Digital enhancement of experimental pipelines

By D. Rossinelli, J. Hu, L. Marshall Jr., N. Kozak, T. Zahtila, D. Brouzet, M. Cutforth, M. G. Mungal AND G. Iaccarino

For over a century, advances in sensing and data acquisition have enabled laboratory experiments to underpin fluid mechanics discoveries. We propose image enhancement strategies to assist in experimentation of complex flow phenomena using a machine learning-based model inversion. The creation of supervised training sets is the key to inversion feasibility, and in turn relies on from-first-principles high-fidelity numerical solution of the governing equations of both the experiment and sensing apparatus. We address three classes of image enhancement problems: experimental signal restoration, novel view synthesis for enhanced optical access, and cross-modality image synthesis to display intrinsic, unmeasured quantities. Our results focus on turbulent combustion flows, specifically rotating Schlieren images around the streamwise axis. The implications of this approach extend considerably beyond this specific modality and flow problem.

1. Introduction

Under the guiding paradigm of "simulate many, test once", CFD has transformed engineering practice by reducing cost and accelerating design iteration. CFD can be used for studying and predicting complex flow phenomena, but for a given problem it remains fundamentally constrained by the spatiotemporal scales it can resolve under a computing budget (Brouzet *et al.* 2026). Subgrid-scale models have proven effective to mitigate these limits, but it is virtually impossible to avoid the introduction of model-form uncertainty.

Over the last decade, the CFD community has proposed the use of machine learning (ML) to achieve better performance in forward modeling (Beck *et al.* 2019; Duraisamy *et al.* 2019; Duraisamy 2021), but the results have been both surprising and mixed. Some brief and pertinent observations include that learned operators are able to capture families of solutions on curated training distributions, yet degrade when changing boundary conditions, geometries, or forcing. Further, long horizon roll-outs may violate conservation or drift (Bar-Sinai *et al.* 2019; Li *et al.* 2020), or have difficulties across operating regimes (Kochkov *et al.* 2021).

Whilst a great deal of attention has gone to forward modeling, a less explored yet important unsolved issue in ML-assisted fluid mechanics is the gap between simulations and experimental observations. The motivation is readily apparent, because while high-fidelity CFD data are abundant, most laboratory measurements are produced by optical instruments, so the simulated fields cannot directly serve as priors that guide and constrain the interpretation of experimental data. A unified space in which CFD outputs and experimental measurements coexist in a consistent fashion may be key to fundamentally reshaping how ML can assist fluid mechanics.

Given that ML-assisted fluid mechanics advances in the context of the present wave of ML revolutionizing many scientific domains, it is worth pointing out that to the best of our knowledge, the current trend was sparked by early successes in solving biomedical image processing tasks (Ciresan *et al.* 2012, 2013). Although somewhat simpler than full

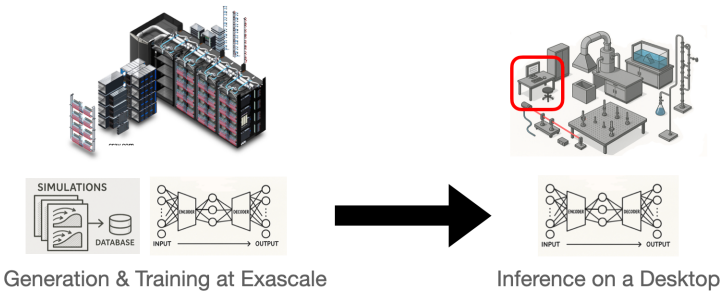

Figure 1. Exascale-to-lab pipeline, from supercomputing-grade set generation and training on flow physics and optics (left), to workstation-grade in-lab inference on experimental imagery (right).

reconstructions, those successes with both image classification and image segmentation represent inverse problems, not forward problems.

The present work realigns with the fundamental assumption that ML is inherently suitable at solving inverse problems. Radically departing from the many attempts to improve CFD with ML, here we enhance real data acquisition in experiments by leveraging a synthetic supervised training set, created with high-fidelity simulation counterparts. We elaborate on three classes of image enhancement:

- Signal restoration. Predict the signal that the same sensor would have acquired for the same experimental realization under ideal measurement conditions.
- Novel view synthesis (NVS). Predict the signal that the same sensor, placed at a different position, would have acquired for the same experimental realization.
- Cross-modality image synthesis. Predict the signal that a different sensor would have acquired for the same experimental realization.

Regarding the last point, hardware constraints, limited optical access, and run-to-run variability make it often impossible to acquire all desired modalities simultaneously or to register them reliably after the fact.

We approach the three problem classes with one fundamental idea: starting with the output of high-fidelity simulations, create a supervised set of synthetic image pairs that faithfully replicates both the physics and the acquisition modality in laboratories, and use this set to learn the inverse relationship via ML training. With the trained model available, the idea is to deploy the inversion on measurements obtained experimentally (Figure 1). Training of the ML model is supercomputing-grade and can harness exascale systems, but once trained, inference can be carried out rapidly on a desktop workstation at the lab near the physical experiment. While our results are centered on NVS, we elaborate on all three problem classes when no experimental ground truth is available, and discuss the applicability of the same approach across diverse modalities and flow problems.

The structure of this work is as follows. Section 2 reviews related work on how ML is currently used to assist experimental flow imaging. Section 3 details the DEEP approach for NVS. Section 4 presents metrics, ablations, and CFD-to-experiment analyses. Section 5 provides insight into findings, limitations, and extends our approach to signal restoration and cross-modality image synthesis. Finally, Section 6 summarizes contributions and outlines directions for future work.

2. Prior work

Research on ML-assisted experimental fluid mechanics has followed three main veins: measurement fidelity improvement including through signal restoration, modeling enhancement through NVS, and cross-modality image synthesis (Vinuesa *et al.* (2023)). Accordingly we now review the progress in each vein.

