Combustion and Flame 278 (2025) 114218

Contents lists available at ScienceDirect

Combustion and Flame

journal homepage: www.sciencedirect.com/journal/combustion-and-flame

Deep learning based combustion chemistry acceleration method for widely applicable NH_3/H_2 turbulent combustion simulations

Sipei Wu [a], Wenkai Liang [a,*], Kai Hong Luo [a,b,**]

[a] Center for Combustion Energy, Key Laboratory for Thermal Science and Power Engineering of Ministry of Education, Department of Energy and Power Engineering, Tsinghua University, Beijing 100084, China

[b] Department of Mechanical Engineering, University College London, Torrington Place, London WC1E 7JE, United Kingdom

ARTICLE INFO

Keywords:
Turbulent premixed combustion
Fuel blends
Detailed chemical mechanism
Combustion chemistry acceleration
Artificial neural network
Manifold sampling

ABSTRACT

Simulating reacting flows with detailed chemistry is often prohibitively expensive due to the complexity of reaction mechanisms and the numerical stiffness arising from disparate chemical time scales. While recent advancements in neural networks offer potential for efficiently capturing the dynamics of stiff chemistry, its application to dual-fuels with drastic differences in reactivity such as ammonia (NH_3) and hydrogen (H_2) remains challenging. In this study, we present a neural network model with variable time steps aimed at enhancing the efficiency of combustion chemistry simulations focusing on the complex dual-fuel NH_3/H_2 under premixed combustion. We improved the "sampling-training" workflow based on previous HFRD method to overcome the challenge of generalizing neural network models to fuel blends under premixed combustion. This workflow involves three improvements: defining the base manifold using unity Lewis number laminar flames, introducing continuously controllable randomization, and employing a training process with mass conservation and heat release rate similarity constraints. Our approach is validated against simulations of planar turbulent premixed flames and temporally-evolving jet flames across various conditions. The model demonstrates high accuracy and consistency, achieving a chemical calculation acceleration of 7 times and an overall simulation acceleration of 5 times using a model with 4 hidden layers and 800 neurons on the same CPU device. When a GPU is adopted, the chemical calculation acceleration increases to 30 times, and the overall simulation acceleration reaches 10 times.

Novelty and Significance Statement

Utilizing detailed chemistry in reacting flow simulations drastically increases computational cost due to numerical stiffness and disparate time scales. A promising approach is to replace the time-consuming ODE solvers with compact neural networks. Despite the rapid development of the neural network approach for accelerating combustion kinetics calculations, the application of this concept to fuel blends with varying mixing ratios and reactivities remains insufficient, particularly in turbulent premixed flames. In this study, we improved the neural network framework that could predict the kinetics of fuel blends of low-reactivity fuel NH_3 and high-reactivity fuel H_2, highlighting the applications to binary fuel with large reactivity differences. Specifically, the unity Lewis number laminar flames are leveraged as an economical thermochemical base manifold, given their close resemblance to turbulent flame profiles. A continuously controllable randomization method is introduced to balance model capacity and computational efficiency by adjusting key parameters. Additionally, a loss function with mass conservation and heat release rate similarity constraints ensures stable long-term predictions. The well-trained neural network model was coupled with CFD codes to simulate two challenging cases across a wide range of turbulent intensities and fuel compositions. The results show visually identical scalar fields and highly accurate statistical outcomes, even under intense turbulence and after $O(10^4)$ neural network model calls, with a O(10) acceleration in computation.

* Corresponding author.
** Corresponding author at: Center for Combustion Energy, Key Laboratory for Thermal Science and Power Engineering of Ministry of Education, Department of Energy and Power Engineering, Tsinghua University, Beijing 100084, China.
E-mail addresses: liangwenkai@tsinghua.edu.cn (W. Liang), k.luo@ucl.ac.uk (K.H. Luo).

https://doi.org/10.1016/j.combustflame.2025.114218
Received 29 September 2024; Received in revised form 20 April 2025; Accepted 3 May 2025
Available online 26 May 2025
0010-2180/© 2025 The Authors. Published by Elsevier Inc. on behalf of The Combustion Institute. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

1. Introduction

Due to the ongoing climate crisis and the increasing levels of atmospheric carbon dioxide (CO_2) [1], it is imperative to explore zero-carbon fuels and carbon-neutral combustion technologies to reduce greenhouse gas emissions and meet stringent pollutant regulations. The dual-fuel engine concept, which utilizes a high-reactivity pilot fuel to ignite low-reactivity premixed charges, has been gaining significant attentions [2]. Among the various alternative dual-fuels, hydrogen (H_2) and ammonia (NH_3) have become important topics in recent combustion research. While H_2 is recognized as a clean fuel with well-understood combustion characteristics, challenges related to its supply chain, economic storage, and secure transportation have hindered widespread adoption [3]. In contrast, NH_3 has emerged as a promising zero-carbon fuel due to its advantages in storage and transportation. Nevertheless, NH_3 presents challenges such as low reactivity, low flame speeds, high autoignition temperature, and a narrow flammability range compared to H_2, making stable ignition and combustion difficult. To enhance the combustion of NH_3/air mixtures, the addition of H_2 is a common approach. The significant differences in reactivity between NH_3 and H_2 complicate the chemical kinetics modeling and reacting flow simulations for NH_3/H_2 mixtures.

To accurately address the chemical kinetics, burning characteristics, and nitric oxide (NO_x) emissions across various combustion devices, such as gas turbines and internal combustion engines, it is essential to employ detailed mechanisms in simulations [4]. However, reacting flow simulations involving detailed chemistry can be prohibitively expensive due to the size of reaction mechanisms and the numerical stiffness arising from the broad range of chemical time scales [5]. Evaluating the chemical source terms at every grid node and for each time step of the simulation is particularly computationally intensive, usually accounting for over 80% of the total computing time [6]. More importantly, the calculation of chemical reactions is a common challenge faced by all combustion simulations, and the acceleration of this submodule can benefit almost all simulation codes, even other broader fields involving detailed chemical kinetic modeling in chemical, biological and atmospheric systems.

The advent of deep learning technologies offers a promising solution to the challenges in simulating stiff chemistry [7,8]. Artificial neural networks (ANNs) have shown the ability to capture the dynamics of stiff chemistry with larger time steps more efficiently than traditional stiff ordinary differential equation (ODE) solvers, which typically require small step sizes and computationally expensive implicit methods (e.g., SEULEX [9] and CVODE [10]) [11–13]. ANNs are trained to map initial states to final states and can be applied in an explicit manner. Replacing these costly solvers with pre-trained neural ODE solvers has emerged as a promising alternative.

ANNs have been utilized in combustion research since at least 1996 [14]. More advanced neural network architectures such as self-organizing maps (SOMs) have been used to cluster the high-dimensional thermochemical space and trained ANN for each cluster separately [15]. This SOM-ANN approach has been extended to various complex multi-dimensional problems [16–19]. Ding et al. [20] proposed a multiple ANN approach to build the ANN for each species, successfully coupled with simulation of CH_4/air piloted flames. Deep neural networks (DNNs) with complex structures comprising large numbers of layers and neurons are gaining increasing attention. Hansinger et al. [21] adopted deep residual networks for flamelet/progress variable tabulation with application to the piloted flame. Recently, neural operators have also been explored. Goswami et al. [22] employed autoencoder-based deep operator networks (AE-DeepONet) for the kinetics of a skeletal syngas model. DeepONet is also applied by Kumar et al. [23] for the complex chemical kinetics of n-dodecane at high- and low-temperatures. Weng et al. [24] evaluated and tested the performance of Fourier neural operators (FNO) in solving stiff chemical kinetics.

One key challenge in training ANNs or DNNs for the tabulation of thermochemistry is preparing the dataset. Early validation efforts trained neural networks with data from a single case and tested them in the same case (*a-priori* validation), which frequently led to poor generalization to new cases (*a-posteriori* validation). This issue arises because thermochemical states sampled from a multi-dimensional case represent only a small subset of the broader thermochemical space and may not be suitable for other multi-dimensional problems. To address this, many studies have focused on *a-posteriori* validation of ANN/DNN models across various cases. Fig. 1 reviews existing sampling methods for preparing the dataset, which are categorized into global sampling and manifold sampling.

Global sampling involves sampling the training data from the entire composition space. The most straightforward method, Monte Carlo (MC) sampling, is impractical for large mechanisms due to the complexity of data distribution and the limited capacity of model parameters. To mitigate this, effective data filtering is essential. Zhang et al. [25] introduced the multi-scale (MS) sampling method to address the challenge of collecting diverse thermochemical states for H_2 combustion. Building on this, Xu et al. [26] proposed a global multi-scale sampling (GMS) method that improves the prediction accuracy for methane and n-heptane kinetics by sampling along the evolution trajectory of Monte Carlo data and filtering based on the temporal gradient range.

Manifold sampling, which is also employed in this work, assumes that thermochemical states in high-dimensional simulations lie nearly on a low-dimensional manifold [27]. This method uses canonical combustion configurations to identify base states and generate realistic states through data randomization algorithms. 0-D homogeneous reactors such as autoignition have been widely applied to prepare dataset due to their simplicity [25,28,29]. Sen et al. [30] trained the ANN on a table extracted from an unsteady flame—vortex interaction (FVI) simulation and applied it to turbulent premixed flames. In their further work [31,32], linear eddy mixing (LEM) simulations are employed

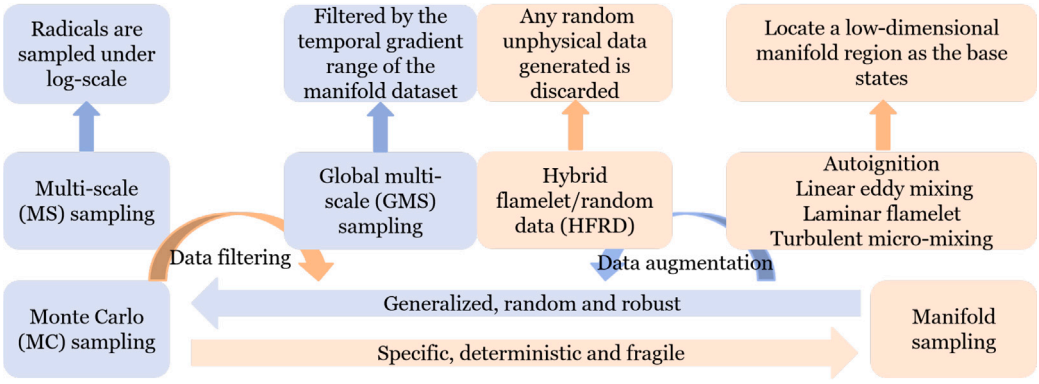

Fig. 1. Review of sampling methods.

for ANN training to include the effect of the eddies on a wide range of scales. Chatzopoulos et al. [16] simulated an ensemble of laminar flamelets to generate training samples, which was also the basis of the approach in Franke et al. [17]. Wan et al. [33] trained the ANN using a turbulent non-adiabatic non-premixed micro-mixing based canonical problem. Nguyen et al. [34] also sampled from a turbulent micro-mixing problem and decomposed the thermochemical hyperspace into clusters to facilitate the training of neural networks. To simulate turbulent premixed flames, Readshaw et al. employed [35] a set of 1-D premixed flame simulations to generate data. To enhance the robustness and generalization capability of the models, data augmentation on the base manifold is necessary. Recently, the hybrid flamelet/random data (HFRD) method [20,36] used unsteady laminar flamelets with ignition and extinction phases to generate training data, with successful application to various turbulent non-premixed flames. HFRD method is further developed by Li et al. [37,38] and applied to hydrogen, ethylene, and Jet-A premixed combustion.

