

RESEARCH ARTICLE 

Hydrogen reaction rate modeling based on convolutional neural network for large eddy simulation

Quentin Malé¹ , Corentin J. Lapeyre² and Nicolas Noiray¹

¹CAPS Laboratory, Department of Mechanical and Process Engineering, ETH Zürich, Zürich, Switzerland

²NVIDIA Corporation, Santa Clara, CA, USA

Corresponding author: Quentin Malé; Email: qumale@ethz.ch

Received: 29 August 2024; **Revised:** 28 November 2024; **Accepted:** 24 December 2024


Keywords: artificial intelligence; computational fluid dynamics; data-driven reacting flow modeling; hydrogen combustion; machine learning

Abstract

This article establishes a data-driven modeling framework for lean hydrogen (H_2)-air reaction rates for the Large Eddy Simulation (LES) of turbulent reactive flows. This is particularly challenging since H_2 molecules diffuse much faster than heat, leading to large variations in burning rates, thermodiffusive instabilities at the subfilter scale, and complex turbulence-chemistry interactions. Our data-driven approach leverages a Convolutional Neural Network (CNN), trained to approximate filtered burning rates from emulated LES data. First, five different lean premixed turbulent H_2 -air flame Direct Numerical Simulations (DNSs) are computed each with a unique global equivalence ratio. Second, DNS snapshots are filtered and downsampled to emulate LES data. Third, a CNN is trained to approximate the filtered burning rates as a function of LES scalar quantities: progress variable, local equivalence ratio, and flame thickening due to filtering. Finally, the performances of the CNN model are assessed on test solutions never seen during training. The model retrieves burning rates with very high accuracy. It is also tested on two filter and downsampling parameters and two global equivalence ratios between those used during training. For these interpolation cases, the model approximates burning rates with low error even though the cases were not included in the training dataset. This a priori study shows that the proposed data-driven machine learning framework is able to address the challenge of modeling lean premixed H_2 -air burning rates. It paves the way for a new modeling paradigm for the simulation of carbon-free hydrogen combustion systems.

Impact Statement

Large eddy simulation (LES) plays a key role in the design of new combustion chambers for robust and safe combustion of carbon-free fuels such as H_2 . It is a very efficient method that consists of filtering and modeling small-scale flow physics that would otherwise be very computationally expensive to resolve. However, the high molecular diffusion of H_2 challenges current subfilter-scale modeling approaches for lean premixed mixtures. By establishing a framework for data-driven machine learning modeling of LES burning rates, this work can enable high-fidelity simulation of hydrogen combustion systems. These combustion systems are essential for achieving a low-carbon economy.

 This research article was awarded Open Materials badge for transparent practices. See the Data Availability Statement for details

© The Author(s), 2025. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

1. Introduction

The urgent need to drastically reduce our greenhouse gas emissions requires us to develop alternative solutions to fossil fuel combustion. Sustainably produced hydrogen (H₂) is a promising energy carrier for the decarbonization of combustion systems where there is no viable technology to be a substitute. This is the case, for example, with gas turbines for power generation or aviation, as well as some industrial burners. However, compared with conventional fuels, H₂-air combustion involves higher flame temperatures, much higher flame speeds, and a very wide flammability range. These characteristics pose considerable challenges to the design of new H₂ combustion devices, ensuring stable, robust, and safe operation. In this context, high-fidelity simulation of turbulent combustion plays a key role. It is essential both for a fundamental understanding of flame behavior and as an aid to the design of industrial combustion devices. It is also crucial for modeling safety scenarios, with hydrogen expected to be widely used to decarbonize the industry.

The community nowadays uses Large Eddy Simulation (LES) methods, which involve simulating only the largest flow structures, to compute real complex configurations at an affordable computing cost. The smaller scales are modeled, according to subfilter-scale models (Poinso and Veynante, 2011; Fiorina et al., 2015). In LES, a filtered quantity φ is defined as

$$\bar{\varphi}(\mathbf{x}, t) = \int_{-\infty}^{\infty} \varphi(\mathbf{x}', t) F(\mathbf{x} - \mathbf{x}') d\mathbf{x}', \quad (1.1)$$

where F is the LES filter. Balance equations are obtained by applying the filter F to the multispecies Navier–Stokes equations. The filtered transport equation of a reactive scalar φ reads

$$\partial_t(\bar{\rho}\tilde{\varphi}) + \nabla \cdot (\bar{\rho}\tilde{\mathbf{u}}\tilde{\varphi}) - \nabla \cdot (\bar{\rho}\tilde{D}_\varphi\nabla\tilde{\varphi}) = \nabla \cdot \Gamma_\varphi + \bar{\omega}_\varphi, \quad (1.2)$$

where $\tilde{\cdot}$ denotes density-weighted Favre filtering (i.e. $\tilde{\cdot} = \bar{\rho}\cdot/\bar{\rho}$), with density ρ , time t , velocity \mathbf{u} . D_φ , Γ_φ and ω_φ are the molecular diffusivity, subfilter-scale flux, and source term of φ . The filtering operation produces unclosed terms located on the RHS of Eq. (1.2) that must be modeled. The subfilter-scale flux comprises the transport of φ by unresolved fluctuations of molecular diffusive flux and momentum

$$\Gamma_\varphi = \Gamma_{\varphi,D} - \Gamma_{\varphi,C} \text{ with } \Gamma_{\varphi,D} = \overline{\rho D_\varphi \nabla \varphi} - \bar{\rho} \tilde{D}_\varphi \nabla \tilde{\varphi} \text{ and } \Gamma_{\varphi,C} = \overline{\rho \mathbf{u} \varphi} - \bar{\rho} \tilde{\mathbf{u}} \tilde{\varphi}. \quad (1.3)$$

It is usually modeled as neglecting $\Gamma_{\varphi,D}$ and involving gradient transport models for $\Gamma_{\varphi,C}$ with a turbulent viscosity hypothesis (Boger et al., 1998). The countergradient phenomenon can be taken into account via more complex modeling (Tullis and Cant, 2002a, 2002b), although a large part will be described at the resolved scale in LES (Boger et al., 1998). Alternatively, Γ_φ can be tabulated within the filtered tabulated chemistry framework (Fiorina et al., 2010) or modeled via Machine Learning (ML) algorithms (Seltz et al., 2019). Modeling of the filtered chemical source term $\bar{\omega}_\varphi$ is extremely challenging because combustion generally occurs within reaction fronts that are smaller than the LES computational mesh. Furthermore, there is a strong interaction between turbulence and chemistry: chemical heat release influences turbulence, while turbulence, in turn, modifies flame structure, which can either enhance or inhibit chemical reactions.

To overcome these difficulties, several approaches have been proposed in the context of premixed combustion, which is the focus of this article. A common technique is to artificially thicken the reaction front, enabling it to be resolved on the given computational mesh. This is the basis of the Thickened Flame (TF) model, which achieves this by increasing all diffusivities by a factor F while decreasing all reaction rates by the same factor to maintain the correct flame propagation speeds (Butler and O'Rourke, 1977). The LES modeling then focuses on an efficiency function E taking into account the effects of thickening on flame propagation, such as the loss of flame wrinkling (Colin et al., 2000; Charlette et al., 2002; Poinso and Veynante, 2011). Another approach consists of filtering the flame front with a LES filter larger than the mesh size (Boger et al., 1998; Duwig, 2007). Filtering is usually done on the equation of a reaction progress variable c representing the evolution of combustion within the flame front. The balance equation for \tilde{c} may be recast as a level set G-equation (Kerstein et al., 1988).

$$\partial_t(\bar{\rho}\tilde{c}) + \nabla \cdot (\bar{\rho}\tilde{u}\tilde{c}) + \nabla \cdot (\bar{\rho}\tilde{u}\tilde{c}) = \nabla \cdot (\bar{\rho}D_c\nabla \cdot c) + \bar{\omega}_c = \overline{\rho s_d|\nabla c|}, \quad (1.4)$$

which gives rise to a flame front displacement term at speed s_d , $\overline{\rho s_d|\nabla c|}$, which requires modeling. Another way to close the equation is to tabulate all the closure terms from filtered laminar premixed flames (Fiorina et al., 2010), assuming the turbulent flame can be represented as an ensemble of locally One-Dimensional (1D) laminar flamelets (Peters, 1988). In all cases, a model to retrieve the subfilter-scale effects on the thickened/filtered flame front is needed.

Modeling subfilter-scale phenomena is particularly challenging for hydrogen combustion. Indeed, H_2 is a very peculiar fuel due to its small Lewis number. Its high molecular diffusion, much faster than heat, gives rise to unique phenomena, especially when burned lean (that is with excess air) and premixed. First, lean H_2 flames are prone to thermodiffusive instabilities. The disparity between heat and H_2 mass fluxes (called differential diffusion) amplify small flame front perturbations, inducing significant flame wrinkling and strong variations of the reaction rates along the flame front (Poinso and Veynante, 2011; Berger et al., 2022a, 2022b). Second, burning rates increase significantly when lean H_2 flame is subjected to a strain (Petrov and Ghoniem, 1994; Vagelopoulos et al., 1994), up to a certain limit. The resulting increase in flame speed can be substantial. For example, the experimental work in Vagelopoulos et al. (1994) reports a three-fold enhancement in flame speed compared to an unstrained flame.

