

Numerical Simulation of NO_x Emissions in a Lean Premixed Hydrogen Combustor Using Different Levels of Fidelity

Original

Numerical Simulation of NO_x Emissions in a Lean Premixed Hydrogen Combustor Using Different Levels of Fidelity / Folcarelli, L., Marci', S., Madonia, V., Ferrero, A., Masseni, F., Pastrone, D.G.. - ELETTRONICO. - Proceedings of the AIAA AVIATION 2026 Forum:(2026), pp. 1-35. (AIAA AVIATION 2026 Forum San Diego, CA 8-12 June 2026) [10.2514/6.2026-4794].

Availability:

This version is available at: 11583/3016367 since: 2026-10-06T16:21:28Z

Publisher:

American Institute of Aeronautics and Astronautics

Published

DOI:10.2514/6.2026-4794

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

AIAA preprint/submitted version e/o postprint/Author's Accepted Manuscript

(Article begins on next page)

Numerical Simulation of NO_x Emissions in a Lean Premixed Hydrogen Combustor Using Different Levels of Fidelity

Lorenzo Folcarelli ^{*}, Sebastiano G. I. Marci [†], Vincenzo Madonia [‡], Andrea Ferrero [§], Filippo Masseni [¶] and Dario Pastrone ^{||}

Politecnico di Torino, Torino, Italy, 10129

This work presents a comprehensive computational framework for the prediction of NO_x emissions in lean premixed hydrogen combustion systems for aviation applications, using the swirl-stabilized AHEAD combustor as the reference test case. Three modeling levels are assessed across inlet air temperatures from 313 to 693 K, mass flow rates from 80 to 295 kg/h, and equivalence ratios from 0.30 to 1.00: a 2D axisymmetric RANS with the $k-\omega$ SST closure, a 3D RANS with the Explicit Algebraic Reynolds Stress Model, and a Chemical Reactor Network calibrated via the Covariance Matrix Adaptation Evolution Strategy. Turbulence–chemistry interactions are described through the Eddy Dissipation Concept, and both the Naik and CRECK kinetic mechanisms are tested. Given the inherent availability of both 2D and 3D CFD datasets, a lightweight modified U-Net is introduced to correct the computationally efficient 2D fields toward their corresponding 3D planar references. Validation across the operating envelope indicates that while the models exhibit differences in predictive accuracy, the Chemical Reactor Network successfully captures the fundamental emission trends, providing NO_x predictions with an accuracy comparable to the CFD models but at a computational cost several orders of magnitude lower. Furthermore, the U-Net effectively bridges the spatial fidelity gap, reducing the 2D–3D discrepancies for velocity, temperature, NO, and NO₂ fields while adding only a marginal computational overhead compared to full 3D computations.

Nomenclature

Roman symbols

$a_{n,j}$	= NASA7 polynomial coefficients of species j , [–]
a_{FI}	= fuel–air heat-exchange correlation constant, W/°C
a_{ij}	= component of the dimensionless anisotropy tensor, [–]
A_j	= EARSM model coefficients, [–]
$A_{\text{mixing tube}}$	= cross-sectional area of the mixing tube, m ²
b_{FI}	= fuel–air heat-exchange correlation constant, °C
c_p	= specific heat capacity at constant pressure, J/(kg·K)
$\hat{c}_p$	= molar heat capacity at constant pressure, J/(mol·K)
c_v	= specific heat capacity at constant volume, J/(kg·K)
C	= number of channels of the U-Net input/output tensor, [–]
$C_{3\text{D}}$	= average computational cost of one baseline 3D RANS–EARSM simulation, CPU h
$C_{\text{refined only}}$	= average computational cost of the refined-mesh iterations only, CPU h
$C_{3\text{D,ref}}$	= effective computational cost of one refined 3D RANS–EARSM simulation, CPU h
C_μ	= $k-\omega$ turbulence model constant, [–]
C_τ	= EDC time-scale constant, [–]
C_ξ	= EDC length-scale constant, [–]
$D_{j,m}$	= molecular mass diffusion coefficient of species j , m ² /s
$D_{T,j}$	= thermal (Soret) diffusion coefficient of species j , kg/(m·s)

^{*}PhD candidate, Department of Mechanical and Aerospace Engineering, Corso Duca degli Abruzzi 24, Torino, Italy, AIAA Member

[†]PhD student, Department of Mechanical and Aerospace Engineering, Corso Duca degli Abruzzi 24, Torino, Italy, AIAA Member

[‡]PhD student, Department of Mechanical and Aerospace Engineering, Corso Duca degli Abruzzi 24, Torino, Italy, AIAA Member

[§]Associate Professor, Department of Mechanical and Aerospace Engineering, Corso Duca degli Abruzzi 24, Torino, Italy

[¶]Assistant Professor, Department of Mechanical and Aerospace Engineering, Corso Duca degli Abruzzi 24, Torino, Italy, AIAA Member

^{||}Full Professor, Department of Mechanical and Aerospace Engineering, Corso Duca degli Abruzzi 24, Torino, Italy, AIAA Associate Fellow

f	= generic scalar, [-]
G	= volumetric heat release of chemical reactions, W/m ³
h	= mixture sensible specific enthalpy, J/kg
$\tilde{h}_j$	= sensible specific enthalpy of species j , J/kg
h_j^0	= molar enthalpy of formation of species j , J/mol
h_{in}	= specific enthalpy of the inlet stream of a CRN reactor, J/kg
H	= height (rows) of the U-Net input/output tensor, [-]
II_{Ω}, IV	= invariants of the strain-rate and vorticity tensors used in the EARSMS expansion, [-]
$J(\vec{\Theta})$	= CRN-calibration objective function, [-]
$\vec{J}_j$	= diffusion mass flux vector of species j , kg/(m ² ·s)
k	= laminar thermal conductivity, W/(m·K)
k_{eff}	= effective thermal conductivity, W/(m·K)
m	= reactor mass content, kg
$\dot{m}$	= mass flow rate, kg/s
M	= total number of species in the mixture, [-]
M_j	= molar mass of species j , kg/mol
n_{cal}	= number of experimental points used for CRN calibration, [-]
N	= EARSMS model coefficient, [-]
p	= static pressure, Pa
P_k	= production rate of turbulent kinetic energy, m ² /s ³
Pr_t	= turbulent Prandtl number, [-]
r	= radial coordinate, m
R	= universal gas constant, J/(mol·K)
$\mathcal{R}_j$	= volumetric net rate of creation of species j , kg/(m ³ ·s)
Sc_t	= turbulent Schmidt number, [-]
S_{ij}	= strain-rate tensor component, 1/s
T	= static temperature, K
T_{AF}	= adiabatic flame temperature, K
$\text{tr}\{\cdot\}$	= trace operator, [-]
u_0	= mean cold-flow velocity in the mixing tube, m/s
$\overline{u_i u_j}$	= component of the Reynolds stress tensor, m ² /s ²
u_k	= specific internal energy of species k , J/kg
U_i	= i -th component of the mean velocity, m/s
$\vec{v}$	= velocity vector, m/s
v_x, v_y, v_z	= Cartesian velocity components, m/s
v_z, v_r, v_{θ}	= axial, radial and tangential velocity components in cylindrical coordinates, m/s
V	= volume, m ³
$\dot{V}$	= volume flow rate, m ³ /s
$\dot{V}_{axial}, \dot{V}_{swirl}$	= volumetric flow rates through the axial orifice and the radial swirler, m ³ /s
w_i	= weighting factor of point i in the CRN objective function, [-]
W	= width of the U-Net input/output tensor, [-]
x_i, x_j, x_k	= Cartesian spatial coordinates, m
y^+	= dimensionless wall distance, [-]
Y_j	= mass fraction of species j , [-]
Y_j^*	= fine-scale mass fraction of species j , [-]
$Y_{k,in}$	= inlet mass fraction of species k in a CRN reactor, [-]
Y_{log}	= log ₁₀ -transformed nitrogen-oxide mole-fraction field used in U-Net pre-processing, [-]
z	= axial coordinate, m
<i>Greek symbols</i>	
β_i	= EARSMS tensor-polynomial coefficients, [-]
χ	= axial-injection ratio, [-]
δ_{ij}	= Kronecker delta, [-]
ϵ	= dissipation rate of turbulent kinetic energy, m ² /s ³

ε	= small positive floor used in the logarithmic transform of NO/NO ₂ fields, [-]
θ	= azimuthal coordinate, rad
$\theta_1, \dots, \theta_{10}$	= CRN calibration parameters, [-]
$\vec{\Theta}$	= CRN calibration vector $[\theta_1, \dots, \theta_{10}]$, [-]
κ	= turbulent kinetic energy, m ² /s ²
κ_{fs}	= volume fraction of the fine-scale (EDC) structures, [-]
μ	= laminar dynamic viscosity, Pa·s
μ_t	= turbulent dynamic viscosity, Pa·s
ν	= kinematic viscosity, m ² /s
ρ	= density, kg/m ³
ρ_{air}	= air density in the mixing tube, kg/m ³
$\hat{\Sigma}, \Sigma_{ij}$	= non-dimensional strain-rate tensor and its components, [-]
$\hat{\Omega}, \Omega_{ij}$	= non-dimensional vorticity tensor and its components, [-]
$\hat{\tau}$	= total stress tensor, Pa
τ	= turbulent time scale (k/ε), s
τ^*	= EDC reaction time scale, s
ϕ	= global equivalence ratio, [-]
ξ	= mass fraction of the fine-structure region, [-]
ω	= specific dissipation rate of turbulent kinetic energy, 1/s

Superscripts and overscripts

0	= formation condition
*	= fine-scale value
+	= wall-distance unit (y^+)
$\overline{(\cdot)}$	= Favre (density-weighted) average
$\overline{(\cdot)}$	= Reynolds (ensemble) average
$(\dot{\cdot})$	= time derivative or flow rate per unit time
$(\vec{\cdot})$	= vector quantity
$(\hat{\cdot})$	= tensor quantity

Acronyms and abbreviations

AHEAD	= Advanced Hybrid Engines for Aircraft Development
CAD	= Computer-Aided Design
CFD	= Computational Fluid Dynamics
CMA-ES	= Covariance Matrix Adaptation Evolution Strategy
CNN	= Convolutional Neural Network
CPU	= Central Processing Unit
CRECK	= Politecnico di Milano CRECK Modeling Group kinetic mechanism
CRN	= Chemical Reactor Network
EARSM	= Explicit Algebraic Reynolds Stress Model
EDC	= Eddy Dissipation Concept
ELU	= Exponential Linear Unit
FL_{AX}	= axial flame zone (CRN reactor)
FL_{SW}	= swirled flame zone (CRN reactor)
GPU	= Graphics Processing Unit
IRZ	= Inner Recirculation Zone
LES	= Large Eddy Simulation
MAE	= Mean Absolute Error
MAPE	= Mean Absolute Percentage Error
MSE	= Mean Squared Error
NASA7	= NASA seven-coefficient thermodynamic polynomial
NO _x	= nitrogen oxides
ORZ	= Outer Recirculation Zone
OUT	= Outlet CRN reactor

PIV	= Particle Image Velocimetry
ppmv	= parts per million by volume
RANS	= Reynolds-Averaged Navier-Stokes
Res ₁ , Res ₂	= CRN inlet reservoirs
RMS	= Root Mean Square
SST	= Shear Stress Transport (Menter $k-\omega$ model)
U-Net	= U-shaped Convolutional Encoder-Decoder Network
Z _{AX}	= axial CRN reactor zone

I. Introduction

THE growing urgency to address environmental degradation and climate change has positioned emission reduction as a crucial global challenge across all industrial sectors. As international agreements such as the Paris Agreement tighten regulatory frameworks and public awareness of environmental impacts increases, industries worldwide are adopting innovative approaches to minimize their ecological footprint [1].

Hydrogen-fueled gas turbine engines represent one of the most promising pathways toward zero-carbon aviation, as hydrogen/air combustion generates carbon-free products, eliminating the problem of CO and CO₂ production [2]. However, transitioning to hydrogen combustion introduces a set of technical challenges that must be addressed. A significant concern is that hydrogen/air mixtures exhibit higher adiabatic flame temperatures than conventional kerosene/air mixtures [3]. This temperature differential is particularly problematic because combustion at elevated temperatures promotes the formation of nitrogen oxides (NO_x) through thermal mechanisms [4]. NO_x emissions are especially concerning in aviation contexts due to their contribution to ozone formation at cruise altitudes and their impact on air quality degradation near airports [5]. Fortunately, hydrogen possesses combustion characteristics that can be leveraged to mitigate NO_x formation. Its wide flammability range enables stable combustion under leaner conditions than conventional fuels [6]. Operating in lean premixed combustion regimes reduces peak flame temperatures, thereby minimizing thermal NO_x formation while maintaining acceptable combustion efficiency and stability.

