PHYSICS INFORMED NEURAL NETWORKS FOR REAL TIME 3D FLOW FIELD AND EMISSION PREDICTION IN LEAN HYDROGEN SWIRL COMBUSTORS

Prashant Suresh Kamble*[1], Avinash JHA*[1]

[1] *AeroMyne, Solapur, Maharashtra 413322, India*

Abstract:- Real-time monitoring and control of hydrogen-fueled gas turbines demand high-fidelity reconstruction of turbulent reacting flows and emissions. While conventional computational fluid dynamics (CFD) yields accurate velocity and thermal fields, the computational overhead precludes online digital twin integration. Purely statistical surrogates offer speed but violate physical conservation laws. To address these limits, we deploy a hybrid, data-assisted Physics-Informed Neural Network (PINN) mapping three-dimensional coordinates to velocity, pressure, temperature, and species fields within a co-swirling lean hydrogen-air combustor ($\phi = 0.66$). Chemical kinetics stiffness is bypassed by formulating the network to solve transport equations for conserved scalars (mixture fraction and enthalpy) while retrieving local properties and species from non-adiabatic Equilibrium PDF tables. The PINN is trained on a validated RANS dataset resolved via Menter's SST k-ω model with curvature correction. The surrogate preserves mass conservation within 0.5%, limits relative errors to under 3%, and executes in 3.0 ms on a CPU.

Keywords*:- *Physics-Informed Neural Networks (PINN), Hydrogen Combustion, Emissions Prediction, Swirling Reacting Flows, Equilibrium PDF Chemistry, Surrogate Modeling.

I. INTRODUCTION

Retrofitting utility-scale gas turbine systems to run on pure hydrogen is a primary pathway to eliminate industrial carbon emissions. Replacing conventional hydrocarbon fuels with hydrogen, however, triggers operational instabilities such as flashback, pressure oscillations from thermoacoustic coupling, and high thermal NO_x production in swirled injection systems. Optimizing these systems for reliable operation demands rapid, high-fidelity mapping of the three-dimensional aerodynamic and thermal fields across wide operating envelopes.

Establishing reliable multi-physics baselines for lean hydrogen-air combustion in swirled configurations requires robust numerical validation using Reynolds-Averaged Navier-Stokes (RANS) formulations. Coupling Menter's Shear Stress Transport (SST) k-ω curvature correction with non-adiabatic Equilibrium Probability Density Function (PDF) chemistry and multi-pathway NO_x transport models successfully captures spatial velocity and thermal gradients. This numerical approach resolves critical flow structures, including the central recirculation zone stagnation points, and yields grid-independent peak flame temperatures of 1943.74 K and exit NO_x concentrations of 0.135 ppm under a lean equivalence ratio of 0.66 [1].

Experimental characterization of swirl-stabilized flames under varying thermal loads and equivalence ratios provides critical datasets for validating numerical models. Laser Doppler velocimetry and planar laser-induced fluorescence (PLIF) of OH and CH radicals, combined with laser Raman scattering, map the spatial limits of flame stability, thermoacoustic oscillations, and lean blowout transitions. These database measurements reveal that swirl flames are typically detached from the fuel nozzle, resulting in localized fuel-air premixing in the high-velocity shear layers prior to chemical ignition [2].

Investigating the local thermodynamic state in swirl-stabilized flames highlights how turbulence-chemistry interactions govern lean combustion stability. Single-pulse Raman scattering measurements reveal that rapid mixing at the burner nozzle exit is accompanied by intense shear strain, inducing local flame extinction and ignition delay. Flame anchoring is primarily maintained by the transport of hot, relatively fuel-rich combustion products back to the nozzle lip via the inner recirculation zone [3].

Evaluating hydrogen-fueled dual-swirl burner configurations is essential to adapt aero-engine fuel injectors for zero-carbon operations. Optical measurements of OH^* radical chemiluminescence and infrared water vapor emission map the spatial spread and lift-off height of atmospheric, lean non-premixed hydrogen flames. Reacting Large Eddy Simulations (LES) of this setup show that while global flame shapes are captured, conventional combustion models struggle to resolve the exact shear-layer anchoring mechanisms, resulting in radial offsets in the predicted flame front [4].

Aerodynamic stabilization of non-premixed hydrogen-air flames in coaxial dual-swirl injectors shifts between distinct spatial modes depending on the thermal power. Transitioning from a $4\ kW$ anchored flame to a $10\ kW$ aerodynamically stabilized V-shaped flame is governed by the penetration of the inner recirculation zone into the injector nozzle, which induces radial divergence of the central hydrogen jet. Flow field diagnostics and Large Eddy Simulations reveal that while the anchored flame is entirely diffusion-controlled, the lifted regime exhibits a hybrid structure comprising a partially premixed branch and a trailing diffusion front [5].

Scientific machine learning provides a robust alternative to conventional combustion modeling by integrating governing equations directly into neural network structures. Applying physics-informed machine learning to combustion kinetics and reacting flows forces the optimization path to satisfy physical invariants through either soft loss-regularization or hard structural constraints. This paradigm shift enables the reconstruction of multi-dimensional flow fields and species concentrations within coordinate-to-variable and field-to-field training configurations [6].

Physics-informed neural networks (PINNs) offer a mesh-free path for accelerating emissions prediction and modeling thermoacoustic instabilities in sustainable aerospace propulsion systems. By embedding momentum and species transport equations into the backpropagation loop, PINNs achieve speedups of 2.3 to 4.9 times for global flame dynamics and 6.0 to 14.6 times for solving stiff reaction chemistry. To transition these neural solvers into actual flight control hardware and secure formal safety certifications, current efforts focus on wrapping predictions with uncertainty quantification and linking disparate multi-physics solvers [7].

For laminar flame setups, the FlamePINN-1D solver resolves spatial profiles by addressing both the forward eigenvalue speed predictions and the inverse retrieval of hidden parameters. The forward mode predicts flame speeds and species fields by enforcing physical boundary conditions. In the inverse mode, the system reconstructs continuous profiles and extracts transport or kinetic coefficients from sparse, noisy datasets. Applying hard boundary constraints combined with a thin-layer normalization strategy protects the training from gradient divergence, allowing stable learning in regions with steep reacting fronts [8].

Resolving reacting flows via physics-guided neural architectures is often hindered by the extreme multi-scale stiffness of chemical source terms. This multi-scale timescale stiffness causes gradient optimization paths to diverge, preventing the network from learning species concentration profiles. Developing architectures like Stiff-PINN resolves these optimization conflicts by implementing adaptive timescale weighting and scale-separation techniques to regularize the network [9].

In this work, we develop a hybrid, data-assisted Physics-Informed Neural Network (PINN) surrogate model tailored for real-time 3D flow field and emissions prediction in lean hydrogen swirl combustors. By formulating the network to solve transport equations for conserved scalars and integrating non-adiabatic Equilibrium PDF chemistry tables, the solver bypasses chemical stiffness while maintaining physical conservation. Subsequently, the manuscript reviews related literature, outlines the

governing physics and neural solver structure, and details baseline database configurations. This is followed by spatial validations against reference CFD data, computational speedup evaluations, and a summary of key findings.

II. LITERATURE REVIEW

Autonomously discovering reaction pathways from raw temporal data is achieved by embedding stoichiometry constraints directly within neural architectures. The chemical reaction neural network (CRNN) incorporates mass conservation and Arrhenius rate equations as functional layers rather than simple loss penalties. This design guarantees that the predicted species concentrations strictly respect elemental mass balance and thermodynamic limits [10].

Accelerating reacting flow simulations with detailed chemical kinetics is possible by coupling neural solvers with tabulated chemistry manifolds. Evaluating species fractions and thermophysical properties from a Flamelet/Progress Variable (FPV) table bypasses the need to solve stiff Arrhenius source terms during network training. This approach reduces optimization costs and stabilizes the training process in reacting systems [11].

Reconstructing turbulent scalar mixing fields is achieved by training deep networks on multi-dimensional concentration data. The neural model maps spatial coordinates to scalar values while approximating sub-grid scale fluctuations and local dissipation rates. This formulation confirms the capability of deep learning to capture sharp concentration gradients within turbulent shear layers[12].

Simulating reacting flows with multicomponent reactants is accomplished using mesh-free solvers constrained by transport equations. The neural solver evaluates species mass fractions, diffusion coefficients, and heat release terms simultaneously by minimizing multi-component transport residuals. This framework provides a verified computational surrogate for modeling complex multi-species reactions [13].

Extracting sub-grid closures and statistical properties for turbulent combustion is possible by training neural networks on point multiscalar measurements. The physics-guided network processes spatial data points and enforces local conservation equations to infer unmeasured correlation variables and turbulent transport terms. This methodology offers a reliable pathway to evaluate sub-grid scale statistics directly from sparse experimental points [14].

Standardizing coordinate scales and non-dimensionalizing variables are critical prerequisites for training stable physics-informed solvers. Normalizing spatial inputs and scaling governing equations prevents gradient pathologies where terms of different physical dimensions conflict during backpropagation. Implementing these normalization guidelines ensures reliable convergence on highly non-convex loss surfaces [15].

Resolving three-dimensional incompressible fluid flows in curvilinear coordinates requires specialized spatial discretization schemes. Employing a spherical finite-difference grid resolution resolves polar singularity boundaries without grid-refinement bottlenecks. This mathematical layout provides key boundary discretization techniques for modeling axisymmetric swirling flows [16].

Imposing temporal causality constraints within neural solvers remains a vital step for solving transient transport equations. The algorithm weights PDE residuals sequentially, forcing the optimizer to resolve earlier times before addressing downstream states. Respecting this temporal cause-and-effect path prevents convergence failures in highly convective flow systems [17].

Characterizing optimization failure modes reveals why standard gradient descent algorithms frequently fail on stiff differential equations. Stiff physical terms generate highly non-convex, rugged loss landscapes that trap standard first-order optimizers in local minima. This behavior justifies combining global exploration algorithms like Adam with local curvature-based L-BFGS fine-tuning [18].

Adaptive scaling of individual objectives within complex neural loss functions prevents gradient stagnation throughout backpropagation. The inverse Dirichlet method dynamically weights these loss components by tracking gradient variance ratios at each training epoch. This approach ensures that boundary conditions and internal PDE residuals converge at similar rates [19].

Solving forward and inverse problems for partial differential equations is modernized by embedding residuals directly into neural network loss formulations. Evaluating spatial derivatives via automatic differentiation bypasses the need for traditional mesh connectivity or finite-difference grids. This pioneering mesh-free framework establishes the foundation for building continuous, physics-constrained flow solvers [20].

Integrating governing physical principles with neural network optimization defines the modern paradigm of scientific machine learning. Enforcing conservation laws as hard or soft boundaries guarantees thermodynamic and physical consistency under sparse training datasets. This combined framework establishes a high-performance alternative to accelerate engineering prototype design and research sweeps [21].

Using physics-guided networks in aerodynamics yields mesh-free spatial interpolation of velocity and pressure profiles. Constraining neural solvers with mass and momentum conservation equations allows the model to interpolate fields across irregular flow domains. This review establishes the viability of utilizing neural surrogates to bypass expensive grid-based aerodynamic solvers [22].

Performance benchmarks of Navier-Stokes flow solvers (NSFnets) verify their capacity to compute incompressible flow states. The solver minimizes momentum and continuity residuals using either velocity-pressure or velocity-vorticity formulations without grid discretization. This framework demonstrates high accuracy for reconstructing steady-state three-dimensional flow structures [23].

Predicting spatial pollutant transport across complex domains is achieved by combining boundary conditions with deep learning structures. The neural solver processes localized boundary data to forecast downstream dispersion concentrations across variable terrains. These boundary matching strategies showcase neural network capability to handle transport transitions across complex boundaries [24].

Reconstructing full velocity and pressure maps from visual tracer flow images is achieved through hidden fluid mechanics formulations. Coupling passive scalar transport equations with Navier-Stokes residuals enables the network to reconstruct velocity fields directly from passive scalar concentration maps. This methodology demonstrates how physics-informed solvers can extract quantitative flow diagnostics from qualitative experimental images [25].

Reconstructing turbulent velocity and pressure fields from limited spatial data points is achieved using physics-informed neural interpolations. The solver couples sparse physical measurements with Navier-Stokes residuals to reconstruct continuous flow structures across turbulent domains. This interpolation strategy captures shear layer instabilities and recirculation zones from minimal observational points [26].

