. Further, if one of the vertices of the triangular beam deflects outside the coverage of the lenses or is blocked by the knife edge, then the contribution of that triangular element is excluded from the final image.

3.5. Sensing

In the experimental setup, sensor arrays probe the irradiance field at the film plane. To computationally emulate sensing, we must solve a resampling problem: the final image has to be represented with pixels, whereas the radiant intensity is represented as a linear superposition of projected triangles, with homogeneous but diverse intensities as described by Eq. (3.1). The main strategy is to employ an accumulation framebuffer. We iterate over each triangle and perform rasterization to obtain pixel fragments, and we reduce these by accumulating intensities in the framebuffer. This approach, however, is vulnerable to aliasing, since both refractive index and optical elements shear triangles considerably. An approach to effectively remove aliasing is to compute the area of a target pixel intersecting with the incident beam triangle, as shown in Figure 4. This is equivalent to convolving the incident beams with a box function, cascaded by an impulse train. Such a low-pass filter can be implemented very efficiently, resulting in comparable computational time to that of a classical rasterization, but results in a considerable improvement in image quality.




Figure 6. Shadowgraphy relying on impulse sampling of a beam's incident radiation (left), and improved approach based on box sampling(right).

4. Verification

Numerically solving the ray equation is the critical step in simulating the propagation of light through the medium of interest. For optimal performance in complex media, it is paramount that the employed numerical scheme is both highly accurate and consistent. In Section 3, two approaches were mentioned that permit recasting of Eq. (3.3) into a system of ODEs, namely, the Euler method and the approach we suggest, a Taylor series method. Both are advanced numerically using an RK4 scheme. We present a 1D verification problem with a pseudo-discontinuity in n featuring a sharp sigmoid function that is perpendicular to the ray paths. This case assesses the ability of the schemes to step through discontinuities, which may arise in practice from shocks or material change. The gradient is analytically defined. Thus, any sources of error in verification are attributable to the employed numerical schemes.

We consider a refractive index field varying in one dimension due to a sharp discontinuity, represented by

$$n(x, y) = \bar{n} - \frac{n_\delta}{1 + \exp(-\sigma(x - \tilde{x}))} \quad x \in [x_L, x_R],$$

where \bar{n} , n_δ , $\sigma \gg 1$, and $\tilde{x} \in [x_L, x_R]$ are constants. Here, the sigmoid function is used as a continuous approximation of the Heaviside step function representing a sharp change in IoR. This sharp change in IoR is practically motivated, specifically arising from physics such as the presence of shocks in high speed flows, flames in turbulent combustion, or the interface between dissimilar fluid phases in multiphase systems. The ray component that is perpendicular to the surface of the discontinuity must traverse through the discontinuity and remain orthogonal to the surface, since the refraction does not change its direction. This test case thus examines whether the numerical schemes are capable of producing ray paths that are consistent with the underlying physics. In our implementation, we set $\bar{n} = 1.5$, $n_\delta = 0.5$, $\sigma = 5000$, and $\tilde{x} = 0.4$. We consider three rays with initial positions of

$$\vec{r}_0^T = [x_0, y_0] = \begin{cases} [0.26, -0.6], \\ [0.26, 0], \\ [0.26, 0.6] \end{cases} \quad (4.1)$$

The initial directions for the three rays are set to $d\vec{r}/ds_0 = [1, 0]^T$, which indicate that the rays are parallel to the x -axis and perpendicular to the sharp step in the sigmoid




Figure 7. Qualitative comparison between real schlieren imagery (a) from Strelau *et al.* (2023), imagery formed with the present approach (b), and simplified in silico imaging widely used in literature (c) from Passiatore *et al.* (2024).

function. The step size for the arc length used in the numerical scheme is set as $\Delta s = 0.01$. The trajectories of the three rays simulated using the Euler system and the Taylor series system are shown alongside the surface plot of the IoR field in Figure 5. The plot demonstrates that our approach successfully computes the trajectory that traverses the sharp change, while the Euler system failed to do so.