2.1. *Signal restoration*

Signal restoration in fluid mechanics has drawn inspiration from pioneering work in computer vision, where learned priors have been shown to recover qualitative detail from noisy or under-sampled data (Dong *et al.* 2015). These reconstruction principles have found remarkable success in biomedical flow imaging (Long *et al.* 2023), where super-resolution has been applied to upscale details of blood flows captured with magnetic resonance imaging (MRI) at coarse resolution. Adapted to address similar challenges in fluid mechanics, notable successes include the enhanced reconstruction of largely under-resolved turbulent flow fields (Fukami *et al.* 2019) where convolutional neural networks (CNNs) were applied to coarse 2D images of flow around a cylinder to recover fine-scale vortical structure below the original resolution, flow reconstruction from limited sensor measurements (Callaham *et al.* 2019) and from experimental PIV data (Wei & Guo 2025; Jaroslawski *et al.* 2025).

2.2. *Novel view synthesis*

NVS is the synthesis of measurements that a sensor would record in different pose or configuration, and in fluid mechanics it has received sparse attention (Yousif *et al.* 2023). The key architecture here has been neural radiance fields (NeRFs) which seek to generate a field representation of the object, relying on line integrals through field quantities to determine novel views (Mildenhall *et al.* 2021). Some efforts have been aimed at NVS for Schlieren-style deflection fields using NeRFs (Li *et al.* 2024; He *et al.* 2025), but NeRFs require re-training and multiple views per scene. More recently, geometry-unaware NVS methods have been introduced: Watson *et al.* (2022) and Liu *et al.* (2023) proposed pose-conditioned diffusion models, performing NVS through single view data. Jin *et al.* (2024) proposed a transformer-based approach for fully data-driven NVS. Moreover, equivariance is expected to play a central role in NVS for physics-based systems, as it enforces that predictions transform consistently under prescribed groups of transformations of the input (Du *et al.* 2022). The need for such transformation-aware architectures has driven recent work in areas such as ray and light-field reconstruction for visualization (Xu *et al.* 2022, 2024), 3D shape reconstruction (Chatzipantazis *et al.* 2022), and molecular reconstruction (Arora *et al.* 2025), yet we are not aware of any explicit treatment in the context of NVS for experimental fluid mechanics.

2.3. *Cross-modality image synthesis*

Measurements of virtual sensors can be predicted by learning how to mix signals acquired from real sensors. In the context of ML, methods for cross-modal synthesis fall into four main classes:

- common latent space,
- knowledge distillation,
- conditional generative adversarial networks (cGANs), and

• mutual information maximization.

Shared latent space approaches encode modalities into a common representation framework and thereafter decode to the target, whereas knowledge distillation methods transfer learned representations from one modality to guide learning another. cGANs learn mappings by conditioning the generator on input data, enabling realistic target generation guided by using source modalities. Finally, mutual information maximization seeks to directly increase the shared information between input and output representations (Azad *et al.* 2025). Biomedical imaging offers extensive literature on missing modality synthesis and information-theoretic limits. These methods can directly be applied to fluid mechanics (McNaughton *et al.* 2023; Dayarathna *et al.* 2024). In fluid applications, physics-informed mappings from background-oriented Schlieren (BOS) density to velocity and pressure demonstrate that fusing measurement operators with governing equations is effective and recent advance using neural implicit fields corroborate this approach (Cai *et al.* 2021; He *et al.* 2025).

The application and improvement of cross-modality synthesis to turbulent combustion remain largely unexplored. One such application is the inference of the heat release rate (HRR) from chemiluminescence. This mapping is non-unique and highly non-linear with functional dependence on numerous flow and state properties such as pressure and species concentration among others. However, as flow becomes highly turbulent and thus highly non-linear, assumptions made in current estimation methods (Lauer & Sattelmayer 2010) no longer hold. One non-linear extension used LES to generate HRR and radical species fields, extracted proper orthogonal decomposition (POD) modes, and trained a multilayer perceptron (MLP) to map the modal coefficients (Li *et al.* 2021).

2.4. *Core idea of this work*

The fundamental strategy proposed herein is aligned with three recent works in particle-image velocimetry (PIV) and four works on 4D-flow MRI reconstruction. The recent PIV works feature a CNN-based dense motion estimator trained on artificial PIV datasets and then applied to laboratory images (Cai *et al.* 2019), a stereoscopic recurrent all-pairs field transform (RAFT) variant trained on multi-fidelity Reynolds-Averaged Navier-Stokes (RANS)/Direct Numerical Simulation (DNS) with synthetic PIV images and evaluated on wind-tunnel measurements (Elrefaie *et al.* 2024), and a deep preprocessing network trained on synthetic particle visualizations to denoise real PIV frames (Fan *et al.* 2023). In all three of these PIV contributions, the imaging model is limited to particle rendering with simple noise/contrast filters. Analogous ideas have been investigated in 4D-flow MRI, where CFD-derived synthetic acquisitions enabled ML-based reconstruction and upscaling of measured blood-flow fields (Long *et al.* 2023; Ferdian *et al.* 2020; Rutkowski *et al.* 2021). Although those studies target flows with comparatively fewer spatiotemporal scales, their success underscores the impact and relevance of simulation-driven, modality-matched training and motivates our exploration on higher-Re $(10,000 - 100,000)$, complex experiments with non-trivial optics.

3. Novel view synthesis for Schlieren imaging

Rotating Schlieren images around the streamwise axis is one instance of NVS, however with peculiarities:

• Single-view synthesis. We have only a single viewpoint, as opposed to classical NVS methods relying on many viewpoints to achieve accurate field reconstructions (Yu *et al.* 2021).