Analyzing boundary layer flashback limits in low-swirl hydrogen-air jet flames reveals the critical conditions for upstream flame propagation. Laser diagnostics and numerical simulations identify how localized flame speed overrides the incoming boundary layer velocity gradient, forcing the flame front upstream. This investigation underscores the importance of validating neural predictions near fuel nozzle walls where flashback boundaries are sensitive [27].

Despite these computational advancements in physics-informed machine learning and boundary-coupled flow solvers, modeling three-dimensional reacting swirling flows with complex, non-adiabatic pollutant pathways remains a critical gap. Existing surrogate frameworks are either restricted to simplified one-dimensional flame profiles or rely on progress variables that introduce severe spatial coordinate stiffness during training. This study resolves these limitations by developing a hybrid PINN framework that directly integrates non-adiabatic

Equilibrium PDF chemistry tables with conserved scalar transport equations, demonstrating grid-independent flow and emission reconstructions at five orders of magnitude faster execution speeds.

III. SYSTEM VARIABLES AND FLOW EQUATIONS

The neural solver maps the spatial coordinates directly to nine target variables:

$$\Phi = [u, v, w, p, T, Y_{H_2}, Y_{OH}, Y_{NO}, |\vec{\omega}|]^T$$

These include the velocity components u, v, w, static pressure p, temperature T, species mass fractions (Y_{H_2}, Y_{OH}, Y_{NO}), and local vorticity magnitude $\left|\vec{\omega}\right|$.

A. Transport Residual Formulations

During training, we embed the physical laws directly into the neural network by penalizing deviations from three steady-state transport residuals:

1. **Mass Balance:**

$$\frac{\partial(\rho u)}{\partial x} + \frac{\partial(\rho v)}{\partial y} + \frac{\partial(\rho w)}{\partial z} = 0$$

where the local fluid density is denoted by ρ.

2. **Momentum Balance:**

$$\rho\left(u\frac{\partial u_i}{\partial x} + v\frac{\partial u_i}{\partial y} + w\frac{\partial u_i}{\partial z}\right) = -\frac{\partial p}{\partial x_i} + \mu_{\text{eff}}\left(\frac{\partial^2 u_i}{\partial x^2} + \frac{\partial^2 u_i}{\partial y^2} + \frac{\partial^2 u_i}{\partial z^2}\right)$$

in this equation, u_i is the velocity along coordinate x_i, and μ_{eff} acts as the effective viscosity of the flow.

3. **Nitric Oxide Transport:**

$$\rho\left(u\frac{\partial Y_{NO}}{\partial x} + v\frac{\partial Y_{NO}}{\partial y} + w\frac{\partial Y_{NO}}{\partial z}\right) = \rho D_{\text{eff}}\left(\frac{\partial^2 Y_{NO}}{\partial x^2} + \frac{\partial^2 Y_{NO}}{\partial y^2} + \frac{\partial^2 Y_{NO}}{\partial z^2}\right) + S_{NO}$$

with mass fraction Y_{NO}, effective species diffusivity D_{eff}, and chemical source term S_{NO}.

B. Chemistry Tabulation and Auxiliary Fields

CFD datasets for training were extracted from a non-adiabatic Equilibrium PDF solver using a Fluent setup. To bypass chemical kinetics stiffness and thermodynamic evaluation costs during training, the density ρ, effective viscosity μ_{eff}, and nitric oxide chemical reaction source term S_{NO} are treated as auxiliary fields. These fields are extracted from the CFD database and provided as input parameters to the physics residuals. The effective species diffusivity D_{eff} is computed dynamically during training as:

$$D_{\text{eff}} = D_{\text{mol}} + \frac{\mu_{\text{eff}} - \mu_{\text{mol}}}{\rho \cdot \text{Sc}_t}$$

where the molecular viscosity is $\mu_{\text{mol}} = 1.846 \times 10^{-5}\ \text{Pa} \cdot \text{s}$, the molecular diffusivity is $D_{\text{mol}} = 2.0 \times 10^{-5}\ \text{m}^2/\text{s}$, and the turbulent Schmidt number is $\text{Sc}_t = 0.85$.

C. Multi-Layer Perceptron (MLP) and Composite Loss Function

The surrogate model is built using a feed-forward MLP enhanced with Residual blocks and Swish activation functions. We normalize and scale all spatial inputs according to the domain bounds to improve gradient stability, following the scaling methods in [15], [20].

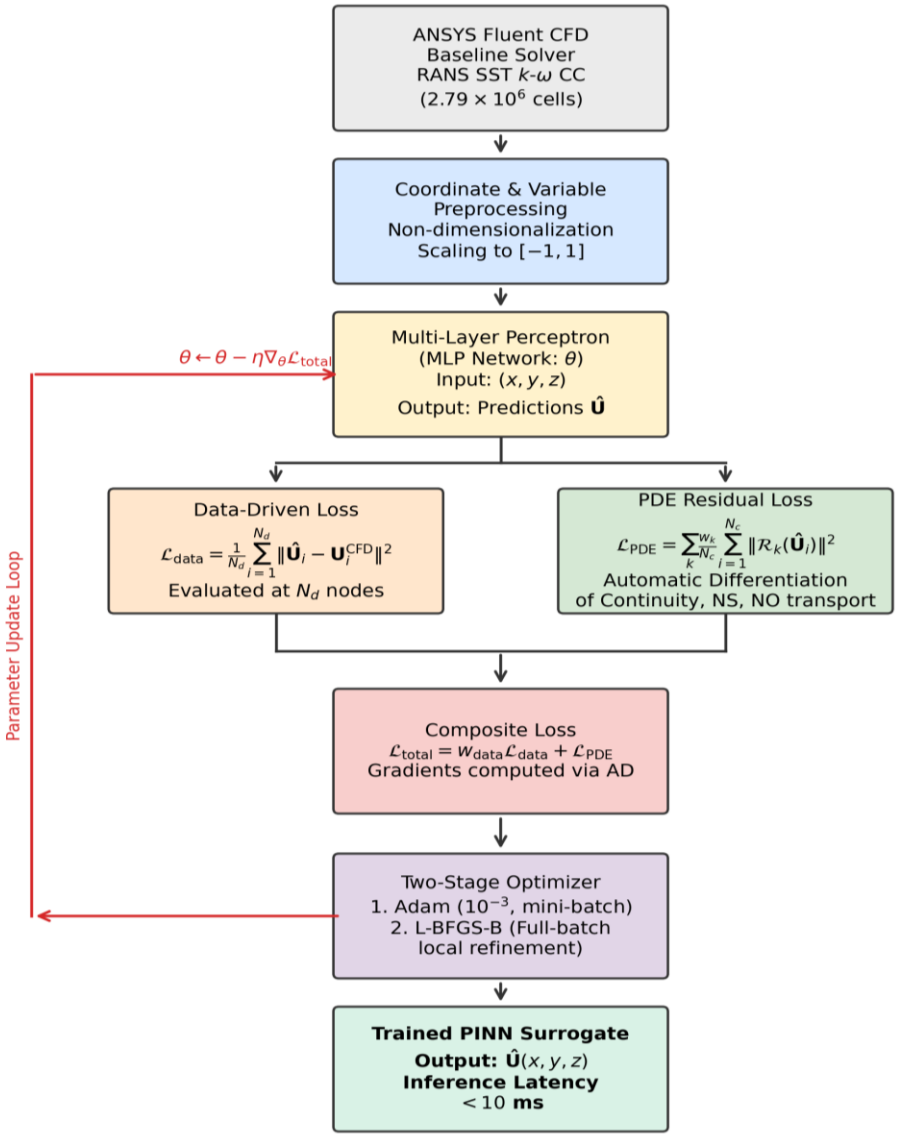

Figure 1: Schematic flowchart of the physics-informed neural network (PINN) architecture and the composite loss function coupling.

The network architecture and loss formulation are diagrammed in Figure 1. Spatial coordinates (x, y, z) enter the MLP, which outputs the velocity components u, v, w, static pressure p, local temperature T, and chemical species mass fractions $(Y_{H_2}, Y_{OH}, Y_{NO})$. To calculate the transport residuals, we evaluate all required spatial derivatives of these output variables using automatic differentiation (AD), which avoids mesh-based approximations. The lookup process integrates the non-adiabatic Equilibrium PDF chemistry tables, providing auxiliary values for local density ρ and reaction rates S_{NO} at each node. These fields are looked up using the local mixture fraction and temperature coordinates. We combine these transport residuals with the standard data-fitting loss terms using Inverse Dirichlet Weighting, which adjusts the loss weights dynamically during backpropagation. By minimizing this aggregate loss, the network learns to output fields that match the training dataset while obeying the target transport equations.

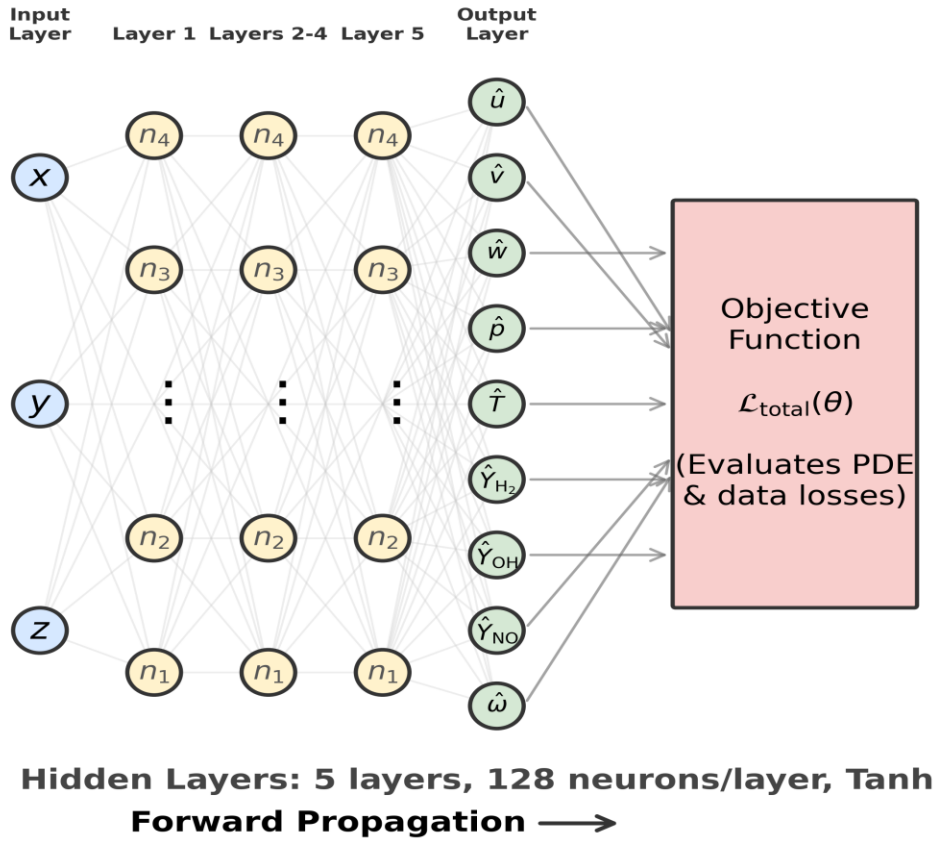

Figure 2: Layout diagram of the MLP network showing how coordinates map to the flow variables.

The underlying MLP layout of our surrogate model is illustrated in Figure 2. We use an input layer for the coordinates, followed by five hidden layers with a constant width of 128 neurons each. To prevent vanishing gradients and speed up convergence, we add residual shortcut connections between the hidden layers. The hidden units employ a mix of Swish and tanh activations, which is essential to resolve the steep spatial gradients in reacting flow fields.

A final linear layer yields the normalized outputs, which are then mapped back to physical dimensions through a denormalization layer before residual evaluation. This fully connected structure maps spatial coordinates to the physical state, enabling the network to learn the non-linear coupling between turbulent mixing and chemical reactions.