Lean hydrogen combustion systems present complex fluid dynamic and thermochemical interactions that require sophisticated computational modeling approaches to accurately predict emissions performance. Computational Fluid Dynamics (CFD) represents a useful tool for investigating these phenomena, offering detailed insights into flow fields, temperature distributions, and species concentrations that directly influence emissions formation [7]. The choice of turbulence modeling approach significantly impacts the accuracy of NO_x predictions in CFD simulations. Reynolds-Averaged Navier-Stokes (RANS) methods offer computational efficiency but require careful selection of turbulence closure models. The Explicit Algebraic Reynolds Stress Model (EARSIM) has demonstrated superior performance in swirling flows characteristic of gas turbine combustors, as it better captures the anisotropic nature of turbulence compared to simpler two-equation models [8]. This capability is particularly important for accurately predicting the mixing processes that influence local equivalence ratios and, consequently, NO_x formation rates. Several researchers have applied CFD methodologies to hydrogen combustion systems with varying levels of fidelity [9–16]. However, the computational cost of CFD simulations with a detailed chemistry mechanism that includes NO_x production remains high.

This study aims to develop and validate a comprehensive computational framework for predicting nitrogen oxide emissions in hydrogen-fueled combustion systems, specifically focusing on the swirl-stabilized AHEAD combustor [17]. Recognizing the distinct computational and physical differences between these CFD fidelity levels, this study naturally extends to testing a U-Net architecture designed to map and correct 2D fields using 3D reference data. The investigation employs a progressive hierarchical approach consisting of:

- Two-dimensional axisymmetric RANS simulations using the $k-\omega$ SST turbulence model for computationally efficient baseline predictions;
- Three-dimensional RANS simulations utilizing the EARSIM turbulence model to accurately capture anisotropic turbulence effects in the swirling flow;
- A Chemical Reactor Network (CRN) methodology implemented in Cantera to allow rapid emissions assessment with a computational cost several orders of magnitude lower than CFD;
- A data-driven correction framework based on a modified U-Net architecture to systematically correct discrepancies between the lower-fidelity 2D and 3D planar reference fields.

The resulting multi-fidelity database, spanning variations in inlet air mass flow, temperature, and equivalence ratio,

serves as a foundation for developing computationally efficient surrogates. Such tools are essential for accelerating the design and optimization of zero-carbon aircraft propulsion systems that must satisfy both stringent performance requirements and environmental emissions targets.

The remainder of this paper is organized as follows: Section II describes the AHEAD combustor geometry and the experimental operating conditions. Section III details the computational fluid dynamics setup, including the EARSM turbulence closure and the reactive flow modeling. Section IV presents the development and calibration of the Chemical Reactor Network. Section V discusses the validation against experimental data and the comparison between the different modeling fidelities. Section VI introduces the U-Net architecture and the data-driven multi-fidelity correction strategy.

II. Test Case Description

The reference configuration adopted in the present work is the hydrogen-fired AHEAD combustor, developed within the “Advanced Hybrid Engines for Aircraft Development” framework and experimentally investigated by Reichel et al. [17, 18]. The burner is a partially premixed, swirl-stabilized configuration designed for low-emission hydrogen combustion. A CAD representation of the geometry is shown in Fig. 1.

Air is supplied to an upstream plenum before entering the premixing section through two distinct paths. The main swirling contribution passes through a radial swirl generator, while a second air stream is admitted along the burner centerline through a variable-diameter axial orifice. The fraction of air entering through this central branch is controlled by the axial-orifice diameter and is quantified by the axial injection ratio χ , defined as

$$\chi = \frac{\dot{V}_{\text{axial}}}{\dot{V}_{\text{axial}} + \dot{V}_{\text{swirl}}} \quad (1)$$

The axial injection ratio χ is not imposed independently. Instead, it results from the pressure-loss ratio between the axial orifice and the radial swirler, and therefore depends on the selected axial-orifice diameter and on the combustor operating conditions [19].

Hydrogen is supplied through a central fuel conduit and injected directly into the premixing section in the axial direction. The fuel injection system consists of 16 circumferentially distributed orifices, each with a diameter of 0.8 mm. Downstream of the air and fuel injection region, the reactants flow through a cylindrical mixing tube of diameter 34 mm, which discharges into a cylindrical combustion chamber with diameter of 105 mm.

A small additional fraction of the plenum air is introduced near the downstream end of the mixing tube through 22 circumferentially arranged dilution holes, each having a diameter of 0.6 mm. These holes are used to reduce the local equivalence ratio in the near-wall region, thereby limiting the susceptibility to boundary-layer flashback. Additional details on the combustor design are provided in Ref. [20].




Fig. 1 CAD drawing of the AHEAD combustor.

The investigated operating conditions are specified by the total inlet air mass flow rate, the air preheating temperature, the global equivalence ratio, the geometrical swirl number and the diameter of the axial air-injection orifice. The geometrical swirl number is defined on the basis of the radial swirl generator only, and therefore does not include the

contribution of the air entering through the axial orifice. This parameter is adjusted by changing the blockage of the radial swirl generator using annular rings. All the conditions considered in the present work correspond to a geometrical swirl number of 0.9 and an axial air orifice diameter of 8.8 mm, which corresponds to a χ equal to 12.5 %.

The hydrogen is stored at ambient temperature, but its temperature at injection is influenced by heat exchange with the preheated-air environment surrounding the fuel supply system. Following the correlation proposed in Ref. [17], the fuel temperature is estimated from the inlet air temperature, hydrogen mass flow rate and equivalence ratio as

$$T_{H_2} = \frac{a_{FI}T_{in} + b_{FI}(\dot{m}_{H_2}c_{p,H_2} - 0.5a_{FI})}{0.5a_{FI} + \dot{m}_{H_2}c_{p,H_2}} \quad (2)$$

where T_{in} is the inlet air temperature, $\dot{m}_{H_2}$ is the hydrogen mass flow rate, and c_{p,H_2} denotes the constant-pressure specific heat capacity of hydrogen. The latter is assumed constant over the temperature interval relevant to the present operating conditions. The coefficients $a_{FI} = 1.40 \text{ W/}^\circ\text{C}$ and $b_{FI} = 34.9 \text{ }^\circ\text{C}$ were obtained experimentally [17].

The complete experimental matrix covers inlet air temperatures between 313 K and 693 K, total inlet air mass flow rates from 116 kg/h to 255 kg/h, and equivalence ratios in the range $0.20 \leq \phi \leq 1.00$. In the present numerical analysis, however, the cases at $\phi = 0.20$ are not included since neither the RANS simulations nor the CRN calculations achieved ignition under these very lean conditions.

The experimental NO_x data used for comparison are summarized in Table 1. The measurements, reported by Reichel et al. and collected for the AHEAD test case at atmospheric pressure [21, 22], are expressed as parts per million by volume on a dry basis, corrected to 15% O_2 concentration. The data show the expected increase of NO_x emissions with equivalence ratio. For lean conditions below approximately $\phi = 0.50$ the measured values remain comparatively low, whereas a rapid growth is observed as the mixture approaches richer operating conditions, especially at elevated inlet air temperature.

Table 1 Experimental NO_x measurements for the investigated operating conditions, extracted from Ref. [22].

T_{air} [K]	$\dot{m}_{\text{air,tot}}$ [kg/h]	ϕ [-]	NO_x [ppmv, dry, 15% O_2]
313	255	0.30	1.83
313	255	0.40	2.32
313	255	0.50	3.61
313	255	0.60	6.22
453	180	0.30	3.56
453	180	0.40	3.69
453	180	0.50	5.94
453	180	0.60	9.39
623	130	0.30	3.93
623	130	0.40	5.28
623	130	0.50	10.12
623	130	0.60	26.13
623	130	0.80	182.33
623	130	0.87	212.07
623	130	0.96	305.54
693	116	0.30	3.90
693	116	0.40	6.75
693	116	0.50	16.64

III. Computational Fluid Dynamics

A. 2D Axisymmetric RANS

The simulations are performed in Ansys® Fluent, using a 2D steady swirl axisymmetric flow approximation, i.e. neglecting every derivative along the azimuthal θ direction while accounting for a tangential component of velocity v_θ variable along the axial z and radial r directions.

$$\frac{\partial (\cdot)}{\partial \theta} = 0 \quad (3)$$

$$\vec{v}(z, r) = [v_z, v_r, v_\theta] \quad (4)$$

Introducing the Favre-averaged variables, denoted by a tilde, obtained for a generic scalar f as:

$$\tilde{f} = \frac{\overline{\rho f}}{\bar{\rho}} \quad (5)$$

where the overbar indicates the typical Reynolds averaging and ρ represents the density, the corresponding Navier-Stokes equations are written in compressible form in Eqs.(6)–(8) [23].

$$\nabla \cdot (\bar{\rho} \vec{v}) = 0 \quad (6)$$

$$\nabla \cdot (\bar{\rho} \vec{v} \vec{v}) = -\nabla \bar{p} + \nabla \cdot (\hat{\tau}) \quad (7)$$

$$\nabla \cdot \left(\bar{\rho} \vec{v} \left(\tilde{h} + \frac{\tilde{v}^2}{2} \right) \right) = \nabla \cdot \left(k_{eff} \nabla \tilde{T} - \sum_j \tilde{h}_j \vec{J}_j + \hat{\tau} \cdot \vec{v} \right) + G \quad (8)$$

In Eqs.(6)–(8) p represents the pressure, $\hat{\tau}$ the total stress tensor, h the sensible enthalpy, k_{eff} the effective conductivity, J_j the diffusion flux of species j , and G the volumetric heat of chemical reactions. The total stress tensor $\hat{\tau}$ is formed by the sum of the laminar and turbulent stresses:

$$\hat{\tau}_{ij} = 2(\mu + \mu_t) \left(\tilde{S}_{ij} - \frac{1}{3} \frac{\partial \tilde{v}_k}{\partial x_k} \delta_{ij} \right) - \frac{2}{3} \bar{\rho} \kappa \delta_{ij} \quad (9)$$

with μ and μ_t indicating the laminar and turbulent dynamic viscosity, respectively, κ the turbulent kinetic energy and S_{ij} the strain rate tensor, i.e.:

$$\tilde{S}_{ij} = \frac{1}{2} \left(\frac{\partial \tilde{v}_i}{\partial x_j} + \frac{\partial \tilde{v}_j}{\partial x_i} \right) \quad (10)$$

The turbulent viscosity is evaluated using the k – ω Shear Stress Transport (SST) two-equations turbulence model developed by Menter [24]. The effective thermal conductivity k_{eff} is defined as:

$$k_{eff} = k + \frac{c_p \mu_t}{Pr_t} \quad (11)$$

where k is the laminar thermal conductivity, c_p is the specific heat capacity at constant pressure and Pr_t is the turbulent Prandtl number. In this work, Pr_t was selected to be 0.85, and the c_p is obtained as the mixture-average of the $c_{p,j}$ of all the species in the mixture. The molar heat capacity at constant pressure $\hat{c}_p$ is approximated by NASA7 polynomials [25] for the single species j , i.e.:

$$\frac{\hat{c}_{p,j}(\tilde{T})}{R} = a_{0,j} + a_{1,j} \tilde{T} + a_{2,j} \tilde{T}^2 + a_{3,j} \tilde{T}^3 + a_{4,j} \tilde{T}^4 \quad (12)$$

with R being the universal gas constant. Thus, the sensible enthalpy for the single species j in the case of ideal gas approximation is calculated as:

$$\tilde{h}_j = \int_{T_{ref}}^{\tilde{T}} c_{p,j} d\tilde{T} \quad (13)$$

with T_{ref} being 298.15 K. Then, the sensible enthalpy of the mixture is evaluated as:

$$\tilde{h} = \sum_j \tilde{Y}_j \tilde{h}_j \quad (14)$$

The diffusion flux of species j was modeled using the full multicomponent formulation including Soret diffusion [23]:

$$\vec{J}_j = -\rho \sum_{k=1}^{M-1} D_{jk} \nabla \tilde{Y}_k - \frac{\mu_t}{Sc_t} \nabla \tilde{Y}_j - D_{T,j} \frac{\nabla \tilde{T}}{\tilde{T}} \quad (15)$$

where Sc_t is the turbulent Schmidt number set to 0.7, Y is the mass fraction, D_{jk} are the generalized Fick's law diffusion coefficients obtained from the Maxwell–Stefan equations [23], and $D_{T,j}$ is the Soret diffusion coefficient.

Finally, the spatial discretization is performed using second-order upwind schemes and the coupled pressure-velocity scheme.

1. Geometry Approximation for 2D simulation

The 2D geometry fluid domain is built with the constraint that the inlets in the axisymmetric geometry have the same total area of the real combustor, to ensure that the fluid enters through the inlets with the same velocity as it would in the real 3D geometry. The 2D geometry CAD with dimensions and boundary conditions used in this work is shown in Fig. 2.




Fig. 2 Simplified 2D axisymmetric geometry of the AHEAD combustor, including the main dimensions and boundary-condition locations used in the RANS simulations.

The 2D geometry of the radial swirler is approximated based on the actual swirler design, which features eight slots. Each slot has a total length of 25 mm, connected by two semicircular arcs with a radius of 2.25 mm. For a configuration with a geometric swirl number of 0.9, the slots are partially blocked by rings, obstructing 7 mm of their length.

In the 2D axisymmetric approximation, the swirler is modeled with an inlet followed by a wall condition to represent the obstruction caused by the blocking rings. Since the continuous cylindrical surface in the 2D model represents the distributed area of the swirler slots, the inlet length in the simulation is shorter than the actual slot length. It is, therefore, necessary to determine the inlet position in the 2D approximation, and it was decided to retain the geometric centroid of the actual slot.