Although many of the aforementioned works have applied ANNs or DNNs to combustion problems with fixed fuel compositions, there are very few examples of ANN/DNN application to fuel blends. An early work on blended fuels with varying mixing ratios is [39], which employed a wide range of mixing ratios and initial temperature for the flamelet simulations to account for fuel blends and radiative heat losses. The trained ANNs are applied to two turbulent flames with different fuel compositions: a pure methane flame, Sandia flame D, and Sydney flame HM1, which is a methane/hydrogen flame. Moreover, compared with non-premixed flames, generating suitable data for ANN training that successfully anticipates a premixed flame composition space is even more difficult [35]. The applicability of ANNs to predict the stiff chemistry dynamics of NH_3/H_2 in turbulent premixed flames, where significantly different kinetics are involved, has yet to be thoroughly examined. To achieve this, the ANN or DNN model must accurately reproduce the predictions of ODE solvers across a wide range of equivalence ratios and mixing ratios.

In this study, we develop a DNN model with variable time steps to enhance the efficiency of combustion chemistry simulations for turbulent premixed NH_3/H_2 flames. Using NH_3/H_2 as the example of fuel blends with significantly different chemical properties, we improved the "sampling-training" workflow based on previous HFRD method. The improvements includes defining the base manifold using unity Lewis number laminar flames, introducing continuously controllable randomization, and employing a training process with mass conservation and heat release rate similarity constraints. We simulate planar turbulent premixed flames with intensive turbulence and temporally-evolving jet flames with strong shear flows using both CVODE and the proposed DNN model. Two key hyperparameters are optimized to balance data coverage and model complexity. Then, scalar distributions and statistical results are compared across various temperatures, time steps, equivalence ratios, and mixing ratios.

The rest of the paper proceeds as follows: Section 2 introduces the chemistry ODE system and the details of the "sampling-training" workflow. Section 3 describes the validation configurations employed in the present study. Qualitative and quantitative analysis of proposed method are provided in Section 4. Finally, Section 5 gives conclusions and suggests future work.

2. Chemistry acceleration framework

2.1. Problem statement

In the combustion simulations, complex interactions between convection, diffusion, and chemical reaction processes are described by equations such as partial differential equations (PDE), ordinary differential equations (ODE), and algebraic equations. For example, the mass fraction advances in time, as shown in the following equation,

$$Y(t+\Delta t) = Y(t) + (-C + D + R)\Delta t, \tag{1}$$

where Y, C, D, and R are the mass fraction, convection term, diffusion term, and chemical reaction term of the mass fraction transport equation. In many reacting flow codes, the operator-splitting strategy [40] is employed where chemistry and transport are decoupled and integrated with distinct numerical methods. This strategy splits the governing equation into sub-steps and integrates each separately to advance to the next time step [41]. The chemistry contribution results in a system of ODEs at each cell of the computational domain, which is typically integrated using a stiff ODE solver. The transport contribution results in a system of PDEs, usually solved by finite difference/volume methods (FDM/FVM). A typical solving process is outlined below,

$$\begin{aligned} Y^*(t+\Delta t) &= Y(t) + (-C + D)\Delta t \\ Y(t+\Delta t) &= Y^*(t+\Delta t) + \int_0^{\Delta t} R\Delta t. \end{aligned} \tag{2}$$

This widely adopted operator-splitting strategy offers opportunities to accelerate each submodule separately. This paper focuses on the accelerating submodule of the chemical reaction source term.

Considering a chemical system of n_s species and n_r reactions,

$$\sum_{i=1}^{n_s} \upsilon'_{ij}\mathcal{M}_i \rightleftharpoons \sum_{i=1}^{n_s} \upsilon''_{ij}\mathcal{M}_i \text{ for } j = 1, \dots, n_r, \tag{3}$$

where $\mathcal{M}_i$ denotes species i, υ'_{ij} and υ''_{ij} represent the molar stoichiometric coefficients of species i in reaction j. For each cell in the computational domain, the change in composition due to chemistry over a computational timestep Δt is the solution of a system of $(n_s + 1)$ ODEs for the species mass fraction and temperature. The species equations and temperature equation for a constant pressure system are

$$\begin{aligned} \frac{dY_i}{dt} &= \frac{\omega_i W_i}{\rho} \text{ for } i = 1, \dots, n_s, \\ \frac{dT}{dt} &= -\frac{\sum_{i=1}^{n_s} \omega_i h_i W_i}{\rho c_p}, \end{aligned} \tag{4}$$

where Y_i, ω_i, W_i, and h_i are the mass fraction, net production rate, molecular weight, and molar enthalpy of species i, and ρ is mixture mass density. T is temperature, and c_p is the heat capacity at constant pressure. The ideal gas law is used as the equation of state.

Here, ω_i, the net production rate of species i, is calculated from

$$\omega_i = \sum_{j=1}^{n_r} \left(\upsilon'_{ij} - \upsilon''_{ij}\right)\left(K_{fj}\prod_{i=1}^{n_s}[X]_i^{\upsilon'_{ij}} - K_{rj}\prod_{i=1}^{n_s}[X]_i^{\upsilon''_{ij}}\right), \tag{5}$$

where K_{fj} and K_{rj} are the forward and reverse rates of reaction j, which are usually modeled using the Arrhenius law with necessary modifications. $[X]_i$ is the molar concentration of species i. $\dot{Q} = -\sum_{i=1}^{n_s} \omega_i h_i$ is defined as heat release rate (HRR).

Considering the requirement to invoke expensive stiff ODE solvers at every time step and on every grid, it is prudent to replace them with compact neural networks. Neural networks eliminate the stiffness differences (i.e., load imbalance) under various thermochemical states and offer the potential for further acceleration with the help of advanced GPUs. The "sampling-training" framework applied in this work is developed based on hybrid flamelet/random data (HFRD) generation approach, as shown in Fig. 2. This framework divided the whole "sampling-training" process into three parts,

1. Calculate the base manifold.
2. Perform randomization to enhance model generalization capability and robustness.
3. Train the neural network model.

After these steps, we can couple the pre-trained neural network into high-dimensional CFD codes to accelerate combustion chemistry.

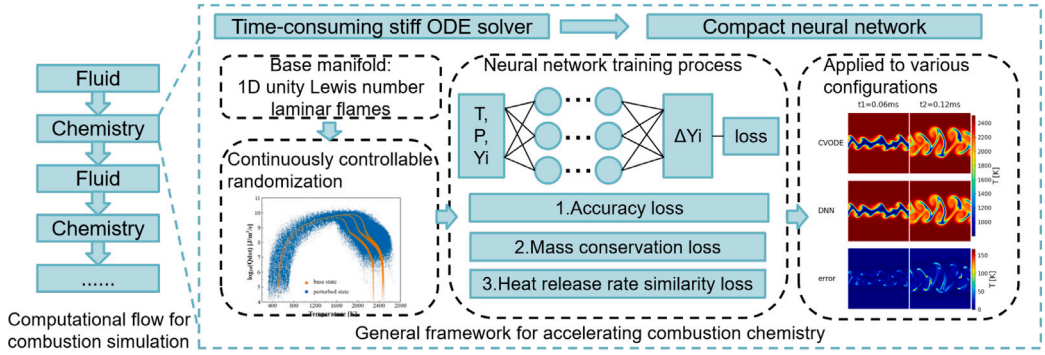

Fig. 2. General framework for combustion chemistry acceleration.

Table 1
Summary of flame simulations used to generate the dataset of base thermochemical states. p is the pressure, T_u is the unburnt gas temperature, ϕ is the equivalence ratio, $X(H_2)$ is the mixing ratio of H_2.

Simulation	Combustion mode	Conditions	Values used
Flame	Premixed	p [Bar]	1
		T_u [K]	800
		ϕ [−]	0.5, 0.6, 0.7, 0.8, 0.9, 1.0, 1.1, 1.2, 1.3, 1.4, 1.5
		$X(H_2)$ [−]	0, 0.05, 0.10, 0.15, 0.20, 0.25, 0.30

2.2. Base manifold

Following Readshaw et al. [35], which used a set of laminar premixed flames as the canonical problem for turbulent premixed flames, we generate a dataset of thermochemical states consisting of temperature, pressure, and species mass fractions $\left[T, p, Y_1, Y_2, \ldots, Y_{n_s}\right]$ using 1-D unity Lewis number laminar premixed flame simulations. It is noted that the unity Lewis number transport model is based on findings that, as turbulence intensity increases, conditioned profiles of various local mixture characteristics tend to resemble those computed for unity Lewis number laminar flames [42–45]. This is true of NH_3 and NH_3/H_2 turbulent flames, as depicted in Fig. 3. The selected conditions span an extensive range of equivalence ratios and mixing ratios for dual-fuels in NH_3/H_2 combustion, as summarized in Table 1. The mixing ratio of H_2 is defined as,

$$X(H_2) = \frac{[H_2]_0}{[NH_3]_0 + [H_2]_0}, \tag{6}$$

where $[H_2]_0$ and $[NH_3]_0$ denote the molar concentration of H_2 and NH_3 in the fuel [46,47]. The equivalence ratio is varied to cover both lean and rich conditions, and the mixing ratio is adjusted to reflect typical experimental and industrial applications [48,49]. 77 1-D transient laminar premixed flames are simulated using OpenFOAM [50]. All simulations utilized the Otomo mechanism [46], which contains 32 species and 213 reactions.

Each flame's computational domain is 20 mm long, discretized into 1000 uniform cells. The domain is initialized with a premixed fuel/air mixture (unburnt gas) on one side and equilibrium states (burnt gas) on the other. The inlet velocity is set to the laminar flame speed calculated by Cantera. A time step of 1×10^{-7} s is used, with each simulation running for 10,000 steps. Thermochemical states are sampled every 10 time steps.

2.3. Randomization method

Randomization is a necessary step to enhance model generalization and robustness. Firstly, high-dimensional reacting flow simulations used for *a-posteriori* validation may produce thermochemical states that deviate from the base manifold. Secondly, studies [20,25,26,29,37] have demonstrated that relying solely on simple, smooth base manifold data can lead to model overfitting, making the trained model susceptible to disturbances and unable to tolerate even 5% noise [25]. Thirdly, temperature, pressure, and species distributions are strongly correlated and exhibit multi-scale features. Leveraging these features can improve sampling efficiency. Therefore, we propose a stratified and continuous randomization strategy that matches the data complexity with the parameters of the neural network model.

As described in Section 2.2, initial states for randomization are sampled from 1-D laminar premixed flames, taking advantage of the strong correlations between temperature, pressure, and various components. Due to the uniform grid, many states are concentrated around the unburnt and burnt states, necessitating data balancing. To address this, we employ a stratified sampling method, selecting 100 reaction progress variables from a uniform distribution over the interval [0,1] for each laminar flame and each saved time step, resulting in 7.7 million initial states. The reaction progress variable c is defined as

$$c = \frac{T - T_u}{T_b - T_u}, \tag{7}$$

where T_u and T_b denote the unburnt and burnt temperatures, calculated according to the fuel composition.

Given that physical quantities exhibit multi-scale distributions, initial thermochemical state $[T_0, p_0, Y_{i,0}]$ is randomly perturbed to $[T', p', Y_i'']$ in different ways. For each initial state, the temperature is perturbed by injecting 10% noise,

$$T' = (1 + \alpha)T_0,\ \alpha \sim \text{Uniform}(-0.1,\ 0.1). \tag{8}$$

Pressure is sampled from a wide uniform distribution,

$$p' \sim \text{Uniform}(0.5,\ 2.0)\ bar. \tag{9}$$

The key developments above the previous randomization work [20,35–39] lie in the randomization of mass fractions and filtering methods.