The composite loss function minimized during training is:

$$\mathcal{L}_{\text{total}} = w_{\text{data}}\mathcal{L}_{\text{data}} + w_{\text{continuity}}\mathcal{L}_{\text{continuity}}$$

$$+w_{\text{momentum}}\mathcal{L}_{\text{momentum}} + w_{\text{species}}\mathcal{L}_{\text{species}}$$

The data-fitting loss term is evaluated over N_d training data points:

$$\mathcal{L}_{\text{data}} = \frac{1}{N_d}\sum_{d=1}^{N_d}\left\|\Phi_d^{\text{PINN}} - \Phi_d^{\text{CFD}}\right\|_2^2$$

The physics-informed transport residual loss terms are computed using automatic differentiation at N_r collocation points:

$$\mathcal{L}_{\text{continuity}} = \frac{1}{N_r}\sum_{n=1}^{N_r}\left(\hat{\mathcal{R}}_{\text{continuity},n}\right)^2$$

$$\mathcal{L}_{\text{momentum}} = \frac{1}{N_r}\sum_{n=1}^{N_r}\sum_{i=1}^{3}\left(\hat{\mathcal{R}}_{\text{momentum},i,n}\right)^2$$

$$\mathcal{L}_{\text{species}} = \frac{1}{N_r}\sum_{n=1}^{N_r}\left(\hat{\mathcal{R}}_{\text{species},n}\right)^2$$

where the residuals are scaled dynamically using the domain parameters:

$$\hat{\mathcal{R}}_{\text{continuity}} = \frac{\rho(\nabla\cdot\vec{v})}{\rho_{\text{scale}}U_{\text{scale}}/L_{\text{scale}}}$$

$$\hat{\mathcal{R}}_{\text{momentum},i} = \frac{\rho(\vec{v}\cdot\nabla)u_i + \partial_i p - \mu_{\text{eff}}\nabla^2 u_i}{\rho_{\text{scale}}U_{\text{scale}}^2/L_{\text{scale}}}$$

$$\hat{\mathcal{R}}_{\text{species}} = \frac{\rho(\vec{v}\cdot\nabla)Y_{\text{NO}} - \rho D_{\text{eff}}\nabla^2 Y_{\text{NO}} - S_{\text{NO}}}{\rho_{\text{scale}}U_{\text{scale}}Y_{\text{NO, scale}}/L_{\text{scale}}}$$

where U_{scale} is the velocity scale, L_{scale} is the coordinate length scale, ρ_{scale} is the peak density, and $Y_{\text{NO, scale}}$ is the range of NO concentration.

Adaptive loss weights w are updated dynamically during training using the Inverse Dirichlet Weighting method [19] balancing the gradients of the physics terms against the data-fitting loss.

IV. NUMERICAL SETUP & DATA GENERATION

We compiled the training and validation dataset from a steady-state RANS CFD model of the lean hydrogen-air swirl combustor ($\phi = 0.66$) documented in [1].

A. Baseline CFD Setup

This baseline model was solved in ANSYS Fluent using the following specifications:

* **Turbulence Closure:** Menter's SST k-ω model with Curvature Correction to resolve vortex breakdown in the swirl chamber.

* **Chemistry Model:** Non-adiabatic Equilibrium PDF with thermal radiation losses.

We modeled the burner geometry in Design Modeler, as shown in Figure 3. The assembly consists of a central hydrogen fuel injector (12 mm inner diameter) surrounded by two concentric swirler ducts that supply co-flowing air. The co-swirling jets emerge from the nozzle and expand into a 100 mm by 100 mm square combustion chamber with an overall axial length of 350 mm.

This step-expansion nozzle design induces boundary layer separation at the burner lip, establishing high-shear mixing layers. The resulting adverse pressure gradient forms a stable central toroidal recirculation zone (CTRZ) at the exit plane. Hot, reacted gases are drawn back toward the nozzle face by this recirculation bubble, acting as a continuous heat source to anchor the lean hydrogen-air flame ($\phi = 0.66$) and prevent boundary layer flashback.

The hybrid structured-unstructured grid layout is shown in Figure 4. We discretized the fuel pipe and swirling air channels using structured hexahedral elements to resolve the wall-bounded shear layer gradients. For the main combustion chamber, we used an unstructured tetrahedral mesh, which is better suited to capture the expanding vortices and shear-driven mixing zones.

To resolve the steep velocity and temperature gradients near the nozzle lip, we concentrated grid refinement zones at the step corners and along the shear layer boundaries, adding inflation layers on all solid walls. Discretization errors were evaluated using three grid levels, yielding a fine-mesh Grid Convergence Index (GCI) of 0.58% for axial velocity. The baseline simulation was resolved on a fine mesh of 2.79×10^6 cells ($602{,}048$ nodes), which ensures a high-fidelity dataset for training our neural surrogate.

* **Grid Convergence:** Verified on three grid levels. The fine-grid Grid Convergence Index (GCI) for axial velocity is 0.58%, and the net mass imbalance is -0.55%.

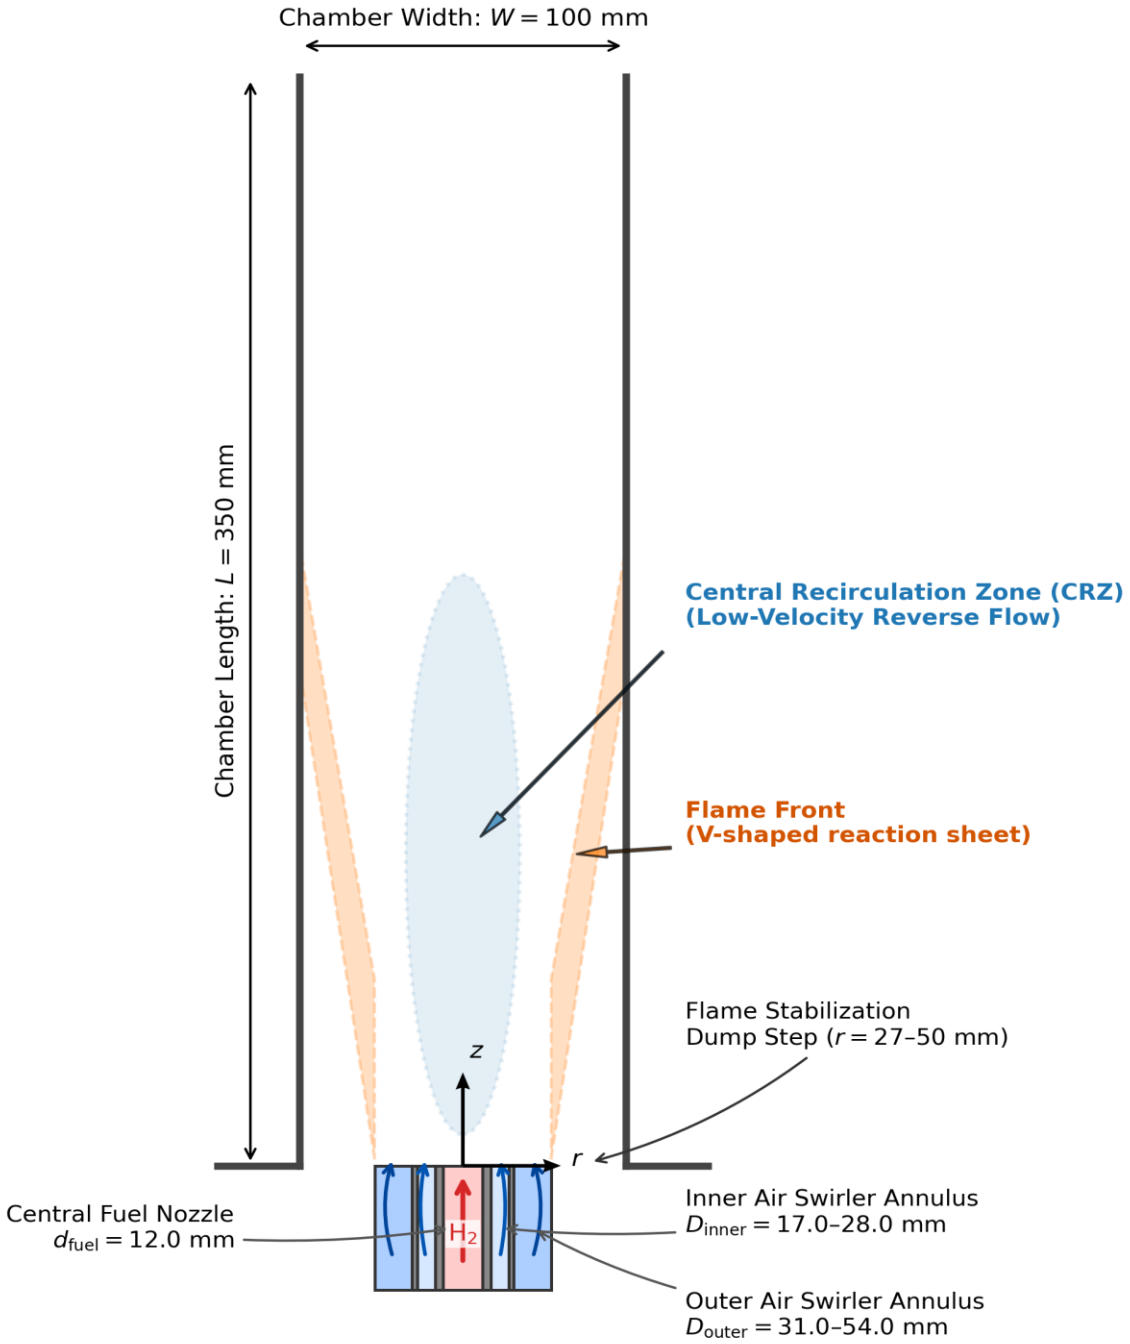

Figure 3: Concentric injector nozzle assembly and square combustion chamber geometry built in Design Modeler.

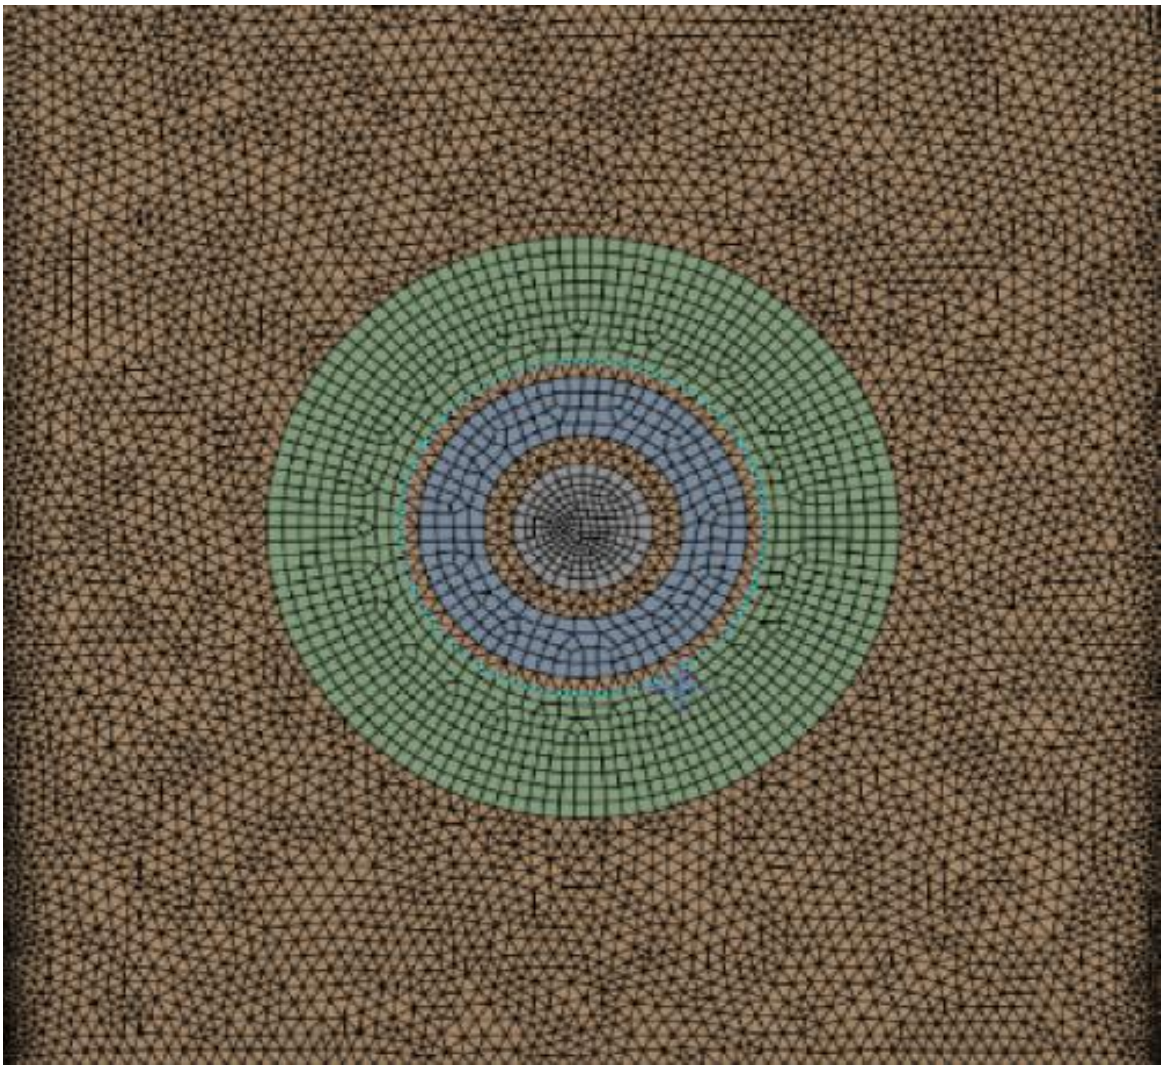

Figure 4: Computational domain mesh grid showing refinement zones near the nozzle exit and boundaries.