The same decision to retain the geometric centroid was applied in approximating the inlet area for both the fuel and the dilution air. In this approximation, the effect of the boundary layer occupying different portions of the injection orifice between the real and approximated geometry was neglected. Accounting for this would have required modifying the geometry of the fuel injection conduit as a function of the equivalence ratio. Additionally, to better approximate the inlet conditions of the hydrogen stream in the premixing tube, a portion of the fuel conduit is also included in the 2D geometry, which can be seen in Detail A in Fig. 2. Assuming that 2.5% of the total inlet air mass flow rate passes through the dilution orifices, the mass flow rates at the boundaries are determined using the definition of the axial injection ratio, χ , as follows:

$$\begin{cases} \dot{m}_{\text{dil}} = 0.025 \dot{m}_{\text{tot}} \\ \dot{m}_{\text{swirl}} = (1 - \chi) (\dot{m}_{\text{tot}} - \dot{m}_{\text{dil}}) \\ \dot{m}_{\text{axial}} = \dot{m}_{\text{tot}} - \dot{m}_{\text{swirl}} - \dot{m}_{\text{dil}} \end{cases} \quad (16)$$

In Eq. 16, $\dot{m}_{\text{tot}}$ represents the total air mass flow rate, $\dot{m}_{\text{dil}}$ is the mass flow rate through the dilution orifices, $\dot{m}_{\text{axial}}$ is the mass flow rate through the axial orifice, and $\dot{m}_{\text{swirl}}$ is the mass flow rate through the swirl orifice.

Additionally, it is necessary to define the tangential velocity component at the swirler inlet. However, due to changes in the inlet geometry in the 2D RANS model, the geometrical swirl number cannot be directly applied to establish boundary conditions. Instead, a tangential-to-radial velocity ratio of 0.8 at the swirler inlet was found to provide good agreement with experimental results, based on comparisons between CFD predictions and experimental Particle Image Velocimetry (PIV) data, extracted from Ref.[22], for a single operating condition.

B. 3D RANS with Explicit Algebraic Reynolds Stress Model

The three-dimensional simulations employ a 3D steady RANS approach with the Explicit Algebraic Reynolds Stress Model in Ansys® Fluent [23]. Due to the complexity of the notation, the Favre averaging symbols are omitted throughout this subsection for clarity, although their application remains mathematically valid. Unlike the 2D axisymmetric approach, the 3D model accounts for variations in all spatial directions:

$$\frac{\partial (\cdot)}{\partial \theta} \neq 0 \quad (17)$$

This yields a full three-dimensional velocity field:

$$\vec{v}(x, y, z) = [v_x, v_y, v_z] \quad (18)$$

The continuity, momentum, and energy conservation equations (Eqs. 6, 7 and 8) remain valid. However, the EARS model approach differs substantially from the k - ω SST model in its treatment of the Reynolds stress tensor because it computes Reynolds stresses using the anisotropy tensor [8]:

$$\overline{u_i u_j} = \kappa \left(a_{ij} + \frac{2}{3} \delta_{ij} \right) \quad (19)$$

where κ is the turbulent kinetic energy, a_{ij} is a component of the dimensionless anisotropy tensor, and δ_{ij} is the Kronecker delta. The anisotropy tensor $\hat{a}$ is determined by solving the following implicit algebraic matrix equation:

$$N \hat{a} = -A_1 \hat{\Sigma} + (\hat{a} \hat{\Omega} - \hat{\Omega} \hat{a}) - A_2 \left(\hat{a} \hat{\Sigma} - \hat{\Sigma} \hat{a} - \frac{2}{3} \text{tr}\{\hat{a} \hat{\Sigma}\} \right) \quad (20)$$

In this equation, $\hat{\Sigma}$ is the strain rate tensor and $\hat{\Omega}$ is the vorticity tensor, and N :

$$N = A_3 + A_4 \left(\frac{P_k}{\epsilon} \right) \quad (21)$$

where P_k is the production rate of turbulent kinetic energy and ϵ is its dissipation rate. The coefficients A_j depend on the coefficients of the pressure-strain term in the Reynolds stress transport model. In the Ansys Fluent implementation, they are set to $A_1=1.245$, $A_2=0$, $A_3=1.8$, and $A_4=2.25$ [23].

The non-dimensional strain-rate and vorticity tensors are defined as:

$$\Sigma_{ij} = \frac{1}{2} \tau \left(\frac{\partial U_i}{\partial x_j} + \frac{\partial U_j}{\partial x_i} \right), \quad \Omega_{ij} = \frac{1}{2} \tau \left(\frac{\partial U_i}{\partial x_j} - \frac{\partial U_j}{\partial x_i} \right) \quad (22)$$

where τ is the turbulent time scale, defined as the ratio of turbulent kinetic energy to its dissipation rate:

$$\tau = \frac{\kappa}{\epsilon} = \frac{1}{C_\mu \omega} \quad (23)$$

with $C_\mu = 0.09$ and ω being the specific dissipation rate.

To obtain an explicit solution, the anisotropy tensor is expressed as a tensor polynomial, represented as a series of products of the tensors $\hat{\Sigma}$ and $\hat{\Omega}$ with scalar coefficients:

$$a_{ij} = \beta_1 \Sigma_{ij} + \beta_3 \left(\Omega_{ik} \Omega_{kj} - \frac{1}{3} II_{\Omega} \delta_{ij} \right) + \beta_4 (\Sigma_{ik} \Omega_{kj} - \Omega_{ik} \Sigma_{kj}) \\ + \beta_6 \left(\Sigma_{ik} \Omega_{kl} \Omega_{lj} + \Omega_{ik} \Omega_{kl} \Sigma_{lj} - \frac{2}{3} IV \delta_{ij} - II_{\Omega} \Sigma_{ij} \right) \quad (24)$$

The β coefficients depend on invariants of $\hat{\Sigma}$ and $\hat{\Omega}$, calculated using formulas provided by Wallin and Johansson [8]. The EARSM implementation in Ansys Fluent differs from the original model in several aspects:

- 1) The fourth-order tensor polynomial contribution is neglected for computational efficiency [23];
- 2) The tensor basis is modified according to Apsley and Leschziner [26], but the model remains algebraically equivalent to the original formulation [23];
- 3) Standard eddy viscosity $\mu_t = \rho \kappa / \omega$ is used for diffusion terms, whereas the original model uses the effective eddy viscosity [23].

For three-dimensional flows, the equation to be solved for the function N is of sixth order, and no direct explicit solution can be derived. However, an analytic solution derived from a cubic equation is implemented in Ansys Fluent as recommended by Wallin and Johansson [8, 27]. The key advantage of EARSM over the 2D $k-\omega$ SST model is its ability to capture anisotropic turbulence characteristics [28]. This is valuable for a combustor study, as the nonlinear formulation better represents complex flow phenomena such as streamline curvature, rotation effects, and secondary flows that influence mixing patterns and reaction zones.

The spatial discretization and chemistry modeling approach maintain consistency with the 2D simulations, utilizing second-order upwind schemes and coupled pressure-velocity algorithm.

C. Reactive Flow modeling in CFD

The volumetric heat of reaction G is calculated as:

$$G = - \sum_j \frac{h_j^0}{M_j} \mathcal{R}_j \quad (25)$$

where h_j^0 , M_j and $\mathcal{R}_j$ are the enthalpy of formation, the molecular weight and the volumetric rate of creation, respectively, of the j^{th} species.

In addition to Eqs. (6)–(8), a conservation equation is considered for $M - 1$ species, where M denotes the total number of species in the mixture. The mass fraction of the excluded species is calculated as 1 minus the sum of the mass fractions of all the other species. To minimize the numerical error, N_2 is chosen to be the species not directly transported in the flow domain, as it is expected to have the largest mass fraction. These conservation equations are then written as:

$$\nabla \cdot (\bar{\rho} \vec{v} \tilde{Y}_j) = -\nabla \cdot \vec{J}_j + \mathcal{R}_j \quad (26)$$

The reaction rates $\mathcal{R}_j$ are modeled using the Eddy Dissipation Concept (EDC) [29], taking into account turbulence and chemistry interactions. The EDC model assumes that the reactions only occur in the fine scales, i.e. in the smallest turbulence scales. It is then necessary to calculate the volume of the computational cell that is occupied by these fine scales, i.e. the volume fraction κ_{fs} :

$$\kappa_{fs} = \frac{\xi^2}{1 - \xi^3} \quad (27)$$

In Eq. 27 the variable ξ represents the fine-structure region mass fraction, which is calculated as:

$$\xi = C_{\xi} \left(\frac{\nu \epsilon}{\kappa^2} \right)^{\frac{1}{4}} \quad (28)$$

where C_{ξ} is a constant value equal to 2.1377, ν is the kinematic viscosity, κ is the turbulent kinetic energy and ϵ is the rate of dissipation of turbulent kinetic energy. Moreover, the reactions are integrated over the *time scale* τ^* , which is defined as:

$$\tau^* = C_\tau \left(\frac{\nu}{\epsilon} \right)^{\frac{1}{2}} \quad (29)$$

where C_τ is a constant value equal to 0.4082. Finally, the reaction rates $\mathcal{R}_j$ can be written as [23]:

$$\mathcal{R}_j = \rho K_{fs} \frac{Y_j^* - \tilde{Y}_j}{\tau^*} \quad (30)$$


where Y_j^* represents the fine-scale species mass fraction after reacting over the time scale τ^* in an adiabatic and isobaric perfectly stirred reactor, whose initial conditions are the species and temperature at the end of the previous iteration. In this work, to guarantee that the methodology proposed is independent of the chemical description considered, two chemical mechanisms have been used: these are the Naik mechanism, available in Ansys Fluent [30] and the CRECK [31–33] mechanism. These two mechanisms include detailed nitrogen-oxidation chemistry and they are summarized in Table 2.

Table 2 Specifications of chemical mechanisms employed in this study. For each mechanism, the table reports the total number of species, the species list, and the total number of reactions.

Mechanism	Species	Species list	Reactions
Naik [30]	21	H, H ₂ , H ₂ O, H ₂ O ₂ , HNO, HNO ₂ , HO ₂ , HONO, HONO ₂ , N, N ₂ , N ₂ O, NH, NH ₂ , NH ₃ , NNH, NO, NO ₂ , O, O ₂ , OH	105
CRECK [31–33]	31	Ar, H, H ₂ , H ₂ NN, H ₂ NO, H ₂ O, H ₂ O ₂ , He, HNO, HNO ₂ , HNOH, HO ₂ , HONO, HONO ₂ , N, N ₂ , N ₂ H ₂ , N ₂ H ₃ , N ₂ H ₄ , N ₂ O, NH, NH ₂ , NH ₂ OH, NH ₃ , NNH, NO, NO ₂ , NO ₃ , O, O ₂ , OH	203

D. Computational Mesh and Grid Sensitivity

The 2D computational domain is discretized using a baseline mesh of approximately 2×10^5 cells. To adequately resolve the near-wall boundary layers, the mean y^+ is below 0.5 across all solid boundaries, with local maxima remaining strictly below 10 across all simulated operating points. The domain is partitioned into three distinct topological regions. The hydrogen injection tube features the highest spatial resolution, with a mean cell area of $3.42 \times 10^{-10} \text{ m}^2$. A primary refinement region extends from the air inlet through the primary reaction zone, utilizing a mean cell size of $2.91 \times 10^{-8} \text{ m}^2$ to capture the steep thermal and species gradients. Finally, a coarser mesh of $1.77 \times 10^{-7} \text{ m}^2$ is deployed in the downstream post-flame region where scalar gradients diminish. A grid independence study was performed by uniformly quadrupling the baseline cell count: the macroscopic flame topology and scalar fields showed negligible deviations, confirming that the 2×10^5 cell grid is fully resolved and numerically asymptotic. The 2D mesh is depicted in Fig. 3.



(a) Mesh at the air inlet and hydrogen injection region.

(b) Mesh size variation in the combustion chamber.

Fig. 3 Computational 2D mesh.

For the 3D simulations, a poly-hexcore meshing strategy was evaluated using two grids of 2.1×10^6 and 3.2×10^6 cells. On the coarser grid, the target volume element size is 2 mm within the flame, mixing tube, swirler, and injection




Fig. 4 Longitudinal section of the 3D computational mesh, including a detailed view of the mixing tube.

plate, bounded between 0.5 mm and 4 mm depending on local geometric features. The surface mesh is 1 mm on the combustion chamber walls and 0.5 mm on the mixing tube and swirler, again bounded between 0.5 and 4 mm. The refined grid halves the target volume element size in the reactive regions to 1 mm and reduces the local volume maximum to 2 mm, while surface specifications remain unchanged. Comparison between the two refinement levels revealed discrepancies in the temperature field, and consequently in NO_x production, particularly within the outer recirculation zone. Therefore, to resolve more accurately these structures, all 3D reactive flow simulations were performed using the refined 3.2×10^6 cell mesh. The refined 3D mesh is presented in Fig. 4.

IV. Chemical Reactor Network

The reduced-order model used in this work is a Chemical Reactor Network (CRN) built according to the methodology introduced in Ref. [34] and in Ref. [35]. The network is implemented in Cantera [36], which is used to evaluate the thermochemical state of each reactor and to integrate the coupled conservation equations governing the system.