These phenomena invalidate usual Flame Surface Density (FSD) approaches (Pope, 1988; Trouvé and Poinso, 1994; Marble and Broadwell, 1997; Boger et al., 1998) based on the assumption that subfilter-scale effects are limited to an increase in flame surface density through wrinkling due to turbulence alone. Aniello et al. (2022) have recently attempted to address this issue by multiplying the filtered chemical source terms by a corrective factor, determined from Direct Numerical Simulation (DNS) (Berger et al., 2022a, 2022b), decoupling the enhancement of combustion rates due to thermodiffusion instabilities and usual turbulence wrinkling. However, this is no longer possible when the wrinkling scales induced by turbulent eddies become smaller than the characteristic lengths at which the instabilities occur. Moreover, Berger et al. (2022c) have reported complex and synergistic interactions between thermodiffusion instabilities and turbulence, casting doubt on the possibility of decoupling these phenomena. Gaucherand et al. (2024b) attempted to correct unresolved straining effects by multiplying the filtered chemical source terms by a strain-to-unstrained flame speed ratio determined using 1D flames. The strain rate was assumed to be constant over the entire domain at a value of 2000 s^{-1} , not taking into account the very wide strain variation along the turbulent flame. Additional efforts must be made to account for local strain effects, which is extremely difficult in LES because the information on the strain that a turbulent flame actually undergoes at the subfilter-scale level is unknown. Subfilter Probability Density Function (PDF) modeling for H_2 turbulent combustion is also being developed (Berger et al., 2022d; Ferrante et al., 2024; Pitsch, 2024; Yao and Blanquart, 2024), where the subfilter-scale distribution of thermochemical variables is usually modeled using β -PDF. The turbulence-chemistry interactions are described through the subfilter PDF while the flamelet assumption (Peters, 1988) links local thermochemical variables to those of the laminar flame elements. Thus, a set of 1D laminar flames is needed to be computed, assuming a specific flame archetype a priori. Defining the subfilter PDF is a particular challenge. For example, Berger et al. (2022d) and Yao and Blanquart (2024) use a Dirac function for the mixture fraction distribution, while it is reported in Pitsch (2024) that it leads to modeling errors. Indeed, preferential diffusion in H_2 -air mixtures can lead to fluctuation of the mixture fraction at the subfilter-scale level. Furthermore, the effect of unstable modes that occur partially at the subfilter scale must be included in the subfilter PDF to avoid underestimating the flame consumption speed (Lapenna et al., 2024; Pitsch, 2024).

The phenomena involved in H_2 turbulent combustion are extremely complex, not well understood, and still subject to discussion, we propose to change the paradigm by developing a data-driven method based on ML instead of trying to model turbulent H_2 burning rates by physical description. This work is motivated by the positive results of such an approach for the approximation of the subfilter flame surface density both a priori and a posteriori for methane-air combustion (Lapeyre et al., 2018, 2019; Ho et al., 2024). The strategy is adapted to hydrogen combustion by training a Neural Network (NN) to estimate the H_2 burning rate directly. In this way, no assumptions are made about the turbulent combustion regime, and

the model can theoretically be applied to any combustion regime encountered during training. Firstly, a database of fully resolved turbulent H₂-air premixed flames are generated using DNS, taking into account preferential and differential diffusion effects. The database includes five different equivalence ratios to assess the ability of the model to account for air/fuel ratio variation. The database is then filtered and downsampled to emulate data from LES. Three different filter sizes are used to assess the ability of the model to account for subfilter effects at different scales. From the emulated LES fields, a Convolutional Neural Network (CNN) is trained to estimate the LES burning rates based on known LES scalar quantities such as the filtered mixture fraction and progress variable. Once trained, the model is found to approximate the burning rates on solutions never seen during training with high accuracy. In addition to this, the model is able to approximate burning rates on filter sizes and equivalence ratios other than those used for training.

First, the problem statement is described, with the definition of the input/output variables of the model (Section 2.1). Then the NN set up to approximate the burning rate function is detailed (Section 2.2). Next, the data generation and training strategy is explained (Section 2.3). Finally, training results are presented and discussed (Section 3).

2. Background and methodology

2.1. Problem statement

The simulation of turbulent flows at high Reynolds numbers cannot generally be performed with sufficient resolution to represent all the scales of the flow. LES aims to simulate only a part of the turbulent spectrum to achieve a reasonable computational cost. Scales are cut by spatial filtering, and unresolved scales are modeled using so-called subfilter-scale models. In the LES framework, the Favre-filtered transport equation for the mass fraction of the k^{th} species Y_k follows Eq. (1.2). As discussed in the introduction, modeling the filtered chemical source terms in LES of reacting flows is extremely challenging because chemistry largely occurs at the subfilter-scale level in most cases. This is further complicated in lean premixed H₂-air mixtures because of thermodiffusive effects, causing additional flame wrinkling and pronounced fluctuations in reaction rates along the flame front. A precise description of the preferential and differential diffusion, turbulence wrinkling, and their interactions (Berger et al., 2022c) at the subfilter-scale level is required.

In this work, we focus on the development of a data-driven model for the filtered chemical source term of H₂ in lean premixed turbulent combustion. ML is used to discover structural relations between resolved scalar quantities and subfilter-scale turbulence-chemistry interactions. We develop this method for the H₂ filtered “burning rate” $\bar{\omega}$, defined as

$$\bar{\omega} = -\bar{\omega}_{\text{H}_2}, \quad (2.1)$$

with the aim of establishing an ML-based modeling framework and a proof of concept for H₂ turbulent combustion. The final implementation will depend on the given LES modeling framework. For example, the ML model could replace the tabulation method proposed in Lapenna et al. (2021, 2024) for the chemical source term of the progress variable. This is what has been done in Seltz et al. (2019) with direct mapping of the closure terms for a methane-air flame. The ML framework could also be based on a database with a one-step chemistry (Millán-Merino and Boivin, 2024; Schiavone et al., 2024) to close the LES multi-species transport equations.

The basic ingredients to describe turbulent flames remain the quantities introduced for laminar flame analysis: the progress variable c and the mixture fraction ζ . The progress variable equals zero in fresh gas, and one in burnt gas, and varies monotonically during combustion. Mixture fraction ζ or the equivalence ratio ϕ is used to inform the ML-based model of the composition of the mixture, which impacts burning rates. The mixture fraction is based on Bilger’s definition (Bilger et al., 1990).

$$\zeta = \frac{\beta - \beta^o}{\beta^f - \beta^o} \text{ with } \beta = \frac{Z_{\text{H}}}{2W_{\text{H}}} - \frac{Z_{\text{O}}}{W_{\text{O}}}, \quad (2.2)$$

where Z_m and W_m are the elemental mass fraction and atomic mass of the element m . The superscripts o and f denote oxidizer and fuel stream. The filtered progress variable \tilde{c} is defined as

$$\tilde{c} = \frac{Y_{\text{H}_2}^u(\tilde{\zeta}) - \tilde{Y}_{\text{H}_2}}{Y_{\text{H}_2}^u(\tilde{\zeta}) - Y_{\text{H}_2}^b(\tilde{\zeta})} = \frac{\tilde{\zeta} - \tilde{Y}_{\text{H}_2}}{\tilde{\zeta} - \max\left(0, \frac{\tilde{\zeta} - \zeta_s}{1 - \zeta_s}\right)}, \quad (2.3)$$

where superscripts u and b denote values in unburnt and burnt gas and ζ_s is the mixture fraction value at stoichiometry. We have arbitrarily chosen to base the progress variable on the H_2 mass fraction. It is shown in [Supplementary Material](#) that it is also possible to base the progress variable on the H_2O mass fraction without any impact on the training and prediction accuracy of the model. In this framework, the challenge is now to model $\bar{\omega}$ as a function of \tilde{c} and $\tilde{\zeta}$ fields.

2.2. Machine learning framework

The aim of the ML-based model is to approximate the function

$$\bar{\omega} = f(\tilde{c}, \tilde{\zeta} \text{ or } \tilde{\phi}, \dots). \quad (2.4)$$

The evolution of the reaction rate depends on the thermochemical state described by the progress variable and mixture fraction, but also on the flow. Indeed, the level of turbulence, flame curvature, and strain, have a major influence on the reaction rate (Poinsot and Veynante, 2011). In this work, we propose to include the topological information by means of spatial convolutions using a CNN. The CNN is found to perform better when the equivalence ratio $\tilde{\phi}$ is used instead of the mixture fraction $\tilde{\zeta}$. The filtered equivalence ratio can be approximated as

$$\tilde{\phi} = \frac{\tilde{\zeta}(1 - \zeta_s)}{\zeta_s(1 - \tilde{\zeta})}. \quad (2.5)$$

The CNN takes as input the field of \tilde{c} and $\tilde{\phi}$ over a domain Ω , performs successive convolutions and operations, and outputs an approximation of $\bar{\omega}$ over the same domain Ω . The function f_{CNN} therefore reads

$$\bar{\omega}^{\text{NN}} = f_{\text{CNN}}(\tilde{c}, \tilde{\phi}) \approx \bar{\omega}. \quad (2.6)$$

CNNs were designed to learn spatial hierarchies of features, from low- to high-level patterns. They are therefore well fitted for recognizing complex flame topology and learning associated burning rates. Furthermore, the convolutions enable to train models on large inputs via parameter sharing. The application of CNNs to the modeling of flame features has already been successfully performed in Lapeyre et al. (2018, 2019), Seltz et al. (2019) and Nikolaou et al. (2019) for fuel-air mixture without significant preferential diffusion effects, unlike the H_2 -air mixtures considered in this article.