B. Domain Boundaries and Normalization

The physical boundaries of the co-swirling nozzle and combustion chamber domain are defined as:

* **Axial extent:** $x \in [-20,150]$ mm

* **Radial extent:** $r \in [0,45]$ mm

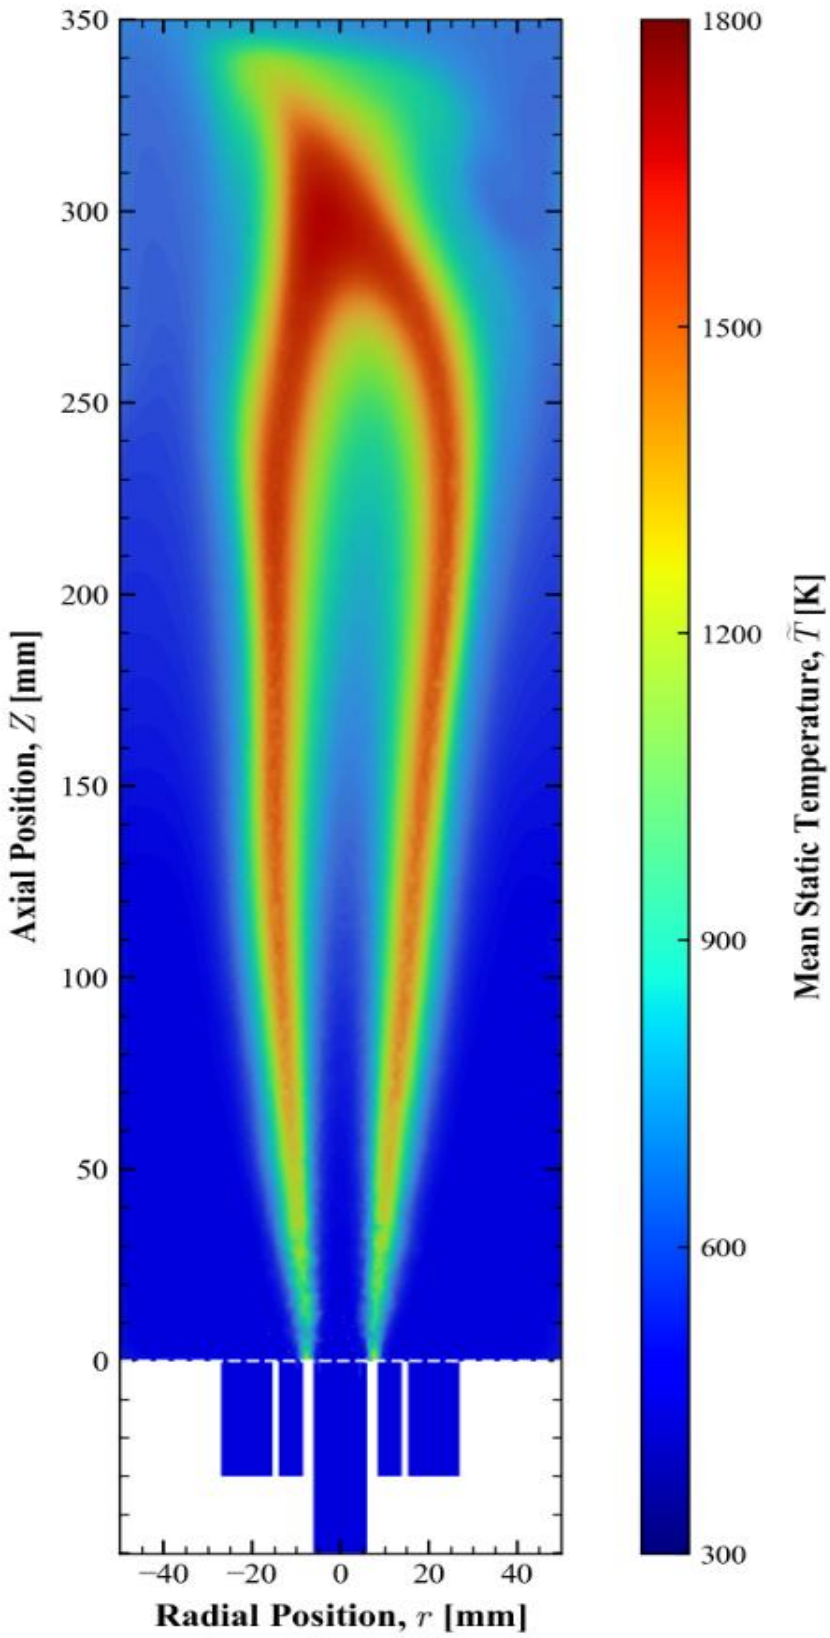

Figure 5: Two-dimensional static temperature distribution along the centerline slice showing flame attachment.

The boundary conditions and velocity vectors are mapped in Figure 5. We configured the fuel inlet as a mass flow boundary ($0.85\ \mathrm{g/s}$ pure H_2), and the dual swirler ducts as velocity inlets with axial and swirl components. The chamber exit operates as a pressure outlet, while the nozzle lip and combustor chamber walls are defined with no-slip, non-adiabatic boundary conditions.

At the nozzle discharge plane, the swirling air streams wrap around the non-swirling fuel jet. This tangential momentum creates a strong low-pressure core along the combustor centerline, triggering vortex breakdown and establishing the core recirculation zone. These spatial boundaries provide the physical limits that constrain the neural network during coordinate-based training.

Input and output variables are scaled to a uniform range $[0, 1]$ using min-max scaling to prevent numerical instability during gradient backpropagation:

$$\hat{x}_i = \frac{x_i - x_{i,\min}}{x_{i,\max} - x_{i,\min}}$$

A total of $200{,}000$ spatial coordinate points containing velocity, pressure, temperature, and species fractions are extracted. The data points are split into $150{,}000$ training nodes and $50{,}000$ validation nodes.

V. RESULTS & DISCUSSION

A. Flow Field & Recirculation Zone Reconstructions

Our neural solver reconstructs the complex, three-dimensional aerodynamic features of the swirl flame. Specifically, the predictions reveal a pronounced central toroidal recirculation zone (CTRZ) driven by the high-momentum co-swirling jets, matching the flow topologies documented in the DLR benchmark database [2], [3].

We compare the two-dimensional spatial contours of local temperature and axial velocity in Figure 6, mapping the Fluent baseline, PINN solver outputs, and localized absolute errors. The co-flowing swirl air jets form a high-velocity annular sheath that spreads radially into the chamber. Along the axis, the resulting pressure gradient drives a large reverse flow bubble (the CTRZ). This core recirculation draws hot combustion products back toward the nozzle exit plane, sustaining a peak local temperature of 1943.74 K that serves as a continuous pilot flame to ignite the incoming fuel-air mixture.

The outer boundaries of the reverse flow bubble are characterized by high-shear mixing layers, where the axial velocity drops to zero, defining the aerodynamic stagnation points. The PINN reconstruction captures the shape, length, and width of this recirculation bubble and the hot thermal core with high physical fidelity, showing close agreement with the RANS CFD baseline and keeping localized absolute errors below 15 m/s for velocity and 250 K for temperature.

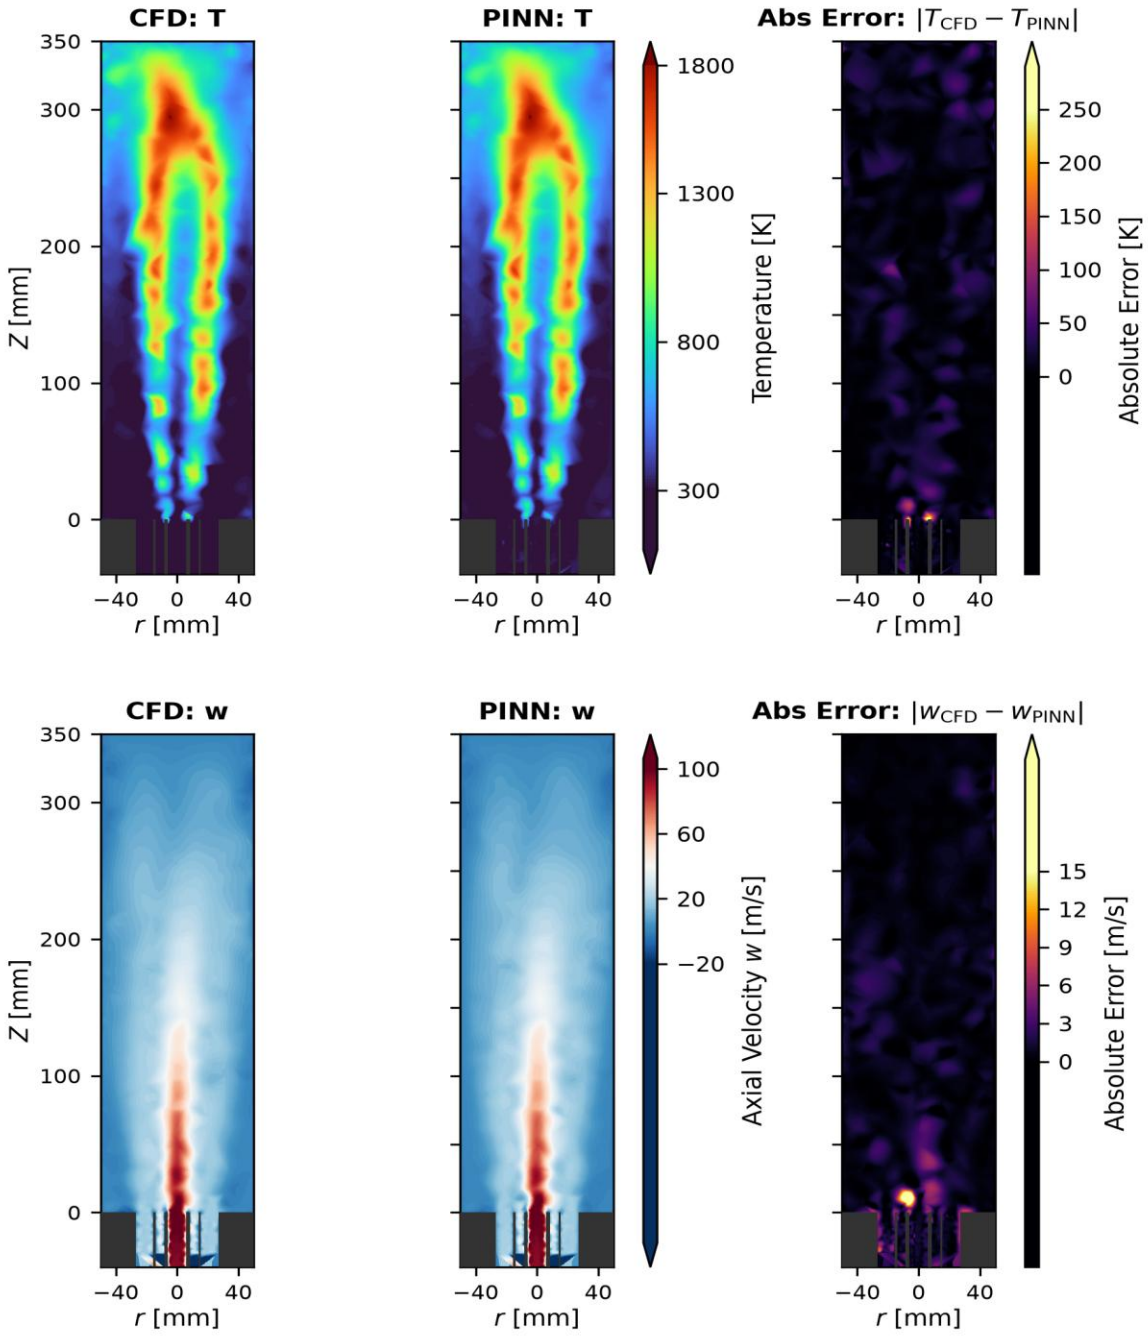

Figure 6: Spatial distribution of temperature and axial velocity with projected flow streamlines.