Considering the significant changes in mass fraction magnitudes, in the HFRD method [20], mass fractions are randomized using an exponent that follows a uniform distribution,

$$\begin{aligned} &Y_i' = Y_{i,0}^b \text{ for } i = 1, \ldots, n_s, \\ &b \sim \text{Uniform}(1 - c,\ 1 + c), \\ &c = 0.2. \end{aligned} \tag{10}$$

Note that parameters b and c are presented in a form different from the original work [20] for ease of comparison. In subsequent studies, Li et al. [37,38] applied this randomization method 10 times and selected $c \in [0.1, 0.15]$. It is observed that this randomization approach generates an exponent following a log-normal-like distribution, which amplifies the magnitude differences between components and ensures sufficient coverage of the radical space. In this work, a larger number of randomization iterations is used to generate a smoother distribution while reducing the degree of randomness in each iteration. Specifically,

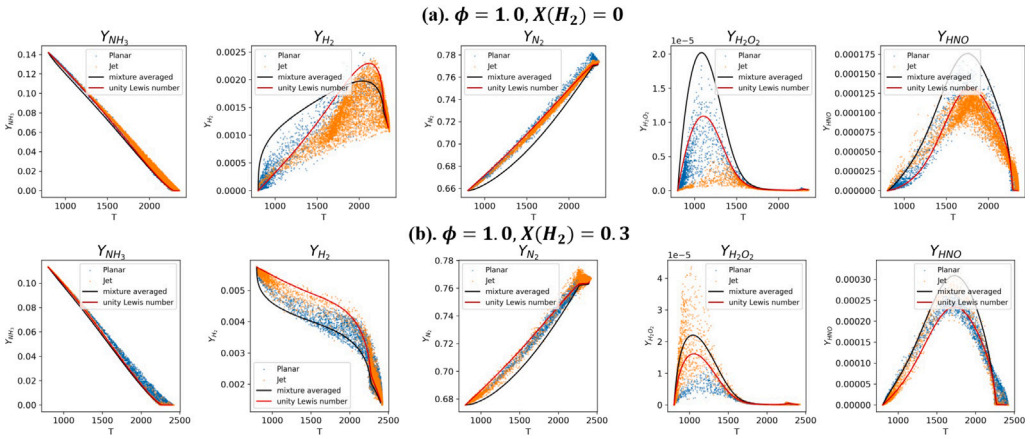

Fig. 3. Conditional profiles of selected mass fractions for (a) pure NH_3 and (b) 30% NH_3/H_2 flames. Blue and orange dots are sampled from the two turbulent flames for validation (Section 3), the black line is the 1-D laminar flame profile using a mixture averaged transport model, and red line is 1-D unity Lewis number profile. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the exponent is generated by multiplying 50 random numbers, each independently sampled from a uniform distribution,

$$\begin{aligned} Y_i' &= Y_{i,0}^{\beta} \text{ for } i = 1, \ldots, n_s, \\ \beta &= \prod_{k=1}^{50} \gamma_k, \\ \gamma_k &\sim \text{Uniform}(1-\xi,\ 1+\xi). \end{aligned} \tag{11}$$

The parameter ξ is crucial to controlling the spread of the data distribution continuously and is typically chosen within the range of 0.005 to 0.02. Finally, the species concentrations are normalized so that the species mass fractions sum to unity,

$$Y_i'' = \frac{Y_i'}{\sum_{i=1}^{n_s} Y_i'}. \tag{12}$$

Randomization based on mathematical formulas can generate numerous unphysical states, which reduce the model's accuracy in the regions of interest. To filter out these unphysical states, the HFRD method [20] discards any data outside the flammability limits and ensures that the generated thermochemical states maintain appropriate molar element ratios. In the context of premixed combustion, Li et al. [37,38] proposed a criterion based on heat release rate change to remove the unphysical perturbed states. In their study, perturbed states with HRR values exceeding 100 times above or below the original HRR values are discarded. In this work, a physical constraint based on both heat release rate and net production rate is proposed as follows,

$$\begin{aligned} &|\dot{Q}'| < max(|2 \times \dot{Q}|, 10^8), \\ &\sum_{i=1}^{n_s} |\omega_i' - \omega_i| < 5 \times \sum_{i=1}^{n_s} |\omega_i|, \end{aligned} \tag{13}$$

where $\dot{Q}$ and ω_i are calculated from the initial state $[T_0, p_0, Y_{i,0}]$, while $\dot{Q}'$ and ω_i' are calculated from the perturbed state $[T', p', Y_i'']$. The randomization process continues until a state that satisfies the physical constraint is found. Based on this data perturbation and filtering process, the training dataset of 7.7 million perturbed states is sampled from 7.7 million initial states. Each perturbed state advanced a variable time step to get the corresponding labeled state by solving the ODE system in Eq. (4) using CVODE in Cantera. The time steps are randomly chosen from 10^{-8} to 10^{-6} second for a wider application range.

2.4. Training details

The goal of the DNN is the same as that of stiff ODE solvers, i.e., predicting the change in mass fraction $[Y_i(t+\Delta t) - Y_i(t)]$ from current thermochemical state $[T(t), p(t), Y_i(t)]$ and given time step Δt. For mass fractions, Box–Cox transformation (BCT) [51] defined as

$$\mathcal{F}(x) = \frac{x^{\lambda} - 1}{\lambda},\ \lambda \neq 0, \tag{14}$$

is applied to emphasize the multi-scale distribution of minor components and avoid the singularity of the log transformation [25]. In this work, we use $\lambda = 0.1$. Thus, the DNN's input is $[T(t), p(t), \mathcal{F}(Y_i(t)), \Delta t]$ and the output is $[\mathcal{F}(Y_i(t+\Delta t)) - \mathcal{F}(Y_i(t))]$. Besides, all the input and output variables are transformed using Z-score normalization [52].

The DNN architecture comprises 4 hidden layers, each with 800 neurons, and uses the Gaussian Error Linear Unit (GELU) [53] as the activation function. Additionally, a multi-objective loss function is proposed, incorporating accuracy, mass conservation, and heat release rate similarity constraints. Let $r_i^{target} = \mathcal{F}(Y_i(t+\Delta t)) - \mathcal{F}(Y_i(t))$ and r_i^{pred} denote the targeted and predicted mass fraction variations before inverse transformation, calculated using CVODE or DNN chemistry ODE solver. Similarly, $Y_i^{target} = Y_i(t+\Delta t)$ and Y_i^{pred} are the targeted and predicted mass fractions, $T^{target} = T(t+\Delta t)$ and $p^{target} = p(t+\Delta t)$ are the temperature and pressure given by the stiff ODE solver. The loss function can be expressed as

$$\begin{aligned} Loss &= Loss_1 + Loss_2 + \eta Loss_3, \\ Loss_1 &= \sum_{i=1}^{n_s} |r_i^{pred} - r_i^{target}|, \\ Loss_2 &= |\sum_{i=1}^{n_s} Y_i^{pred} - 1|, \\ Loss_3 &= |\mathcal{H}(T^{target}, p^{target}, Y_i^{pred}) - \mathcal{H}(T^{target}, p^{target}, Y_i^{target})|. \end{aligned} \tag{15}$$

$Loss_1$ ensures accuracy in mass fractions, $Loss_2$ enforces that the mass fractions sum to unity, and $Loss_3$ ensures accuracy in heat release rates. $\mathcal{H}$ maps the thermochemical state to the heat release rate. Due to the complexity and non-differentiability of $\mathcal{H}$, a pre-trained and differentiable neural network, DNN_2, is used for parameterization. During the training of the chemical acceleration network, DNN_1, the parameters of DNN_2 are frozen and remain unchanged, serving solely as a well-fitted function. Specifically, DNN_1 first predicts the thermochemical state for the next time step. The HRRs are then estimated using the frozen DNN_2 for both the target and predicted thermochemical states. These two HRRs are required to be as close as possible, and their discrepancy is incorporated as an additional term in the loss function when optimizing DNN_1.

DNN_2 has the same hidden layers and neurons as DNN_1, with its output neurons adjusted from n_s to 1. Testing confirms that this architecture accurately predicts the HRR. Due to the typically large values of HRR discrepancy, the trade-off hyperparameter η is set to

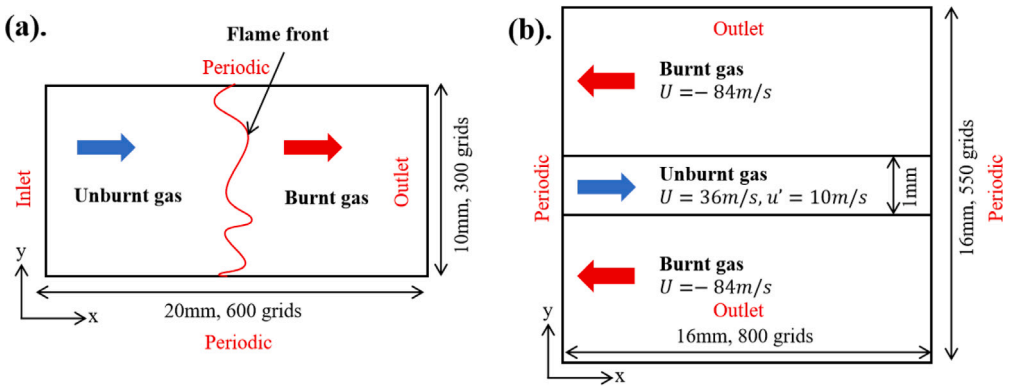

Fig. 4. Computational domain, initial conditions, and boundary conditions for (a) planar turbulent premixed flame and (b) temporally-evolving jet flame.

10^{-10} to balance the different components of the loss function and ensure they are on the same order of magnitude.

Optimization is performed for 1500 epochs using the Adam optimizer [54], with a batch size of 20000 and the initial learning rate of 10^{-3} gradually reduced to 10^{-6} using the cosine annealing learning rate scheduler [55]. The training process for the whole dataset requires about 1.5 h on one NVIDIA RTX 4090 GPU.

3. Validation configurations

Validating the pre-trained DNN using completely unseen high-dimensional configurations is vital in checking the model's generalization capability. In this work, we adopt several DNS cases to avoid the uncertainties due to turbulent and combustion models. Cantera [56], pybind11 [57], and pytorch [58] have been integrated into OpenFOAM [50] to enable a detailed transport model (mixture averaged transport model) and the call of user-defined neural networks within the turbulent reacting flow simulation framework, as recommended by the DeepFlame software [6].

3.1. Configuration of planar turbulent premixed flames

First, planar turbulent premixed flames are simulated using CVODE and DNN to validate model performance in intensive turbulence and across a wide range of equivalence ratios and mixing ratios for dual-fuels.

Fig. 4(a) illustrates the computational domain, initial conditions, and boundary conditions for simulating planar premixed NH_3 and NH_3/H_2 turbulent flames. The computational domain is $L_x \times L_y = 2L \times L$ in the streamwise direction x and transverse direction y, respectively, where L is 10 mm. The simulation is initialized with the corresponding laminar premixed flame solution and a homogeneous isotropic turbulence field based on a prescribed von-Karman Pao energy spectrum [59]. The inlet velocity is chosen to keep the whole flame structure statistically stationary in the computational domain. The inflow temperature and mass fraction boundary conditions are chosen as the corresponding unburnt state, and the outflow is non-reflecting (waveTransmissive in OpenFOAM). Other boundary conditions are periodic. The domain is discretized by a uniform grid of 600 × 300 for all the cases with a grid size Δx of 33.3 μm, which is smaller than the Kolmogorov length scale η, as shown in Table 2.

Three cases with different turbulent intensities are simulated, denoted as case L, case M, and case H, respectively, where "L", "M" and "H" refer to the level of turbulence intensity (low, medium, and high). The reactant is comprised of NH_3/air mixture with an equivalence ratio of 1.0 and a temperature of 800 K. Besides, 9 cases including three different equivalence ratios (0.5, 1.0, and 1.5) and three different mixing ratios of H_2 (0, 0.05 and 0.30) are also validated with medium turbulent intensity, denoted as case M1 ∼ case M9. The simulation parameters of these 11 DNS cases are listed in Table 2, where S_L is the laminar flame speed, δ_L is the laminar flame thickness, u' is the root-mean-square turbulent velocity fluctuation, and l_t is the turbulence integral length scale. The turbulent Reynolds number (Re) is defined as $Re = u' l_t / \nu$, and the Karlovitz number (Ka) is defined as $Ka^2 = ((u'/S_L)^3(\delta_L/l_t))$. The Kolmogorov length scale is defined by $\eta = l_t Re^{-3/4}$.

3.2. Configuration of temporally-evolving jet flames

Next, a more complex configuration of temporally-evolving jet flames of premixed NH_3 and NH_3/H_2 is considered to examine model generalization capability in turbulent flames with strong shear flows.