The combustor is represented as a set of ideal-gas control volumes. Each zone is modeled as a constant-volume reactor with adiabatic and chemically inert walls. For each reactor, the governing equations are the mass balance, the species transport equations, and the energy balance, written as

$$\frac{dm}{dt} = \sum_{\text{in}} \dot{m}_{\text{in}} - \sum_{\text{out}} \dot{m}_{\text{out}}, \quad (31)$$

$$m \frac{dY_k}{dt} = \sum_{\text{in}} \dot{m}_{\text{in}} (Y_{k,\text{in}} - Y_k) + \dot{m}_{k,\text{gen}}, \quad (32)$$

$$mc_v \frac{dT}{dt} = \sum_{\text{in}} \dot{m}_{\text{in}} \left(h_{\text{in}} - \sum_k u_k Y_{k,\text{in}} \right) - \frac{pV}{m} \sum_{\text{out}} \dot{m}_{\text{out}} - \sum_k \dot{m}_{k,\text{gen}} u_k. \quad (33)$$

Here, m is the mass contained in the reactor, while $\dot{m}_{\text{in}}$ and $\dot{m}_{\text{out}}$ are the inlet and outlet mass flow rates. The quantity Y_k denotes the mass fraction of species k , and $\dot{m}_{k,\text{gen}}$ is the net production rate of that species due to chemical

reactions. The reactor temperature, pressure, volume, and constant-volume specific heat are indicated by T , p , V , and c_v , respectively. The incoming specific enthalpy is denoted by h_{in} , whereas u_k is the specific internal energy of species k .

Figure 5 shows how the main flow regions of the combustor are represented within the CRN. The structures identified from the experimental velocity field are reduced to a limited number of reactor zones. In particular, the chamber is discretized into six reactor volumes. The axial region immediately downstream of the mixing tube exit is represented by Z_{AX} , while the main reaction region is modeled by the swirled flame zone FL_{SW} . The central and peripheral recirculating structures are represented by the inner and outer recirculation zones, denoted as IRZ and ORZ, respectively. Downstream of the primary flame region, the predominantly axial portion of the flame is described by FL_{AX} , while the OUT reactor collects the flow leaving the IRZ and FL_{AX} before discharge to the exhaust.




Fig. 5 Mapping of the swirl-stabilized combustor onto the CRN domain. The upper panel shows the experimental axial velocity field reconstructed from Particle Image Velocimetry measurements [22]. The lower panel shows the corresponding representation in terms of inlet streams and reactor zones.

The resulting network topology is shown in Fig. 6. The two inlet reservoirs represent the mixing tube and define two premixed streams with independent equivalence ratios and air mass flow rates. Downstream, the reactor zones describe the combustion chamber. The figure also reports the mass-flow connections between the different regions, which determine how mass, chemical species, and enthalpy are redistributed within the network.




Fig. 6 Topology of the Chemical Reactor Network developed for the AHEAD combustor. The diagram reports the inlet reservoirs, the reactor zones associated with the combustion chamber, and the inter-reactor mass-flow connections.

The inter-reactor mass flow rates are obtained by imposing the mass-flow balance on each reactor of the network.

In addition, a control equation is introduced to enforce the equivalence ratio in the FL_{SW} reactor to match the global equivalence ratio of the operating condition, expressed according to the Bilger formulation based on elemental mass fractions [37]. Since these equations do not uniquely determine all the internal exchanges, eight calibration parameters are introduced to define the relevant mass-flow ratios between the connected regions.

The reactor volumes are estimated through a geometric closure procedure based on the combustor dimensions and on the swirl number at the mixing tube exit. The swirl number is approximated using the Beér–Chigier formulation [38] and is then used to define the deflection angle of the swirling jet. The Z_{AX} volume is approximated as a cylindrical region downstream of the mixing tube, while the ORZ is reconstructed as a solid of revolution whose extent is governed by the resulting jet deflection. The FL_{SW} volume is determined from the interaction between the swirling jet and the chamber wall, distinguishing whether the jet remains confined, reaches the wall at the reference axial distance, or impinges upstream of it. The FL_{AX} and OUT volumes are prescribed as fractions of the total combustion chamber volume through two additional calibration parameters, which are added to those used to define the mass-flow distribution. The IRZ volume is then obtained as the residual volume after subtracting all other zones from the chamber volume [35].

Therefore, the internal flow allocation and volume partitioning are controlled by a ten-dimensional calibration vector

$$\vec{\Theta} = [\theta_1, \theta_2, \theta_3, \theta_4, \theta_5, \theta_6, \theta_7, \theta_8, \theta_9, \theta_{10}] . \quad (34)$$

The meaning and definition of the parameters are summarized in Table 3. The first two parameters define the relative stoichiometry and air-flow partitioning of the two inlet reservoirs. Parameters θ_3 – θ_8 control the mass-flow redistribution among the reactor zones, while θ_9 and θ_{10} define the volume fractions assigned to FL_{AX} and OUT.

Table 3 Definition of the CRN calibration parameters.

Parameter	Physical meaning	Definition
θ_1	Relative stoichiometry of the inlet streams	ϕ_1 / ϕ_2
θ_2	Air mass-flow split between inlet reservoirs	$\dot{m}_{\text{air},1} / \dot{m}_{\text{air},2}$
θ_3	Fraction of Res2 directed to ORZ	$\dot{m}_{2,\text{ORZ}} / \dot{m}_2$
θ_4	$FL_{SW} \rightarrow \text{ORZ}$ split	$\dot{m}_B / (\dot{m}_{2,FL_{SW}} + \dot{m}_C + \dot{m}_E + \dot{m}_A)$
θ_5	IRZ $\rightarrow Z_{AX}$ fraction	$\dot{m}_G / (\dot{m}_F + \dot{m}_H)$
θ_6	IRZ $\rightarrow FL_{SW}$ fraction	$\dot{m}_E / [(\dot{m}_F + \dot{m}_H) (1 - \theta_5)]$
θ_7	IRZ $\rightarrow FL_{AX}$ fraction	$\dot{m}_I / [(\dot{m}_F + \dot{m}_H) (1 - \theta_5) (1 - \theta_6)]$
θ_8	$FL_{AX} \rightarrow \text{IRZ}$ recirculation	$\dot{m}_H / (\dot{m}_D + \dot{m}_I)$
θ_9	FL_{AX} volume fraction	$V_{FL_{AX}} / V_{CC}$
θ_{10}	OUT volume fraction	V_{OUT} / V_{CC}

The CRN calibration is performed by tuning $\vec{\Theta}$ to match the experimental NO_x measurements reported by Reichel et al. [21, 22]. The optimization is carried out using the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [39], selected after a comparison among several optimization algorithms [40]. Four experimental points at the reference operating condition are used for calibration, corresponding to $\phi = 0.40, 0.60, 0.87$, and 0.96 . The objective function is a weighted root-mean-square relative error between measured and predicted NO_x emissions:

$$J(\vec{\Theta}) = \left[\sum_{i=1}^{n_{\text{cal}}} w_i \left(\frac{\text{NO}_{x,\text{sim},i}(\vec{\Theta}) - \text{NO}_{x,\text{exp},i}}{\text{NO}_{x,\text{exp},i}} \right)^2 \right]^{1/2}, \quad w_i = \frac{\text{NO}_{x,\text{exp},i}^{0.85}}{\sum_{j=1}^{n_{\text{cal}}} \text{NO}_{x,\text{exp},j}^{0.85}} \quad (35)$$

where $n_{\text{cal}} = 4$ is the number of experimental points used for calibration. The weighting gives more relevance to operating points with larger absolute NO_x levels, while avoiding an excessive contribution from very lean cases where small absolute deviations can lead to large relative errors [35].

The CRN simulations were performed using both the Naik chemical mechanism and the more comprehensive CRECK mechanism, whose main features are summarized in Table 2. For the CRECK mechanism, the CRN results presented in the following sections were obtained through a subsequent calibration performed on the same calibration points adopted for the Naik mechanism, by perturbing the calibration parameters in the neighborhood of the Naik optimum according to the procedure reported in [35].

The ability of the CRN to reproduce the measured nitrogen oxide emissions is shown in Fig. 7, where the results obtained with the Naik chemical mechanism are reported as an example. The reference operating condition, reported in Fig. 7a, was used for calibration. In this subfigure, the four filled squares indicate the experimental data included in the optimization, while the remaining points are used for validation. The same calibrated parameter vector is then applied to the other operating conditions shown in Figs. 7b–7d, without any further recalibration. These results show that the model is able to predict NO_x emissions and reproduce the correct order of magnitude, with a computational cost several orders of magnitude lower than CFD-based simulations [35]. The model results are discussed in detail in Section V.




Fig. 7 Comparison between CRN predictions and experimental NO_x emissions as a function of the equivalence ratio for four inlet operating conditions. The reference condition in panel (a) is used for calibration. The same calibrated parameter vector is applied to panels (b)–(d) without recalibration.

An uncertainty quantification analysis was also carried out to assess the robustness of the calibrated CRN predictions [40]. The analysis considered separate blocks of uncertain inputs associated with operating conditions, calibration parameters, and geometrical or closure assumptions. The operating-parameter block was found to have the strongest influence on the predicted NO_x . By contrast, the contributions associated with the geometrical and calibration parameters remain limited in the lean and intermediate regimes and become relevant only at high equivalence ratios.

V. NO_x Emissions Predictions

In this section, the NO_x emissions predicted by the modeling hierarchy described in Sec. III and Sec. IV are presented and discussed. Five prediction sets are considered: the 2D axisymmetric RANS model with the Naik mechanism (RANS-2D, Naik), the 2D axisymmetric RANS model with the CRECK mechanism (RANS-2D, CRECK), the 3D RANS model with EARSM turbulence closure and the Naik mechanism (RANS-3D EARSM, Naik), and the Chemical Reactor Network with both the Naik and the CRECK mechanisms (CRN, Naik and CRN, CRECK). Owing to the substantial computational cost associated with the 3D RANS EARSM simulations, this approach was applied only with

Table 4 NO_x emissions at operating points with constant $u_0 = 70$ m/s. Values are reported in ppmv, dry, corrected to 15% O₂.

T_{in} [K]	$\dot{m}_{air}$ [kg/h]	ϕ	Exp.	RANS-2D Naik	RANS-2D CRECK	RANS-3D EARS Naik	CRN Naik	CRN CRECK
313	255	0.30	1.83	0.38	0.42	-	0.18	0.07
313	255	0.40	2.32	1.70	0.63	0.86	1.16	0.48
313	255	0.50	3.61	3.75	1.71	2.27	3.24	1.47
313	255	0.60	6.22	7.62	5.90	4.79	6.91	3.70
313	255	0.70	–	11.40	7.71	9.53	13.88	9.27
453	180	0.30	3.56	0.87	0.33	0.39	0.50	0.21
453	180	0.40	3.69	2.60	0.66	1.48	2.07	0.92
453	180	0.50	5.94	5.29	2.42	3.42	5.27	2.71
453	180	0.60	9.39	10.26	6.55	7.60	11.53	7.57
453	180	0.80	–	44.06	39.46	50.52	54.76	50.43
453	180	1.00	–	84.47	89.26	143.87	131.85	130.31
623	130	0.30	3.93	1.87	0.73	1.03	1.37	0.61
623	130	0.40	5.28	4.90	2.30	2.77	4.01	2.00
623	130	0.50	10.12	9.79	5.63	6.33	10.16	6.51
623	130	0.60	26.13	21.70	16.21	16.80	26.08	21.47
623	130	0.80	182.33	127.61	127.35	146.32	140.68	141.04
623	130	0.87	212.07	195.04	199.12	233.18	218.25	220.04
623	130	0.96	305.54	258.15	266.33	333.04	299.87	301.11
623	130	1.00	–	268.77	282.15	338.93	328.41	329.90
693	116	0.30	3.90	2.54	1.02	1.43	1.94	0.90
693	116	0.40	6.75	6.51	3.23	3.67	5.54	3.00
693	116	0.50	16.64	13.36	8.45	8.66	14.48	10.29
693	116	0.60	–	31.51	25.76	25.86	39.11	34.67
693	116	0.80	–	201.47	204.77	232.25	217.41	222.02
693	116	1.00	–	400.18	416.14	472.49	464.08	466.26

the Naik mechanism.

The complete numerical database is reported in Tables 4 and 5. Table 4 contains the operating points at constant u_0 , corresponding to the four reference inlet air temperatures $T_{in} = 313, 453, 623$ and 693 K and the baseline air mass flow rates of 255, 180, 130 and 116 kg/h, respectively. The parameter u_0 represents the mean velocity in the mixing tube when only air is present (i.e., equivalence ratio = 0) and is calculated as [17]:

$$u_0 = \frac{\dot{m}_{air}}{A_{mixing\ tube} \rho_{air}} \quad (36)$$

Table 5 reports the parametric sweeps in which the inlet air mass flow rate is varied at fixed equivalence ratio, namely $\phi = 0.40$ and $\phi = 0.60$. All emission values are expressed as parts per million by volume on a dry basis, corrected to 15% O₂. The 3D EARS CFD reported the extinction of the flame at the operating point at $T_{in} = 313$ K, $\dot{m}_{air} = 255$ kg/h and $\phi = 0.30$: the corresponding CFD entry is reported as "-".