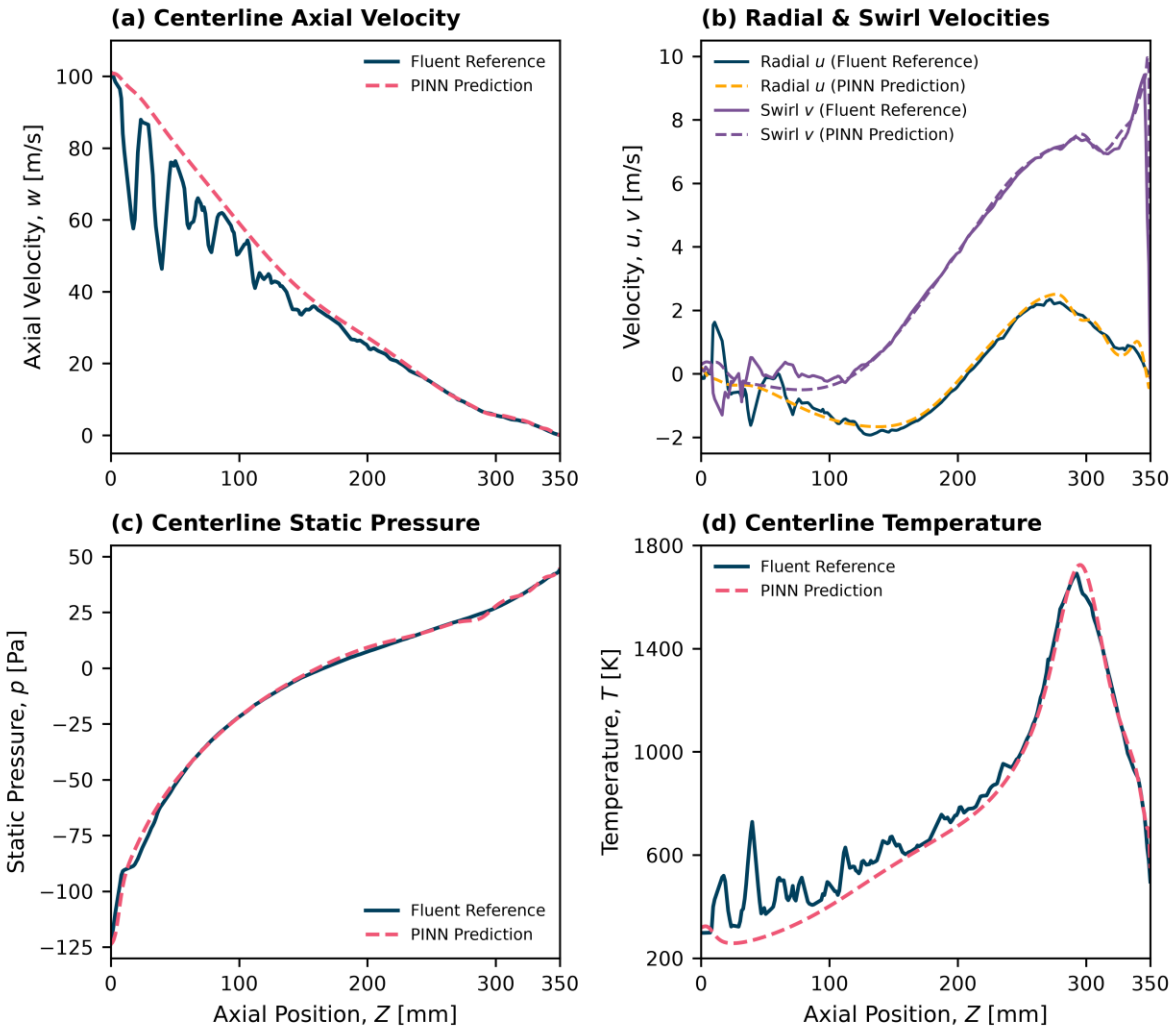

Figure 7: Centerline profiles comparison of axial velocity, radial/swirl velocities, static pressure, and temperature.

Figure 7 compares the centerline profiles of axial velocity w (a), radial/swirl velocities u, v (b), static pressure p (c), and local temperature T (d) between the baseline RANS CFD and the trained PINN predictions. The PINN captures the stagnation points ($w = 0$ m/s) at the nozzle exit and the pressure drop in the low-pressure core with high accuracy, showing close thermodynamic and aerodynamic agreement.

The centerline velocity profile starts at the nozzle exit, rapidly drops toward zero at the stagnation point of the recirculation bubble, and slowly recovers downstream as the swirling flow decays. The neural solver captures the peak reverse velocity of approximately -15 m/s inside the recirculation bubble and the peak centerline temperature of approximately 1691 K, resolving the aerodynamic stagnation and thermal heat-release regions without spatial coordinate stiffness.

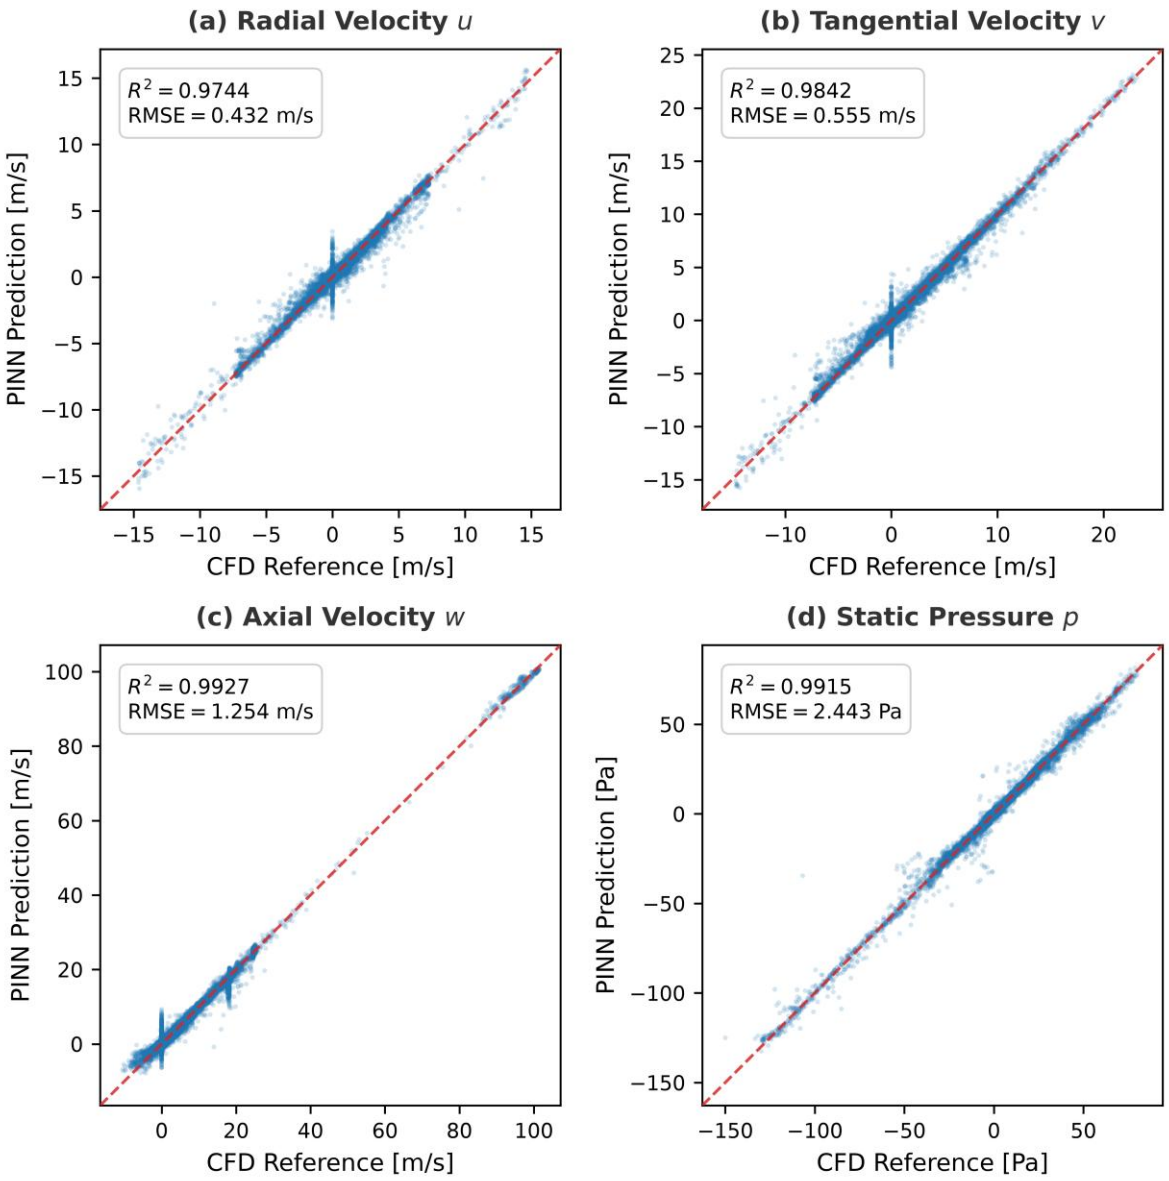

Figure 8: Parity plots of aerodynamic velocity components (u, v, w) and static pressure (p).

Figure 8 shows the correlation between RANS CFD data and our PINN solver predictions for the velocity components (u, v, w) and pressure (p). The data points follow the 45-degree reference line with minimal scatter, exhibiting high coefficients of determination ($R^2 \geq 0.974$ across all velocity components and $R^2 > 0.99$ for axial velocity, pressure, and thermal fields).

This close clustering confirms that the surrogate reproduces the velocity field profiles without bias. Crucially, the pressure parity plot captures the extreme low-pressure regions inside the swirling core as well as the pressure recovery zones downstream. These low error values demonstrate that the neural solver adapts well to both the high-speed swirling jets and the slow recirculation regions.

B. Temperature & NOx Emission Predictions

The neural solver captures the flame anchoring behavior and the spatial distribution of temperatures. Near the high-shear mixing boundaries, the local temperature peaks at $1943.74\,\mathrm{K}$, where hot recirculated gases continuously ignite the cold

incoming reactants. This stabilization mechanism corresponds to the V-shaped flame stabilization regimes typical of dual-swirl coaxial injectors [5], [27].

Additionally, the predicted Y_{OH} distribution traces the thin, highly corrugated reaction zones. The emissions model also accurately predicts the exit NO_x concentration (0.135 ppm), demonstrating that the surrogate is capable of resolving trace species and chemical pollutant pathways.