The CNN architecture used is inspired by the U-net architecture originally developed for image segmentation (Ronneberger et al., 2015) (Figure 1). However, the output activation function is replaced by a linear activation function for the regression task considered in this work. A similar architecture for the approximation of flame surface density has already been used successfully (Lapeyre et al., 2018, 2019). Using 64 channels in the first multi-channel feature map instead of 32 in Lapeyre et al. (2018, 2019) increased the accuracy of the approximation of $\bar{\omega}$ (Eq. (2.6)) in the present work. In total, the network comprises 5,424,193 trainable parameters.

2.3. Data generation and training strategy

The data generation strategy consists of filtering and downsampling data from DNS to emulate LES data with known filtered quantities. The configuration chosen to generate training data is a slot burner (Figure 2) at constant pressure $P = 1$ atm and fresh gas temperature $T^u = 300$ K, as in Lapeyre et al. (2018, 2019). The physical domain consists of a central inlet where a premixed H_2 -air mixture flows at a




Figure 1. U-net-type architecture is used in the present work. Each black box corresponds to a multi-channel feature map. The number of channels is denoted on top of the box. The input sample has three channels: \tilde{c} , $\tilde{\phi}$ and δ_L^0/δ_L^1 . The ratio δ_L^0/δ_L^1 is the inverse of the laminar flame thickening due to filtering (Section 2.3.2). The original size of the cubic sample is N^3 . It is then reduced to $(N/2)^3$ and $(N/4)^3$ during the contracting path before going back to the original size during the expansive path. Gray boxes represent copied feature maps. The arrows denote the different operations.




Figure 2. Diagram of the slot burner configuration used to generate the DNS database. The flame is depicted by an iso-surface at a progress variable $c = 0.5$ colored by $\dot{\omega} = -\dot{\omega}_{H_2}$. The dimensions of the domain are annotated on the right. The length L_x is adapted to the length of the turbulent flame brush, which is a function of the global equivalence ratio ϕ_g .

bulk velocity $U_b = 24$ m/s with velocity fluctuation $u' = 2.4$ m/s, surrounded by two laminar coflows where burnt gas flows at a bulk velocity $U_c = 3.6$ m/s. The injection of turbulence at the central inlet corresponds to homogeneous and isotropic turbulence using a Passot-Pouquet turbulence spectrum (Passot and Pouquet, 1987) with an integral length scale $l_t = 2$ mm. The domain is rectangular with periodic boundary conditions in the z -direction. Adiabatic walls are specified in the y -direction. Both inlets and outlets are specified in the x -direction. Exact dimensions are annotated on Figure 2. The configuration is very similar to that used in Berger et al. (2022c), Coulon et al. (2023), and Gaucherand et al. (2024a) to study H_2 -air and H_2/NH_3 -air premixed turbulent flames. This configuration is computed for five different equivalence ratios (Table 1) to assess the ability of the CNN to approximate the reaction

Table 1. Global equivalence ratio ϕ_g for the five different DNS cases

Case	$\phi_g = 0.35$	$\phi_g = 0.4$	$\phi_g = 0.5$	$\phi_g = 0.6$	$\phi_g = 0.7$
S_L^0 [m/s]	$9.453 \cdot 10^{-2}$	$1.925 \cdot 10^{-1}$	$4.972 \cdot 10^{-1}$	$8.845 \cdot 10^{-1}$	1.302
$\delta_{th,L}^0$ [m]	$1.146 \cdot 10^{-3}$	$7.176 \cdot 10^{-4}$	$4.809 \cdot 10^{-4}$	$4.339 \cdot 10^{-4}$	$4.171 \cdot 10^{-4}$
T^{ad} [K]	1311	1428	1646	1844	2020

The laminar flame characteristic values are also indicated. S_L^0 is the laminar flame speed, $\delta_{th,L}^0$ the laminar thermal flame thickness, T^{ad} the adiabatic flame temperature.

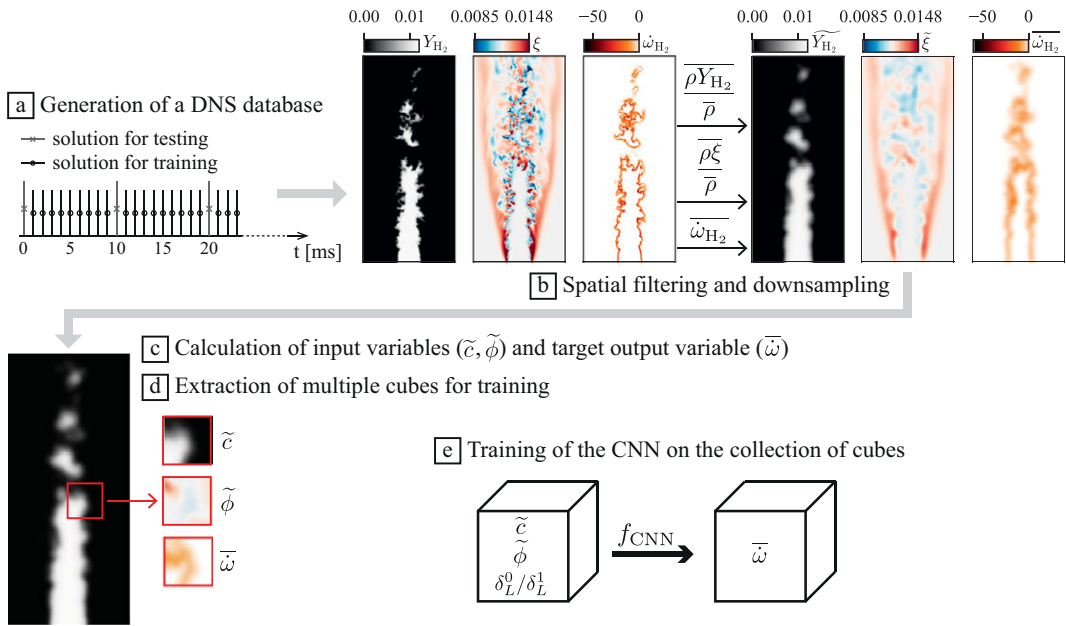


Figure 3. Diagram of the strategy used to generate data and train the CNN. δ_L^0/δ_L^1 is the inverse of the laminar flame thickening due to filtering (Section 2.3.2). $\tilde{\phi}$ is calculated from $\tilde{\xi}$ using Eq. (2.5).

rate over a range of equivalence ratios. All other parameters are kept constant. The Reynolds number of the central inlet is about 10,000 for all cases.

The data generation and training strategy is based on processing multiple snapshots from the five DNS (Figure 3):

- The five DNS cases (Table 1) are computed for 59 ms from a statistically steady state. Three-Dimensional (3D) instantaneous solutions are saved every ms. Solutions with an index multiple of 10 are set aside for testing after training. All other solutions are used for training and validation data. This generates 6 solutions for testing and 54 solutions for training and validation data for each DNS case.
- A spatial filtering is performed on each solution to obtain $\overline{\rho}$, $\overline{\rho Y_{H_2}}$, $\overline{\rho \xi}$ and $\overline{\omega_{H_2}}$. The filtered variables are downsampled to a coarser grid, representing an LES grid. The filtering and downsampling operation is described in Section 2.3.2.
- The filtered progress variable \tilde{c} is calculated using Eq. (2.3) with $\tilde{Y}_{H_2} = \overline{\rho Y_{H_2}}/\overline{\rho}$ and $\tilde{\xi} = \overline{\rho \xi}/\overline{\rho}$. The filtered equivalence ratio $\tilde{\phi}$ is calculated using Eq. (2.5) with $\tilde{\xi} = \overline{\rho \xi}/\overline{\rho}$. The filtered burning rate $\overline{\omega}$ is calculated using Eq. (2.1) with $\overline{\omega_{H_2}}$.

- d. It is important that the CNN does not learn the specific configuration of the slot burner, but only learns generic flame elements. Depending on the effective receptive field of the CNN, training using the emulated LES solutions directly could encode in the network information on the distance to the inlets, outlet, or the wall for example. This would drastically reduce CNN’s ability to generalize and must be avoided. To address this, training is performed on randomly selected cubes in the domain as in Lapeyre et al. (2018, 2019). This also enables the CNN to be invariant to translation (Biscione and Bowers, 2021).
- e. The CNN is trained to approximate $\overline{\omega}$ (Eq. (2.6)) using data from the cubic samples. 10% of samples are randomly selected and set aside for validation, the rest are used for training. The complete description of the training is given in Section 2.3.3.

2.3.1. Direct numerical simulation

DNS of the slot burner cases (Table 1) are performed using the AVBP (Schonfeld and Rudgyard, 1999; Gicquel et al., 2011) massively parallel code solving the compressible multi-species Navier–Stokes equations. A third-order accurate Taylor–Galerkin scheme is adopted for the discretization of the convective terms (Colin and Rudgyard, 2000). Navier–Stokes Characteristic Boundary Conditions (NSCBCs) (Poinsot and Lelef, 1992) are imposed at the inlets (relaxation factor of 1000 s⁻¹) and at the outlet (relaxation factor of 200 s⁻¹). Dynamic viscosity μ follows a power law function of temperature T .

$$\mu = \mu_0 \left(\frac{T}{T_0} \right)^\gamma, \tag{2.7}$$

with $\mu_0 = 8.062 \times 10^{-5}$ kg/m/s, $T_0 = 2.645 \times 10^3$ K and $\gamma = 6.481 \times 10^{-1}$. Thermal diffusivity is computed from the viscosity using a constant Prandtl number: Pr = 0.66. Species diffusivities are computed using a constant Schmidt number specific for each species (Table 2). This approach takes into account non-unity Lewis numbers and preferential diffusion between the different species. It is verified that the errors made by the simplified transport description are negligible by comparing the results with simulations using a mixture-averaged transport model, see [Supplementary Material](#). Soret and Dufour transport processes are ignored in the simulations of the present work. Including these transport, processes could enhance thermodiffusive instabilities and increase source terms and flame surface curvature (Schlup and Blanquart, 2018). Although this would require another training on the new database, it does not affect the conclusions of this work. Indeed, it has been shown that neglecting these effects still leads to simulations that capture much of the peculiar behavior of lean-premixed H₂-air flames including thermodiffusive instabilities and enhanced burning rates (Aspden et al., 2015; Aspden, 2017; Schlup and Blanquart, 2018; Pitsch, 2024). Hydrogen chemical kinetics relies on the San Diego mechanism (Saxena and Williams, 2006), already successfully used for H₂-air premixed combustion in Coulon et al. (2023). This mechanism comprises 9 species and 21 reactions.