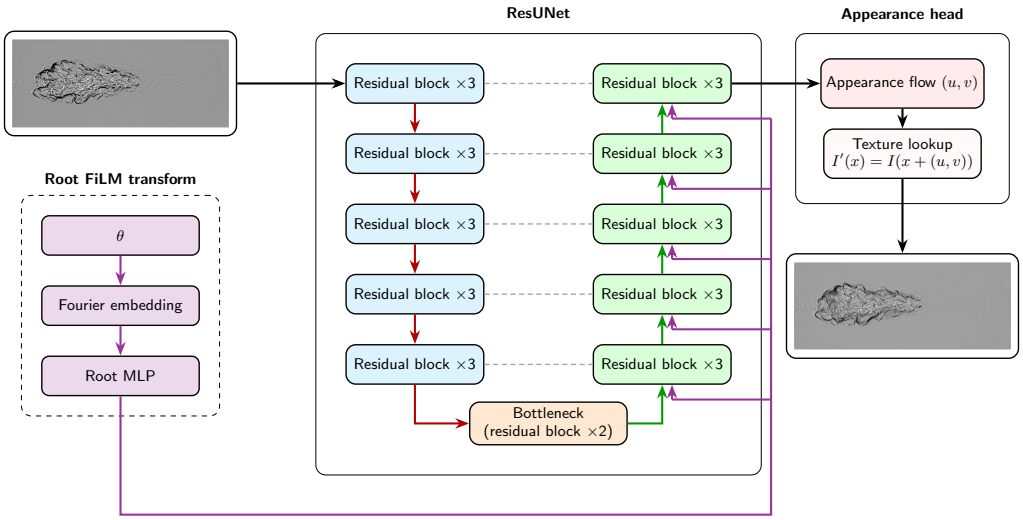

Figure 2. ResUNet with appearance flow and FiLM conditioning. The blue blocks denote the encoder, the orange block denotes the bottleneck, and the green blocks denote the decoder. The desired streamwise axis of rotation θ is processed through a root FiLM transform. Red arrows denote 2×2 max-pooling, and green arrows denote 2×2 upsampling by transpose convolutions. The dotted lines denote skip connections, and the purple lines denote linear projections yielding elementwise FiLM conditioning. As shown, the base image is processed through the ResUNet, with the FiLM conditioning applied in each decoder block. Finally, the appearance head yields a predicted rotated image through a texture lookup.

• Refractive physics. Schlieren imaging is governed by light deflection, whereas traditional NVS is often based on light interacting with Lambertian surfaces or absorption fields.

• Harder inverse reconstruction. Deflection-dominated line-of-sight measurements are highly nonlinear and severely ill posed. Recovering the 3D flow field from Schlieren images would be even harder than in absorption-based modalities.

• Generalization costs. Both NeRFs and Gaussian splatting are largely scene dependent, i.e., they still require some scene-specific re-training. The target use case herein is to process unseen real imagery of experiments, and we thus do not re-train.

To address the challenges our problem poses to conventional NVS methods, we instead propose to train direct image-to-image rotation, foregoing the intermediate field representation. Our model is constructed as follows:

• Appearance flow. We utilize the Ansatz that for small rotations, we can approximate the Schlieren rotation problem as a texture lookup problem, sampling pixels from the original image to construct our rotated view. Inspired by Zhou *et al.* (2016), where appearance flow was used for NVS, we train our model to perform texture lookup through appearance flow.

• ResUNet. The backbone of our network is a ResUNet. We utilize the UNet because of its well-studied performance on image segmentation. Residual blocks are introduced for deeper stabilized training (He *et al.* 2016).

• FiLM conditioning. To incorporate the desired angular rotations around the stream-

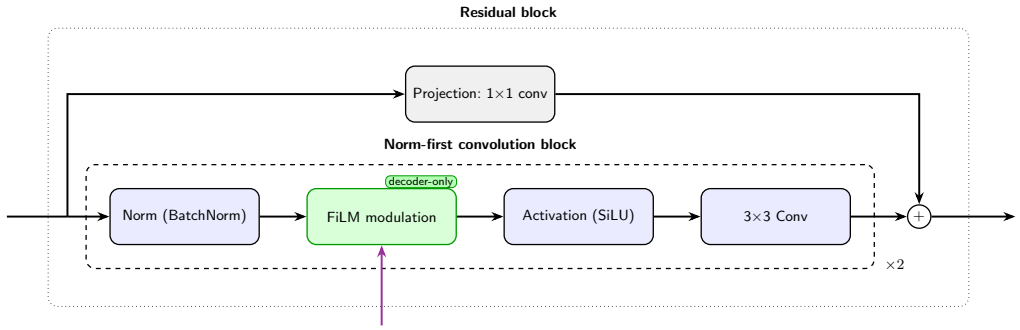

Figure 3. The residual block architecture, consisting of a residual projection around two norm-first convolution blocks. Note that the FiLM modulation is applied only as an affine transform on top of the BatchNorm in the decoder. The FiLM step is skipped in both the encoder and bottleneck.

wise axis into our network, we utilize FiLM (feature-wise linear modulation) conditioning (Perez *et al.* 2018) throughout the decoder to modulate the appearance flow. FiLM provides a simple way to incorporate the viewing angle into the network.

An overview of our network architecture can be found in Figures 2 and 3. The input image is processed through the ResUNet, while the desired rotation θ is processed through an embedding layer and an MLP before being injected into the ResUNet decoder. Finally, the appearance flow head performs texture lookup to accomplish NVS.

We train with a minimum reprojection loss, originally proposed in Godard *et al.* (2019), due to its stable training in the presence of occlusions and disocclusions when warping images. The exact loss function, where α is set to $\alpha = 0.85$, is

$$L(I_{gt}, I_{pred}) = \frac{1}{2}\alpha(1 - \mathrm{SSIM}(I_{gt}, I_{pred})) + (1 - \alpha)||I_{gt} - I_{pred}||_1, \tag{3.1}$$

where I_{gt} is the ground truth image representing an ideal NVS, I_{pred} is the actual NVS, and SSIM denotes the structural similarity index measure.