As shown in Fig. 4(b), the computational domain is a square with a side length of 16 mm, and the central 1 mm along the y-axis is initialized as unburnt gas, while the surrounding region is set as burnt gas. The average velocities of the unburnt and burnt gases are 36 m/s and −84 m/s, respectively. To promote instability and mixing in the boundary between burnt and unburnt gas, turbulent velocity fluctuations of 10 m/s are imposed on the unburnt gas using a prescribed von-Karman Pao energy spectrum [59] with a turbulence integral length scale of 1 mm. The boundary conditions are periodic in the streamwise direction (x) and non-reflecting in the transverse direction (y). The domain consists of 800 × 550 grid points. The grid is uniform in the x-direction and stretched at both ends in the y-direction. The minimum grid spacing used in the DNS grid is 20 μm in each direction, which is sufficient to reasonably resolve the smallest scales of the turbulent flow in the flame region.

Similarly, 9 cases are validated, including three different equivalence ratios (0.5, 1.0, and 1.5) and three different mixing ratios of H_2 (0, 0.05, and 0.30), and are denoted as case J1 ∼ case J9. The parameters of case J1 ∼ case J9 are the same as those of case M1 ∼ case M9, shown in Table 2. A similar configuration was also utilized by Satio et al. [28] to validate their DNN model for NH_3 combustion. Their model considers only 7 species with maximum mass fractions above 0.01 in the training data ($NH_3, O_2, H_2, OH, H_2O, N_2$, and NO). In contrast, DNN model in current work includes all species involved in detailed chemical kinetics. Additionally, Saito et al.'s validation involves a complex procedure to perform DNN model selection using an ODE solver, with the model applied only at temperatures above 1400 K. Our approach, however, utilizes the same DNN model consistently for all cases.

4. Results and discussion

4.1. Key hyperparameter selection

Given adequate training data to mitigate overfitting, increasing the number of parameters in a neural network can enhance its performance. However, for the purpose of accelerating combustion chemistry

Table 2
Parameters of the planar turbulent premixed flames.

Case	L	M(M2)	H	M1	M3	M4	M5	M6	M7	M8	M9
ϕ	1.0	1.0	1.0	0.5	1.5	0.5	1.0	1.5	0.5	1.0	1.5
$X(H_2)$	0	0	0	0	0	0.05	0.05	0.05	0.3	0.3	0.3
S_L [m/s]	0.56	0.56	0.56	0.21	0.41	0.24	0.62	0.46	0.54	1.2	0.85
δ_L [mm]	0.9	0.9	0.9	1.67	1.43	1.5	0.83	1.32	0.78	0.51	0.85
u' [m/s]	5	10	30	10	10	10	10	10	10	10	10
l_t [mm]	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
u'/S_L	8.9	17.9	53.6	48.6	24.6	41.9	16.1	22.0	18.4	8.3	11.8
l_t/δ_L	1.1	1.1	1.1	0.6	0.7	0.7	1.2	0.8	1.3	1.9	1.2
η [μm]	199	118	52	82	112	82	119	113	84	123	119
$\Delta x/\eta$	0.17	0.28	0.64	0.41	0.3	0.41	0.28	0.3	0.39	0.27	0.28
Re	9	17	52	28	19	28	17	18	27	16	17
Ka	25	72	373	437	146	332	59	118	70	17	37

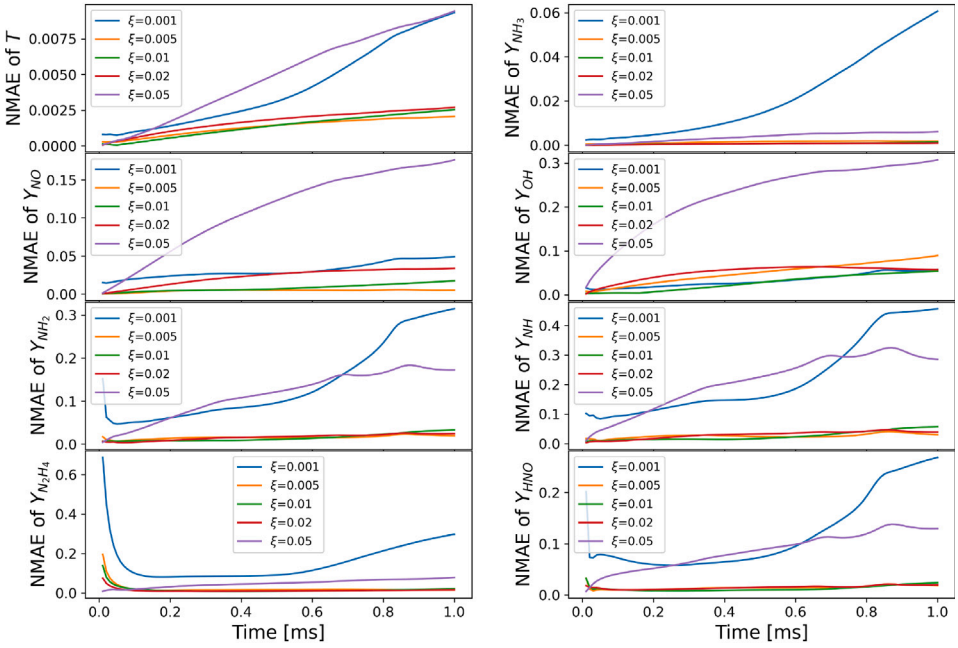

Fig. 5. NMAE of T, NH_3, NO, OH, NH_2, NH, N_2H_4, and HNO versus simulated time using the same DNN model (4 × 800) trained with datasets of different complexity (ξ = 0.001, 0.005, 0.01, 0.02, 0.05).

simulations, the parameter count must be carefully managed to maintain computational efficiency. This work aims to achieve several-fold acceleration on the same hardware by adapting the parameter count to the complexity of the data distribution.

To assess the impact of key hyperparameters that control data and model complexity on both prediction accuracy and computational efficiency, we analyze the error evolution curves for Case M, a NH_3/air premixed flame with an equivalence ratio of 1.0. Model performance is quantified by the normalized mean absolute error (NMAE) across the computational domain, defined as

$$\text{NMAE}_{u,t} = \frac{\sum_{k=1}^{N} \left| u_k^{target} - u_k^{pred} \right|}{\sum_{k=1}^{N} u_k^{target}}, \tag{16}$$

where the subscript u refers to any calculated field (T, p, or Y_i), and t denotes the tth time step. u_k^{target} and u_k^{pred} represent the targeted and predicted value at the kth cell in the computational domain, calculated using CFD code coupled with CVODE or DNN chemistry ODE solver. N is the total grid number.

4.1.1. Data complexity

Randomization techniques are essential for enhancing model generalization, but they can increase data distribution complexity, making neural network fitting more challenging. In this work, we propose a parameter ξ to control the spread of the data distribution continuously. To evaluate the performance of a neural network model across training data of varying complexity, ξ is set to 0.001, 0.005, 0.01, 0.02, or 0.05 for comparison. We employ a DNN model with four hidden layers, each containing 800 neurons in this analysis.

Fig. 5 presents the NMAE for T, NH_3, NO, OH, NH_2, NH, N_2H_4, and HNO as a function of simulated time. These results compare the performance of the same DNN model (4 × 800) trained on datasets with varying complexities (ξ = 0.001, 0.005, 0.01, 0.02, 0.05) against the CVODE solver. The model demonstrates rapid performance improvement for moderate dataset complexities (ξ = 0.005 to 0.02), neither too complex (ξ = 0.05) nor too simple (ξ = 0.001). The NMAE of Y_{OH} decreased from 30% to 6% and the NMAE of Y_{NO} decreased from 15% to 3% compared with the worst case with highest dataset complexity. Additionally, the larger NMAE in OH compared to NO suggests better predictive performance for ammonia chemistry than for hydrogen chemistry. This observation aligns with the statistical results presented later (Figs. 10 and 11). To enhance the accuracy of OH predictions, future work could explore a component-weighted accuracy loss. Consequently, we select ξ = 0.01 for generating training dataset in all subsequent calculations.

4.1.2. Model complexity

Ideally, we aim to develop a DNN model with as few parameters as possible to decrease computation time of chemistry evaluation.

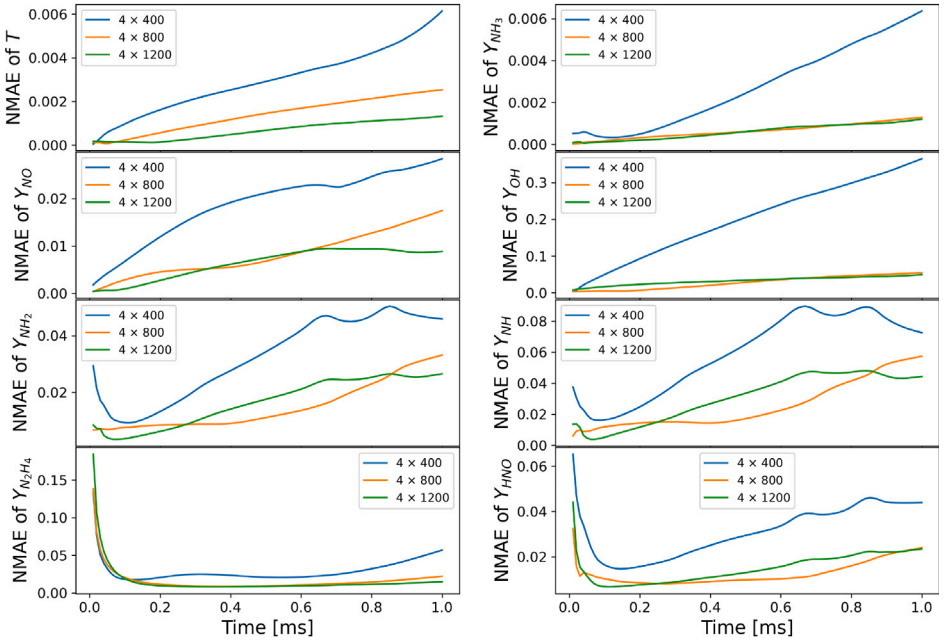

Fig. 6. NMAE of T, NH_3, NO, OH, NH_2, NH, N_2H_4, and HNO versus simulated time using DNN models of different complexity (4 × 400, 4 × 800, 4 × 1200) trained with the same dataset ($\xi = 0.01$).

Table 3
The performance comparison between DNNs and CVODE (baseline) on architectures w/o GPU on case M using 8 processors. The results represent a step-wise average of the chemical and total times.

	Chemical (s)	Total (s)	Chemical speed up	Total speed up
CVODE (baseline, 8 CPUs)	11.51	12.47	1.0	1.0
DNN (4 × 400, 8 CPUs)	0.85	1.76	13.5	7.1
DNN (4 × 800, 8 CPUs)	1.63	2.55	7.0	4.9
DNN (4 × 1200, 8 CPUs)	4.37	5.29	2.6	2.3
DNN (4 × 400, 8 CPUs + 1 GPU)	0.34	1.24	33.8	10.0
DNN (4 × 800, 8 CPUs + 1 GPU)	0.38	1.28	30.3	9.7
DNN (4 × 1200, 8 CPUs + 1 GPU)	0.44	1.35	26.1	9.2

However, the data complexity is considerable since the sampling space encompasses various conditions (temperature, pressure, equivalence ratios, and mixing ratios of dual-fuels). An interesting finding is that, after a certain threshold, further reducing the number of model parameters may lead to a significant drop in prediction accuracy. We evaluate models with four hidden layers, each containing 400, 800, and 1200 neurons, respectively. The prediction accuracy and inference time for various configurations are shown in Fig. 6 and Table 3. As illustrated, models with 800 and 1200 neurons achieved similar accuracies, both significantly outperforming the model with 400 neurons. Too few neurons can lead to incomplete coverage of thermochemical space, resulting in significant deviations of the model in certain states.

Furthermore, as the number of neurons decreases from 1200 to 800 to 400, the chemical and overall speed-up on the CPU increases linearly. On the GPU, there are fewer disparities in chemistry and overall speed-up. Based on the trade-off between speed-up and accuracy, we select a model architecture with four hidden layers, each with 800 neurons, as the optimized model. This model achieves a chemical calculation acceleration of 7 times and an overall simulation acceleration of 5 times using a model with 4 hidden layers and 800 neurons on the same CPU device. When a GPU is adopted, the chemical calculation acceleration increases to 30 times, and the overall simulation acceleration reaches 10 times.