The mesh is a homogeneous cartesian grid with constant element size $\Delta_x = 0.1$ mm. This gives a resolution index $RI = \delta_{th,L}^0 / \Delta_x + 1$ that varies between 12.5 and 5.2 for ϕ_g between 0.35 and 0.7. $\delta_{th,L}^0$ is the laminar thermal flame thickness. A mesh sensitivity study was carried out where we verified that turbulent combustion statistics are not affected by mesh refinements. This study is included as [Supplementary Material](#). The length of the domain in the x -direction L_x is adapted to the length of the turbulent flame brush. It varies from 76 mm for $\phi_g = 0.35$ to 36 mm for $\phi_g = 0.7$ (Figure 2). The number of grid points

Table 2. Schmidt Sc_k and Lewis Le_k numbers for the species used for the DNS of H₂-air flames

Species k	H ₂	H	O ₂	OH	O	H ₂ O	HO ₂	H ₂ O ₂	N ₂
Sc_k	0.23	0.14	0.80	0.53	0.52	0.58	0.80	0.81	0.91
Le_k	0.34	0.21	1.20	0.80	0.79	0.88	1.21	1.22	1.37

accordingly varies from 39 to 19 million. Calculations were carried out on the new Alps Research Infrastructure at the Swiss National Supercomputing Centre (CSCS) using AMD EPYC 7742 CPUs. A total of 1.6 million core hours were used to generate the database. Depending on the grid size, computations are performed using between 5,120 and 2,560 computing cores with good efficiency thanks to the massively parallel computing ability of the AVBP code.

2.3.2. Filtering and downsampling

A spatial filtering is performed to obtain $\bar{\rho}$, $\overline{\rho Y_{H_2}}$, $\overline{\rho \zeta}$ and $\overline{\omega_{H_2}}$, in order to calculate \tilde{c} , $\tilde{\phi}$ and $\overline{\omega}$ (Eqs (2.3), (2.5), and ((2.1)). The filtered quantity of a variable φ from the DNS field is defined by Eq. (1.1). The filtered 3D fields are obtained by successively filtering in each direction x , y and z using a 1D Gaussian filter. The 1D filter function is written in discrete form as

$$F(u) = \begin{cases} e^{-\frac{1}{2}(\frac{u}{\sigma})^2} & \text{if } u \in [-z.. +z] \text{ with } z = 4\sigma, \\ 0 & \text{otherwise} \end{cases} \tag{2.8}$$

and then normalized by its sum $\sum_{u=-z}^z F(u)$. The filtered fields are then downsampled to a coarser grid to emulate LES fields. Three different standard deviations σ are used in this work to emulate three different LES resolutions (Table 3). For each σ , a downsampling factor DSF is associated. It is chosen so as to have a LES grid with a resolution index $RI = \delta_{c,L}^1 / \Delta_x + 1$ of about 6. $\delta_{c,L}^0$ and $\delta_{c,L}^1$ are the laminar flame thickness based on the progress variable of a non-filtered and filtered 1D flame

$$\delta_{c,L}^0 = (\max|\nabla c|)^{-1} \text{ and } \delta_{c,L}^1 = (\max|\nabla \tilde{c}|)^{-1}. \tag{2.9}$$

2.3.3. Training method

The generated cubic samples (Figure 3) of size $16 \times 16 \times 16$ points are used to train the CNN to approximate $\overline{\omega}$ (Eq. (2.6)). A set of 10% of the samples is randomly selected to form a validation dataset and monitor overfitting. The remaining samples are used for training. At each epoch, each cubic sample in the training dataset has a 25% probability of undergoing between one and three 90° rotations along any axis, or of being flipped in any direction. It is done to add variability to the orientation of the flame front in the dataset. The model must have no preferential orientation and learn an isotropic function.

Training is accomplished by backpropagation. The trainable parameters of the CNN are updated iteratively on batches of data from the training dataset by stochastic gradient descent with the AdamW optimizer (Kingma and Ba, 2017; Loshchilov and Hutter, 2019). The parameters are adjusted to minimize the Mean Squared Error (MSE) loss function overall output pixels of the prediction compared to the target. The batch size per GPU is 200 and the training is terminated after 1,500 epochs. The base learning rate

Table 3. Flame characteristics for the three sets of LES parameters (σ , DSF) used in this work

Case	$\phi_g = 0.35$	$\phi_g = 0.4$	$\phi_g = 0.5$	$\phi_g = 0.6$	$\phi_g = 0.7$	
$\sigma = 4$	$\delta_{c,L}^1$ [m]	$1.41 \cdot 10^{-3}$	$1.06 \cdot 10^{-3}$	$9.28 \cdot 10^{-4}$	$9.01 \cdot 10^{-4}$	$9.09 \cdot 10^{-4}$
DSF = 2	$\delta_{c,L}^0 / \delta_{c,L}^1$	0.81	0.59	0.41	0.36	0.33
$\Delta_x = 0.2$ mm	RI	8.03	6.30	5.64	5.51	5.55
$\sigma = 8$	$\delta_{c,L}^1$ [m]	$2.06 \cdot 10^{-3}$	$1.83 \cdot 10^{-3}$	$1.75 \cdot 10^{-3}$	$1.75 \cdot 10^{-3}$	$1.75 \cdot 10^{-3}$
DSF = 4	$\delta_{c,L}^0 / \delta_{c,L}^1$	0.55	0.34	0.21	0.18	0.17
$\Delta_x = 0.4$ mm	RI	6.16	5.58	5.38	5.36	5.37
$\sigma = 16$	$\delta_{c,L}^1$ [m]	$3.66 \cdot 10^{-3}$	$3.53 \cdot 10^{-3}$	$3.47 \cdot 10^{-3}$	$3.45 \cdot 10^{-3}$	$3.45 \cdot 10^{-3}$
DSF = 8	$\delta_{c,L}^0 / \delta_{c,L}^1$	0.31	0.18	0.11	0.09	0.09
$\Delta_x = 0.8$ mm	RI	5.58	5.41	5.34	5.32	5.31

Δ_x is the grid cell size, equal to DSF times the original cell size of 0.1 mm. $\delta_{c,L}^0$ and $\delta_{c,L}^1$ are the laminar flame thickness based on the progress variable of a non-filtered and filtered 1D flame (Eq. (2.9)). $RI = \delta_{c,L}^1 / \Delta_x + 1$ is the resolution index of the filtered flame on the LES grid.

equals 10^{-3} . The evolution of the loss function on the training and validation datasets during training is given in Section 3.

The PyTorch (Ansel et al., 2024) Python library is used to implement and train the CNN. The Distributed Data-Parallel (DDP) feature in PyTorch is employed to parallelize the training process across four NVIDIA H100 GPUs. Each GPU performs its own forward and backward pass on a subset of the training dataset. The gradients computed on each GPU are then synchronized and averaged to update the model parameters. Training performances are reported in Section 3.

3. Results and discussion

3.1. Training and testing

The training and validation database comprises $40 \times 54 \times 5 \times 3 = 32,400$ cubic samples (40 random cubes per solution, 54 solutions per case, 5 different equivalence ratios, 3 different sets of LES parameters (filtering and downsampling)). 3,240 (10%) of these cubes are randomly set aside for validation. Training the CNN model for 1,500 epochs on this database takes 47 minutes. 1,500 epochs are sufficient because the loss function evaluated on the validation dataset no longer decreases significantly (Figure 4). The set of model parameters giving rise to the lowest loss is selected as the best model (black circle in Figure 4). This best model is used to assess the ability of the CNN to approximate burning rates on the test solutions (full domain, never seen during training). The Normalized Mean Absolute Error (NMAE) is used to estimate the accuracy of the CNN model (Figure 4). It is defined as follows:

$$\text{NMAE} = \frac{1}{\text{mean}(\bar{\omega}^*)} \frac{1}{N} \sum_{i=1}^N \left| \bar{\omega}_i^* - \bar{\omega}_i^{\text{NN}} \right|, \tag{3.1}$$

where $\bar{\omega}^*$ is the filtered burning rate from the DNS dataset and $\bar{\omega}^{\text{NN}}$ is the filtered burning rate modeled by the CNN. The mean absolute error is only about 5% of the average burning rate for all cases. The very high accuracy of the CNN model is confirmed by the distribution of CNN-modeled burning rate $\bar{\omega}^{\text{NN}}$ versus ground-truth filtered burning rate $\bar{\omega}^*$ (Figure 5). The log-normalized Two-Dimensional (2D) histograms show a large majority of points on the $x = y$ axis (that is perfect agreement). Figure 6 shows a visualization of the result using a planar cut through a test solution for the case $\phi_g = 0.4$ and for the three sets of LES parameters (σ , DSF). One can see the ability of the CNN to retrieve the extremely complex turbulent flame