The training set for our model relies exclusively on forward models of Schlieren synthesis (Rossinelli *et al.* 2025) generated from simulations of the laser ignition of a gaseous methane–oxygen combustor (Brouzet *et al.* 2026) within the PSAAP III Center at Stanford, consuming about 3 TB of CFD data to produce roughly 1.2 TB of images. For each time snapshot, we generate 360 frames by rotating the Schlieren setup around the streamwise axis in 1° increments, so that all views for a given snapshot form a group of discrete rotations and any frame can be used to predict any other frame at the same time instant. Our full training dataset consists of three separate simulations with 215 total snapshots, and our evaluation dataset consists of a separate simulation with 29 snapshots.

4. Results

To get a practical sense of the use case and technical challenges, we encourage the reader to consult the online demo (`http://rossinelli.gitlab.io/rotating-schlieren`). Figure 4 shows ideal NVS (a) and achieved NVS (b) for different degree rotations of a Schlieren image taken from a simulation unseen during learning. The appearance flow

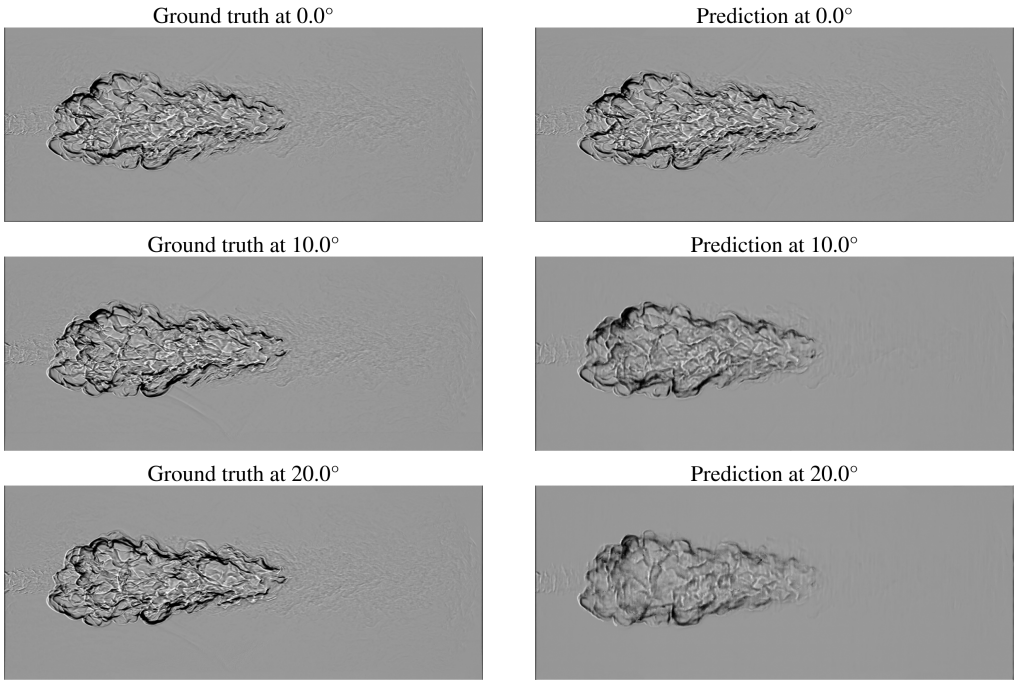

Figure 4. (a) Ground truth and (b) prediction target rotations for an unseen high-fidelity simulation. As shown, the predictions are able to capture the large-scale features of the target Schlieren image, although the smaller features begin smearing as we increase the rotation angle. See the interactive demo for a clearer demonstration of how the angular rotations appear smooth on inference.

allows for accurate reconstruction at low-rotation angles as it resamples the reference input. The angular rotations appear smooth, too, showing that the model is correctly utilizing FiLM to synthesize smooth rotation. The model performance decreases for larger angles, as small-scale features are smeared. Nonetheless, maintaining large-scale features allows a volumetric perception of the Schlieren image, giving insight into the underlying geometry of the refractive field.

Figure 5 displays the PSNR (peak signal to noise ratio) and SSIM against the rotation angle for the unseen simulation experiment, with the middle 90% represented through the bands. The SSIM at each pixel position is computed as in Wang *et al.* (2004) using a Gaussian filter of width $\sigma = 1.5$ over a window of width 11. Canonical stabilization coefficients $k_1 = 0.01$ and $k_2 = 0.03$ are used. At 0° rotation, the NVS task collapses to reproducing the original view, and the network matches the ground truth within practical accuracy. As the rotation angle deviates from 0°, the quality of the synthesized views deteriorates rapidly, with clearly visible image distortions. We attribute the large variance in PSNR and SSIM to the increasing reconstruction complexity for different timepoints of the simulation, yielding larger errors as the timestep, and thus task complexity, increases.

Following prior work that evaluates reconstruction fidelity in the frequency domain by comparing radially averaged power spectra (Durall *et al.* 2020; Dzanic *et al.* 2020), we also compute a radial power-retention score as

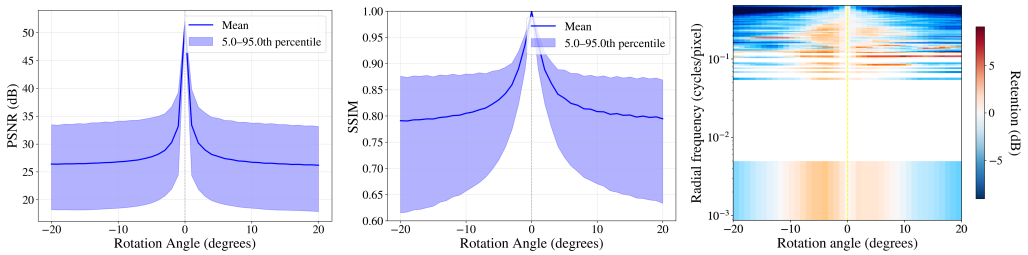

(a) (b) (c)