Although direct integration of chemistry on GPUs achieves a speedup comparable to DNN inference on GPUs [60–62], our method significantly accelerates combustion simulations on the same CPU device. This advantage is particularly beneficial for supercomputers where expensive GPUs are not available.

4.2. Planar turbulent premixed flames

4.2.1. Comparisons of scalar distributions

Fig. 7 visualizes the instantaneous snapshots of heat release rate, temperature, and species mass fraction from CVODE, DNN, and absolute error for Case M and H, with Ka of 72 and 373, respectively. Case M belongs to the thin reaction zone regime, where the smallest eddy cannot disrupt the inner layer. In contrast, Case H, with a Ka exceeding 100, features small-scale vortices to disrupt the inner reaction zone. Here, the turbulent flame is expected to be in the distributed reaction zone [63]. In Case H, high turbulence intensity leads to unexpected gas composition, resulting in more accumulated errors in the DNN case. The absolute error for temperature and most radicals in Case H is an order of magnitude higher than in Case M at 1.0 ms (after 10,000 time steps using DNN-based calculations). For example, the temperature error is consistently less than 80 K in Case M, while in Case H, there are substantial pockets with errors reaching up to 750 K. This discrepancy is due to the flame position deviation and the rapid accumulation of errors under high turbulence intensity. However, the statistical results remain accurate for both Case M and Case H, as presented later (Figs. 10 and 11).

Fig. 8 compares the instantaneous snapshots for different equivalence ratios with 30% blending of hydrogen. The NO concentration accumulated in the rich flame is significantly lower than that of lean or stoichiometric flames. Meanwhile, the concentrations of NH, NH_2, and N_2H_4 radicals in the reaction zone of the rich flame are notably higher compared to lean or stoichiometric flames. Additionally, the

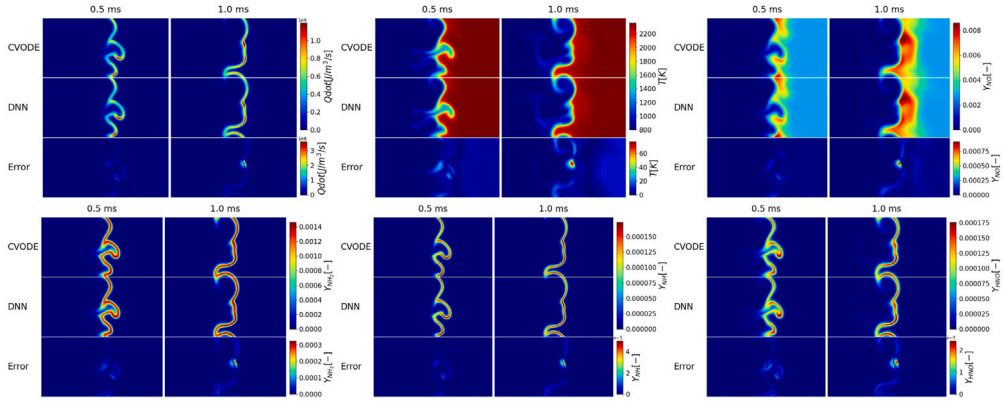

(a) Case M, $\phi = 1.0, X(H_2) = 0, u' = 10\text{m/s}$

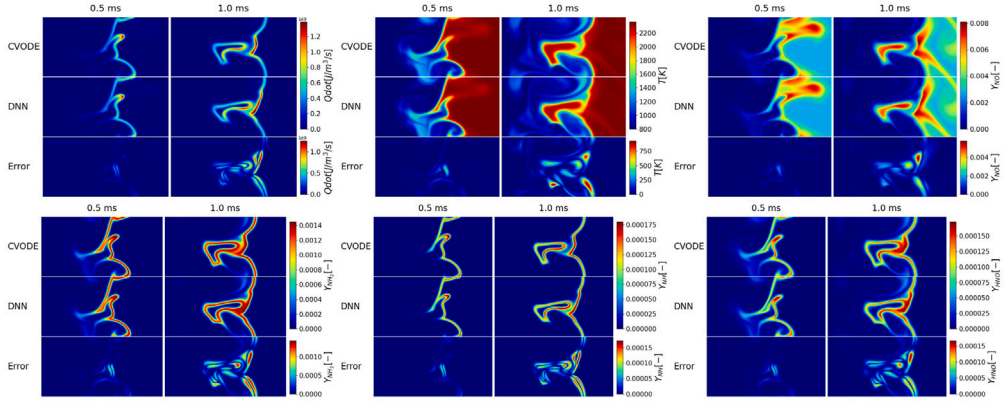

(b) Case H, $\phi = 1.0, X(H_2) = 0, u' = 30\text{m/s}$

Fig. 7. Comparisons of the distributions of heat release rate, temperature, and species mass fraction from CVODE, DNN, and absolute error for Case M and H (pure NH_3 combustion with different turbulent intensities).

lean flames are more wrinkled and stretched than the rich flames, which could due to the thermal-diffusive instability caused by hydrogen. The DNN accurately reproduces these phenomena, demonstrating its ability to capture the ODE behaviors over a wide range of conditions with high fidelity. Overall, the DNN model effectively predicts scalar distributions across various equivalence ratios and H_2 blendings, even under extreme turbulent conditions.

4.2.2. Turbulent flame speed comparison

In this article, we apply the evolution curve of turbulent flame speed S_T, which is one of the most important global parameters for characterizing turbulent premixed flames, as a comprehensive metric for comparing CVODE and DNN, defined as

$$S_T = \frac{1}{A}\int_V \frac{\dot{Q}}{\rho c_p \left(T_b - T_u\right)} dV, \tag{17}$$

where $\dot{Q}$ is the heat release rate, ρ is the density, c_p is the specific heat capacity at constant pressure, T_b and T_u are the burnt and unburnt gas temperature respectively, and A is the equivalent flame front area.

Fig. 9 compares various simulation parameters, including different equivalence ratios, mixing ratios, time steps, and unburnt gas temperatures, to demonstrate the applicability of the trained DNN model.

Subfigures (a), (b), and (c) illustrate the DNN's performance across different equivalence ratios and mixing ratios. The DNN accurately predicts both lean (a) and rich (c) combustion scenarios, capturing variations in turbulent flame speed from pure NH_3 to mixtures with up to 30% H_2. Despite the significant chemical property differences in dual-fuel mixtures, the DNN effectively handles a broad range of conditions.

Subfigure (d) demonstrates the DNN's accuracy under varying turbulence intensities, characterized by Ka of 25, 72, and 373. The randomized data distribution, based on a unity Lewis number manifold, closely resembles distributions of intense turbulence, enabling the DNN to predict chemical reactions accurately under high turbulence intensity.

Subfigure (e) displays the DNN's capability to handle variable time steps, accurately adapting to 1x, 5x, and 10x the original time step. This flexibility eliminates the need to train the model for each time step, making it suitable for multi-scale simulations, such as large eddy simulation (LES) and DNS, and for scenarios with variable time steps while maintaining a fixed Courant number.

Finally, Subfigure (f) highlights the DNN's accuracy across different initial unburnt gas temperatures, T_u. Although the model was trained on data sampled from a base manifold at $T_u = 800\,\text{K}$ for 1D flames, the incorporation of randomized temperature and pressure in the training data enables the DNN to accurately predict turbulent flames with T_u values of 700 K and 900 K.

4.2.3. Comparisons of statistical results

This subsection performs statistical analyses to examine the model performance further. Fig. 10 presents the profiles of the mean streamwise temperature T and mass fractions Y_{NH_3}, Y_{NO}, Y_{OH}, Y_{NH_2}, Y_{NH}, $Y_{N_2H_4}$ and Y_{HNO}, for Case L, M, H, M7, M8, and M9. The curves of these parameters demonstrate good quantitative agreements between DNN and CVODE across low-to-high turbulence intensity, lean and rich combustion, and different fuel compositions.

To further compare the flame structure between CVODE-based calculation and DNN-based calculation, Fig. 11 shows conditionally averaged mass fractions Y_{NH_3}, Y_{NO}, Y_{OH}, Y_{NH_2}, Y_{NH}, $Y_{N_2H_4}$ and Y_{HNO} at

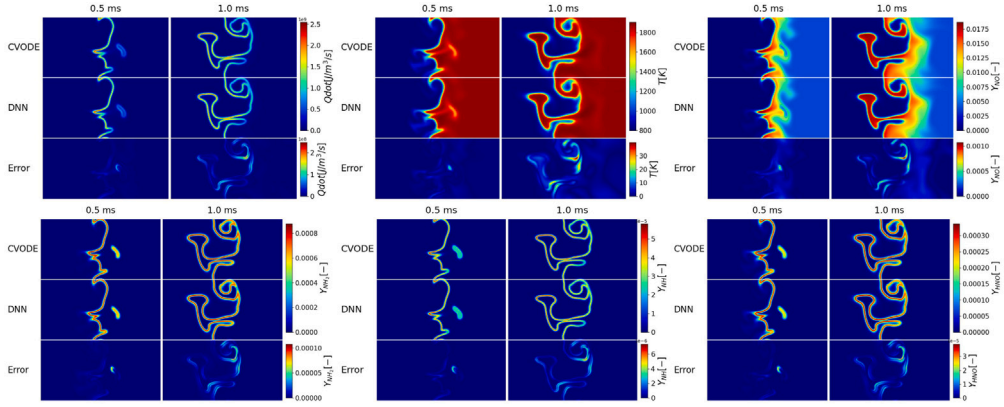

(a) Case M7, $\phi = 0.5, X(H_2) = 0.3, u' = 10m/s$

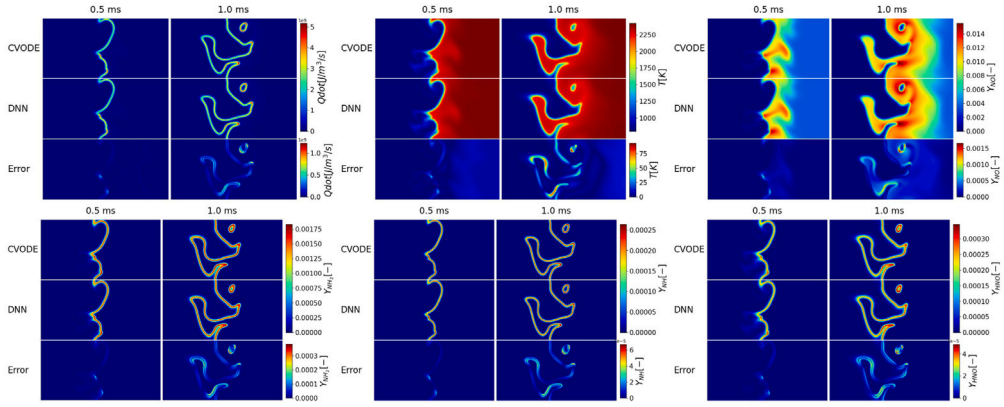

(b) Case M8, $\phi = 1.0, X(H_2) = 0.3, u' = 10m/s$

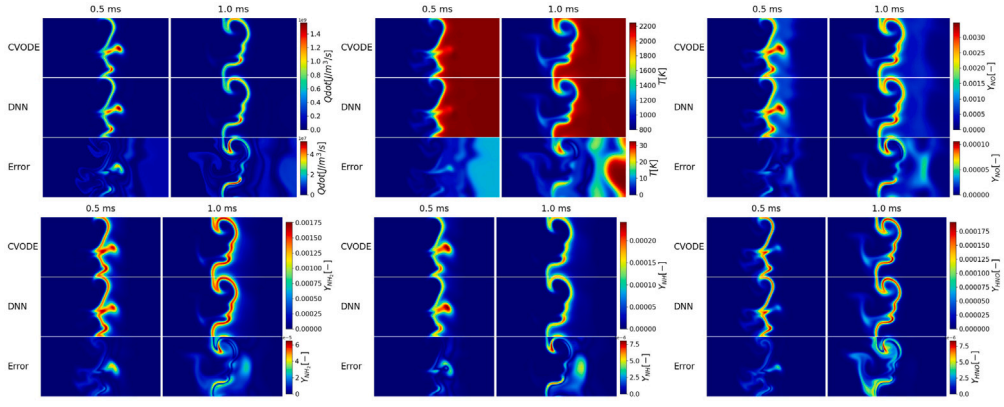

(c) Case M9, $\phi = 1.5, X(H_2) = 0.3, u' = 10m/s$

Fig. 8. Comparisons of the distributions of heat release rate, temperature, and species mass fraction from CVODE, DNN, and absolute error for Case M7, M8, and M9 (NH_3/H_2 combustion with different equivalence ratios).