The eighteen operating points of Table 4 for which experimental NO_x measurements are available span nearly two orders of magnitude, from ~1.8 to ~305 ppmv (dry, 15% O₂). To prevent the high-emission points from dominating the aggregated error metrics, the dataset has been partitioned into a low-NO_x regime ($\phi \in [0.30, 0.50]$, 12 points, NO_x ≤ 17 ppmv) and a high-NO_x regime ($\phi > 0.50$, 6 points, up to ~305 ppmv). Table 6 reports the Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAPE) for each model in the two regimes. The RANS-3D simulation that did not achieve ignition was counted as 0 ppmv for the calculation of MAE and MAPE in Table 6.

A. Validation and Split-Regime Error Trade-off

The two regimes display the expected opposite error structure. In the low-NO_x regime, the absolute error is bounded (MAE ≤ 3.4 ppmv for every model) but the relative error is large (29–68%); in the high-NO_x regime the absolute error grows by roughly an order of magnitude (9–21 ppmv) while the MAPE drops to 10–21%. Across the full envelope,

Table 5 NO_x emissions for the inlet air mass-flow-rate parametric sweeps at fixed equivalence ratio. Values are reported in ppmv, dry, corrected to 15% O₂.

T_{in} [K]	$\dot{m}_{air}$ [kg/h]	ϕ	RANS-2D Naik	RANS-2D CRECK	RANS-3D EARS Naik	CRN Naik	CRN CRECK
313	130	0.40	3.09	1.47	1.01	1.12	0.48
313	184	0.40	2.18	0.91	0.92	1.14	0.48
313	220	0.40	1.88	0.74	0.90	1.16	0.48
313	295	0.40	1.56	0.58	0.84	1.18	0.49
453	80	0.40	6.44	4.27	1.60	2.16	1.05
453	130	0.40	3.69	1.63	1.55	2.09	0.96
453	150	0.40	3.14	1.33	1.52	2.08	0.94
453	205	0.40	2.34	0.95	1.47	2.07	0.91
453	230	0.40	2.17	0.87	1.47	2.07	0.90
623	93	0.40	7.23	4.06	2.83	4.34	2.39
623	110	0.40	5.83	2.96	2.80	4.35	2.25
623	148	0.40	4.31	1.92	2.75	4.15	2.06
693	82	0.40	10.23	6.23	3.75	6.04	3.46
693	100	0.40	7.72	3.80	3.69	5.75	3.19
693	133	0.40	5.67	2.62	3.62	5.41	2.87
313	130	0.60	13.52	11.10	5.28	8.07	4.86
313	184	0.60	10.78	7.74	4.90	7.44	4.22
313	220	0.60	8.88	5.83	4.75	7.15	3.93
313	295	0.60	6.49	3.60	4.53	6.73	3.53
453	80	0.60	25.30	22.55	9.44	15.76	12.03
453	130	0.60	15.35	10.04	8.07	12.99	9.08
453	150	0.60	12.27	8.53	7.76	12.32	8.39
453	205	0.60	9.77	5.57	7.24	11.08	7.28
453	230	0.60	9.02	5.38	7.09	10.68	6.85
623	93	0.60	31.13	26.30	20.39	32.04	27.80
623	110	0.60	26.20	20.55	18.39	28.72	24.31
623	148	0.60	18.48	13.75	15.82	24.60	19.86
693	82	0.60	47.96	42.64	33.88	49.53	45.95
693	100	0.60	37.50	31.53	28.81	43.26	39.16
693	133	0.60	26.86	21.94	23.73	36.03	31.25

CRN Naik returns the lowest MAE (9.40 ppmv) and the lowest MAPE (10.3%) in the high-NO_x regime; however, RANS-2D Naik is the most accurate in the low-NO_x regime in both absolute (MAE = 1.19 ppmv) and relative terms (MAPE = 28.9%). All five formulations reproduce the near-exponential rise of NO_x with ϕ seen at $T_{in} = 623$ K, where the experimental data move from 3.93 ppmv ($\phi = 0.30$) to 305.54 ppmv ($\phi = 0.96$).

B. Mechanism Comparison: Naik vs. CRECK

Within identical architectures, the Naik mechanism systematically produces more emissions than CRECK in the lean regime, as also shown in Ref. [34]. Averaged over the twelve low-NO_x points, the predicted NO_x ratio Naik/CRECK is 2.26 for RANS-2D and 2.08 for the CRN; over the six high-NO_x points, the same ratio narrows to 1.19 (RANS-2D) and 1.27 (CRN). Notably, at $\phi = 0.87$ ($T_{in} = 623$ K) the ranking inverts with CRECK mechanism predicting more emissions: CRECK predicts 199.12 against 195.04 ppmv for Naik in RANS-2D. The spread between mechanisms is essentially independent of architectural fidelity: it does not widen from CRN to RANS-2D. The robustness of this observation is supported by the fact that the CRECK-based CRN was not obtained through a direct substitution of the

Table 6 Split-regime validation metrics computed against the experimental NO_x data of Table 4.

Model	Low- NO_x ($\phi \leq 0.50$)		High- NO_x ($\phi > 0.50$)	
	MAE [ppmv]	MAPE [%]	MAE [ppmv]	MAPE [%]
RANS-2D Naik	1.19	28.9	20.97	17.0
RANS-2D CRECK	3.34	66.0	20.04	20.4
RANS-3D EARSMS Naik	2.94	58.9	16.19	19.4
CRN Naik	1.48	38.5	9.40	10.3
CRN CRECK	3.20	67.6	10.45	17.6

kinetic mechanism in the Naik-calibrated network, but by applying the recalibration procedure described in Sec. IV. The persistence of a comparable Naik/CRECK separation after recalibration further suggests that the Naik/CRECK discrepancy is dominated by the kinetic description itself, although a more systematic investigation across architectures would be needed to fully decouple kinetic and turbulence–chemistry effects.

C. Computational Fidelity

Holding the mechanism fixed (Naik), RANS-3D EARSMS does not deliver a systematic accuracy benefit over the lower-fidelity formulations. In the low- NO_x regime its MAE = 2.94 ppmv and MAPE = 58.9% are between RANS-2D Naik and the CRECK-based predictions; in the high- NO_x regime it improves on RANS-2D in MAE (16.19 vs. 20.97 ppmv) but is outperformed by CRN Naik on both metrics. The three Naik-based predictions remain coherent in trend across the entire ϕ sweep, but cross over near $\phi = 0.80$ – 1.00 : RANS-3D EARSMS over-predicts the near-stoichiometric experimental point (333.04 vs. 305.54 ppmv), CRN Naik provides the closest match (299.87 ppmv), also benefiting from the fact that this operating condition was included in the calibration set, while RANS-2D Naik under-predicts (258.15 ppmv). The error in the low equivalence ratio region could be attributed mainly to the lack of a transient formulation in the CFD and in the modeling of turbulence, resulting in neglecting thermoacoustic instabilities and the formation of hot spots characteristic of lean hydrogen flames. In the high equivalence ratio region, instead, the production of NO_x pollutants is highly correlated to local equivalence ratio fluctuations, as the production of pollutants is non-linearly related to temperature; thus, a better description of mixing and turbulence is beneficial to the 3D CFD simulation.

D. Parametric Sensitivities

Table 5 isolates the inlet-temperature effect from the residence-time effect. At fixed mass flow ($\dot{m}_{air} \approx 130$ kg/h) and $\phi = 0.40$, RANS-2D Naik increases from 3.09 ppmv ($T_{in} = 313$ K) to 5.67 ppmv ($T_{in} = 693$ K), a factor ~ 1.8 ; at $\phi = 0.60$ the same temperature span yields a factor ~ 2.0 (from 13.52 to 26.86 ppmv). The residence-time effect over the surveyed mass-flow ranges is of comparable or larger magnitude: at $T_{in} = 453$ K and $\phi = 0.60$, RANS-2D Naik decreases by a factor ~ 2.8 from 25.30 to 9.02 ppmv as $\dot{m}_{air}$ rises from 80 to 230 kg/h. The two effects therefore compete on a similar quantitative footing in this combustor: at the nominal operating points of Table 4 they reinforce one another, since T_{in} is increased while $\dot{m}_{air}$ is reduced to keep u_0 constant, so that the apparent thermal sensitivity in Table 4 is the coupling of two distinct physical mechanisms.

A further qualitative difference between architectures emerges from the parametric sweeps. RANS-2D shows the strongest dependence on $\dot{m}_{air}$, while RANS-3D EARSMS is markedly less sensitive: at $T_{in} = 453$ K and $\phi = 0.60$, for example, the 3D model drops by only 25% (from 9.44 to 7.09 ppmv) over the same mass-flow span where RANS-2D drops by 64%. The CRN response is intermediate with 32 and 43% for Naik and CRECK, respectively. The inlet-temperature response shows the same architectural ordering. At $\phi = 0.60$ and $\dot{m}_{air} \approx 130$ kg/h, NO_x grows by a factor ~ 4.5 between $T_{in} = 313$ and 693 K in both RANS-3D EARSMS (from 5.28 to 23.73 ppmv) and the CRN (from 8.07 to 36.03 ppmv, Naik), whereas RANS-2D Naik yields only a factor ~ 2 (from 13.52 to 26.86 ppmv); the same ordering is recovered at $\phi = 0.40$ (factors ~ 3.6 and ~ 4.8 for RANS-3D EARSMS and CRN, against ~ 1.8 for RANS-2D). With the kinetic source term held fixed across the three architectures (Naik), the spread in both the $\dot{m}_{air}$ and the T_{in} sensitivities can be traced to the predicted residence-time and temperature fields rather than to the chemistry. For design exploration this implies that RANS-2D should be interpreted as an upper-bound estimator of NO_x at low $\dot{m}_{air}$.

within the operating conditions investigated here, and that RANS–3D EARSMS and the CRN are the more reliable tools for extrapolation across the mass-flow envelope. Overall, the results support the construction of a multi-fidelity surrogate model for rapid NO_x assessment. However, the data also show that the fidelity discrepancy is regime-dependent: the correction from 2D RANS to 3D EARSMS or CRN predictions cannot be assumed constant, but must account explicitly for equivalence ratio, inlet temperature and mass flow rate.

E. Computational Cost

The computational cost associated with the different fidelity levels was evaluated in terms of CPU hours. The values reported in this section refer to the average computational cost over the corresponding simulation dataset. All CFD simulations were performed on Intel(R) Xeon(R) Gold 6338 CPUs at 2.00 GHz. As a representative example, only the results obtained with the Naik chemical mechanism are considered. This choice avoids mixing the effect of model fidelity with the additional computational cost associated with changing the chemical mechanism. The computational cost is therefore compared across the different fidelity levels considered in this work: CRN, 2D CFD, 3D CFD and refined 3D CFD calculations.

The CRN represents the lowest-cost level of the hierarchy, requiring approximately 18 s on a single core, corresponding to $5.0 \cdot 10^{-3}$ CPU h per operating point. The 2D axisymmetric RANS simulations required, on average, 13.85 CPU h per case. The transition from 2D to 3D RANS with EARSMS increased the average cost to 247.67 CPU h per case, i.e. approximately 17.9 times the cost of a 2D simulation. This increase is mainly associated with the larger number of cells, the fully three-dimensional flow description and the more expensive turbulence closure required to represent anisotropic swirling-flow effects.

The refined 3D simulations require a separate interpretation. Rather than being started from a uniform or preliminary initial field, the baseline 3D solution was interpolated onto the refined mesh and used as the initial condition. For this reason, the effective cost of one refined 3D data point was evaluated as the sum of the average baseline 3D cost and the cost of the corresponding refined-mesh iterations. This gives $C_{3D,ref} = C_{3D} + C_{refined\ only} = 247.67 + 286.27 = 533.94$ CPU h.

The refined-only contribution is therefore slightly larger than the full baseline 3D cost; however, the baseline 3D simulations required $1.42 \cdot 10^{-1}$ CPU h/iteration, whereas the refined 3D simulations required $5.71 \cdot 10^{-1}$ CPU h/iteration. Hence, each refined-mesh iteration was approximately four times more expensive than one baseline 3D iteration. Consequently, the values reported here represent the effective cost of the refined-mesh procedure adopted in this work. If the refined 3D simulations had been started directly from an initial field, rather than from an interpolated baseline 3D solution, the total computational cost would have been substantially higher.

Table 7 summarizes the average computational cost of the four fidelity levels. The cost increase across the hierarchy spans more than five orders of magnitude, from $5.0 \cdot 10^{-3}$ CPU h for the CRN to 533.94 CPU h for the refined 3D CFD level. This large separation in cost confirms the relevance of a multi-fidelity strategy: the CRN and 2D simulations enable broad parametric exploration, whereas the 3D and refined 3D simulations provide higher-fidelity information at a much higher computational expense.

Table 7 Average computational cost of the different fidelity levels using the Naik chemical mechanism.