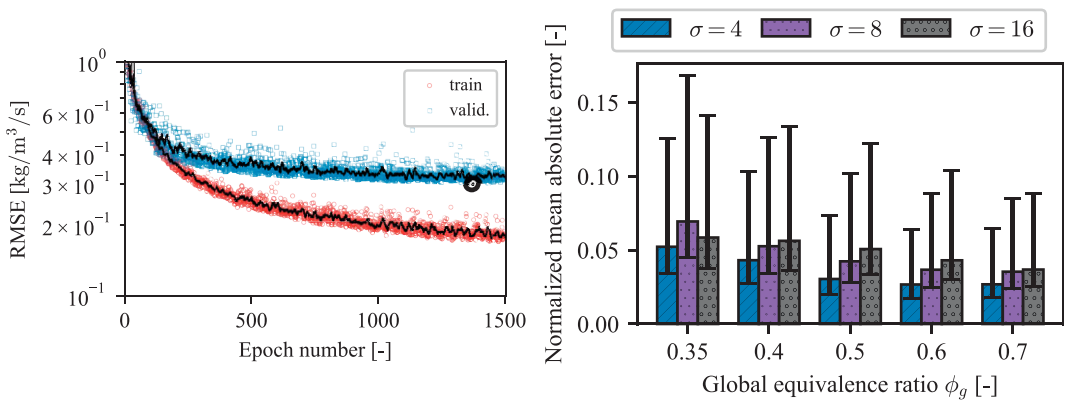


Figure 4. Left: Evolution of the RMSE during training, evaluated over the training dataset (red circles) and the validation dataset (blue squares). Black solid lines are moving averages of the RMSE. The black circle shows the lowest RMSE over the validation dataset. The model parameters at this specific epoch are selected. Right: Normalized mean absolute error over the testing solutions (Eq. (3.1)) for the different equivalence ratios and LES parameters used for building the training dataset. Error bars show the first and third quartiles of the data points.

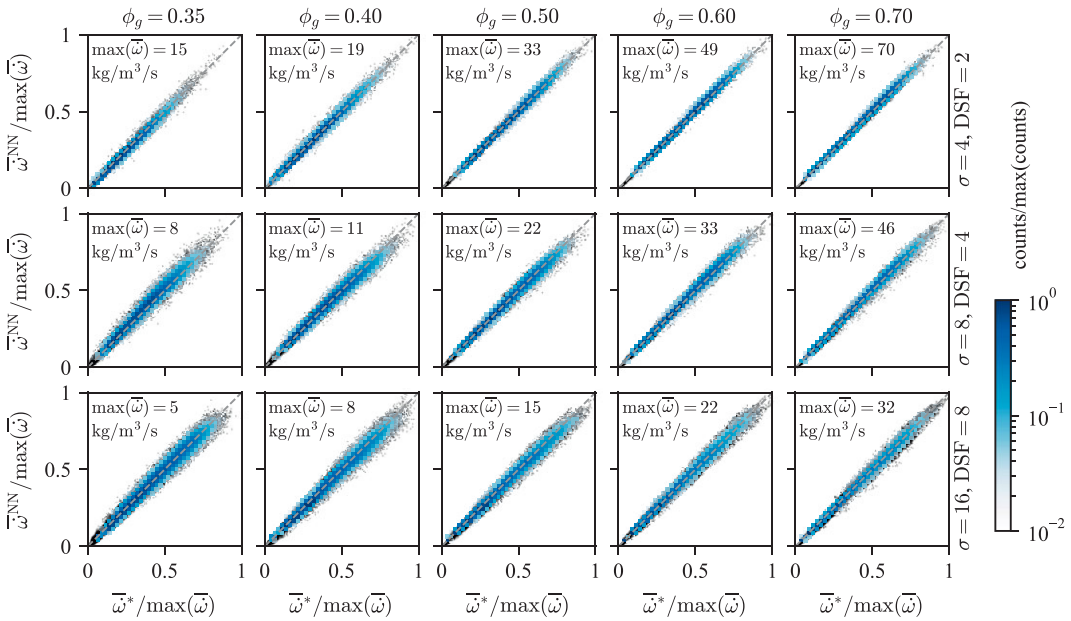


Figure 5. Scatter plots with 2D histograms: CNN-modeled burning rate $\bar{\omega}^{\text{NN}}$ versus ground-truth filtered burning rate $\bar{\omega}^*$. Individual values are normalized by the maximum burning rate in the datasets. The points used for the histograms have a progress variable c : $0.05 \leq c \leq 0.95$. Histogram values below the color scale are transparent. Gray dashed line indicates $x = y$ (that is zero error). Each column corresponds to a global equivalence ratio. Each row corresponds to a set of LES parameters (filtering and downsampling). Data are collected from the testing solutions.

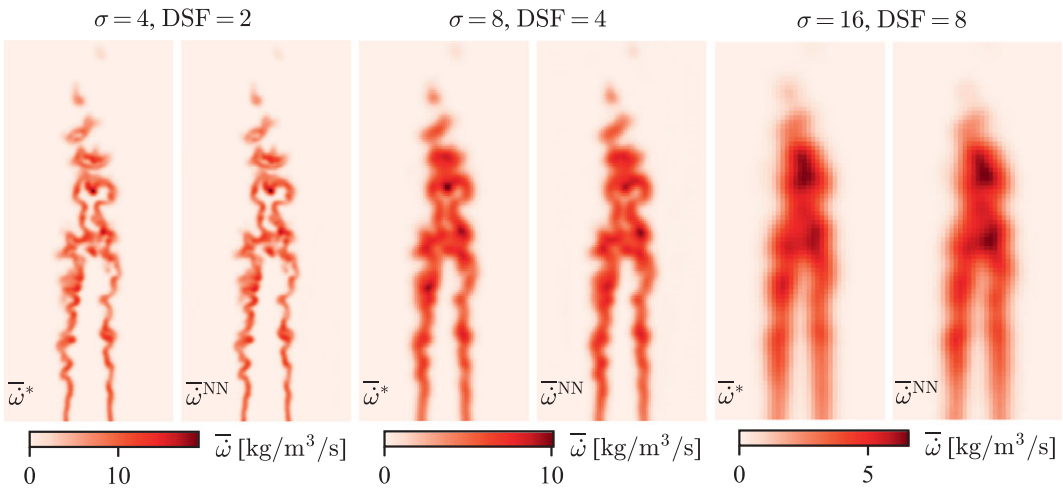


Figure 6. Planar cut normal to the z -axis, in the middle of the domain, colored by the ground-truth filtered burning rate $\bar{\omega}^*$ and the CNN-modeled burning rate $\bar{\omega}^{\text{NN}}$ for three sets of LES parameters. The global equivalence ratio is $\phi_g = 0.4$. The complex turbulent flame topology of the burning rates is remarkably well reproduced by the CNN.

topology for all filter sizes. Even the burning pockets detached from the main flame brush are correctly modeled by the CNN. This is true for all other equivalence ratios. The visualization of the cases $\phi_g = 0.35, 0.5, 0.6$ and 0.7 are given as [Supplementary Material](#) to avoid overloading the manuscript.

3.2. Generalization ability

The ability of the CNN model to generalize outside the cases for which it has been trained is now assessed. For this, two new sets of LES parameters ($\sigma = 6, \text{DSF} = 3$) and ($\sigma = 12, \text{DSF} = 6$), and two new equivalence ratios $\phi_g = 0.45$ and 0.55 are used.

Filtering and downsampling with the two new sets of LES parameters are performed on test solutions of the existing DNS with global equivalence ratios $\phi_g = 0.35, 0.4, 0.5, 0.6$ and 0.7 . Although the distribution of points $\bar{\omega}^{\text{NN}}$ versus $\bar{\omega}^*$ is a little more scattered ([Figure 7](#)), the majority of the CNN-modeled burning rates are in very good agreement with their ground-truth counterparts. This demonstrates the ability of the model to generalize between two LES filter sizes for which it has been trained. It is important to note, however, that training with all three filter sizes ([Table 3](#)) is necessary to achieve this result. We have tried to train the CNN with only two filter sizes: $\sigma = 4$ and 16 , and found that the resulting model has a significant bias for $\sigma = 8$, see [Supplementary Material](#).

The ability of the model to approximate burning rates for other global equivalence ratios is now assessed. For this, two new DNSs are performed with a global equivalence ratio $\phi_g = 0.45$ and 0.55 and six solutions per DNS (snapshot every 1 ms) are used to evaluate the model on the filter sizes $\sigma = 4, 8$ and 16 ([Figure 8](#)). The distribution of points $\bar{\omega}^{\text{NN}}$ versus $\bar{\omega}^*$ is a little more scattered compared to inference on solutions with a global equivalence ratio used for training. The majority of the points are nevertheless clustered along the $x = y$ axis: the approximation of the burning rates is still very good. The planar cuts show a flame topology very well retrieved by the CNN model. These excellent results are due to the large number of equivalence ratios included in the training dataset. We have trained the CNN with only $\phi_g = 0.35, 0.4, 0.6$ and 0.7 (that is without $\phi_g = 0.5$) and that model is flawed for the global equivalence ratio $\phi_g = 0.5$, see [Supplementary Material](#).