$t = 1.0$ ms for Case L, M, H, M7, M8, and M9. Conditionally averaged Y_{NH_2} and Y_{NH} reach their peaks at relatively high temperatures, while $Y_{N_2H_4}$ peaks at relatively low temperatures. This indicates that NH_3 gradually decomposes into NH_2, and then into NH, while the rate of the addition reactions of NH_i, such as $NH + NH + M \rightleftarrows N_2H_2 + M$, and $NH_2 + NH_2 + M \rightleftarrows N_2H_4 + M$, are enhanced with relatively low temperatures. All the conditional profiles given by the DNN agree well with the CVODE calculations. The only notable discrepancy is in the species OH at high temperatures in Cases M8 and M9, which feature high H_2 blending. The accurate predictions for N-containing radicals by the neural network suggest superior prediction for NH_3 chemistry compared to H_2 chemistry. This may be due to insufficient radical concentration coverage for H_2 chemistry in the dataset, highlighting the need for more balanced data distribution in future studies.

4.3. Temporally-evolving jet flames

4.3.1. Comparisons of scalar distributions

Fig. 12 illustrates the snapshot comparisons of pure NH_3 (Case J2) and 30% H_2 blending (Case J8). In Case J8, both the heat release rate

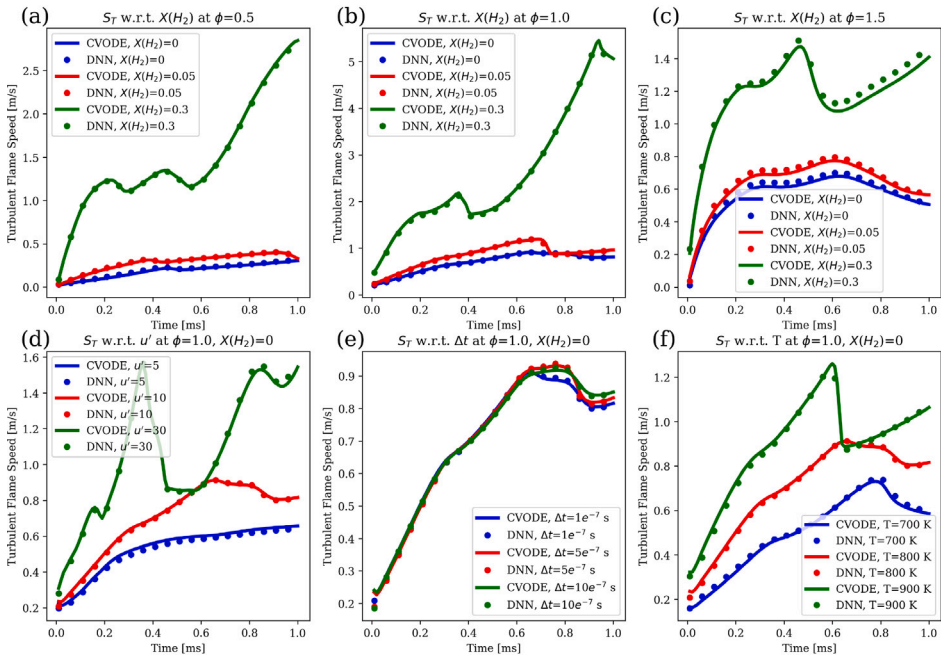

Fig. 9. Turbulent flame velocity of planar premixed NH_3/H_2 turbulent flames versus simulated time (lines from CVODE, symbols from DNN). Subfigures (a), (b), and (c) show validations under various equivalence ratios ϕ and mixing ratios $X(H_2)$. Subfigure (d) shows validations for different turbulence intensities u', subfigure (e) for different time steps Δt, and subfigure (f) for different unburnt gas temperatures T_u.

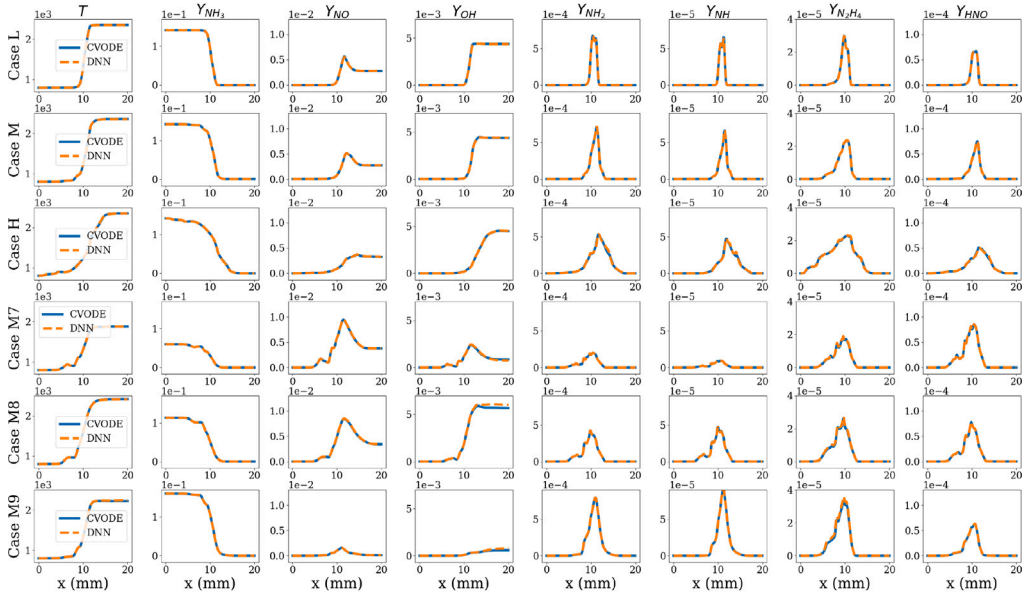

Fig. 10. The mean profile of streamwise temperature T and mass fractions Y_{NH_3}, Y_{NO}, Y_{OH}, Y_{NH_2}, Y_{NH}, $Y_{N_2H_4}$ and Y_{HNO}, for Case L, M, H, M7, M8, and M9.

and temperature are significantly higher, demonstrating the enhanced reactivity due to H_2 addition. The NO concentration also doubles, and the marked differences in N-containing radicals indicate distinct chemical reaction pathways.

After 0.1 ms, the mixing of unburnt and burnt gases intensifies, suggesting that the advection and diffusion of chemical species become more pronounced. Despite the intense mixing and broader gas composition, the accumulated error remains minimal, indicating that the effective randomization strategy has covered a sufficiently broad thermochemical space. The DNN provides highly accurate predictions for key distributions such as heat release rate, temperature, NO, and various intermediate radicals, demonstrating its ability to accurately simulate strong shearing turbulent flames across a wide range of operating conditions.

4.3.2. Turbulent flame speed comparison

Fig. 13 compares turbulent flame velocity of temporally-evolving jet flames versus simulated time (lines from CVODE, symbols from DNN) across various equivalence ratios ϕ and mixing ratios $X(H_2)$. Unlike statistically steady flames, temporally-evolving jet flames capture the changes in flow and composition conditions as the observation window shifts from upstream to downstream along the actual jet flame. Each subplot illustrates the rapid increase in fuel consumption rate with higher H_2 volume percentage, emphasizing the significantly different chemical properties of NH_3 and H_2. The temporal evolution of turbulent flame velocity exhibits similar trends between DNN-based and CVODE-based calculations, indicating that the DNN model accurately captures these variations across a wide range of equivalence and mixing

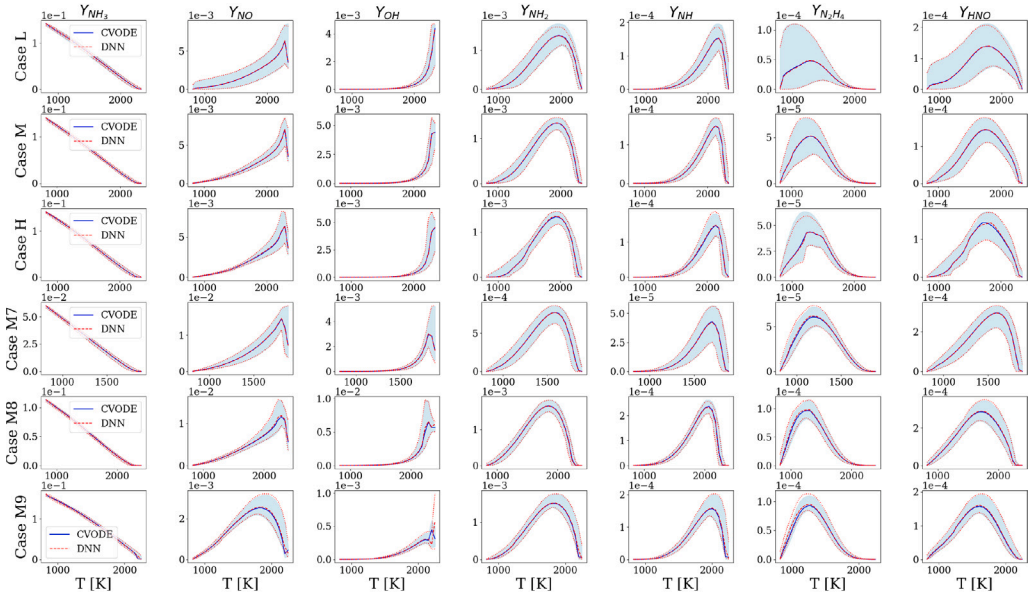

Fig. 11. Conditionally averaged mass fractions Y_{NH_3}, Y_{NO}, Y_{OH}, Y_{NH_2}, Y_{NH}, $Y_{N_2H_4}$ and Y_{HNO} at $t = 1.0$ ms for Case L, M, H, M7, M8, and M9. Blue solid lines and red dashed lines represent conditionally averaged values calculated using CVODE and the DNN, respectively. Blue shaded regions and red dotted lines denote the mass fraction bounds in the temperature interval.

ratios. This consistency further demonstrates the DNN model's broad adaptability in turbulent flames with strong shear flows.

4.3.3. Comparisons of statistical results

Similarly, the transverse mean profiles (Fig. 14) and conditionally averaged profiles (Fig. 15) are plotted for Case J1, J2, J3, J7, J8, and J9 at $t = 0.1$ ms. The results indicate that the DNN model reproduces the overall flame structure with high fidelity.

For the transverse mean profiles (Fig. 14), despite significant differences in radical distributions across turbulent flames with varying fuel compositions and equivalence ratios, the DNN achieves good agreement in temperature and species profiles for all cases. In the conditionally averaged profiles (Fig. 15), slight deviations are observed only for Y_{NH_2}, Y_{NH}, and Y_{HNO} in Case J1. Given that Case J1 is near the boundary in the training set (extremely lean, pure NH_3), and that the mass fractions of these radicals are an order of magnitude lower compared to other cases, this level of prediction accuracy is satisfactory.

5. Conclusion and further scope

To address the high computational costs associated with using detailed chemical kinetics in combustion simulations, we develop a DNN model with variable time steps to enhance the efficiency of combustion chemistry simulations for turbulent reacting flows. Using NH_3/H_2 as an example of complex dual-fuels with significantly different chemical properties, we improved the "sampling-training" workflow based on previous HFRD method, which includes calculating the base manifold, performing randomization, and conducting training. Unity Lewis number laminar flames are chosen as the base manifold because turbulent diffusion overwhelms molecular diffusion under large turbulent intensities. To enhance randomization, we introduce a method to continuously control data complexity and a filtering criterion for precise coverage of the thermochemical space in turbulent flames. The model training process is carefully designed by adding mass fraction conservation and heat release rate constraints to the loss function.