Figure 5. (a) Mean PSNR and (b) mean SSIM for a full unseen test simulation, and the (c) radial power retention for a single representative timestep, all plotted against rotation angle for model predictions compared to the ground truth. The sharp spike around 0° represents no rotation. The bands for the PSNR and SSIM plots denote the middle 90% samples.

$$R(\theta, k) = \frac{\mathrm{PSD}_{\mathrm{pred}}(\theta, k)}{\mathrm{PSD}_{\mathrm{ref}}(\theta, k)}, \tag{4.1}$$

where PSD is the radially averaged power spectral density (RAPSD). We compute the RAPSD using the mean over an annulus in the frequency plane. Figure 5 (c) shows a representative radial power-retention score for a single timestep as a heatmap between radial frequency and angular rotation, using the same simulation data as in Figure 4. Masked regions (white) indicate bins where the retention ratio is unreliable because the reference power is near zero or the annulus contains too few samples. The model exhibits traits of a low-pass filter: increasing $|\theta|$ attenuates high frequencies and preserves the amplitude of low frequencies. This result corresponds to a loss in fine-scaled features during rotation.

5. Discussion

Our results foresee the opportunity of using ML to assist the image enhancement of experimental data, showing a promising direction for future research. Our current NVS model for Schlieren images, despite its limitations, is able to provide some level of depth perception, giving insight into the three-dimensionality of the experimental data (see the online demo for a clearer view of these capabilities).

We believe there are several avenues along which we can make current modeling better:

- Snapshot NVS from time series. The current model makes predictions based only on single snapshots. Experimental Schlieren videos include temporal information, which might help the model. Understanding how to transform temporal data to extract information and reconstruct spatial data is an exciting avenue of future research.
- S^2 Rotations. Currently, we are training the model to rotate only along one axis. Training rotations across different axes is expected to improve not only the volumetric perception of the image, but also the model capability in understanding the true underlying geometry, alleviating some of the depth-perception issues seen during our experimental data validation.
- Heavy-duty architectures. Our current proof-of-concept model relies on the ResUNet backbone. However, recent years have seen heavier model architectures such as diffusion

(Ho *et al.* 2020; Song *et al.* 2020) and ViT (Vision Transformers) (Dosovitskiy 2020) are now central to image synthesis. Such models come with increased computational costs, but have been shown to generalize even with low geometric biasing (Jin *et al.* 2024). Under an image-to-image reconstruction framework, these appear promising future directions to further model capabilities.

• Spectral bias. The suppression of high-frequency components is strongly noticeable throughout our model validations. An explanation is that the appearance flow hypothesis breaks down at larger angles. Moreover, low-frequency bias is a well-reported phenomenon within the ML literature (Rahaman *et al.* 2019; Sitzmann *et al.* 2020). This, in turn, affects our benchmark, as the model must segment foreground and background flow features to reconstruct the optical flow. Alleviating the frequency bias (Khodakarami *et al.* 2026) might improve model performance.

• Deep perceptual metrics. Zhang *et al.* (2018) showed how using the deep features of VGG CNN networks trained on ImageNet as a similarity metric outperforms classical metrics in agreeing with human visual perceptions.

While our results focused on a specific use case, this approach may be applicable to modalities and flow problems well beyond Schlieren images for internal combustion flows. We briefly elaborate on two very different applications.

5.1. *Signal restoration & turbulence*

Significantly more challenging and equally if not more impactful, is the signal restoration of outdoor optical observations, such as ground-based observatories, where optical rays are distorted by refractive index fluctuations due to turbulence in the atmosphere. How can we correct for these effects? State-of-the-art approaches rely on temporal averages with ML-based dynamic weights (DeTurNet/DeTurb-style restorers and the ATSyn family) (Su *et al.* 2022; Li *et al.* 2023; Zhang *et al.* 2024). The issue is how to train these models since no ground truth is available. The common practice is CFD-unaware, with two procedural alternatives: (i) Zernike-expansion models that capture global aberrations but miss strongly local, anisoplanatic distortions (Noll 1976); and (ii) phase-screen simulators that draw 2D Gaussian random fields with Kolmogorov spectra (Roggemann *et al.* 1995). Because these generators are not grounded in flow physics, their biases propagate into the learned inverse. Both approaches fail to reproduce hidden-phase phenomena that emerge at long range with smoothly varying refractive indices, where distinct ray bundles hit the same detector pixel. In contrast, simulations that couple flow and optics reproduce this effect. Our approach assists these ML-based models by creating a CFD + optics synthetic dataset, like the example shown in Figure 6 (a).

A related opportunity arises in combustion diagnostics, where PIV is notoriously difficult to deploy in high-temperature reacting flows. Except for specialized and very costly setups, strong refractive-index gradients in the flame front warp and decorrelate the particle images, invalidating standard reconstruction pipelines. Within the DEEP framework, flow- and optics-consistent simulations can generate paired distorted and distortion-free realizations of flame PIV, providing the synthetic supervision needed to unlock end-to-end PIV reconstruction in regimes where conventional PIV processing is currently unfeasible.