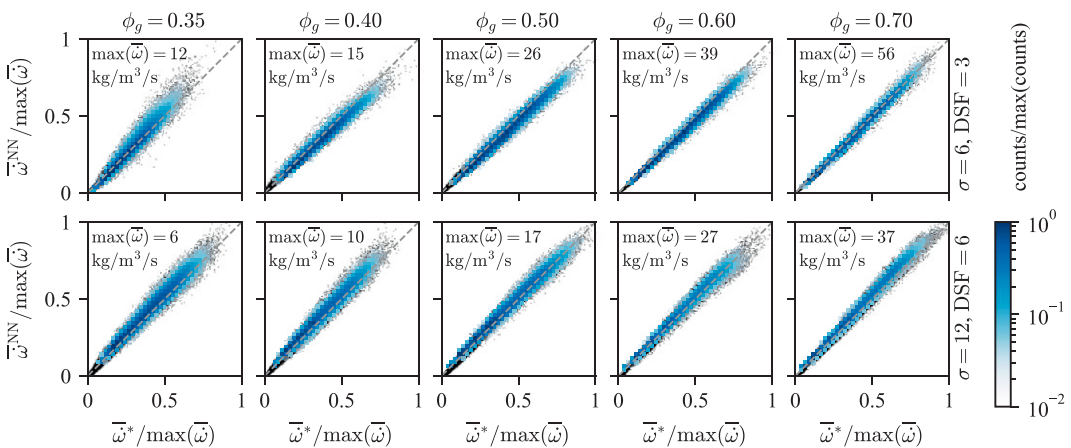


Figure 7. Scatter plots with 2D histograms: CNN-modeled burning rate $\bar{\omega}^{\text{NN}}$ versus ground-truth filtered burning rate $\bar{\omega}^*$. Individual values are normalized by the maximum burning rate in the datasets. The points used for the histograms have a progress variable c : $0.05 \leq c \leq 0.95$. Histogram values below the color scale are transparent. Gray dashed line indicates $x = y$ (that is zero error). Each column corresponds to a global equivalence ratio. Each row corresponds to a set of LES parameters (filtering and downsampling). Data are collected from solutions with two sets of LES parameters that were not included in the training dataset. The CNN approximates the burning rates with good accuracy, demonstrating the ability to generalize to other LES parameters from which it has been trained.

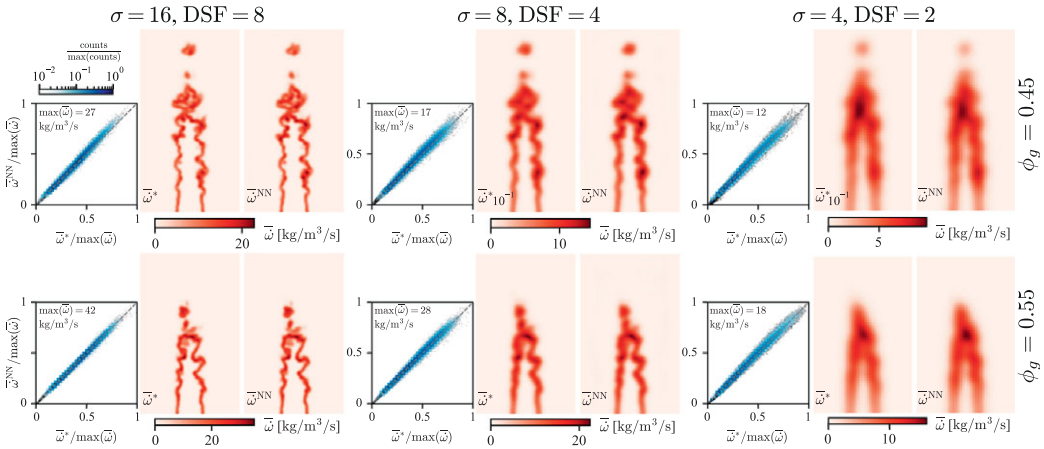


Figure 8. Scatter plots and 2D histograms together with planar cuts comparing CNN-modeled burning rate $\bar{\omega}^{NN}$ to ground-truth filtered burning rate $\bar{\omega}^*$. See caption of Figure 5 for a full description of how the histograms are constructed. The two global equivalence ratios $\phi_g = 0.45$ and 0.55 were not included in the training dataset. The CNN approximates the burning rates with high accuracy, demonstrating the ability to generalize to other equivalence ratios from which it has been trained.

These generalization results are very promising and show that a CNN model can generalize with a reasonable amount of training cases. This is particularly important with regard to the global equivalence ratio ϕ_g . Computing a DNS for each global equivalence ratio has a significant computing cost. In contrast, the computational cost of a filtering and downsampling operation to produce a new set of LES parameters (σ, DSF) is negligible.

3.3. Comparison with a filtered tabulated chemistry approach

A filtered tabulated chemistry modeling is implemented to compare CNN results with an established approach for modeling filtered burning rates. In the filtered tabulated chemistry framework, the flame structure in the direction normal to the flame front is assumed identical to the structure of a planar 1D freely propagating premixed laminar flame (Fiorina et al., 2010). For a given filter size, it is then possible to apply the filter operation introduced in Section 2.3.2 (Eqs (1.1) and (2.8)) to a precomputed 1D flame at a given global equivalence ratio ϕ_g and extract the burning rates directly as

$$\bar{\omega} \approx \bar{\omega}^+(\tilde{c}, \phi_g), \tag{3.2}$$

where $\bar{\omega}^+$ is the filtered burning rate of the 1D flame, mapped as a function of the filtered progress variable \tilde{c} for a given filter size σ . To account for the local fluctuation of the equivalence ratio in the domain, it is possible to retrieve the reaction rate in a 1D flame computed with the local equivalence ratio ϕ used to define the fresh gas composition. In this case, we write

$$\bar{\omega} \approx \bar{\omega}^+(\tilde{c}, \tilde{\phi}). \tag{3.3}$$

Flame wrinkling at the subfilter-scale level is lost in LES. A subfilter-scale wrinkling factor Ξ needs to be applied to compensate for the loss in flame surface (Poinso and Veynante, 2011). The filtered burning rate for filtered tabulated chemistry modeling therefore reads

$$\bar{\omega}^F = \Xi \bar{\omega}^+(\tilde{c}, \phi_g) \text{ or } \bar{\omega}^F = \Xi \bar{\omega}^+(\tilde{c}, \tilde{\phi}), \tag{3.4}$$

depending on whether the global or local equivalence ratio is used to define the fresh gas composition in the corresponding 1D flame. The wrinkling factor Ξ is usually modeled by algebraic expressions. Models

based on a fractal description of the flame front (Gouldin, 1987; Gouldin et al., 1989) are very common. They assume that in a range of physical scales bounded by an inner cutoff η and an outer cutoff L , the flame front is a fractal surface of dimension $2 \leq D_f \leq 3$. The wrinkling factor is given by

$$\Xi = \left(\frac{L}{\eta}\right)^{D_f-2}. \tag{3.5}$$

L corresponds to the size of the largest unresolved wrinkles, which is of the order of the thickness of the LES flame estimated as $\delta_{c,L}^1$ (Eq. (2.9)). η is the size of the smallest wrinkles and varies from $\eta = L$ (that is no unresolved wrinkling, $\Xi = 1$) to $\eta = \delta_{c,L}^0$ (that is filling of the space by a flame surface of fractal dimension D_f between cut-off scales $\delta_{c,L}^0$ and L , $\Xi = (L/\delta_{c,L}^0)^{D_f-2}$). The latter situation is reached when turbulence intensity is sufficiently high, which is often the case in practice (Veynante et al., 2012; Veynante and Moureau, 2015). We then assume $\eta = \delta_{c,L}^0$ in Eq. (3.5), maximizing the flame surface density at the subfilter-scale level. The fractal dimension is estimated as $D_f = 2.5$ following the work of Charlette et al. (2002). The set of 1D flames is computed using the San Diego mechanism (Saxena and Williams, 2006) for the chemical kinetics (same as the training database for the CNN) and a mixture-averaged transport model for the diffusivities.

CNN-based modeling significantly outperforms the filtered tabulated chemistry one, especially for the lean cases (e.g. $\phi_g = 0.35$ in Figure 9). Due to the extremely high thermodiffusive effects in very lean cases, the burning rates of the turbulent 3D field can be much higher than those in a freely propagating planar flame at the same equivalence ratio. For example, for the case $\phi_g = 0.35$, $\sigma = 4$, DSF = 2, the maximum burning rate in the 3D field $\max(\bar{\omega}^*)$ is about 17 times greater than that in a laminar planar flame $\max(\bar{\omega}^+)$. This is due to the fact that flame strain is not considered in the tabulation method applied and that the subfilter wrinkling factor Ξ only takes into account turbulence-induced wrinkling, ignoring that induced by thermodiffusive instabilities and their interactions with turbulence (Berger et al., 2022c). This $\phi_g = 0.35$ case is a truly extreme test case to illustrate the strength of the CNN model. Tabulated chemistry modeling does not fail so spectacularly for equivalence ratios close to stoichiometry, where preferential and differential diffusion effects are less pronounced (Berger et al., 2022b) (e.g. $\phi_g = 0.7$ in Figure 9). The CNN results are still much more accurate thanks to the consideration of the 3D flame structure via successive convolutions. To further demonstrate the CNN’s ability to learn from flame