Fidelity level	Mean cost [CPU h]	Mean cost/iteration [CPU h/iter]
CRN, Naik	5.00×10^{-3}	–
RANS–2D, Naik	13.85	9.70×10^{-3}
RANS–3D EARSMS, Naik	247.67	1.42×10^{-1}
RANS–3D EARSMS refined, Naik	533.94	5.71×10^{-1}

VI. U-Net

Within multi-fidelity CFD frameworks, convolutional neural networks are widely used to reconstruct or correct flow fields obtained from lower-fidelity simulations. In previous work [41], a U-Net architecture was employed for a super-resolution task between two mesh resolutions. In the present study, a lightweight modified U-Net is adopted for a different multi-fidelity objective. Here, the fidelity gap does not stem from different mesh resolutions, but from the different physical and numerical descriptions of the flow, namely the relation between a two-dimensional simulation and the corresponding planar field extracted from a three-dimensional computation. Therefore, the purpose of the network

is not only to adapt the spatial resolution of the solution, but also to learn a data-driven correction of the 2D prediction by approximating the systematic discrepancy between the 2D solution and the corresponding 3D planar field.

A. Methodology

The adopted model is inspired by the U-shaped encoder–decoder architecture proposed in [42]. Ref [43] has shown that this architecture has many applications in CFD. This architecture is suitable for field-to-field regression because it combines multiscale feature extraction with the preservation of local spatial information through skip connections. In the present framework, the input to the network is the flow field obtained from the 2D simulation, while the target is the corresponding field extracted on a selected plane of the 3D solution (Section A–A, Fig. 1). The model is therefore trained to approximate the mapping between the 2D field and the associated 3D planar field, thereby providing a corrected solution that accounts, within the limits of the training dataset, for the systematic contribution of three-dimensional effects missing from the 2D model.



Fig. 8 Computational domain and selected region of interest used for the U-Net-based multi-fidelity correction. The red box identifies the portion of the combustor in which the 2D flow fields are mapped to the corresponding planar fields extracted from the 3D simulation.

The U-Net correction is applied to a selected region of interest starting from the end of the mixing tube at $z = 0.127$ m, as shown in Fig. 8. This region corresponds to the portion of the combustor covered by the PIV measurements [22], where the comparison with the numerical results is performed. Although the same methodology could in principle be extended to the full computational field, restricting the learning problem to this box is more consistent with the objectives of the present study, since it focuses the correction on the combustion chamber region used for validation and avoids allocating model capacity to parts of the domain that are not directly involved in the experimental comparison.

B. Network architecture

As explained in Section VI.C, the network is split into two branches, each based on the same lightweight modified U-Net topology. The input tensor has size $H \times W \times C = 48 \times 128 \times 3$, where the three channels correspond either to the velocity components or to the thermochemical variables. The output tensor has the same spatial size and number of channels and represents the corrected 3D-planar fields.

The encoder consists of four resolution levels with channel widths 16, 32, 64, and 128. Each level includes two 3×3 convolutional layers with same padding, each followed by batch normalization and ELU activation. Downsampling is performed using 2×2 max pooling, reducing the spatial resolution from 48×128 to 3×8 at the bottleneck. The bottleneck uses 256 filters in the first convolution and 128 filters in the second convolution.

The decoder mirrors the encoder through bilinear upsampling, skip concatenation with the corresponding encoder feature maps, and two convolutional layers per level. The decoder output channel widths are 64, 32, 16, and 16. A final 1×1 convolution maps the decoder features to the three output channels. Unlike the original U-Net implementation, same-padding convolutions are used; therefore, no cropping operation is required before concatenating encoder and decoder feature maps.

Dropout is applied only in the deeper part of the network, with rates 0.1 in the last two encoder levels, in the bottleneck, and in the first two decoder levels. This provides mild regularization while preserving the high-resolution details reconstructed in the final decoder stages.

To improve the training process and the prediction accuracy, a single-channel version of the present network ($H \times W \times C = 48 \times 128 \times 1$) was considered. In this configuration, convolutional kernels are allowed to specialize on a single target feature, avoiding the possible interference associated with the simultaneous learning of multiple fields. This strategy was tested specifically on the NO field, whose prediction proved to be particularly challenging in this test case.




Fig. 9 Lightweight modified U-Net architecture used for multi-fidelity field correction. The same topology is used independently for the velocity and thermochemical branches.

C. Field decomposition

A key design choice concerns the decomposition of the physical fields used for training. Instead of employing a single six-channel network, the variables are split into two physically related groups and two independent U-Net models are trained. The first network is dedicated to the velocity field and takes as input the axial, radial, and tangential velocity components, producing the corresponding corrected 3D-planar velocity components. The second network is dedicated to the thermochemical variables and processes temperature, NO mole fraction, and NO₂ mole fraction.

This separation is introduced because the two groups of variables are characterized by different physical mechanisms, magnitudes, spatial supports, and degrees of regularity. The velocity components are mainly associated with the global flow topology and recirculation structures, whereas the thermochemical quantities are more sensitive to reaction zones, scalar transport, and localized gradients. A single six-channel network would require a common set of weights and optimization dynamics for variables with substantially different statistical properties. Even after normalization, this may lead to an unbalanced training process in which one group of variables dominates the gradient updates. By splitting the problem into two branches, each network specializes in a more coherent mapping, improving training stability and interpretability.

D. Pre-processing and normalization

Before training, all input and target fields are scaled consistently using normalization coefficients computed only from the training set. The same coefficients are then applied to the validation data in order to avoid information leakage. Velocity components and temperature are normalized directly. The NO and NO₂ mole-fraction fields, instead, are first transformed using logarithmic scaling. This choice is motivated by their trace-species character: variations in these quantities are often small in absolute magnitude but significant in relative terms. The logarithmic transformation therefore improves the representation of low-magnitude differences before normalization.

For each nitrogen oxide mole fraction X , the transformed variable is defined as

$$X_{\log} = \log_{10}(X + \varepsilon), \quad (37)$$

where $\varepsilon = 1 \times 10^{-12}$ is a small positive floor introduced to avoid singular values in regions where the species concentration is zero or numerically negligible. The logarithmic transformation compresses the dynamic range of the scalar fields and

increases the relative importance of low-concentration regions during training.

The transformed fields are subsequently normalized to the interval $[0, 1]$. This scaling is consistent with the sigmoid activation used in the final layer of the network. During post-processing, the inverse normalization and inverse logarithmic transformation are applied to recover the physical mole-fraction values.

E. Hyperparameter setting

Both subnetworks adopt the same topology but are trained independently with separate weights and branch-specific optimization settings. The optimization is carried out using the Adam optimizer [44] with mean squared error as the loss function. The loss is evaluated in the normalized variable space. This corresponds for NO and NO₂ species to the logarithmically transformed and normalized mole fractions, mitigating the dominance of high-concentration peaks, improving the representation of variations over several orders of magnitude.

After excluding the two cases used for the testing phase presented in Sec. VI.F, the dataset, comprising the simulations in Tables 4 and 5, is randomly partitioned into training (80%) and validation (20%) subsets. To ensure reproducibility, the split is performed using a fixed random seed. A batch size of 6 is used for the training phase.

Different initial learning rates are assigned to the two branches. The velocity network is trained with an initial learning rate of 5×10^{-5} , whereas the thermochemical network uses a slightly smaller value of 3×10^{-5} . This more conservative setting is adopted because the thermochemical branch includes trace-species fields, namely NO and NO₂, whose relevant variations occur over very small absolute values and can be more sensitive to optimization noise. Training is allowed to proceed up to a maximum of 10 000 epochs. This value is used as an upper bound rather than as a stopping criterion in itself, since the final retained model is selected through a checkpoint strategy based on the minimum validation loss.

A mild scheduled learning-rate reduction is applied during training, with a lower bound of 10^{-6} . Overall, the adopted hyperparameter setting aims at balancing convergence stability and predictive accuracy while accounting for the different numerical and physical characteristics of the two subsets of output variables.

F. Results

The performance of the proposed multi-fidelity correction strategy is evaluated on two representative operating conditions that are not used in the training or validation subsets, referred to as Case 1 and Case 2. Both cases correspond to the same inlet air mass flow rate and inlet temperature, namely $\dot{m}_{air} = 130$ kg/h and $T_{in} = 623$ K, while differing in equivalence ratio. Case 1 corresponds to $\phi = 0.40$, whereas Case 2 corresponds to $\phi = 0.60$. This pair of cases is therefore useful to assess whether the U-Net correction remains effective when the thermochemical regime is modified while keeping the main operating conditions fixed.

1. Case 1

The correction strategy is initially assessed using a reference operating condition defined by $\dot{m}_{air} = 130$ kg/h, $\phi = 0.40$, and $T_{in} = 623$ K. Figure 10 reports the comparison for the three velocity components. The original 2D simulation already reproduces the main large-scale organization of the flow, but it exhibits non-negligible deviations from the 3D reference, especially in the shear-layer region and in the lower part of the selected domain. These discrepancies are particularly evident for the axial and radial velocity components, where the transition between the high-velocity jet region and the surrounding flow is locally shifted or intensified. The U-Net correction preserves the global structure of the 2D field while moving the solution closer to the 3D planar reference.

This is confirmed by the residual error maps, where both the magnitude and the spatial extension of the error are reduced compared with the original 2D–3D discrepancy.

The corresponding thermochemical fields are shown in Fig. 11. For temperature, the U-Net prediction corrects most of the discrepancy observed in the 2D simulation and provides a field that is much closer to the 3D planar solution, both in terms of spatial distribution and absolute level. The residual error after correction is strongly reduced, indicating that the network is able to learn the systematic thermal discrepancy between the two fidelity levels.

The nitrogen oxide fields are more challenging. NO and NO₂ are trace species, and their relevant variations occur over very small absolute values. Their prediction is therefore more sensitive to local gradients, small spatial shifts, and the adopted normalization strategy. The U-Net captures the main spatial organization of both species and reproduces the dominant structures observed in the 3D reference. Nevertheless, the residual maps show that local errors remain in regions characterized by sharp scalar transitions. This behavior is expected, since even small displacements of localized




Fig. 10 Case 1: comparison between the 2D simulation, the U-Net prediction, and the corresponding 3D planar reference for the three velocity components. For each component, the first three columns report the 2D field, the corrected field predicted by the U-Net, and the 3D reference field, respectively. The last two columns show the absolute discrepancies between the original 2D solution and the 3D reference, and between the U-Net prediction and the 3D reference.




Fig. 11 Case 1: comparison between the 2D simulation, the U-Net prediction, and the corresponding 3D planar reference for the thermochemical variables. The rows show temperature, NO mole fraction, and NO₂ mole fraction, respectively. For each variable, the first three columns report the 2D field, the U-Net prediction, and the 3D reference, while the last two columns show the absolute discrepancies with respect to the 3D reference.

trace-species structures can generate visible discrepancies in the error maps.




Fig. 12 Case 1: radial profiles extracted at $z = 0.140$ m for the velocity components and thermochemical variables. The 2D simulation, U-Net prediction, and 3D planar reference are compared along the radial coordinate.




Fig. 13 Case 1: radial profiles extracted at $z = 0.200$ m for the velocity components and thermochemical variables. The 2D simulation, U-Net prediction, and 3D planar reference are compared along the radial coordinate.

A more quantitative local comparison is provided in Figs. 12 and 13. The velocity profiles confirm the trends observed in the two-dimensional maps. For the axial velocity, the U-Net prediction closely follows the 3D reference in the main peak region and improves the agreement with respect to the original 2D simulation. A similar behavior is observed for the radial velocity, where the network reduces the overprediction of the 2D solution around the main positive peak and better reproduces the subsequent negative region. The tangential velocity is also well corrected, with the predicted profile remaining close to the 3D reference over most of the radial coordinate.

For temperature, the correction is also substantial. The U-Net shifts the high-temperature region toward the 3D reference and reduces the systematic discrepancy observed in the 2D simulation. The agreement is improved both in the transition region and in the high-temperature plateau.

The behavior of the nitrogen oxide species is less uniform. While NO_2 is reasonably improved by the U-Net correction, the NO profile remains more difficult to reconstruct when it is learned jointly with temperature and NO_2 in the same thermochemical branch. In particular, as shown in Fig. 12, for radial coordinates larger than 0.02 m, the predicted NO profile does not fully recover the 3D reference. This suggests that NO is more sensitive to local thermochemical discrepancies and may require a more specialized mapping than the other thermochemical variables.

This observation motivated a refinement of the training strategy, in which the NO field was treated with a dedicated model rather than being learned jointly with temperature and NO_2 . The aim of this separate training is to focus the optimization process specifically on the spatial and magnitude variations of NO, avoiding competition with variables characterized by different scales, smoother distributions, or stronger global correlations. Figure 14 reports the radial NO mole-fraction profiles obtained after the dedicated NO training. At $z = 0.140$ m, the corrected profile reproduces the main radial structure of the 3D reference, including the low-NO region around the centerline and the recovery toward the outer plateau. Compared with the original 2D solution, the U-Net prediction substantially reduces the overestimation in the transition region and shifts the profile toward the 3D planar reference.




Fig. 14 Case 1: radial profiles of NO mole fraction extracted at $z = 0.140$ m and $z = 0.200$ m after dedicated NO training.

At $z = 0.200$ m, the correction remains more limited: the prediction decreases the discrepancy with respect to the 2D baseline, but it still underpredicts the 3D reference over most of the radial coordinate. This indicates that the dedicated NO training improves the qualitative radial trend, although some residual error persists. The benefit of the dedicated NO training is even clearer at $z = 0.134$ m, where the comparison between the joint thermochemical model and the dedicated NO model shows a more pronounced improvement of the radial profile, as reported in Fig. 15.