Ultimately, without imposing an excessively small arc-length step size Δs , the Euler system may not be able to traverse through sharp changes for two reasons. The first is the lack of projection consideration in the ODE system (i.e. $\nabla n \cdot \frac{d\vec{r}}{ds}$ is not removed from the direction update), and the second is the lack of enforcement of the unity condition of the direction vector.

4.1. Sensing

Comparison of impulse sampling and box sampling approaches is shown in Figure 6, demonstrating the superiority of the latter.

5. Results

5.1. Laser ignition of gaseous methane/oxygen combustor

The proposed schlieren simulation technique was applied to obtain numerical schlieren visualization snapshots of the laser ignition process of gaseous methane/oxygen combustion, conducted in Passiatore *et al.* (2024). Crucially, experimental data is also available for the same setup in Strelau *et al.* (2023), enabling both a qualitative and a quantitative assessment of the performance of the proposed method. Figure 7 depicts the evolution of the flame kernel at different time instances. Experimental results are shown in the top row, results from the present method in the middle, and pseudo-schlieren from Passiatore *et al.* (2024) in the bottom row (consisting of slices of overlaid density gradient magni-

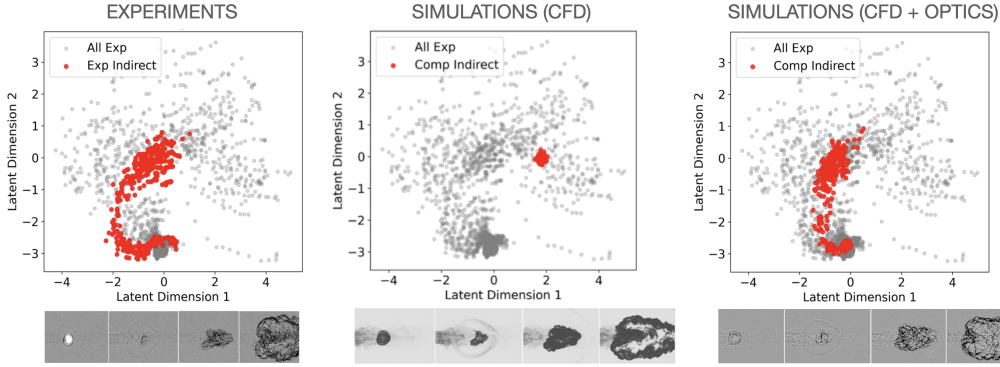


Figure 8. Trajectories of the turbulent combustion process in a rocket combustor represented in the latent space. The latent space is created by examining the experimental schlieren images with a GMVAE (left). CFD output followed by simplified modeling of the schlieren imaging is incompatible with the generated latent space (center), whereas CFD-based images formed with the present approach enable a unified investigation of the underlying dynamics (right).

tude). The superior quality of the current visualization is observed both in terms of the resolved scales and perceptual resemblance to experimental images. It is essential to note that the same source numerical fields are used to obtain results both for Passiatore *et al.* (2024) and the present work.

A further important assessment of the proposed technique was conducted to better characterize the underlying structure of the perceived flow field and its applicability in downstream machine learning tasks. To this end, a Gaussian mixture variational autoencoder (GMVAE) was trained (Fan *et al.* 2025). GMVAE extends the standard VAE by incorporating a mixture of Gaussian priors in the latent space, to learn disentangled, cluster-structured representations from experimental schlieren images. The GMVAE model employs a U-Net-based convolutional architecture for both encoder and decoder modules, enabling hierarchical feature learning across spatial scales. We first fit the GMVAE model to a dataset consisting only of experimental schlieren images. The trained model is then used to extract latent representations from both the numerical schlieren images and the CFD-based images formed with the present approach. Comparing the distributions of latent representations corresponding to ignition experiments and simulations, we observe that the CFD-based images formed with the present approach produce considerably more realistic and physically consistent features in the schlieren image space. The GMVAE was specifically chosen for its ability to extract meaningful latent representations while simultaneously preserving physical structure and clustering experiments based on physical states. The model, serving as a dimensionality reduction technique, projected the time series trajectories from the numerical schlieren images onto a latent manifold. Our objective was to use these latent representations to test whether the proposed schlieren method yields a more accurate representation of the underlying latent dynamics when compared to previous non-optics-based methods and experimental observations. As shown in Figure 8, the trajectories of the proposed method demonstrate significantly better agreement with experimental data and capture underlying dynamics uncaptured by images directly from CFD. Most importantly, this implies that available




Figure 9. Simulated hypersonic boundary layer visualized through, (a) shadowgraphy, (b) schlieren obtained with a horizontal knife edge, and (c,d) vertical knife edges (blocker on the left, and on the right, respectively).

simulation data can be used in conjunction with existing experimental schlieren for enhanced flow diagnostics and access to full three-dimensional virtual flow fields.