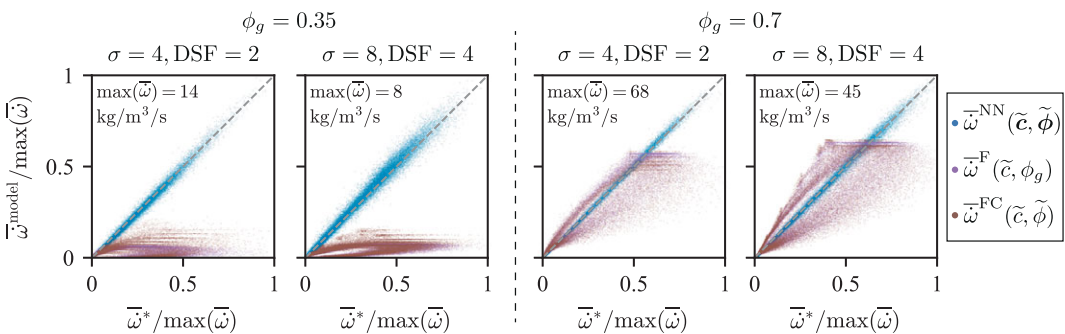


Figure 9. Scatter plots of the CNN-modeled burning rate $\bar{\omega}^{\text{NN}}$ and the filtered tabulated chemistry burning rate $\bar{\omega}^{\text{F}}$ or $\bar{\omega}^{\text{FC}}$ (Eq. (3.4)) versus ground-truth filtered burning rate $\bar{\omega}^*$. Individual values are normalized by the maximum burning rate in the datasets. Gray dashed line indicates $x = y$ (that is zero error). Two cases of LES parameter sets are presented for two global equivalence ratios: the lowest equivalence ratio $\phi_g = 0.35$ for which the filtered tabulated chemistry is significantly flawed due to very strong thermodiffusive effects; the highest equivalence ratio $\phi_g = 0.7$ for which thermodiffusive effects are less important.

topology, training has been performed without information on the equivalence ratio for the case $\phi_g = 0.35$. This study is given as [Supplementary Material](#). It is found that the burning rate approximation is still very accurate, proving our initial claim that a CNN is capable of recognizing complex flame topology and inferring associated burning rates by learning spatial hierarchies of features.

4. Conclusion and future work

Trustworthy LES of lean premixed hydrogen combustion requires advanced modeling methods considering thermodiffusive effects on burning rates at the subfilter-scale level. These effects are extremely complex and their coupling with turbulence is controversial, making the derivation of analytical models very difficult. We therefore propose a pure ML approach where a CNN is trained to approximate LES burning rates from emulated LES fields, generated by explicitly filtering many DNS snapshots. The configuration used to generate the training data is a slot burner where a turbulent flow of premixed H₂-air mixture is placed in between two coflows of burnt gas burnt at the same equivalence ratio. The training dataset is built up by performing five DNSs of this configuration, each with a unique global equivalence ratio. The DNS solutions are then filtered and downsampled using three different sizes to emulate LES solutions. Many small cubes are extracted from the solutions to train the CNN model so that it does not learn the specific configuration but rather generic flame elements, improving its generalization ability later on.

The trained CNN model retrieves burning rates with very high accuracy on the 3D full-scale test solutions. It is able to model burning rates in a very complex turbulent flame topology with flame elements detached from the main flame brush. In addition, the CNN approximates burning rates with low errors for filter and downsampling parameters and global equivalence ratios other than those for which it has been trained. The ML-based modeling outperforms a priori a classical filtered tabulated chemistry approach with correction for flame surface density due to subfilter-scale flame wrinkling. This is particularly true for very lean flames, where differential diffusion effects are major and undermine modeling based on flame surface density.

These results are extremely promising. They demonstrate that it is possible to build an ML-based model of the burning rates of turbulent lean premixed H₂-air flames—an extremely challenging task—that can generalize outside the scope for which it has been trained. Generalization is very important if we are really going to consider implementing this type of model in engineering simulation software. This work paves the way for a new modeling approach to turbulent H₂ combustion, which we want to pursue with the scientific community in the very near future.

We can then draw up a roadmap for integration into an LES code. First, the modeling of the subfilter-scale fluxes in Eq. (1.2) remains to be addressed. It could be done through gradient transport models with a turbulent viscosity hypothesis (Boger et al., 1998). If this closure fails to retrieve correct flame propagation speeds, we could consider training the CNN to approximate the fluxes using several output channels. Second, the evolution of the entire chemical system must be described. This can be done by training the CNN on the source term of a progress variable, transported instead of the individual species that make up the reactive mixture. This is done, for example, in Lapenna et al. (2021, 2024) with a tabulation method for the closure of the transport equation of the progress variable. However, the transport of the mixture fraction must be modeled to account for equivalence ratio fluctuations. We could also train the model again on a database generated with a single global chemical reaction (Millán-Merino and Boivin, 2024; Schiavone et al., 2024). This has the advantage of being able to transport physical species instead of a progress variable, accounting for preferential diffusion through the multi-species Navier–Stokes equations directly.

Second, there are a number of practical integration challenges to be tackled. In a massively parallel HPC code, the fluid domain is decomposed in many partitions distributed over the MPI tasks (typically one MPI per GPU). When inference is performed on each MPI, artefacts may appear at the boundaries of the partitions (Serhani et al., 2024). To avoid this, the inference must be given its neighboring context to

produce overlapping predictions. However, this incurs additional MPI communication costs that must be evaluated on a case-by-case basis. Another major challenge is related to the voxel grid requirement for the CNN. If the CFD code uses an irregular grid that does not match the voxel/pixel structure, two grids have to be used: one for the fluid solver, and one for the CNN inference. This requires the development of efficient methods for interpolating on the fly and transferring data between the two grids. The PhyDLL¹ (Serhani et al., 2024) open-source library is one effort that could help to address these issues. These questions will be the focus of a future project, now that we have shown that CNN-based modeling of the burning rates is possible with this work.

Supplementary material. The supplementary material for this article can be found at <http://doi.org/10.1017/dce.2025.1>.

Data availability statement. Code for training and inference is available via GitLab at https://gitlab.com/male.quentin/cnn_h2flame. A set of solutions from DNS is distributed via Kaggle as part of BLASTNet² (<https://www.kaggle.com/organizations/blastnet>) with the following identifiers:

- blastnet/premixed-flame-slot-burner-dns-h2air-phi035
- blastnet/premixed-flame-slot-burner-dns-h2air-phi04
- blastnet/premixed-flame-slot-burner-dns-h2air-phi05
- blastnet/premixed-flame-slot-burner-dns-h2air-phi06
- blastnet/premixed-flame-slot-burner-dns-h2air-phi07

Acknowledgments. The authors gratefully acknowledge CERFACS for providing the LES solver AVBP and the Antares and ARCANE libraries.

Author contribution. Q.M.: conceived the research project, performed the DNS, generated the database, trained the CNN, post-processed the results, performed the analysis, discussed the results, and wrote the article. C.J.L.: discussed the results, and reviewed and edited the article. N.N.: discussed the results, and reviewed and edited the article.

Funding statement. This research was funded in part by the Swiss National Science Foundation (SNSF) under grant agreement No. 219938. This work was supported by a grant from the Swiss National Supercomputing Centre (CSCS) under project ID s1262.

Competing interest. The authors declare none.

Ethical standard. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

References

- Aniello A, Laera D, Berger L, Attili A and Poinso T (2022) Introducing thermodiffusive effects in large-eddy simulation of turbulent combustion for lean hydrogen-air flames. In *Center for Turbulence Research, Proceedings of the Summer Program*, pp. 267–277.
- Ansel J, Yang E, He H, Gimelshein N, Jain A, Voznesensky M, Bao B, Bell P, Berard D, Burovski E, Chauhan G, Chourdia A, Constable W, Desmaison A, DeVito Z, Ellison E, Feng W, Gong J, Gschwind M, Hirsh B, Huang S, Kalambarkar K, Kirsch L, Lazos M, Lezcano M, Liang Y, Liang J, Lu Y, Luk CK, Maher B, Pan Y, Puhersch C, Reso M, Saroufim M, Siraichi MY, Suk H, Zhang S, Suo M, Tillet P, Zhao X, Wang E, Zhou K, Zou R, Wang X, Mathews A, Wen W, Chanan G, Wu P and Chintala S (2024) PyTorch 2: Faster machine learning through dynamic python bytecode transformation and graph compilation. In *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, pp. 929–947, La Jolla, CA, USA. ACM.
- Aspden A (2017) A numerical study of diffusive effects in turbulent lean premixed hydrogen flames. *Proceedings of the Combustion Institute* 36(2), 1997–2004.
- Aspden A, Day M and Bell J (2015) Turbulence-chemistry interaction in lean premixed hydrogen combustion. *Proceedings of the Combustion Institute* 35(2), 1321–1329.
- Berger L, Attili A and Pitsch H (2022a) Intrinsic instabilities in premixed hydrogen flames: Parametric variation of pressure, equivalence ratio, and temperature. Part 1 - Dispersion relations in the linear regime. *Combustion and Flame* 240, 111935.
- Berger L, Attili A and Pitsch H (2022b) Intrinsic instabilities in premixed hydrogen flames: Parametric variation of pressure, equivalence ratio, and temperature. Part 2 – Non-linear regime and flame speed enhancement. *Combustion and Flame* 240, 111936.