5.2. *Cross-modality image synthesis (chemiluminescence)*

We cast the inference of instantaneous HRR from single-view OH^* chemiluminescence as an image-to-image mapping trained on CFD+optics pairs. The literature shows that OH^* is widely used as qualitative HRR marker, but proportionality is conditional and

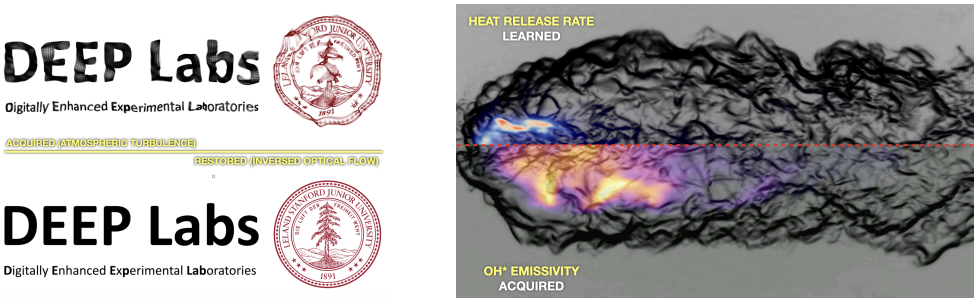

(a) (b)

Figure 6. (a) Long-range imaging through atmospheric turbulence, depicting warping/blur from the refractive-index fluctuations in the acquired frame (top) and learning by inversion in the restored frame (bottom); (b) cross-modality synthesis for combustion, depicting acquired OH* chemiluminescence in the acquired frame (bottom) and the learned heat-release-rate field (top), using training pairs generated from a CFD and optics synthetic dataset.

degrades with pressure, mixture, strain, and line-of-sight. Model-based corrections or ancillary information are often required (Docquier & Candel 2002; Hardalupas & Orain 2004; Nori & Seitzman 2009; Lauer *et al.* 2011; Sardeshmukh *et al.* 2017; Shen *et al.* 2022; Feng *et al.* 2023; Schiavone *et al.* 2024). We therefore synthesize camera-space OH^* measurements from reacting CFD with emissivity kinetics and spectral filtering and pair each frame with HRR to supervise the inverse learning. We then apply the learned map to acquired OH^* images. Figure 6 (b) illustrates a training sample pair: acquired OH^* (bottom) versus the learned HRR (top), trained on modality-matched synthetic data.

6. Conclusions

We introduced a modality-matched approach that couples high-fidelity CFD with explicit imaging forward models to synthesize training data and learn inverse maps in camera space. We elaborated on three problem classes: signal restoration, NVS, and cross-modality image synthesis. We demonstrated the idea on rotating Schlieren images as an example of NVS. By training and evaluating in the measurement modality, like-for-like comparison to experiments becomes amenable, and preliminary results indicate that high-fidelity forward modeling can be inverted via ML to gain insight into experimental data.

Ablation studies, as well as uncertainty quantification, are needed to assess the model robustness and sample complexity.

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

Arora, S., Pan, S.-H., Kumar, S., Singh, S. K., Chauhan, U., Alhalabi, W., Arya, V., Alsulami, B. S., Hsu, C.-H., Chui, K. T. & Gupta, B. B. 2025 CryoEMNet driven symmetry-aware molecular reconstruction through deep learning enhanced electron microscopy. *Sci. Rep.* **15**, 34584.

Azad, R., Dehghanmanshadi, M., Khosravi, N., Cohen-Adad, J. & Merhof, D. 2025 Addressing missing modality challenges in MRI images: a comprehensive review. *Comput. Visual Media* **11**, 241–268.

Bar-Sinai, Y., Hoyer, S., Hickey, J. & Brenner, M. P. 2019 Learning Data-Driven Discretizations for Partial Differential Equations. *Proc. Natl. Acad. Sci.* **116**, 15344–15349.

Beck, A., Flad, D. & Munz, C.-D. 2019 Deep neural networks for data-driven LES closure models. *J. Comput. Phys.* **398**, 108910.

Brouzet, D., Rossinelli, D., Zahtila, T., Voci, A., Strelau, R., Warner, A., Gejji, R. M., Slabaugh, C. D. & Iaccarino, G. 2026 Thermal expansion-driven laser ignition in a gas subscale rocket combustor. *Combust. Flame* **284**.

Cai, S., Wang, Z., Fuest, F., Jeon, Y. J., Gray, C. & Karniadakis, G. E. 2021 Flow over an espresso cup: inferring 3D velocity and pressure fields from tomographic background-oriented Schlieren via physics-informed neural networks. *J. Fluid Mech.* **915**, A102.

Cai, S., Zhou, S., Xu, C. & Gao, Q. 2019 Dense motion estimation of particle images via a convolutional neural network. *Exp. Fluids* **60**, 73.

Callaham, J. L., Maeda, K. & Brunton, S. L. 2019 Robust flow reconstruction from limited measurements via sparse representation. *Phys. Rev. Fluids* **4**, 103907.

Chatzipantazis, E., Pertigkiozoglou, S., Dobriban, E. & Daniilidis, K. 2022 SE(3)-equivariant attention networks for shape reconstruction in function space. *arXiv preprint 2204.02394* .

Ciresan, D. C., Giusti, A., Gambardella, L. M. & Schmidhuber, J. 2012 Deep neural networks segment neuronal membranes in electron microscopy images. In P. L. Bartlett, F. C. N. Pereira, C. J. C. Burges, L. Bottou & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012.* (pp. 2852–2860). Curran Associates Inc.

Ciresan, D. C., Giusti, A., Gambardella, L. M. & Schmidhuber, J. 2013 Mitosis detection in breast cancer histology images with deep neural networks. In K. Mori, I. Sakuma, Y. Sato, C. Barillot & N. Navab (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2013, Proceedings, Part II*, *Lecture Notes in Computer Science*, vol. 8150 (pp. 411–418). Springer.

Dayarathna, S., Islam, K. T., Uribe, S., Yang, G., Hayat, M. & Chen, Z. 2024 Deep learning based synthesis of MRI, CT and PET: review and analysis. *Med. Image Anal.* **92**, 103046.

Docquier, N. & Candel, S. 2002 Combustion control and sensors: a review. *Prog. Energy Combust. Sc.* **28**, 107–150.

Dong, C., Loy, C. C., He, K. & Tang, X. 2015 Image super-resolution using deep convolutional networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **38**, 295–307.