2. Case 2

This second case is considered to evaluate the behavior of the same correction strategy at a higher equivalence ratio. The operating point is thus defined by $\dot{m}_{air} = 130$ kg/h, $\phi = 0.60$, and $T_{in} = 623$ K. Since the qualitative trends are similar to those observed for Case 1, the discussion focuses on the consistency of the correction and on the additional difficulties associated with the thermochemical fields.

As for Case 1, the original 2D simulation captures the dominant flow organization but shows localized discrepancies with respect to the 3D planar reference. The U-Net prediction reduces these discrepancies while preserving the main structure of the flow. The residual maps confirm that the correction is not limited to a single operating condition, but remains effective when the equivalence ratio is increased.




Fig. 15 Case 1: comparison of NO radial profiles extracted at $z = 0.134$ m before and after dedicated NO training. Panel (a) shows the prediction obtained from the joint thermochemical model, whereas panel (b) shows the result obtained with the dedicated NO model.

The temperature field follows the same behavior observed for Case 1: the U-Net prediction reduces the discrepancy between the 2D solution and the 3D planar reference and better reproduces the thermal structure inside the selected region. The nitrogen oxide fields remain more demanding because they are trace species with localized and low-magnitude variations. Nevertheless, the predicted fields reproduce the main spatial patterns of the 3D reference and reduce part of the original model discrepancy.




Fig. 16 Case 2: radial profiles extracted at $z = 0.140$ m for the velocity components and thermochemical variables. The 2D simulation, U-Net prediction, and 3D planar reference are compared along the radial coordinate.

The radial profiles in Fig. 16 confirm the trends observed in the maps. The velocity components and temperature are generally shifted toward the 3D reference by the U-Net correction. For the nitrogen oxide species, the profiles again reveal a different level of difficulty between NO and NO₂. The NO prediction is less accurate than the other fields when obtained from the joint thermochemical model, confirming the need for a dedicated treatment of this variable.

The effect of the dedicated NO training is particularly evident for Case 2, as shown in Fig. 17. The original 2D




Fig. 17 Case 2: radial profile of NO mole fraction extracted at $z = 0.140$ m after dedicated NO training.

simulation strongly underpredicts the NO mole fraction over most of the radial coordinate and fails to reproduce the correct order of magnitude in the high-NO region. The dedicated U-Net correction recovers the main radial transition and predicts values much closer to the 3D reference. Although the profile remains underestimated at large radial coordinates, the corrected solution captures the dominant radial trend and substantially reduces the discrepancy with respect to the 2D baseline.

Overall, the analysis of Cases 1 and 2 shows that the proposed U-Net correction is effective in reducing the discrepancy between the 2D simulations and the corresponding 3D planar references. The improvement is clear for the velocity components and temperature, while the trace-species fields remain more challenging. The radial comparisons reveal that NO is less accurately reconstructed than NO_2 when the two species are learned within the same thermochemical branch. This motivated the introduction of a dedicated NO model, which improves the reconstruction of the radial NO profiles and confirms the benefit of a more specialized training strategy for this field.

3. PIV comparison

The previous analysis assessed the U-Net correction with respect to the 3D numerical planar reference, which represents the high-fidelity target used during training. A further comparison is carried out against the experimental PIV profiles [22] in order to evaluate the physical consistency of the corrected velocity fields. This comparison is not equivalent to a direct training validation, since the U-Net is not optimized against the PIV measurements, but against the 3D CFD planar fields. Therefore, the PIV data provide an independent reference to assess whether the numerical correction moves the solution toward experimentally observed trends.


Figures 18 and 19 show the radial profiles of axial and radial velocity at two axial locations, $z = 0.140$ m and $z = 0.200$ m, respectively. For each location, the original 2D simulation, the U-Net prediction, and the PIV measurements are compared for the selected operating cases. The left panel of each subfigure reports the axial velocity, while the right panel reports the radial velocity.

At $z = 0.140$ m, both axial and radial velocity profiles show that the U-Net prediction does not systematically approach the PIV measurements over the entire radial coordinate. In the peak region, the corrected profile remains close to the original 2D solution, while the PIV profile indicates a different radial distribution and, in some cases, a different peak level. This suggests that the 3D CFD target used for training does not fully reproduce the experimental axial and radial velocity profile in this section, and consequently the U-Net correction cannot be expected to recover the PIV behavior if this information is not present in the training target.


The profiles at $z = 0.200$ m confirm the same general trend. The axial and radial velocity remains difficult to improve with respect to the PIV data: the U-Net prediction tends to follow the numerical 2D/3D trend and does not fully recover the experimental profile shape, particularly around the main peak and in the outer radial region. Conversely, the radial velocity comparison again shows a closer predicted profile to the PIV measurements than the original 2D simulation in the regions where the 2D solution exhibits the largest deviation.

Overall, the PIV comparison highlights an important limitation and, at the same time, a useful outcome of the proposed approach. The U-Net correction is effective in reducing the discrepancy between 2D and 3D CFD fields, but its agreement with the experiment is bounded by the accuracy of the 3D numerical reference used as training target. As a result, the method should be interpreted as a multi-fidelity numerical correction rather than as a direct experimental calibration. However, both radial and axial velocity profiles reveal remaining discrepancies.

An additional point emerging from the comparison is that the disagreement with the PIV data is not specific to




(a) Case 1




(b) Case 2

Fig. 18 Comparison between PIV [22], the original 2D simulation, and the U-Net prediction at $z = 0.140$ m.



(a) Case 1



(b) Case 2

Fig. 19 Comparison between PIV [22], the original 2D simulation, and the U-Net prediction at $z = 0.200$ m.

the U-Net-corrected fields, as even higher-fidelity LES results of Ref. [45] show non-negligible deviations from the experimental points. As shown in Fig. 20, these curves are clipped to a radius of 50 millimeters due to the difficulty of extracting the original LES data near the wall. Therefore, the mismatch between U-net prediction and PIV should not be interpreted solely as a limitation of the data-driven correction, but more generally as evidence of a residual CFD-to-experiment gap that is already present at the LES level. This interpretation is consistent with the discussion reported in the reference LES study [45], where residual discrepancies with the experiment are still acknowledged, in particular in the extension of the inner recirculation zone and in the peak velocity of the swirling jet, even though the overall agreement is considered improved with respect to lower-fidelity approaches. For this reason, the comparison against PIV should be regarded as physically informative but not as an absolute ground truth in a strict validation sense.

4. Mean error metrics over the test set

To quantify the accuracy of the neural-network correction, the predicted profiles are compared against the corresponding 3D reference solutions and contrasted with the original 2D axisymmetric baseline. For a generic scalar field f , the relative L_2 and L_∞ errors are defined as in Eqs. (38)-(39).

$$\text{Rel}L_2 = 100 \frac{\sqrt{\sum_{i=1}^{N_x} \sum_{j=1}^{N_y} (f_{ij} - f_{ij}^{\text{ref}})^2}}{\sqrt{\sum_{i=1}^{N_x} \sum_{j=1}^{N_y} (f_{ij}^{\text{ref}})^2}} \quad (38)$$

$$\text{Rel}L_\infty = 100 \frac{\max_{i,j} |f_{ij} - f_{ij}^{\text{ref}}|}{\max_{i,j} |f_{ij}^{\text{ref}}|} \quad (39)$$

Table 8 reports the mean relative L_2 and L_∞ errors over the full test set. The results show that the proposed correction consistently reduces the discrepancy with respect to the 3D reference for all investigated quantities.

Table 8 Mean relative errors over the full test set. Both the 2D baseline and the neural-network prediction are compared against the 3D reference solution.

Quantity	Rel. L_2 2D [%]	Rel. L_2 Pred. [%]	Rel. L_∞ 2D [%]	Rel. L_∞ Pred. [%]
u_z [m/s]	29.42	5.20	39.78	16.60
u_r [m/s]	38.20	5.52	57.15	15.04
u_θ [m/s]	29.60	3.83	16.96	7.28
T [K]	9.73	1.35	44.47	12.78
X_{NO} [-]	116.94	16.30	313.78	49.87
X_{NO_2} [-]	85.21	12.26	75.84	40.52

5. Computational cost

The U-Net training was performed on an NVIDIA RTX A5000 GPU using TensorFlow [46] with cuDNN [47] acceleration. The computational cost reported in Table 9 is estimated from representative TensorFlow training logs and should be interpreted as an indicative wall-clock cost rather than as a hardware-independent benchmark. When two U-Net branches are trained concurrently on the same GPU, that is, the first used approach, the effective time increases by 66% per epoch because of shared GPU resources.

Table 9 Indicative U-Net training cost on an NVIDIA RTX A5000 GPU.

Training configuration	Time/epoch	Time for 3000 epochs	Time for 10000 epochs
Single U-Net branch	$O(1 \text{ s})$	$O(1 \text{ h})$	$O(3 \text{ h})$




Fig. 20 Comparison between experimental data and numerical predictions on Case 2 for the velocity fields along the radial coordinate. LES results are extracted from Ref. [45].

The maximum number of epochs is therefore computationally affordable compared with the cost of the CFD simulations used to generate the training targets. Moreover, the retained model is selected according to the minimum validation loss through a checkpoint strategy; consequently, the maximum epoch count should be interpreted as an upper training limit rather than as the effective optimal training duration. The dominant computational cost of the overall workflow remains the generation of the high-fidelity 3D reference fields, while the trained U-Net provides a comparatively inexpensive correction of the 2D solution toward the corresponding 3D planar representation.

VII. Conclusions

This work presented a multi-fidelity computational framework for the prediction of NO_x emissions in the lean premixed, swirl-stabilized AHEAD hydrogen combustor. Three modeling levels were assessed across inlet temperatures from 313 to 693 K, air mass flow rates from 80 to 295 kg/h, and equivalence ratios from 0.30 to 1.00: a 2D axisymmetric RANS with the $k-\omega$ SST closure, a 3D RANS with the EARSM closure, and a Chemical Reactor Network calibrated via CMA-ES. Each level was tested with both the Naik and CRECK kinetic mechanisms, with the exception of the 3D EARSM case. Because these models present a clear gap between computational efficiency and predictive depth, the framework was also extended to data-driven techniques to bridge them. A lightweight modified U-Net was trained to correct 2D fields toward the corresponding 3D planar reference, effectively infusing the computationally efficient 2D simulations with higher-fidelity 3D flow features.

The split-regime validation showed that no single architecture dominates across the operating envelope: RANS–2D Naik is the most accurate in the low- NO_x regime ($\text{MAE} = 1.19$ ppmv), while CRN Naik provides the lowest error in the high- NO_x regime ($\text{MAE} = 9.40$ ppmv, $\text{MAPE} = 10.3\%$). The Naik mechanism systematically yields higher emissions than CRECK in lean conditions, with the discrepancy remaining largely independent of architectural fidelity. The parametric sweeps further revealed that inlet temperature and mass flow rate exert competing effects of comparable magnitude on NO_x emissions, and that the sensitivity to both parameters is architecture-dependent: RANS–2D tends to overestimate NO_x at low mass flow rates, while RANS–3D EARSM and the CRN show a more consistent response across the operating envelope. Furthermore, the CRN proves to be more than five orders of magnitude faster than the highest-fidelity CFD simulations.

The U-Net correction substantially reduced the 2D–3D discrepancy for the velocity components and temperature, whereas trace species, particularly NO, required a dedicated training branch to recover the correct radial trend. The agreement with the experiment remains constrained by the fidelity of the 3D reference solution and by the limited experimental information available for the present benchmark. The U-Net correction should therefore be interpreted as a numerical consistency improvement rather than as a direct experimental calibration.

Future work will focus on two directions. First, the construction of a fully integrated multi-fidelity framework leveraging the established hierarchy (CRN, 2D RANS, 3D EARSM) to balance computational speed and predictive accuracy across design exploration tasks. Second, a refinement of the U-Net loss formulation: while the mean squared error proved effective for velocity field reconstruction, specialized loss functions should be explored to better capture the nonlinearities governing chemical species predictions.