We simulate planar turbulent premixed flames with high turbulence intensity and temporally-evolving jet flames with strong shear using both CVODE and the DNN model. Two key hyperparameters are optimized to balance data coverage and model complexity, highlighting the trade-off between data coverage and computational efficiency. The optimized DNN accurately reproduces scalar distributions and statistical results, even for radicals, across a diverse range of unburnt temperatures, pressures, time steps, equivalence ratios, and mixing ratios. Our approach achieves a chemical calculation acceleration of 7 times and an overall simulation acceleration of 5 times using a model with 4 hidden layers and 800 neurons on the same CPU device. When a GPU is adopted, the chemical calculation acceleration increases to 30 times, and the overall simulation acceleration reaches 10 times.

The complexity of DNN models is constrained by the need to accelerate combustion chemistry, especially when dealing with complex dual-fuels with significantly different chemical properties. Training a universal model for the entire thermochemical space with limited parameters remains challenging. Future research needs to explore more representative training state distributions, and develop more efficient model architectures and training processes.

CRediT authorship contribution statement

Sipei Wu: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation. **Wenkai Liang:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Funding acquisition, Formal analysis, Conceptualization. **Kai Hong Luo:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

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

Acknowledgments

The authors gratefully acknowledge the support from the National Natural Science Foundation of China (Grant No. 52250710681 and 52406144). Support from the UK Engineering and Physical Sciences Research Council under the project "UK Consortium on Mesoscale Engineering Sciences (UKCOMES)" (Grant No. EP/X035875/1) is also acknowledged. This work made use of computational support by CoSeC, the Computational Science Centre for Research Communities, through UKCOMES.

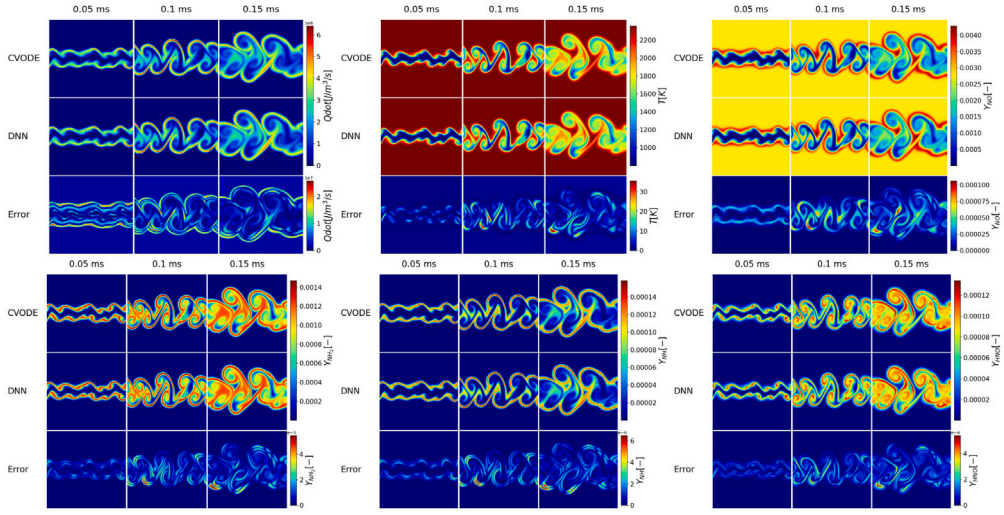

(a) Case J2, $\phi = 1.0, X(\mathrm{H_2}) = 0$

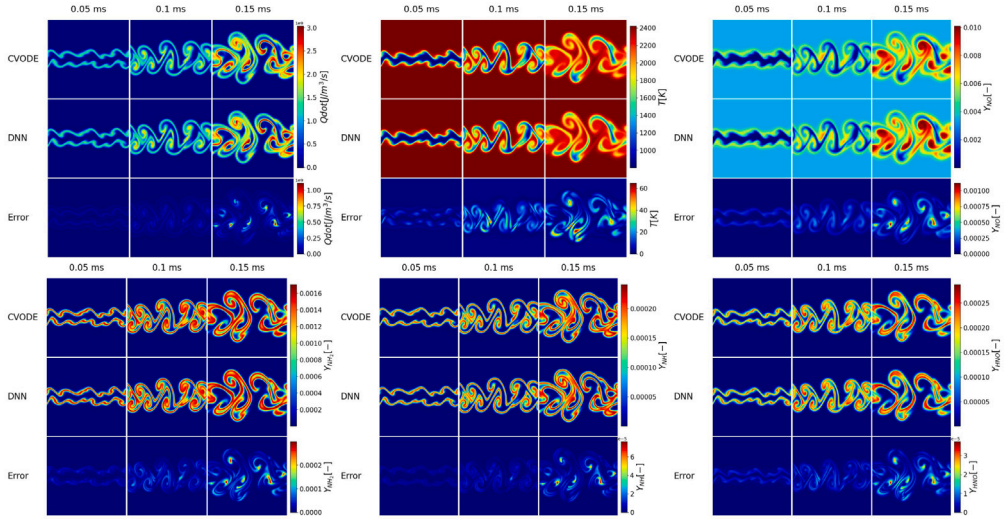

(b) Case J8, $\phi = 1.0, X(\mathrm{H_2}) = 0.3$

Fig. 12. Comparisons of the distributions of heat release rate, temperature, and species mass fraction from CVODE, DNN, and absolute error for Case J2 and J8 (NH_3 and NH_3/H_2 combustion with stoichiometric equivalent ratio).

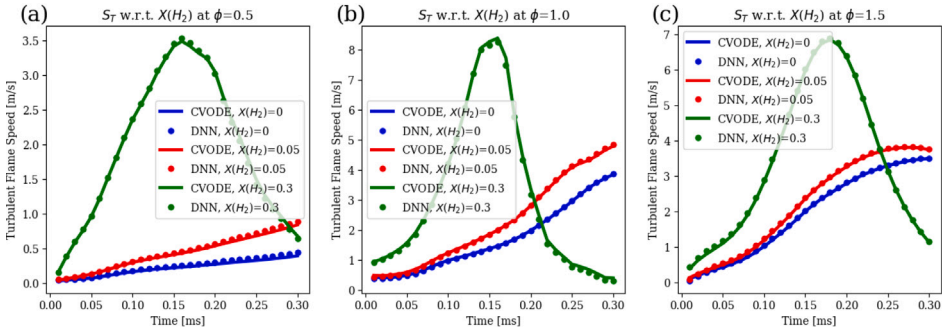

Fig. 13. Turbulent flame velocity of temporally-evolving jet flames versus simulated time (lines from CVODE, symbols from DNN). Subfigures (a), (b), and (c) show validations under various equivalence ratios ϕ and mixing ratios $X(\mathrm{H_2})$.

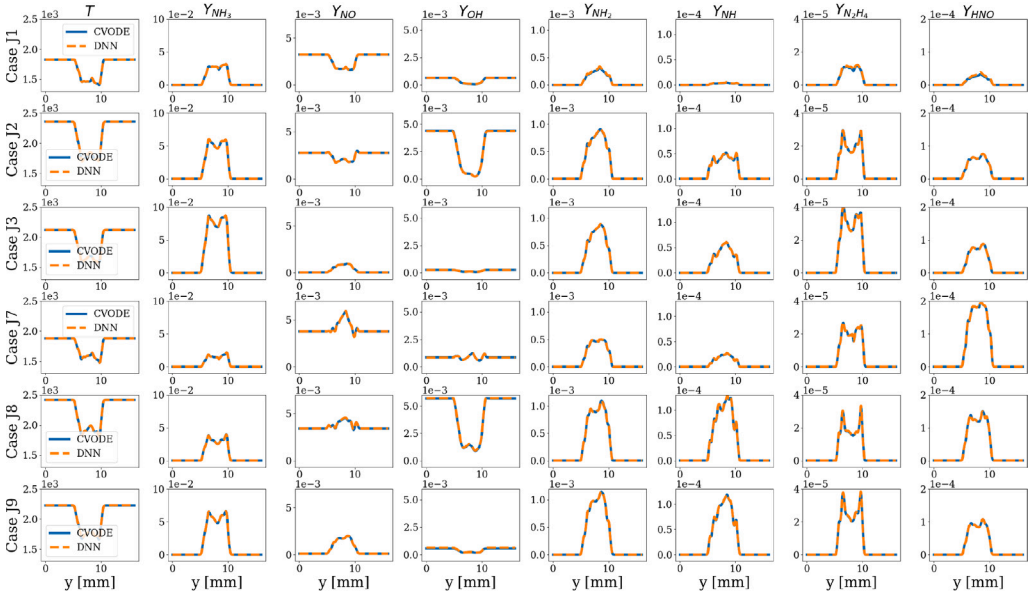

Fig. 14. The mean profile of transverse temperature T and mass fractions Y_{NH_3}, Y_{NO}, Y_{OH}, Y_{NH_2}, Y_{NH}, $Y_{N_2H_4}$ and Y_{HNO} at $t = 0.1$ ms for Case J1, J2, J3, J7, J8, and J9.

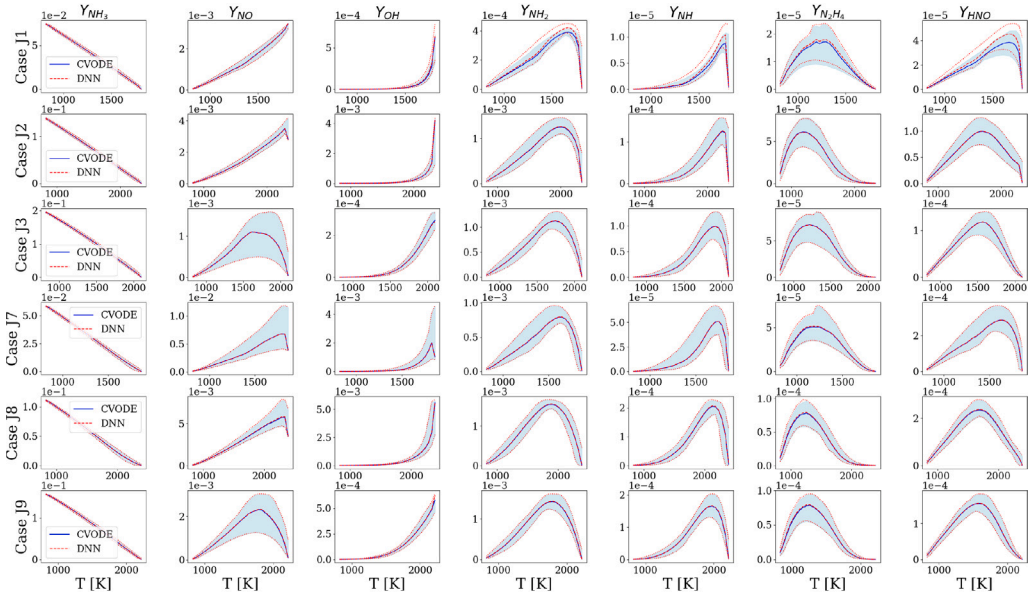

Fig. 15. Conditionally averaged mass fractions Y_{NH_3}, Y_{NO}, Y_{OH}, Y_{NH_2}, Y_{NH}, $Y_{N_2H_4}$ and Y_{HNO} at $t = 0.1$ ms for Case J1, J2, J3, J7, J8, and J9. Blue solid lines and red dashed lines represent conditionally averaged values calculated using CVODE and the DNN, respectively. Blue shaded regions and red dotted lines denote the mass fraction bounds in the temperature interval.

References

[1] S. Solomon, G.-K. Plattner, R. Knutti, P. Friedlingstein, Irreversible climate change due to carbon dioxide emissions, Proc. Natl. Acad. Sci. 106 (6) (2009) 1704–1709.