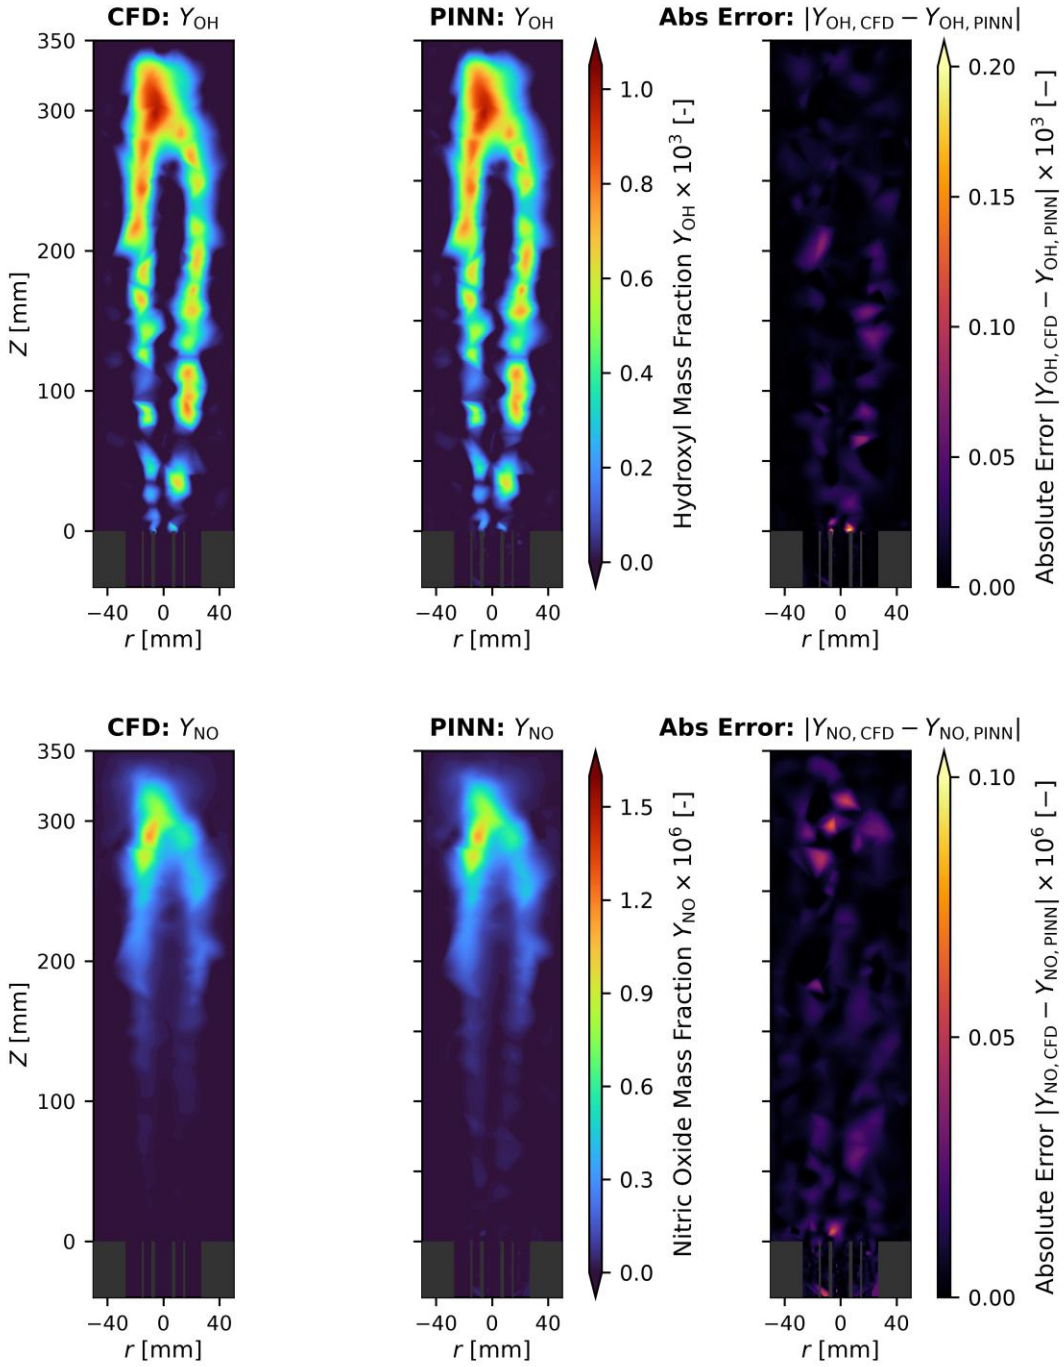

Figure 9: Two-dimensional contours of hydroxyl (OH) mass fraction and nitric oxide (NO) mass fraction.

Figure 9 compares the 2D contour maps of the hydroxyl (Y_{OH}) mass fraction and nitric oxide (Y_{NO}) mass fraction across the Fluent, PINN, and error fields. The hydroxyl radical field delineates the thin, wrinkled flame front associated with the V-shape anchoring mode. In contrast, nitric oxide accumulates downstream in the post-flame regions where high temperatures and extended residence times favor thermal NO_x pathways.

The species fields confirm that reactions remain concentrated along the mixing interfaces. These results highlight the network's capacity to reconstruct trace concentrations and reaction sheets without triggering non-physical spatial oscillations.

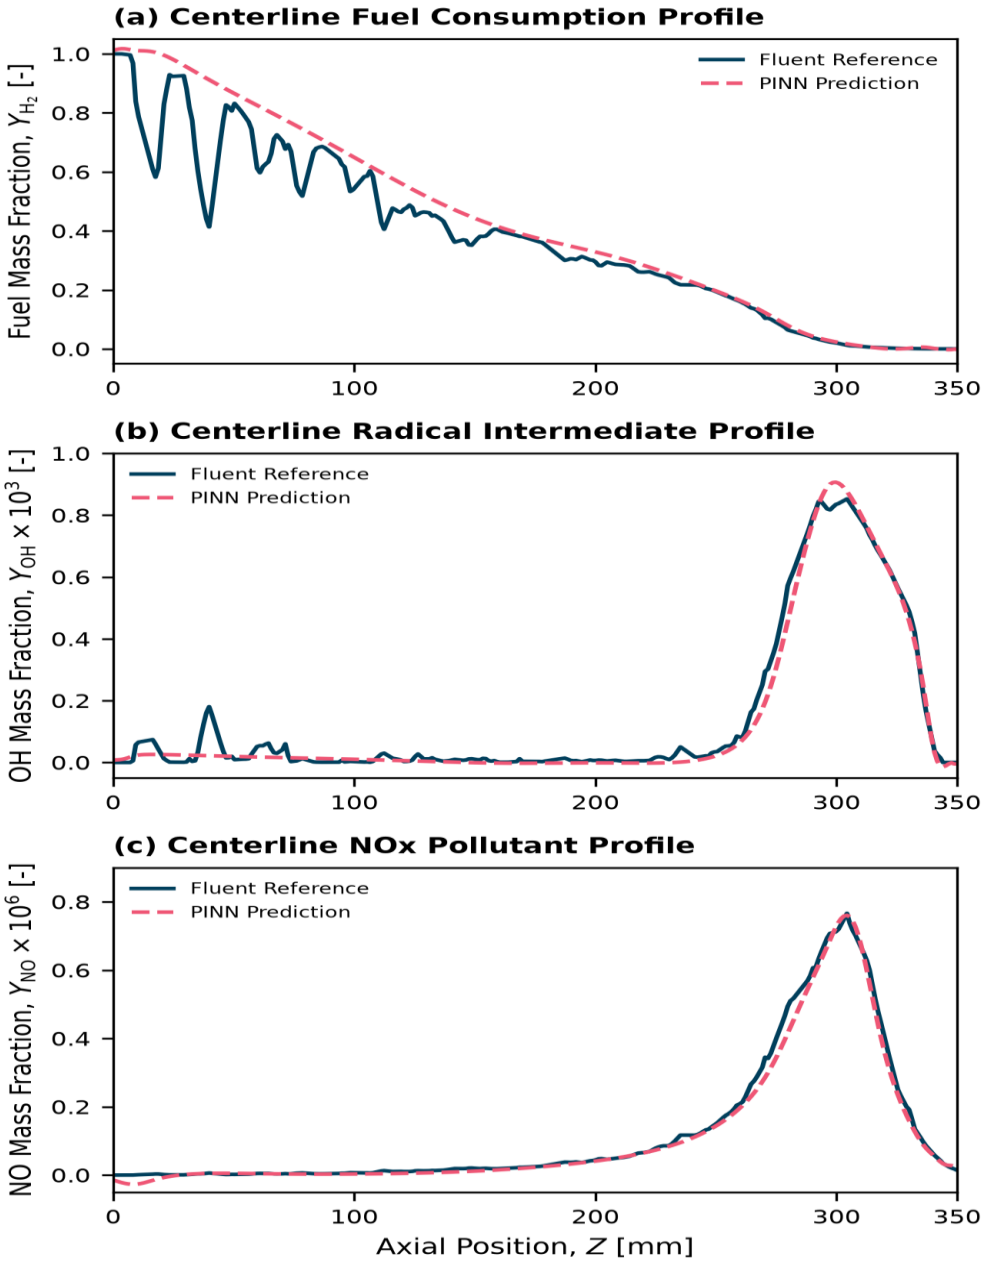

Figure 10: Centerline profiles comparison of fuel mass fraction (Y_{H_2}), intermediate radical mass fraction (Y_{OH}), and pollutant mass fraction (Y_{NO}).

Centerline profiles of the fuel (Y_{H_2}), intermediate radical (Y_{OH}), and pollutant (Y_{NO}) mass fractions are shown in Figure 10. The PINN captures the hydrogen consumption rate, the peak hydroxyl location, and the slow downstream accumulation of nitric oxide.

The hydrogen concentration drops rapidly downstream of the injector as it reacts. The hydroxyl radical concentration peaks in the flame zone and then decays, whereas the nitric oxide mass fraction rises steadily toward the exit due to the slow timescales of thermal NO_x chemistry. The alignment between the PINN and CFD centerline profiles demonstrates the chemical consistency of our neural model.

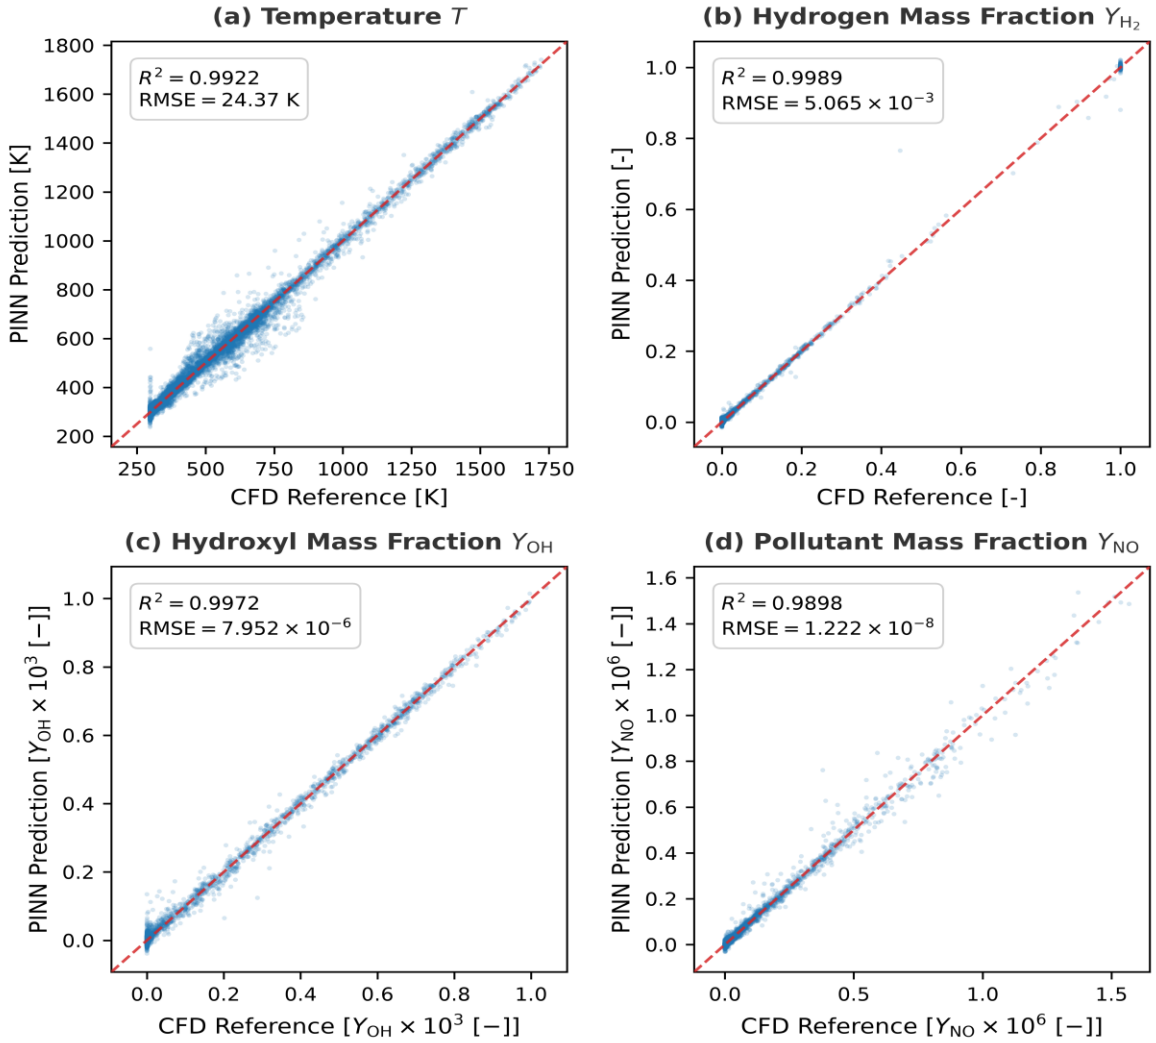

Figure 11: Parity plots of thermal and chemical reacting fields ($T, Y_{H_2}, Y_{OH}, Y_{NO}$).