Dosovitskiy, A. 2020 An image is worth 16x16 words: transformers for image recognition at scale. *arXiv preprint 2010.11929* .

Du, W., Zhang, H., Du, Y., Meng, Q., Chen, W., Zheng, N., Shao, B. & Liu, T.-Y. 2022 SE(3) equivariant graph neural networks with complete local frames. In *Proc. Int. Conf. Mach. Learn.* (pp. 5583–5608). PMLR.

Duraisamy, K. 2021 Perspectives on machine learning-augmented Reynolds-averaged and large eddy simulation models of turbulence. *Phys. Rev. Fluids* **6**, 050504.

Duraisamy, K., Iaccarino, G. & Xiao, H. 2019 Turbulence modeling in the age of data. *Annu. Rev. Fluid Mech.* **51**, 357–377.

Durall, R., Keuper, M. & Keuper, J. 2020 Watch your up-convolution: CNN based generative deep neural networks are failing to reproduce spectral distributions. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 7890–7899).

Dzanic, T., Shah, K. & Witherden, F. 2020 Fourier spectrum discrepancies in deep network generated images. *Adv. Neural Inf. Process. Syst.* **33**, 3022–3032.

Elrefaie, M., Hüttig, S., Gladkova, M., Gericke, T., Cremers, D. & Breitsamter, C. 2024 On-site aerodynamics using stereoscopic PIV and deep optical flow learning. *Exp. Fluids* **65**, 183.

Fan, Y., Guo, C., Han, Y. & Qiao, W. 2023 Deep-learning-based image preprocessing for particle image velocimetry. *Appl. Ocean Res.* **130**, 103406.

Feng, X., Liang, H., Suo, J. & Zheng, L. 2023 OH* chemiluminescence characteristics of the RP-3 fueled dual-swirl direct-mixing combustor. *ACS Omega* **8**, 30716–30726.

Ferdian, E., Suinesiaputra, A., Dubowitz, D. J., Zhao, D., Wang, A., Cowan, B. & Young, A. A. 2020 4D Flow Net: super-Resolution 4D Flow MRI Using Deep Learning and Computational Fluid Dynamics. *Front. Phys.* **8**, 138.

Fukami, K., Fukagata, K. & Taira, K. 2019 Super-resolution reconstruction of turbulent flows with machine learning. *J. Fluid Mech.* **870**, 106–120.

Godard, C., Mac Aodha, O., Firman, M. & Brostow, G. J. 2019 Digging into self-supervised monocular depth estimation. In *IEEE/CVF International Conference on Computer Vision (ICCV)* (pp. 3828–3838).

Hardalupas, Y. & Orain, M. 2004 Local measurements of the time-dependent heat release rate and equivalence ratio using chemiluminescent emission from a flame. *Combust. Flame* **139**, 188–207.

He, K., Zhang, X., Ren, S. & Sun, J. 2016 Deep residual learning for image recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 770–778).

He, Y., Zheng, Y., Xu, S., Liu, C., Peng, D., Liu, Y. & Cai, W. 2025 Neural refractive index field: Unlocking the potential of background-oriented schlieren tomography in volumetric flow visualization. *Phys. Fluids* **37**, 017143.

Ho, J., Jain, A. & Abbeel, P. 2020 Denoising Diffusion Probabilistic Models.

Jaroslawski, T., Patil, A. & McKeon, B. J. 2025 Predicting turbulence structure in street-canyon flows using deep generative modeling. *Int. J. Heat Fluid Flow* **115**, 109849.

Jin, H., Jiang, H., Tan, H., Zhang, K., Bi, S., Zhang, T., Luan, F., Snavely, N. & Xu, Z. 2024 LVSM: a Large View Synthesis Model with Minimal 3D Inductive Bias. *arXiv preprint 2410.17242* .

Khodakarami, S., Oommen, V., Bora, A. & Karniadakis, G. E. 2026 Mitigating

spectral bias in neural operators via high-frequency scaling for physical systems. *Neural Netw.* **193**, 108027.

Kochkov, D., Smith, J. A., Alieva, A., Wang, Q., Brenner, M. P. & Hoyer, S. 2021 Machine Learning–Accelerated Computational Fluid Dynamics. *Proc. Natl. Acad. Sci.* **118**, e2101784118.

Lauer, M. & Sattelmayer, T. 2010 On the adequacy of chemiluminescence as a measure for heat release in turbulent flames with mixture gradients. *J. Eng. Gas Turbines Power* **132**, 061502.

Lauer, M., Zellhuber, M., Sattelmayer, T. & Aul, C. J. 2011 Determination of the heat release distribution in turbulent flames by a model based correction of OH* chemiluminescence. *J. Eng. Gas Turbines Power* **133**, 121501.

Li, B., Yao, W., Lee, Y. & Fan, X. 2021 Reconstruction model for heat release rate based on artificial neural network. *Int. J. Hydrogen Energy* **46**, 19599–19616.

Li, J., Meng, X., Xiong, Y., Jia, T., Pan, C. & Wang, J. 2024 Neural deflection field for sparse-view tomographic background oriented Schlieren. *Phys. Fluids* **36**.

Li, X., Chen, J., Ma, H., Gao, Q., Chen, Y., Li, J., Mao, H., Piao, Y., Gao, Y. & Xiao, G. 2023 A DeturNet-Based Method for Recovering Images Degraded by Atmospheric Turbulence. *Remote Sens.* **15**, 5071.

Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A. & Anandkumar, A. 2020 Fourier neural operator for parametric partial differential equations. *arXiv preprint 2010.08895* .

Liu, R., Wu, R., Van Hoorick, B., Tokmakov, P., Zakharov, S. & Vondrick, C. 2023 Zero-1-to-3: zero-shot one image to 3D object. In *IEEE/CVF International Conference on Computer Vision (ICCV)* (pp. 9298–9309).