5.2. High-enthalpy turbulent hypersonic boundary layer

Figure 9 depicts computational schlieren of a turbulent hypersonic boundary layer (Williams *et al.* 2023), for which the hydrodynamic fields are obtained from high-fidelity direct numerical simulation (DNS) accounting for finite-rate chemistry and multicomponent diffusion. Representative of atmospheric re-entry at a freestream Mach number $Ma_\infty = 25$, the flow exhibits high temperatures ($T > 5300\text{ K}$). As such, nearly all molecular oxygen dissociates within the near-wall region, generating significant densities of radical species, including atomic oxygen, atomic nitrogen, and nitric oxide. Owing to the activation of both these complex aerothermochemical processes and multiscale turbulent mixing, as well as their mutual interaction (Williams *et al.* 2024), replication and optical inspection of these re-entry conditions in controlled ground experiments prove highly intractable. The combination of numerical simulation of light transport with fully scale-resolved hydrodynamics calculations presented here therefore provides a predictive complement to ground-test investigations of hypersonic flight, which may be hindered by a combination of limited test durations, optical access, or stagnation enthalpies.

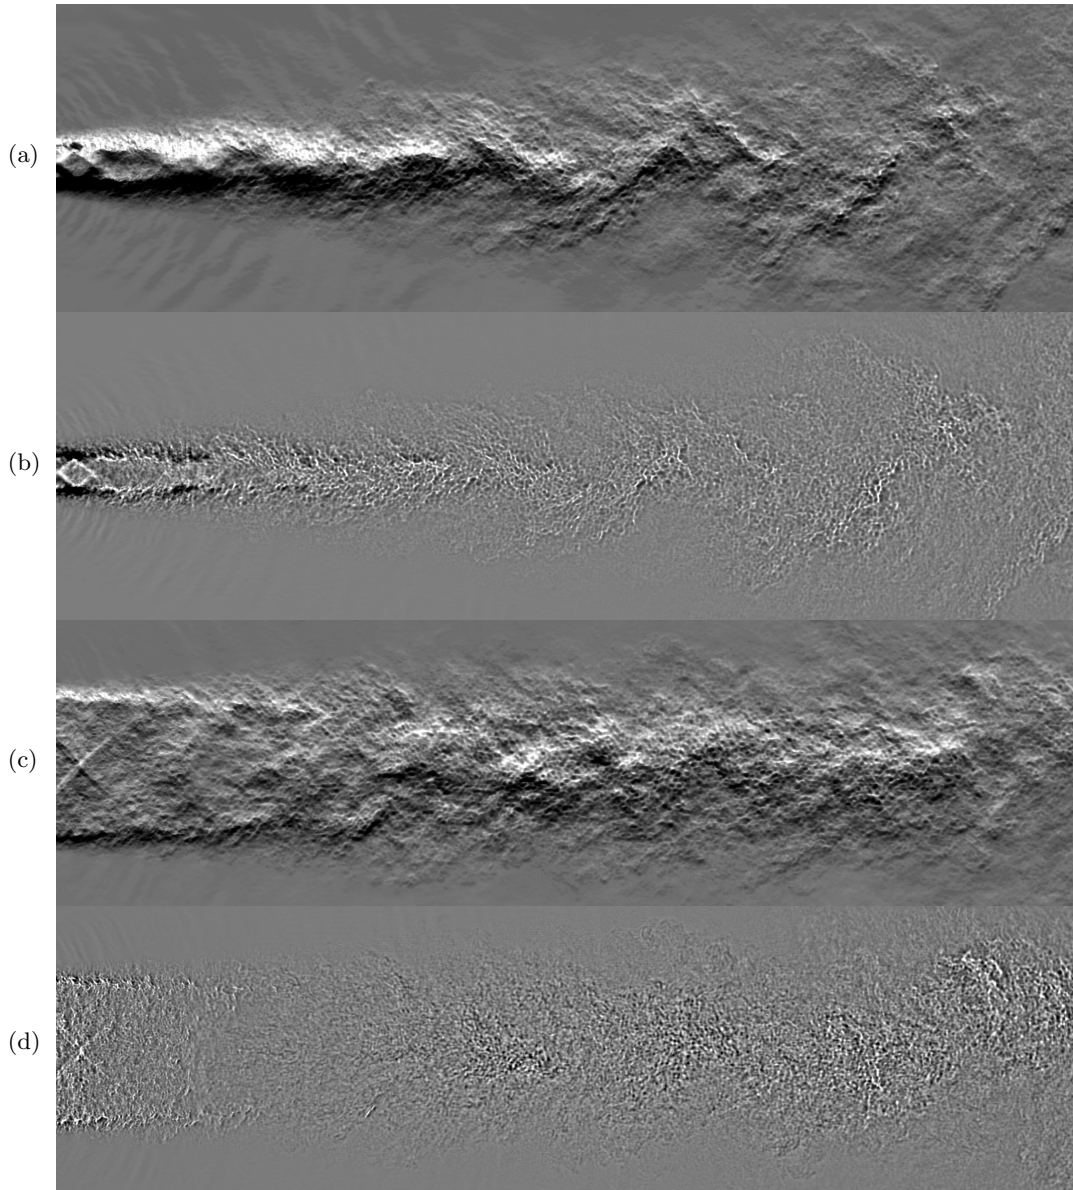


Figure 10. (a) schlieren and (b) shadowgraph visualizations of the supersonic rectangular jet, from the short nozzle side; (c) schlieren and (d) shadowgraph of the same flow problem but from the long nozzle side.

5.3. Supersonic rectangular jet

Figure 10 shows an instantaneous computational schlieren image of an aspect ratio 4:1 supersonic rectangular jet. The flow field originates from the high-fidelity large-eddy simulation (LES) performed by Wu *et al.* (2023). When operating at off-design conditions, a system of shock cells develops in the jet plume. Interaction between the shock cells and downstream-traveling Kelvin-Helmholtz (KH) waves gives rise to broadband shock-