¹ <https://phydll.readthedocs.io>

² <https://blastnet.github.io>

- Berger L, Attili A and Pitsch H** (2022c) Synergistic interactions of thermodiffusive instabilities and turbulence in lean hydrogen flames. *Combustion and Flame* 244, 112254.
- Berger L, Attili A, Wang J, Maeda, K., and Pitsch, H.** (2022d). Development of large-eddy simulation combustion models for thermodiffusive instabilities in turbulent hydrogen flames. In *Center for Turbulence Research, Proceedings of the Summer Program*. 247–256
- Bilger R, Stårner S and Kee R** (1990) On reduced mechanisms for methane-air combustion in nonpremixed flames. *Combustion and Flame* 80(2), 135–149.
- Biscione V and Bowers JS** (2021) Convolutional neural networks are not invariant to translation, but they can learn to be. *Journal of Machine Learning Research* 22(229), 1–28.
- Boger M, Veynante D, Boughanem H and Trouvé A** (1998) Direct numerical simulation analysis of flame surface density concept for large eddy simulation of turbulent premixed combustion. *Symposium (International) on Combustion* 27(1), 917–925.
- Butler T and O'Rourke P** (1977) A numerical method for two dimensional unsteady reacting flows. *Symposium (International) on Combustion* 16(1), 1503–1515.
- Charlette F, Meneveau C and Veynante D** (2002) A power-law flame wrinkling model for LES of premixed turbulent combustion part I: Non-dynamic formulation and initial tests. *Combustion and Flame* 131(1–2), 159–180.
- Colin O, Ducros F, Veynante D and Poinso T** (2000) A thickened flame model for large eddy simulations of turbulent premixed combustion. *Physics of Fluids* 12(7), 1843–1863.
- Colin O and Rudgyard M** (2000) Development of high-order Taylor–Galerkin schemes for LES. *Journal of Computational Physics* 162(2), 338–371.
- Coulon V, Gaucherand J, Xing V, Laera D, Lapeyre C and Poinso T** (2023) Direct numerical simulations of methane, ammonia-hydrogen and hydrogen turbulent premixed flames. *Combustion and Flame* 256, 112933.
- Duwig C** (2007) Study of a filtered Flamelet formulation for large Eddy simulation of premixed turbulent flames. *Flow, Turbulence and Combustion* 79(4), 433–454.
- Ferrante G, Eitelberg G and Langella I** (2024) Differential diffusion modelling for LES of premixed and partially premixed flames with presumed FDF. *Combustion Theory and Modelling* 28(6), 695–730.
- Fiorina B, Veynante D and Candel S** (2015) Modeling combustion chemistry in large Eddy simulation of turbulent flames. *Flow, Turbulence and Combustion* 94(1), 3–42.
- Fiorina B, Vicquelin R, Auzillon P, Darabiha N, Gicquel O and Veynante D** (2010) A filtered tabulated chemistry model for LES of premixed combustion. *Combustion and Flame* 157(3), 465–475.
- Gaucherand J, Laera D, Schulze-Netzer C and Poinso T** (2024a) DNS of turbulent premixed ammonia/hydrogen flames: The impact of thermo-diffusive effects. *Flow, Turbulence and Combustion* 112(2), 587–614.
- Gaucherand J, Schulze-Netzer C, Laera D and Poinso T** (2024b) A subgrid-scale model to account for thermo-diffusive effects in artificially thickened LES models for lean turbulent premixed ammonia/hydrogen flames. *Proceedings of the Combustion Institute* 40(1–4), 105198.
- Gicquel LY, Gourdain N, Bousse J-F, Deniau H, Staffelbach G, Wolf P and Poinso T** (2011) High performance parallel computing of flows in complex geometries. *Comptes Rendus Mécanique* 339(2–3), 104–124.
- Gouldin F** (1987) An application of fractals to modeling premixed turbulent flames. *Combustion and Flame* 68(3), 249–266.
- Gouldin F, Bray K and Chen J-Y** (1989) Chemical closure model for fractal flamelets. *Combustion and Flame* 77(3–4), 241–259.
- Ho JZ, Talei M, Brouzet D, Chung WT, Sharma P and Ihme M** (2024) Augmenting filtered flame front displacement models for LES using machine learning with a posteriori simulations. *Proceedings of the Combustion Institute* 40(1–4), 105311.
- Kerstein AR, Ashurst WT and Williams FA** (1988) Field equation for interface propagation in an unsteady homogeneous flow field. *Physical Review A* 37(7), 2728–2731.
- Kingma, D. P. and Ba, J.** (2017). Adam: A method for stochastic optimization. [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) [cs].
- Lapenna PE, Berger L, Attili A, Lamioni R, Fogla N, Pitsch H and Creta F** (2021) Data-driven subfilter modelling of thermo-diffusively unstable hydrogen–air premixed flames. *Combustion Theory and Modelling* 25(6), 1064–1085.
- Lapenna PE, Remiddi A, Molinaro D, Indelicato G and Creta F** (2024) A-posteriori analysis of a data-driven filtered wrinkled flamelet model for thermodiffusively unstable premixed flames. *Combustion and Flame* 259, 113126.
- Lapeyre CJ, Misdariis A, Cazard N and Poinso T** (2018) A-posteriori evaluation of a deep convolutional neural network approach to subgrid-scale flame surface estimation. In *Center for Turbulence Research, Proceedings of the Summer Program*, pp. 349–358.
- Lapeyre CJ, Misdariis A, Cazard N, Veynante D and Poinso T** (2019) Training convolutional neural networks to estimate turbulent sub-grid scale reaction rates. *Combustion and Flame* 203, 255–264.
- Loshchilov, I. and Hutter, F.** (2019). Decoupled weight decay regularization. [arXiv:1711.05101](https://arxiv.org/abs/1711.05101) [cs, math].
- Marble, F. E. and Broadwell, J. E.** (1997). The coherent flame model for turbulent chemical reactions. Technical Report TRW-9-PU.
- Millán-Merino A and Boivin P** (2024) A new single-step mechanism for hydrogen combustion. *Combustion and Flame* 268, 113641.
- Nikolaou ZM, Chrysostomou C, Vervisch L and Cant S** (2019) Progress variable variance and filtered rate modelling using convolutional neural networks and flamelet methods. *Flow, Turbulence and Combustion* 103(2), 485–501.
- Passot T and Pouquet A** (1987) Numerical simulation of compressible homogeneous flows in the turbulent regime. *Journal of Fluid Mechanics* 181, 441–466.

- Peters N** (1988) Laminar flamelet concepts in turbulent combustion. *Symposium (International) on Combustion* 21(1), 1231–1250.
- Petrov, C. and Ghoniem, A.** (1994). An unsteady strained flame model for turbulent combustion simulations. In *32nd Aerospace Sciences Meeting and Exhibit*, Reno, NV, USA. American Institute of Aeronautics and Astronautics.
- Pitsch H** (2024) The transition to sustainable combustion: Hydrogen- and carbon-based future fuels and methods for dealing with their challenges. *Proceedings of the Combustion Institute* 40(1–4), 105638.
- Poinsot T and Lele S** (1992) Boundary conditions for direct simulations of compressible viscous flows. *Journal of Computational Physics* 101(1), 104–129.
- Poinsot, T. and Veynante, D.** (2011). *Theoretical and Numerical Combustion*, 3rd Edn. Toulouse: CNRS.
- Pope S** (1988) The evolution of surfaces in turbulence. *International Journal of Engineering Science* 26(5), 445–469.
- Ronneberger O, Fischer P and Brox T** (2015) U-net: Convolutional networks for biomedical image segmentation. In Navab N, Hornegger J, Wells WM and Frangi AF (eds), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. Cham: Springer International Publishing, pp. 234–241.
- Saxena P and Williams FA** (2006) Testing a small detailed chemical-kinetic mechanism for the combustion of hydrogen and carbon monoxide. *Combustion and Flame* 145(1–2), 316–323.
- Schiavone FG, Detomaso N, Torresi M and Laera D** (2024) An Arrhenius-based one-step reaction mechanism for hydrogen-air flames simulations in an extended range of operating conditions. *International Journal of Hydrogen Energy* 57, 1229–1243.
- Schlup J and Blanquart G** (2018) Validation of a mixture-averaged thermal diffusion model for premixed lean hydrogen flames. *Combustion Theory and Modelling* 22(2), 264–290.
- Schonfeld T and Rudgyard M** (1999) Steady and unsteady flow simulations using the hybrid flow solver AVBP. *AIAA Journal* 37(11), 1378–1385.
- Seltz A, Domingo P, Vervisch L and Nikolaou ZM** (2019) Direct mapping from LES resolved scales to filtered-flame generated manifolds using convolutional neural networks. *Combustion and Flame* 210, 71–82.
- Serhani A, Xing V, Dupuy D, Lapeyre C and Staffelbach G** (2024) Graph and convolutional neural network coupling with a high-performance large-eddy simulation solver. *Computers & Fluids* 278, 106306.
- Trouvé A and Poinsot T** (1994) The evolution equation for the flame surface density in turbulent premixed combustion. *Journal of Fluid Mechanics* 278, 1–31.
- Tullis S and Cant R** (2002a) Scalar transport modeling in large eddy simulation of turbulent premixed flames. *Proceedings of the Combustion Institute* 29(2), 2097–2104.
- Tullis S and Cant RS** (2002b) Counter-gradient scalar transport in large Eddy simulation of turbulent premixed flames. In Moreau R, Pollard A and Candel S (eds.) *IUTAM Symposium on Turbulent Mixing and Combustion*, vol. 70. Series Title: Fluid Mechanics and its Applications. Netherlands, Dordrecht: Springer, pp. 259–267.
- Vagelopoulos C, Egolfopoulos F and Law C** (1994) Further considerations on the determination of laminar flame speeds with the counterflow twin-flame technique. *Symposium (International) on Combustion* 25(1), 1341–1347.
- Veynante D and Moureau V** (2015) Analysis of dynamic models for large eddy simulations of turbulent premixed combustion. *Combustion and Flame* 162(12), 4622–4642.
- Veynante D, Schmitt T, Boileau M and Moureau V** (2012) Analysis of dynamic models for turbulent premixed combustion. In *Center for Turbulence Research, Proceedings of the Summer Program*, pp. 387–396.
- Yao MX and Blanquart G** (2024) Capturing differential diffusion effects in large eddy simulation of turbulent premixed flames. *Proceedings of the Combustion Institute* 40(1–4), 105500.