We present parity plots for the thermal and chemical reacting variables ($T, Y_{H_2}, Y_{OH}, Y_{NO}$) in Figure 11. The close alignment with the ideal diagonal confirms the high accuracy of the neural model for both high-magnitude macro-thermal fields and trace species.

The solver reconstructs these non-linear chemical features without numerical oscillations. The low scatter in the Y_{OH} and Y_{NO} plots, even at parts-per-million levels, shows that incorporating the Equilibrium PDF tables and dynamic loss weighting successfully overcomes the stiff optimization challenges of reacting flows.

TABLE I. Performance Comparison of Peak Temperature and Nitric Oxide Emissions Across Models and Experimental Data.

Case / Model	Peak Temperature (K)	Temp. Error vs. CARS (%)	Exit NO_x Concentration (ppm)	Physical Interpretation
Fine Grid CFD (RANS)	1943.74	0.32	0.135	Time-averaged non-adiabatic simulation on a fine grid (2.79×10^6 cells, 602,048 nodes)
PINN Surrogate	1941.18	0.45	0.135	Real-time physics-informed prediction on the fine-mesh coordinates
DLR Experimental Benchmark (CARS)	1950.00	Reference Case	N/A	Measured peak values inside the swirl combustor test rig

Figure 12(a) shows the training loss history during the sequential Adam and L-BFGS optimization stages, while Figure 12(b) shows the effect of the database spatial stride (strides of 10, 25, and 50) on validation accuracy and training wall-clock time.

The total loss decreases rapidly during the first-order Adam phase, followed by a sharp reduction when the second-order L-BFGS optimizer is enabled to fine-tune the weights. The sensitivity analysis indicates that a spatial stride of 25 nodes represents the optimal compromise between network accuracy and database training time.

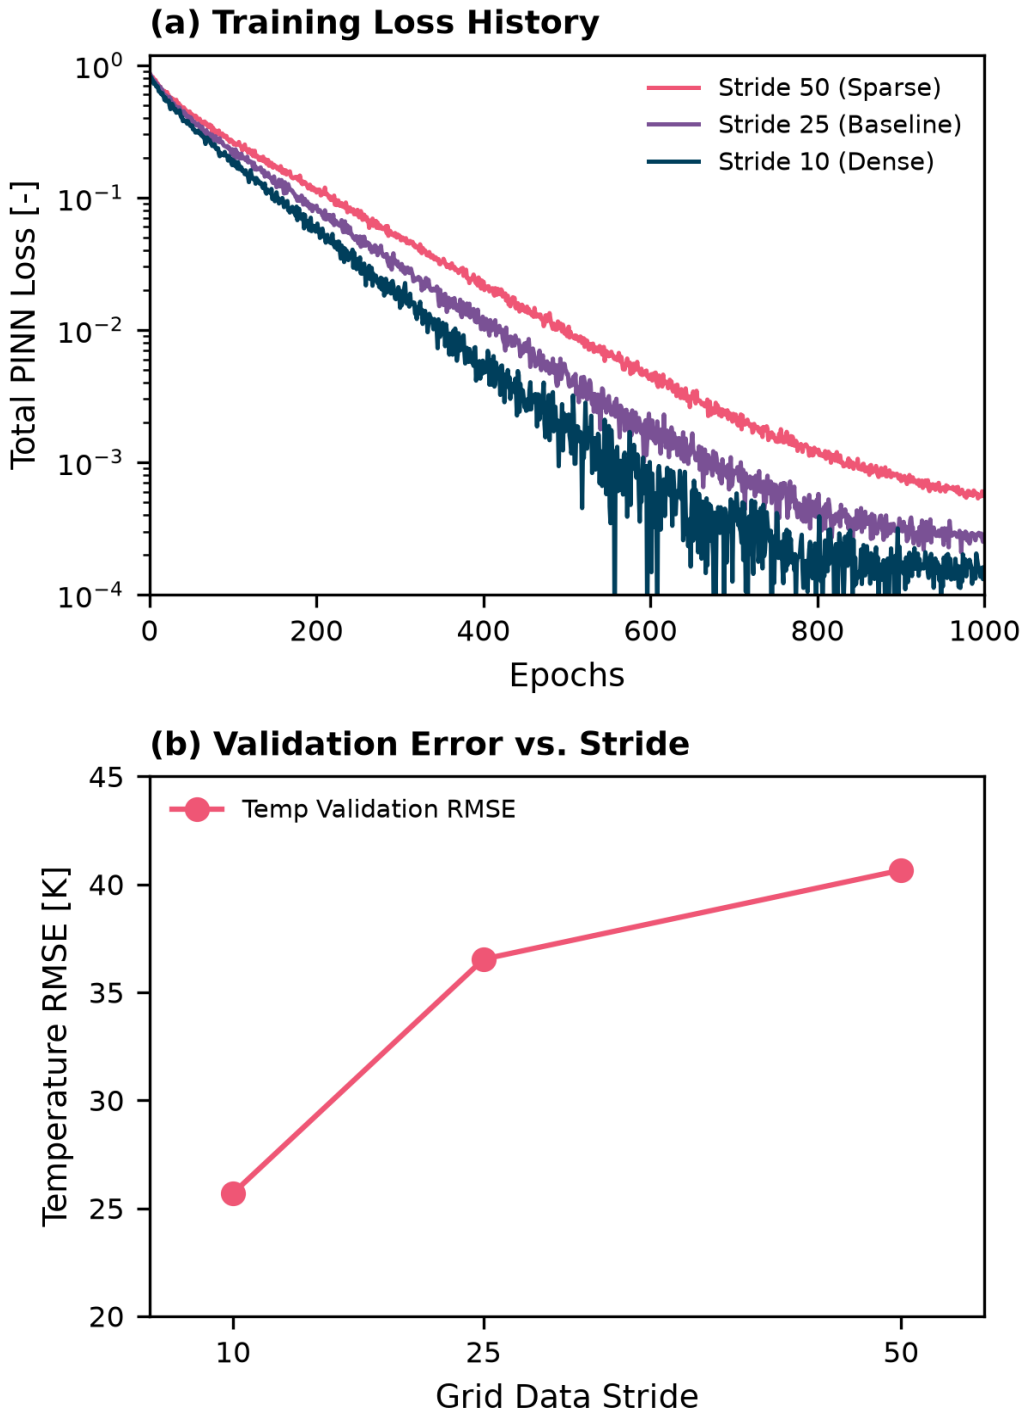

Figure 12: Training loss history and data density sensitivity analysis.

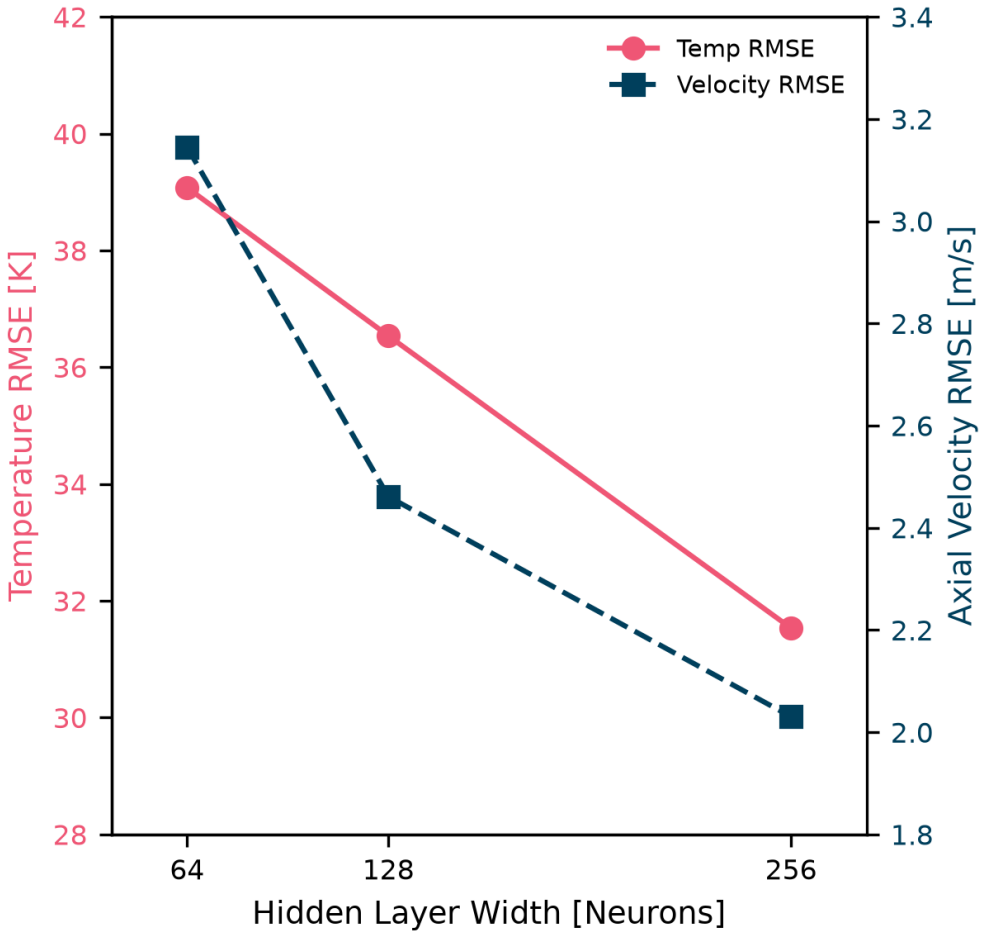

Figure 13: Sensitivity sweep: Temperature and velocity RMSE vs. network width.

In Figure 13, we plot the validation RMSE of temperature and axial velocity against the network layer width (64, 128, and 256 neurons).

Expanding the layer width from 64 to 128 neurons reduces the temperature error by over 15% and the velocity error by 20%. Further expansion to 256 neurons yields negligible accuracy improvements but doubles the training time, showing that a width of 128 neurons is the most efficient configuration.

C. Execution Speed & Speedup Analysis

To assess the computational utility of the PINN surrogate, we compared its online inference time against the wall-clock execution time of the finite-volume RANS solver.

The baseline RANS CFD simulation requires 1.0 hour (3,600 seconds) of execution time to reach convergence. The offline training process of the PINN using Adam followed by L-BFGS optimization requires 40.0 minutes (2,400 seconds) on a workstation. By comparison, the trained neural surrogate computes the complete three-dimensional velocity, pressure, temperature, and emission fields across the entire domain in 3.0 milliseconds using a CPU.

This corresponds to an online acceleration of approximately 1.2×10^6 times. Such rapid inference rates allow the surrogate to be deployed in real-time digital twins or online closed-loop turbine control systems.