Long, D., McMurdo, C., Ferdian, E., Mauger, C. A., Marlevi, D. & Nash, Martyn P.and Young, A. A. 2023 Super-resolution 4D flow MRI to quantify aortic regurgitation using computational fluid dynamics and deep learning. *Int. J. Cardiovasc. Imaging* **39**, 1189–1202.

McNaughton, J., Fernandez, J., Holdsworth, S., Chong, B., Shim, V. & Wang, A. 2023 Machine learning for medical image translation: a systematic review. *Bioengineering* **10**, 1078.

Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R. & Ng, R. 2021 NeRF: representing scenes as neural radiance fields for view synthesis. *Commun. ACM* **65**, 99–106.

Noll, R. J. 1976 Zernike polynomials and atmospheric turbulence. *J. Opt. Soc. Am.* **66**, 207–211.

Nori, V. N. & Seitzman, J. M. 2009 CH* chemiluminescence modeling for combustion diagnostics. *Proc. Combust. Inst.* **32**, 895–903.

Perez, E., Strub, F., De Vries, H., Dumoulin, V. & Courville, A. 2018 FiLM: Visual reasoning with a general conditioning layer. In *Thirty-Second AAAI Conference on Artificial Intelligence*. AAAI Press.

Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F., Bengio, Y. & Courville, A. 2019 On the spectral bias of neural networks. In *International Conference on Machine Learning (ICML)* (pp. 5301–5310). PMLR.

Roggemann, M. C., Welsh, B. M., Montera, D. & Rhoadarmer, T. A. 1995 Method for simulating atmospheric turbulence phase effects for multiple time slices and anisoplanatic conditions. *Appl. Opt.* **34**, 4037–4051.

Rossinelli, D., Li, G., Voci, A., Marshall Jr., L., Hu, J., Brouzet, D., Fan, T., Williams, C., Martin, O., Khanwale, M., Vignat, G., Zahtila, T., Cutforth, M. & Iaccarino, G. 2025 High-fidelity modeling of shadowgraphy and schlieren imaging. *Annual Research Briefs*, Center for Turbulence Research, Stanford University, pp. 89-104 .

Rutkowski, D. R., Roldan-Alzate, A. & Johnson, K. M. 2021 Enhancement of cerebrovascular 4D flow MRI velocity fields using machine learning and computational fluid dynamics simulation data. *Sci. Rep.* **11**, 10240.

Sardeshmukh, S., Lieuwen, T. & Seitzman, J. M. 2017 The use of OH* and CH* as heat release markers in combustion dynamics. *Int. J. Spray Combust. Dyn.* **9**, 317–333.

Schiavone, F. G., Aniello, A., Riber, E., Schuller, T. & Laera, D. 2024 On the adequacy of OH* as heat release marker for hydrogen—air flames. *Proc. Combust. Inst.* **40**, 105248.

Shen, W., Xing, C., Liu, L., Hu, Q. *et al.* 2022 Chemiluminescence-based characterization of heat release rate dynamic in a micro gas turbine combustion chamber. *Energy Inst.* **102**, 32–41.

Sitzmann, V., Martel, J., Bergman, A., Lindell, D. & Wetzstein, G. 2020 Implicit neural representations with periodic activation functions. *Adv. Neural Inf. Process. Syst.* **33**, 7462–7473.

Song, J., Meng, C. & Ermon, S. 2020 Denoising diffusion implicit models. *arXiv preprint 2010.02502* .

Su, C., Wu, X., Guo, Y., Zhang, S., Wang, Z. & Shi, D. 2022 Atmospheric Turbulence Degraded Image Restoration Using a Modified Dilated Convolutional Network. *IET Image Process.* **16**, 3507–3517.

Vinuesa, R., Brunton, S. L. & McKeon, B. J. 2023 The transformative potential of machine learning for experiments in fluid mechanics. *Nat. Rev. Phys.* **5**, 536–545.

Wang, Z., Bovik, A., Sheikh, H. & Simoncelli, E. 2004 Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* **13**, 600–612.

Watson, D., Chan, W., Martin-Brualla, R., Ho, J., Tagliasacchi, A. & Norouzi, M. 2022 Novel view synthesis with diffusion models. *arXiv preprint 2210.04628* .

Wei, L. & Guo, X. 2025 Deep learning framework for velocity field reconstruction from low-cost particle image velocimetry measurements. *Phys. Fluids* **37**.

Xu, Y., Chen, D., Liu, K., Zakharov, S., Ambruș, R., Daniilidis, K. & Guizilini, V. 2024 $SE(3)$ Equivariant Ray Embeddings for Implicit Multi-View Depth Estimation. *Adv. Neural Inf. Process. Syst.* **37**, 13627–13659.

Xu, Y., Lei, J. & Daniilidis, K. 2022 Equivariant light field convolution and transformer. *arXiv preprint 2212.14871* .

Yousif, M. Z., Yu, L., Hoyas, S., Vinuesa, R. & Lim, H. 2023 A deep-learning approach for reconstructing 3D turbulent flows from 2D observation data. *Sci. Rep.* **13**, 2529.

Yu, A., Ye, V., Tancik, M. & Kanazawa, A. 2021 pixelNeRF: Neural radiance fields from one or few images. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 4578–4587).

Zhang, R., Isola, P., Efros, A. A., Shechtman, E. & Wang, O. 2018 The unrea-

sonable effectiveness of deep features as a perceptual metric. In *IEEE/CVF conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 586–595).

Zhang, X., Chimitt, N., Chi, Y., Mao, Z. & Chan, S. H. 2024 Spatio-Temporal Turbulence Mitigation: a Translational Perspective. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.

Zhou, T., Tulsiani, S., Sun, W., Malik, J. & Efros, A. A. 2016 View synthesis by appearance flow. In *European Conference on Computer Vision (ECCV)* (pp. 286–301). Springer.