associated noise (BBSAN) as well as intense tonal noise, known as jet screech (Tam 1995). In Figure 10, we observe intense flapping motions downstream of the jet, associated with the most unstable Kelvin-Helmholtz wave. The numerical schlieren can be used to directly visualize interactions between downstream-traveling instability waves and shocks producing freestream acoustic waves (Manning & Lele 1998). Additionally, numerical schlieren images provide a valuable means to validate the shock structure in the LES directly against experimental schlieren measurements, which is crucial for accurately predicting BBSAN and screech tones. In future work, we plan to apply the numerical schlieren to validate and analyze LES of aspect ratio 2:1 rectangular jets with secondary fluidic injection (Dai *et al.* 2025) as well as dual-stream jets with internal lobed mixers and plugs (Martin *et al.* 2024).

6. Conclusions and outlook

We presented a computational framework that maps the refractive index field from CFD to numerical schlieren and shadowgraph images through a vertical pathway that models light sourcing, light propagation in the test section, optics, and sensing. Three design choices were central to fidelity and robustness: representing a collimated source with a mesh of beams to determine irradiance by area ratios, solving for the ray equation without bounce-back inconsistencies, and modeling analog-to-digital sampling conversion at the detector. The results demonstrate superior fidelity image synthesis, a common image space for consistent comparison between CFD simulations and experimental observations, and a forward operator to assist data augmentation and model inversions through machine learning.

We carried out verifications of the approach against problems featuring invariants with exact closed expressions, and validations on CFD datasets against experiments of turbulent combustion flow problems. Verification confirmed conservation properties and the expected behavior of the propagation scheme. Validation supported quantitative agreement in image statistics and spatial structure. We also tested our approach against hypersonic boundary layers and supersonic jets.