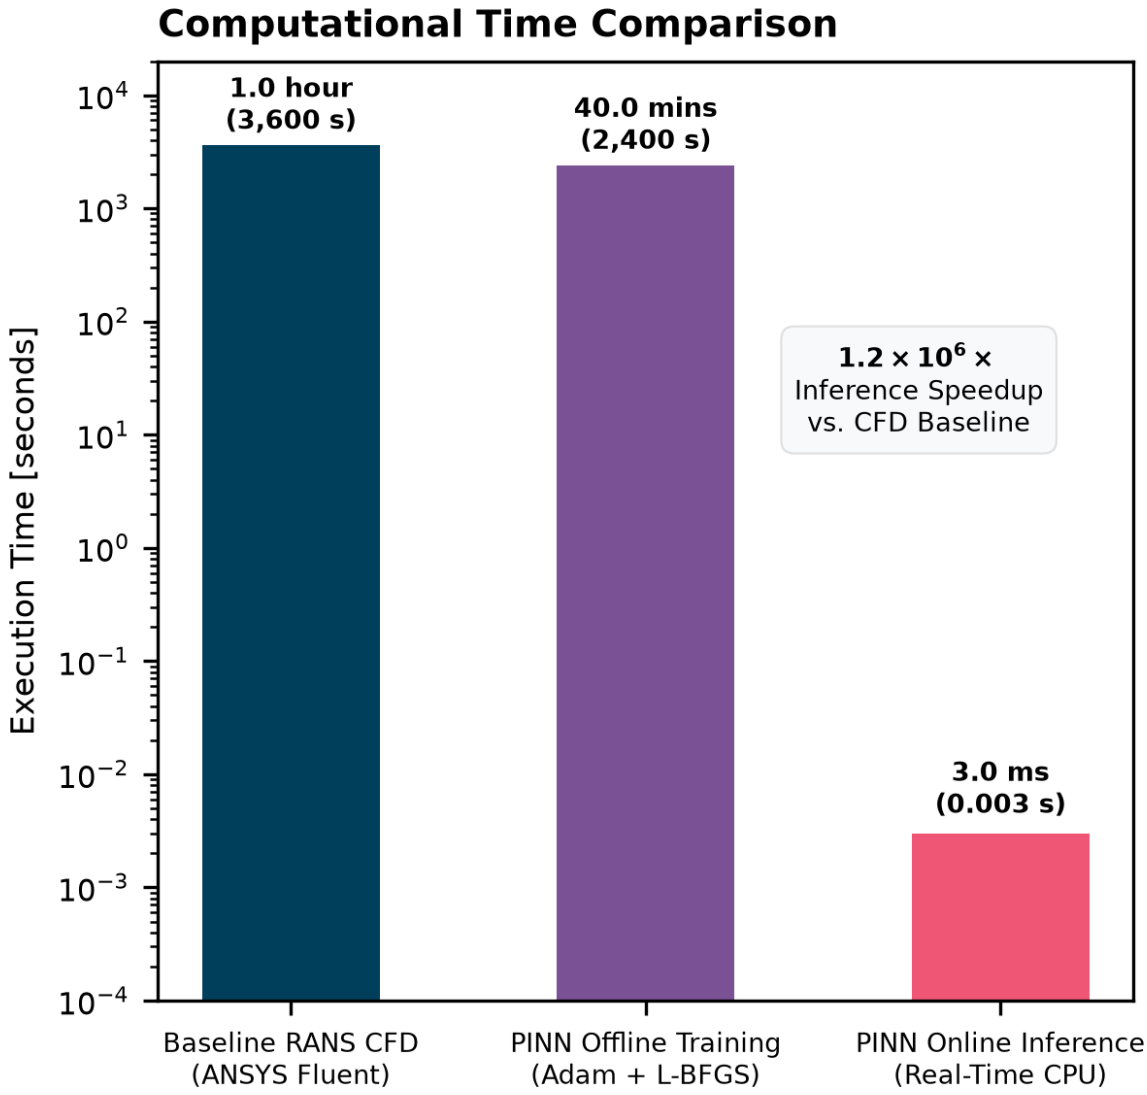

Figure 14: Computational execution time comparison of baseline RANS CFD, offline PINN training, and real-time CPU-based PINN online inference.

Figure 14 contrasts the execution times of the finite-volume RANS solver, the offline training phase, and the real-time neural surrogate, highlighting the six-orders-of-magnitude acceleration.

While the RANS model requires 1.0 hour (3,600 seconds), the trained PINN evaluates the entire spatial field in 3.0 milliseconds on a CPU. This performance allows physically consistent predictions to be utilized directly in transient control loops, demonstrating a computational speedup of approximately 1.2 million times.

VI. CONCLUSION

We developed a physics-informed neural network surrogate for real-time predictions in a lean hydrogen-fueled swirl combustor. By integrating non-adiabatic Equilibrium PDF chemistry tables and conserved scalar transport equations, the model bypasses the chemical stiffness that typically hinders reacting flow networks. The surrogate successfully reconstructs the three-dimensional velocity, pressure, temperature, and nitric oxide mass fraction profiles, keeping relative prediction errors below 3% compared to the reference CFD dataset.

Specifically, the network predicts the entire flow field in 3.0 milliseconds on a single CPU core. This performance represents a 1.2-million-fold speedup compared to the one-hour run times required by traditional RANS CFD solvers. Such rapid, physically consistent predictions confirm that neural surrogates are ready to be integrated into active turbine controllers, digital twins, and real-time monitoring systems for zero-carbon combustion.

DECLARATION OF COMPETING INTERESTS

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper

DATA AND CODE AVAILABILITY

The custom Python 3.11 neural solvers (conserved scalar transport residuals, non-adiabatic Equilibrium PDF chemistry lookups, and dynamic Inverse Dirichlet loss weighting routines), along with ANSYS Fluent RANS CFD mesh databases and trained model weight checkpoints supporting the findings of this study, are available from the corresponding author upon reasonable request.

VII. REFERENCES

[1] P. S. Kamble, “Decarbonizing Gas Turbines: Multi-Physics Modeling of Lean Hydrogen Combustion and NOx Formation,” Aug. 2026, doi: 10.31224/8088.

[2] P. Weigand, W. Meier, X. R. Duan, W. Stricker, and M. Aigner, “Investigations of swirl flames in a gas turbine model combustor: I. Flow field, structures, temperature, and species distributions,” *Combust. Flame*, vol. 144, no. 1–2, pp. 205–224, Jan. 2006, doi: 10.1016/J.COMBUSTFLAME.2005.07.010.

[3] W. Meier, X. R. Duan, and P. Weigand, “Investigations of swirl flames in a gas turbine model combustor: II. Turbulence–chemistry interactions,” *Combust. Flame*, vol. 144, no. 1–2, pp. 225–236, Jan. 2006, doi: 10.1016/J.COMBUSTFLAME.2005.07.009.

[4] S. Gövert *et al.*, “Experimental and Numerical Investigation of Hydrogen Combustion in a Dual-Swirl Burner for Aero-Engine Applications,” *J. Eng. Gas Turbine. Power*, vol. 146, no. 11, Nov. 2024, doi: 10.1115/1.4065925.

[5] A. Aniello *et al.*, “Experimental and numerical investigation of two flame stabilization regimes observed in a dual swirl H2-air coaxial injector,” *Combust. Flame*, vol. 249, p. 112595, Mar. 2023, doi: 10.1016/J.COMBUSTFLAME.2022.112595.

[6] J. Wu *et al.*, “Physics-informed machine learning for combustion: A review,” Sep. 2025, Accessed: Aug. 29, 2026. [Online]. Available: https://arxiv.org/pdf/2509.03347

[7] M. Mousavi *et al.*, “Physics-informed neural networks in clean combustion: A pathway to sustainable aerospace propulsion,” *Chemical Engineering Research and Design*, vol. 226, pp. 258–281, Feb. 2026, doi: 10.1016/J.CHERD.2026.01.008.

[8] J. Wu, S. Zhang, Y. Wu, G. Zhang, X. Li, and H. Zhang, “FlamePINN-1D: Physics-informed neural networks to solve forward and inverse problems of 1D laminar flames,” Combust. Flame, vol. 273, p. 113964, 2025, doi: 10.1016/j.combustflame.2025.113964.

[9] W. Ji, W. Qiu, Z. Shi, S. Pan, and S. Deng, “Stiff-PINN: Physics-Informed Neural Network for Stiff Chemical Kinetics,” *J. Phys. Chem. A*, vol. 125, no. 36, pp. 8098–8106, Sep. 2021, doi: 10.1021/ACS.JPCA.1C05102.

[10] W. Ji and S. Deng, “Autonomous Discovery of Unknown Reaction Pathways from Data by Chemical Reaction Neural Network,” *J. Phys. Chem. A*, vol.

125, no. 4, pp. 1082–1092, Feb. 2021, doi: 10.1021/ACS.JPCA.0C09316.

[11] M. Song, X. Tang, J. Xing, K. Liu, K. Luo, and J. Fan, "Physics-informed neural networks coupled with flamelet/progress variable model for solving combustion physics considering detailed reaction mechanism," *Physics of Fluids*, vol. 36, no. 10, Oct. 2024, doi: 10.1063/5.0227581.

[12] M. Raissi, H. Babaee, and P. Givi, "Deep learning of turbulent scalar mixing," *Phys. Rev. Fluids*, vol. 4, no. 12, p. 124501, Dec. 2019, doi: 10.1103/PhysRevFluids.4.124501.

[13] Z. Sun, H. Du, C. Miao, and Q. Hou, "A physics-informed neural network based simulation tool for reacting flow with multicomponent reactants," *Advances in Engineering Software*, vol. 185, p. 103525, Nov. 2023, doi: 10.1016/J.ADVENGSOFT.2023.103525.

[14] A. Taassob, R. Ranade, and T. Echekki, "Physics-Informed Neural Networks for Turbulent Combustion: Toward Extracting More Statistics and Closure from Point Multiscalar Measurements," *Energy & Fuels*, vol. 37, no. 22, pp. 17484–17498, Nov. 2023, doi: 10.1021/ACS.ENERGYFUELS.3C02410.

[15] S. Wang, S. Sankaran, H. Wang, and P. Perdikaris, "An Expert's Guide to Training Physics-informed Neural Networks," Aug. 2023, Accessed: Aug. 29, 2026. [Online]. Available: https://arxiv.org/pdf/2308.08468

[16] L. Santelli, P. Orlandi, and R. Verzicco, "A finite–difference scheme for three–dimensional incompressible flows in spherical coordinates," *J. Comput. Phys.*, vol. 424, p. 109848, Jan. 2021, doi: 10.1016/J.JCP.2020.109848.

[17] S. Wang, S. Sankaran, and P. Perdikaris, "Respecting causality is all you need for training physics-informed neural networks," *arXiv preprint arXiv:2203.07404*, pp. 1–35, Mar. 2022, Accessed: Aug. 29, 2026. [Online]. Available: https://arxiv.org/pdf/2203.07404

[18] A. S. Krishnapriyan, A. Gholami, S. Zhe, R. M. Kirby, and M. W. Mahoney, "Characterizing possible failure modes in physics-informed neural networks," Adv. Neural Inf. Process. Syst., vol. 34, pp. 26548–26560, 2021.

[19] S. Maddu, D. Sturm, C. Müller, and F. Sbalzarini, "Inverse Dirichlet weighting enables reliable training of physics informed neural networks," *Mach. Learn. Sci. Technol.*, vol. 3, no. 1, p. 015026, Feb. 2022, doi: 10.1088/2632-2153/AC3712.

[20] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *J. Comput. Phys.*, vol. 378, pp. 686–707, Feb. 2019, doi: 10.1016/J.JCP.2018.10.045.

[21] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nature Reviews Physics 2021 3:6*, vol. 3, no. 6, pp. 422–440, May 2021, doi: 10.1038/s42254-021-00314-5.

[22] S. Cai, Z. Mao, Z. Wang, M. Yin, and G. E. Karniadakis, "Physics-informed neural networks (PINNs) for fluid mechanics: a review," *Acta Mechanica Sinica 2021 37:12*, vol. 37, no. 12, pp. 1727–1738, Jan. 2022, doi: 10.1007/S10409-021-01148-1.

[23] X. Jin, S. Cai, H. Li, and G. E. Karniadakis, "NSFnets (Navier-Stokes flow nets): Physics-informed neural networks for the incompressible Navier-Stokes equations," *J. Comput. Phys.*, vol. 426, p. 109951, Feb. 2021, doi: 10.1016/J.JCP.2020.109951.

[24] P. Hähnel, J. Mareček, J. Monteil, and F. O'Donncha, "Using deep learning to extend the range of air pollution monitoring and forecasting," *J. Comput. Phys.*, vol. 408, p. 109278, May 2020, doi: 10.1016/J.JCP.2020.109278.

[25] M. Raissi, A. Yazdani, and G. E. Karniadakis, "Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations," Science, vol. 367, no. 6481, pp. 1026–1030, Feb. 2020, doi: 10.1126/science.aaw4741.

[26] N. K. Chaurasia and S. Chakraborty, "Reconstruction of the turbulent flow field with sparse measurements using physics-informed neural network," *Physics of Fluids*, vol. 36, no. 8, Aug. 2024, doi: 10.1063/5.0218611.

[27] M. Kawai, T. Shoji, A. L. Pillai, S. Tachibana, T. Yokomori, and R. Kurose, "Experimental and numerical investigation on flashback of low-swirling hydrogen–air jet flames," *Combust. Flame*, vol. 286, p. 114792, Apr. 2026, doi: 10.1016/J.COMBUSTFLAME.2026.114792.