References

- [1] United Nations Framework Convention on Climate Change, “Paris Agreement,” UNFCCC, 2015. URL https://unfccc.int/sites/default/files/english_paris_agreement.pdf.
- [2] Verstraete, D., “The Potential of Liquid Hydrogen for long range aircraft propulsion,” Ph.D. thesis, Cranfield University, 2009.
- [3] Lefebvre, A. H., and Ballal, D. R., *Gas Turbine Combustion: Alternative Fuels and Emissions*, 3rd ed., CRC press, 2010. <https://doi.org/10.1201/9781420086058>.
- [4] Turns, S. R., *An Introduction to Combustion: Concepts and Applications*, 3rd ed., McGraw-Hill Education, 2011.
- [5] Lee, D. S., Fahey, D. W., Skowron, A., Allen, M. R., Burkhardt, U., Chen, Q., Doherty, S. J., Freeman, S., Forster, P. M., Fuglestad, J., et al., “The contribution of global aviation to anthropogenic climate forcing for 2000 to 2018,” *Atmospheric Environment*, Vol. 244, 2021, p. 117834. <https://doi.org/10.1016/j.atmosenv.2020.117834>.
- [6] Schefer, R., Wicksall, D., and Agrawal, A., “Hydrogen enrichment for improved lean flame stability,” *International journal of hydrogen energy*, Vol. 28, No. 10, 2003, pp. 1131–1140. [https://doi.org/10.1016/S0360-3199\(02\)00199-4](https://doi.org/10.1016/S0360-3199(02)00199-4).
- [7] Peters, N., *Turbulent combustion*, Cambridge university press, 2000. <https://doi.org/10.1017/CBO9780511612701>.
- [8] Wallin, S., and Johansson, A. V., “An explicit algebraic Reynolds stress model for incompressible and compressible turbulent flows,” *Journal of Fluid Mechanics*, Vol. 403, 2000, pp. 89–132. <https://doi.org/10.1017/S0022112099007004>.
- [9] Capurso, T., Laera, D., Riber, E., and Cuenot, B., “NO_x pathways in lean partially premixed swirling H₂-air turbulent flame,” *Combustion and Flame*, Vol. 248, 2023, p. 112581. <https://doi.org/10.1016/j.combustflame.2022.112581>.
- [10] Mira, D., Lehmkuhl, O., Both, A., Stathopoulos, P., Tanneberger, T., Reichel, T. G., Paschereit, C. O., Vázquez, M., and Houzeaux, G., “Numerical Characterization of a Premixed Hydrogen Flame Under Conditions Close to Flashback,” *Flow, Turbulence and Combustion*, Vol. 104, No. 2, 2020, pp. 479–507. <https://doi.org/10.1007/s10494-019-00106-z>.
- [11] Folcarelli, L., Spagnolo, A., Dicech, F., Ferrero, A., Masseni, F., and Pastrone, D., “Multi-Fidelity Modeling of a Lean Premixed Swirl-Stabilized Hydrogen Burner With Axial Air Injection,” *SCITECH 2025 Forum*, AIAA, 2025. <https://doi.org/10.2514/6.2025-0941>.
- [12] Pashchenko, D., “Comparative analysis of hydrogen/air combustion CFD-modeling for 3D and 2D computational domain of micro-cylindrical combustor,” *International Journal of Hydrogen Energy*, Vol. 42, No. 49, 2017, pp. 29545–29556. <https://doi.org/10.1016/j.ijhydene.2017.10.070>.
- [13] Benim, A. C., and Pfeiffelmann, B., “Comparison of Combustion Models for Lifted Hydrogen Flames within RANS Framework,” *Energies*, Vol. 13, No. 1, 2020. <https://doi.org/10.3390/en13010152>.
- [14] Abdel-Raheem, M., Ibrahim, S., Malalasekera, W., and Masri, A., “Large eddy simulation of hydrogen–air premixed flames in a small scale combustion chamber,” *International Journal of Hydrogen Energy*, Vol. 40, No. 7, 2015, pp. 3098–3109. <https://doi.org/10.1016/j.ijhydene.2014.12.042>.
- [15] Jiang, J., Jiang, X., and Zhu, M., “A computational study of preferential diffusion and scalar transport in nonpremixed hydrogen-air flames,” *International Journal of Hydrogen Energy*, Vol. 40, No. 45, 2015, pp. 15709–15722. <https://doi.org/10.1016/j.ijhydene.2015.08.112>.
- [16] Berger, L., Attili, A., Gauding, M., and Pitsch, H., “LES combustion model for premixed turbulent hydrogen flames with thermodiffusive instabilities: a priori and a posteriori analysis,” *Journal of Fluid Mechanics*, Vol. 1003, 2025, p. A33. <https://doi.org/10.1017/jfm.2024.1201>.
- [17] Reichel, T. G., Goeckeler, K., and Paschereit, O., “Investigation of Lean Premixed Swirl-Stabilized Hydrogen Burner With Axial Air Injection Using OH-PLIF Imaging,” *Journal of Engineering for Gas Turbines and Power*, Vol. 137, No. 11, 2015, p. 111513. <https://doi.org/10.1115/1.4031181>.
- [18] Reichel, T. G., and Paschereit, C. O., “Interaction mechanisms of fuel momentum with flashback limits in lean-premixed combustion of hydrogen,” *International Journal of Hydrogen Energy*, Vol. 42, No. 7, 2017, pp. 4518–4529. <https://doi.org/10.1016/j.ijhydene.2016.11.018>.
- [19] Tanneberger, T., Reichel, T. G., Krüger, O., Terhaar, S., and Paschereit, C. O., “Numerical Investigation of the Flow Field and Mixing in a Swirl-Stabilized Burner With a Non-Swirling Axial Jet,” *Turbo Expo: Power for Land, Sea, and Air, Volume 4B: Combustion, Fuels and Emissions*, 2015, p. V04BT04A026. <https://doi.org/10.1115/GT2015-43382>.

- [20] Reichel, T. G., Terhaar, S., and Paschereit, O., “Increasing Flashback Resistance in Lean Premixed Swirl-Stabilized Hydrogen Combustion by Axial Air Injection,” *Journal of Engineering for Gas Turbines and Power*, Vol. 137, No. 7, 2015, p. 071503. <https://doi.org/10.1115/1.4029119>.
- [21] Reichel, T. G., Terhaar, S., and Paschereit, C. O., “Flashback resistance and fuel–air mixing in lean premixed hydrogen combustion,” *Journal of Propulsion and Power*, Vol. 34, No. 3, 2018, pp. 690–701. <https://doi.org/10.2514/1.B36646>.
- [22] Paschereit, C. O., and Reichel, T. G., “AHEAD Lean-Premixed Swirl-Stabilized Hydrogen Burner TU Berlin Dataset,” <https://tnfworkshop.org/ahead-lean-premixed-swirl-stabilized-hydrogen-burnertu-berlin/>, 2023. Accessed: May 21, 2025.
- [23] Ansys[®] Fluent, Release 2024 R1 - Theory Guide, ANSYS, Inc., 2024.
- [24] Menter, F. R., “Two-equation eddy-viscosity turbulence models for engineering applications,” *AIAA Journal*, Vol. 32, No. 8, 1994, pp. 1598–1605. <https://doi.org/10.2514/3.12149>.
- [25] McBride, B. J., Gordon, S., and Reno, M. A., “Coefficients for Calculating Thermodynamic and Transport Properties of Individual Species,” NASA Technical Memorandum 4513, National Aeronautics and Space Administration, 1993. URL <https://ntrs.nasa.gov/citations/19940013151>.
- [26] Apsley, D. D., and Leschziner, M. A., “Advanced Turbulence Modelling of Separated Flow in a Diffuser,” *Flow, Turbulence and Combustion*, Vol. 63, No. 1, 2000, pp. 81–112. <https://doi.org/10.1023/A:1009930107544>.
- [27] Wallin, S., and Johansson, A. V., “Modelling streamline curvature effects in explicit algebraic Reynolds stress turbulence models,” *International Journal of Heat and Fluid Flow*, Vol. 23, No. 5, 2002, pp. 721–730. [https://doi.org/10.1016/S0142-727X\(02\)00168-6](https://doi.org/10.1016/S0142-727X(02)00168-6).
- [28] Hellsten, A., “New Advanced k- ω Turbulence Model for High-Lift Aerodynamics,” *AIAA Journal*, Vol. 43, No. 9, 2005, pp. 1857–1869. <https://doi.org/10.2514/1.13754>.
- [29] Magnussen, B. F., “The eddy dissipation concept: A bridge between science and technology,” *ECCOMAS thematic conference on computational combustion*, Lisbon, Portugal, 2005, p. 24. URL https://folk.ntnu.no/ivar/edc/BFM_ECOMAS2005_Lisboa.pdf.
- [30] Naik, C. V., “Detailed Mechanism for Hydrogen Combustion with NO_x,” Ansys, Inc., 2009. Originally developed by Reaction Design, now part of Ansys.
- [31] Song, Y., Marrodán, L., Vin, N., Herbinet, O., Assaf, E., Fittschen, C., Stagni, A., Faravelli, T., Alzueta, M. U., and Battin-Leclerc, F., “The sensitizing effects of NO₂ and NO on methane low temperature oxidation in a jet stirred reactor,” *Proceedings of the Combustion Institute*, Vol. 37, No. 1, 2019, pp. 667–675. <https://doi.org/10.1016/j.proci.2018.06.115>.
- [32] Stagni, A., Cavallotti, C., Arunthanayothin, S., Song, Y., Herbinet, O., Battin-Leclerc, F., and Faravelli, T., “An experimental, theoretical and kinetic-modeling study of the gas-phase oxidation of ammonia,” *Reaction Chemistry & Engineering*, Vol. 5, No. 4, 2020, pp. 696–711. <https://doi.org/10.1039/C9RE00429G>.
- [33] Stagni, A., Arunthanayothin, S., Dehue, M., Herbinet, O., Battin-Leclerc, F., Bréquigny, P., Mounaïm-Rousselle, C., and Faravelli, T., “Low- and intermediate-temperature ammonia/hydrogen oxidation in a flow reactor: Experiments and a wide-range kinetic modeling,” *Chemical Engineering Journal*, Vol. 471, 2023, p. 144577. <https://doi.org/10.1016/j.cej.2023.144577>.
- [34] Folcarelli, L., Concetti, R., Cozzi, F., Ferrero, A., Masseni, F., Pastrone, D., Hasslberger, J., and Klein, M., “A Chemical Reactor Network methodology for estimating NO_x emissions under non-premixed lean hydrogen combustion with water injection,” *International Journal of Hydrogen Energy*, Vol. 168, 2025, p. 150561. <https://doi.org/10.1016/j.ijhydene.2025.150561>.
- [35] Madonia, V., Folcarelli, L., Ferrero, A., Pastrone, D. G., and Masseni, F., “NO_x Emission Prediction of a Lean Premixed Hydrogen Combustor via Reactor Network Analysis,” *SCITECH 2026 Forum*, AIAA, 2026. <https://doi.org/10.2514/6.2026-1186>.
- [36] Goodwin, D. G., Speth, R. L., Moffat, H. K., and Weber, B. W., “Cantera: An Object-Oriented Software Toolkit for Chemical Kinetics, Thermodynamics, and Transport Processes,” <https://www.cantera.org>, 2024. Version 3.1.0. Available online at: <https://www.cantera.org>.
- [37] Bilger, R., “The structure of turbulent nonpremixed flames,” *Symposium (International) on Combustion*, Vol. 22, No. 1, 1989, pp. 475–488. [https://doi.org/10.1016/S0082-0784\(89\)80054-2](https://doi.org/10.1016/S0082-0784(89)80054-2).
- [38] Beér, J. M., and Chigier, N. A., *Combustion Aerodynamics*, Applied Science Publishers Limited, 1972.

- [39] Hansen, N., and Ostermeier, A., “Completely Derandomized Self-Adaptation in Evolution Strategies,” *Evolutionary Computation*, Vol. 9, No. 2, 2001, pp. 159–195. <https://doi.org/10.1162/106365601750190398>.
- [40] Madonia, V., Folcarelli, L., Ferrero, A., Masseni, F., and Pastrone, D. G., “Uncertainty Quantification and Computational Cost Reduction of a Chemical Reactor Network for a Hydrogen Combustor,” *Aviation 2026 Forum*, AIAA, 2026.
- [41] Marcì, S. G. I., Muscarà, L., Folcarelli, L., Ferrero, A., Masseni, F., and Pastrone, D. G., “Use of Convolutional Neural Networks for Multifidelity CFD Problems,” *Proceedings of the CEAS AIDAA Conference 2025*, Materials Research Forum LLC, 2026. To appear.
- [42] Ronneberger, O., Fischer, P., and Brox, T., “U-Net: Convolutional Networks for Biomedical Image Segmentation,” *International Conference on Medical image computing and computer-assisted intervention*, Springer, 2015, pp. 234–241. https://doi.org/10.1007/978-3-319-24574-4_28.
- [43] Thuerey, N., Weißenow, K., Prantl, L., and Hu, X., “Deep Learning Methods for Reynolds-Averaged Navier–Stokes Simulations of Airfoil Flows,” *AIAA Journal*, Vol. 58, No. 1, 2020, pp. 25–36. <https://doi.org/10.2514/1.J058291>.
- [44] Kingma, D. P., and Ba, J. L., “Adam: A Method for Stochastic Optimization,” *International Conference on Learning Representations (ICLR)*, 2015. <https://doi.org/10.48550/arXiv.1412.6980>.
- [45] Amerighi, M., Andreini, A., Reichel, T., Tanneberger, T., and Paschereit, C. O., “LES investigation of a swirl stabilized technically premixed hydrogen flame with FGM and TFM models,” *Applied Thermal Engineering*, Vol. 247, 2024, p. 122944. <https://doi.org/10.1016/j.applthermaleng.2024.122944>.
- [46] Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X., “TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems,” , 2015. URL <https://www.tensorflow.org/>, software available from tensorflow.org.
- [47] Chetlur, S., Woolley, C., Vandermersch, P., Cohen, J., Tran, J., Catanzaro, B., and Shelhamer, E., “cuDNN: Efficient Primitives for Deep Learning,” , 2014. <https://doi.org/10.48550/arXiv.1410.0759>.