Limitations of the present work include the application scope, restricted to weakly participating media and collimated illumination. It also assumes grids compatible with the differentiation and interpolation operators used herein.

We are already using such images for large-scale model inversions. Future work will extend the formulation to a broader class of media and optical configurations, including participating media, and will apply the method across diverse flows and sensing setups.

6.1. Acknowledgements

This investigation was funded by the U.S. Department of Energy’s National Nuclear Security Administration (NNSA) via the Stanford PSAAP-III Center for prediction of laser-induced ignition of rocket engines (Grant Number DE-NA0003968).

REFERENCES

- BROWNLEE, C., PEGORARO, V., SHANKAR, S., McCORMICK, P. S. & HANSEN, C. D. 2010 Physically based interactive schlieren flow visualization. In *IEEE Pacific Visualization Symposium (PacificVis)* (pp. 145–152).
- CHEN, M., LIU, B., QIN, F., LIU, X. & ZHOU, B. 2024 Physically based simulation

- of focusing schlieren imaging for a hypersonic boundary layer flow. *Appl. Opt.* **63**, A44–A51.
- DAI, S., GAUTAM, K., YAMASAKI, J., MOHAMMED, A., MARTIN, O. G., CRAMER, J., LELE, S. K. & GUTMARK, E. J. 2025 Exploration of supersonic jet noise mitigation for a rectangular jet using fluidic injection. *AIAA Paper* 2025-3286.
- DIAS PAULINO DA COSTA, F., GAITONDE, A. & JONES, D. 2021 Virtual experiments of light and shock wave interaction using nonlinear ray tracing and photon mapping. *AIAA Paper* 2021-1483.
- DUPONT, T. F. & LIU, Y. 2003 Back and forth error compensation and correction methods for removing errors induced by uneven gradients of the level set function. *J. Comput. Phys.* **190**, 311–324.
- FAN, T., CUTFORTH, M., D'ELIA, M., CORTIELLA, A., DOOSTAN, A. & DARVE, E. 2025 Physically interpretable representation and controlled generation for turbulence data. *arXiv preprint 2502.02605*.
- GAO, P., LI, T., YUAN, Y. & DONG, S. 2022 Numerical approaches and analysis of optical measurements of laser radar cross-sections affected by aero-optical transmission. *Infrared Phys. Technol.* **121**, 104011.
- GONZALEZ-ACUNA, R. G. & CHAPARRO-ROMO, H. A. 2024 The eikonal equation. In *Stigmatic Optics (2nd ed.)*, 2053-2563 (pp. 2–1 to 2–14). IOP Publishing.
- GORI, G. & GUARDONE, A. 2018 VirtuaSchlieren: A hybrid GPU/CPU-based schlieren simulator for ideal and non-ideal compressible-fluid flows. *Appl. Math. Comput.* **319**, 647–661.
- HAVENER, G. 1994 Computational flow imaging: fundamentals and history. *AIAA Paper* 1994-0259.
- HECKBERT, P. & HANRAHAN, P. 1984 Beam tracing polygonal objects. In *ACM SIGGRAPH Comput. Graph.* (pp. 119–127).
- HUANG, Y., SHI, G.-D. & ZHU, K.-Y. 2016 Runge–Kutta ray tracing technique for solving radiative heat transfer in a two-dimensional graded-index medium. *J. Quant. Spectrosc. Radiat. Transf.* **176**, 24–33.
- IHRKE, I., ZIEGLER, G., TEVS, A., THEOBALT, C., MAGNOR, M. & SEIDEL, H.-P. 2007 Eikonal rendering: efficient light transport in refractive objects. *ACM Trans. Graph.* **26**, 59.
- JIANG, Z. & TAKAYAMA, K. 1999 An investigation into the validation of numerical solutions of complex flowfields. *J. Comput. Phys.* **151**, 479–497.
- KOUCHI, T., KODERA, N., TANAKA, K. & SUZUKI, H. 2025 Investigation of instantaneous velocity-fields by schlieren image velocimetry using ray-traced simulated focusing-schlieren images. *AIAA Paper* 2025-0238.
- LEHUJEUR, M. & CHEVROT, S. 2020 On the validity of the eikonal equation for surface-wave phase-velocity tomography. *Geophys. J. Int.* **223**, 908–914.
- LIU, Y., YUAN, Y., GUO, X., SUO, T., LI, Y. & YU, Q. 2020 Numerical investigation of the error caused by the aero-optical environment around an in-flight wing in optically measuring the wing deformation. *Aerosp. Sci. Technol.* **98**, 105663.
- LUTHMAN, E., CYMBALIST, N., LANG, D., CANDLER, G. V. & DIMOTAKIS, P. E. 2019 Simulating schlieren and shadowgraph images from LES data. *Exp. Fluids* **60**.
- MAGNANINI, R. 2022 Explicit complex-valued solutions of the 2D eikonal equation. *Appl. Anal.* **101**, 3934–3946.

- MANNING, T. & LELE, S. 1998 Numerical simulations of shock-vortex interactions in supersonic jet screech. *AIAA Paper* 1998-0282.
- MARTIN, O. G., WU, G. J. & LELE, S. 2024 Noise predictions of a dual-stream jet with forced internal mixing. *AIAA Paper* 2024-4464.
- PASSIATORE, D., WANG, J. M., ROSSINELLI, D., DI RENZO, M. & IACCARINO, G. 2024 Computational study of laser-induced modes of ignition in a coflow combustor. *Flow Turbul. Combust.* **113**, 1055–1079.
- SAKAMOTO, T. 1987 Ray trace algorithms for GRIN media. *Appl. Opt.* **26**, 2943–2946.
- SAMTANEY, R. & ZABUSKY, N. J. 2012 Visualization, feature extraction, and quantification of numerical visualizations of high-gradient compressible flows. In *Flow Visualization* (pp. 339–366). Imperial College Press.
- SETTLES, G. 2001 *Schlieren and shadowgraph techniques: visualizing phenomena in transparent media*. Springer.
- SETTLES, G. S. & HARGATHER, M. J. 2017 A review of recent developments in schlieren and shadowgraph techniques. *Meas. Sci. Technol.* **28**, 042001.
- SHARMA, A., KUMAR, D. V. & GHATAK, A. K. 1982 Tracing rays through graded-index media: a new method. *Appl. Opt.* **21**, 984–987.
- STRELAU, R., FREDERICK, M., SENIOR, W. C., GEJJI, R. & SLABAUGH, C. D. 2023 Modes of laser spark ignition of a model rocket combustor. *AIAA Paper* 2023-2377.
- TAM, C. K. W. 1995 Supersonic jet noise. *Annu. Rev. Fluid Mech.* **27**, 17–43.
- TAMURA, Y. & FUJII, K. 1990 Visualization for computational fluid dynamics and the comparison with experiments. *AIAA Paper* 1990-3031.
- TRALDI, E., BOSELLI, M., SIMONCELLI, E., STANCAMPIANO, A., GHERARDI, M., COLOMBO, V. & SETTLES, G. S. 2018 Schlieren imaging: a powerful tool for atmospheric plasma diagnostic. *EPJ Tech. Instrum.* **5**.
- UNSER, M., ALDROUBI, A. & EDEN, M. 2002 B-spline signal processing. I. Theory. *IEEE Trans. Signal Process.* **41**, 821–833.
- WILLIAMS, C. T., DI RENZO, M. & MOIN, P. 2023 Direct simulation of turbulence-chemistry interaction in a strongly reacting turbulent hypersonic boundary layer. *Annual Research Briefs*, Center for Turbulence Research, Stanford University, pp. 173-184.
- WILLIAMS, C. T., DI RENZO, M. & MOIN, P. 2024 Statistical decomposition of turbulence-chemistry interactions in a hypersonic boundary layer. *Annual Research Briefs*, Center for Turbulence Research, Stanford University, pp. 79-92.
- WU, J. G., LELE, S. K., JEUN, J., GUSTAVSSON, J., SELLAPPAN, P. & KUMAR, R. 2023 Unstructured large-eddy simulations of rectangular jet screech: Assessment and validation. *AIAA J.* **61**, 1224–1237.
- YATES, L. A. 1993 Images constructed from computed flowfields. *AIAA J.* **31**, 1877–1884.