[2] W. Han, Z. Sun, A. Scholtissek, C. Hasse, Machine learning of ignition delay times under dual-fuel engine conditions, Fuel 288 (2021) 119650.

[3] H. Kobayashi, A. Hayakawa, K.K.A. Somarathne, E.C. Okafor, Science and technology of ammonia combustion, Proc. Combust. Inst. 37 (1) (2019) 109–133.

[4] R.C. da Rocha, M. Costa, X.-S. Bai, Chemical kinetic modelling of ammonia/hydrogen/air ignition, premixed flame propagation and NO emission, Fuel 246 (2019) 24–33.

[5] S. Lapointe, S. Mondal, R.A. Whitesides, Data-driven selection of stiff chemistry ODE solver in operator-splitting schemes, Combust. Flame (ISSN: 00102180) 220 (2020) 133–143.

[6] R. Mao, M. Lin, Y. Zhang, T. Zhang, Z.-Q.J. Xu, Z.X. Chen, DeepFlame: A deep learning empowered open-source platform for reacting flow simulations, Comput. Phys. Comm. (ISSN: 00104655) 291 (2023) 108842.

[7] L. Zhou, Y. Song, W. Ji, H. Wei, Machine learning for combustion, Energy AI (ISSN: 2666-5468) 7 (2022) 100128.

[8] M. Ihme, W.T. Chung, A.A. Mishra, Combustion machine learning: Principles, progress and prospects, Prog. Energy Combust. Sci. 91 (2022) 101010.

[9] A. Imren, D. Haworth, On the merits of extrapolation-based stiff ODE solvers for combustion CFD, Combust. Flame 174 (2016) 1–15.

[10] A. Hindmarsh, User Documentation for CVODES v2. 1.0, Tech. Rep., Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States), 2004.

[11] S. Kim, W. Ji, S. Deng, Y. Ma, C. Rackauckas, Stiff neural ordinary differential equations, Chaos: Interdiscip. J. Nonlinear Sci. 31 (9) (2021).

[12] W. Ji, W. Qiu, Z. Shi, S. Pan, S. Deng, Stiff-pinn: Physics-informed neural network for stiff chemical kinetics, J. Phys. Chem. A 125 (36) (2021) 8098–8106.

[13] M. De Florio, E. Schiassi, R. Furfaro, Physics-informed neural networks and functional interpolation for stiff chemical kinetics, Chaos: Interdiscip. J. Nonlinear Sci. 32 (6) (2022).

[14] F. Christo, A. Masri, E. Nebot, Artificial neural network implementation of chemistry with pdf simulation of H_2/CO_2 flames, Combust. Flame 106 (4) (1996) 406–427.

[15] J.A. Blasco, N. Fueyo, C. Dopazo, J. Chen, A self-organizing-map approach to chemistry representation in combustion applications, Combust. Theory Model. 4 (1) (2000) 61.

[16] A. Chatzopoulos, S. Rigopoulos, A chemistry tabulation approach via rate-controlled constrained equilibrium (RCCE) and artificial neural networks (ANNs), with application to turbulent non-premixed $CH_4/H_2/N_2$ flames, Proc. Combust. Inst. 34 (1) (2013) 1465–1473.

[17] L.L. Franke, A.K. Chatzopoulos, S. Rigopoulos, Tabulation of combustion chemistry via Artificial Neural Networks (ANNs): Methodology and application to LES-PDF simulation of Sydney flame L, Combust. Flame 185 (2017) 245–260.

[18] J. An, G. He, K. Luo, F. Qin, B. Liu, Artificial neural network based chemical mechanisms for computationally efficient modeling of hydrogen/carbon monoxide/kerosene combustion, Int. J. Hydrog. Energy 45 (53) (2020) 29594–29605.

[19] J. An, F. Qin, J. Zhang, Z. Ren, Explore artificial neural networks for solving complex hydrocarbon chemistry in turbulent reactive flows, Fundam. Res. 2 (4) (2022) 595–603.

[20] T. Ding, T. Readshaw, S. Rigopoulos, W. Jones, Machine learning tabulation of thermochemistry in turbulent combustion: An approach based on hybrid flamelet/random data and multiple multilayer perceptrons, Combust. Flame (ISSN: 00102180) 231 (2021) 111493.

[21] M. Hansinger, Y. Ge, M. Pfitzner, Deep residual networks for flamelet/progress variable tabulation with application to a piloted flame with inhomogeneous inlet, Combust. Sci. Technol. 194 (8) (2022) 1587–1613.

[22] S. Goswami, A.D. Jagtap, H. Babaee, B.T. Susi, G.E. Karniadakis, Learning stiff chemical kinetics using extended deep neural operators, Comput. Methods Appl. Mech. Engrg. (ISSN: 00457825) 419 (2024) 116674.

[23] A. Kumar, T. Echekki, Combustion chemistry acceleration with deepONets, Fuel 365 (2024) 131212.

[24] Y. Weng, H. Li, H. Zhang, Z.X. Chen, D. Zhou, Extended Fourier Neural Operators to learn stiff chemical kinetics under unseen conditions, Combust. Flame 272 (2025) 113847.

[25] T. Zhang, Y. Yi, Y. Xu, Z.X. Chen, Y. Zhang, W. E, Z.-Q.J. Xu, A multi-scale sampling method for accurate and robust deep neural network to predict combustion chemical kinetics, Combust. Flame (ISSN: 00102180) 245 (2022) 112319.

[26] Z.-Q.J. Xu, J. Yao, Y. Yi, L. Hang, W. E, Y. Zhang, T. Zhang, Solving multiscale dynamical systems by deep learning, 2024, URL http://arxiv.org/abs/2401.01220.

[27] S.B. Pope, Small scales, many species and the manifold challenges of turbulent combustion, Proc. Combust. Inst. 34 (1) (2013) 1–31.

[28] M. Saito, J. Xing, J. Nagao, R. Kurose, Data-driven simulation of ammonia combustion using neural ordinary differential equations (NODE), Appl. Energy Combust. Sci. (ISSN: 2666352X) 16 (2023) 100196.

[29] X. Han, M. Jia, Y. Chang, Y. Li, An improved approach towards more robust deep learning models for chemical kinetics, Combust. Flame (ISSN: 00102180) 238 (2022) 111934.

[30] B.A. Sen, S. Menon, Turbulent premixed flame modeling using artificial neural networks based chemical kinetics, Proc. Combust. Inst. 32 (1) (2009) 1605–1611.

[31] B.A. Sen, S. Menon, Linear eddy mixing based tabulation and artificial neural networks for large eddy simulations of turbulent flames, Combust. Flame 157 (1) (2010) 62–74.

[32] B.A. Sen, E.R. Hawkes, S. Menon, Large eddy simulation of extinction and reignition with artificial neural networks based chemical kinetics, Combust. Flame 157 (3) (2010) 566–578.

[33] K. Wan, C. Barnaud, L. Vervisch, P. Domingo, Chemistry reduction using machine learning trained from non-premixed micro-mixing modeling: Application to DNS of a syngas turbulent oxy-flame with side-wall effects, Combust. Flame 220 (2020) 119–129.

[34] H.-T. Nguyen, P. Domingo, L. Vervisch, P.-D. Nguyen, Machine learning for integrating combustion chemistry in numerical simulations, Energy AI 5 (2021) 100082.

[35] T. Readshaw, L.L. Franke, W. Jones, S. Rigopoulos, Simulation of turbulent premixed flames with machine learning-tabulated thermochemistry, Combust. Flame 258 (2023) 113058.

[36] T. Readshaw, T. Ding, S. Rigopoulos, W. Jones, Modeling of turbulent flames with the large eddy simulation–probability density function (LES–PDF) approach, stochastic fields, and artificial neural networks, Phys. Fluids 33 (3) (2021).

[37] H. Li, R. Yang, M. Zhang, R. Mao, Z.X. Chen, A comprehensive study on the accuracy and generalization of deep learning-generated chemical ODE integrators, 2023, URL http://arxiv.org/abs/2312.16387.

[38] H. Li, R. Yang, Y. Xu, M. Zhang, R. Mao, Z.X. Chen, Comprehensive deep learning for combustion chemistry integration: Multi-fuel generalization and a posteriori validation in reacting flow, Phys. Fluids 37 (1) (2025).

[39] T. Ding, S. Rigopoulos, W. Jones, Machine learning tabulation of thermochemistry of fuel blends, Appl. Energy Combust. Sci. 12 (2022) 100086.

[40] M.A. Singer, S.B. Pope, H.N. Najm, Operator-splitting with ISAT to model reacting flow with detailed chemistry, Combust. Theory Model. 10 (2) (2006) 199–217.

[41] Z. Ren, S.B. Pope, Second-order splitting schemes for a class of reactive systems, J. Comput. Phys. (ISSN: 00219991) 227 (17) (2008) 8165–8176.

[42] A. Aspden, M. Day, J. Bell, Lewis number effects in distributed flames, Proc. Combust. Inst. 33 (1) (2011) 1473–1480.

[43] B. Savard, B. Bobbitt, G. Blanquart, Structure of a high karlovitz n-C_7H_{16} premixed turbulent flame, Proc. Combust. Inst. 35 (2) (2015) 1377–1384.

[44] A. Aspden, M. Day, J. Bell, Towards the distributed burning regime in turbulent premixed flames, J. Fluid Mech. 871 (2019) 1–21.

[45] H. Lee, P. Dai, M. Wan, A. Lipatnikov, Influence of molecular transport on burning rate and conditioned species concentrations in highly turbulent premixed flames, J. Fluid Mech. 928 (2021) A5.

[46] J. Otomo, M. Koshi, T. Mitsumori, H. Iwasaki, K. Yamada, Chemical kinetic modeling of ammonia oxidation with improved reaction mechanism for ammonia/air and ammonia/hydrogen/air combustion, Int. J. Hydrog. Energy 43 (5) (2018) 3004–3014.

[47] A. Alnasif, S. Mashruk, H. Shi, M. Alnajideen, P. Wang, D. Pugh, A. Valera-Medina, Evolution of ammonia reaction mechanisms and modeling parameters: A review, Appl. Energy Combust. Sci. 15 (2023) 100175.

[48] O.I. Awad, B. Zhou, K. Harrath, K. Kadirgama, Characteristics of NH_3/H_2 blend as carbon-free fuels: A review, Int. J. Hydrog. Energy 48 (96) (2023) 38077–38100.

[49] C. Lhuillier, P. Brequigny, N. Lamoureux, F. Contino, C. Mounaïm-Rousselle, Experimental investigation on laminar burning velocities of ammonia/hydrogen/air mixtures at elevated temperatures, Fuel 263 (2020) 116653.

[50] H.G. Weller, G. Tabor, H. Jasak, C. Fureby, A tensorial approach to computational continuum mechanics using object-oriented techniques, Comput. Phys. 12 (6) (1998) 620–631.

[51] G.E. Box, D.R. Cox, An analysis of transformations, J. R. Stat. Soc. Ser. B Stat. Methodol. 26 (2) (1964) 211–243.

[52] S. Patro, Normalization: A preprocessing stage, 2015, arXiv preprint arXiv:1503.06462.

[53] D. Hendrycks, K. Gimpel, Gaussian error linear units (gelus), 2016, arXiv preprint arXiv:1606.08415.

[54] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, 2014, arXiv preprint arXiv:1412.6980.

[55] I. Loshchilov, F. Hutter, Sgdr: Stochastic gradient descent with warm restarts, 2016, arXiv preprint arXiv:1608.03983.

[56] D.G. Goodwin, Cantera c++ user's guide, Calif. Inst. Technol. 32 (2002) 358.

[57] W. Jakob, J. Rhinelander, D. Moldovan, Pybind11 – Seamless operability between C++11 and Python, 2017, https://github.com/pybind/pybind11.

[58] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al., Pytorch: An imperative style, high-performance deep learning library, Adv. Neural Inf. Process. Syst. 32 (2019).

[59] T. Saad, D. Cline, R. Stoll, J.C. Sutherland, Scalable tools for generating synthetic isotropic turbulence with arbitrary spectra, AIAA J. 55 (1) (2017) 327–331.

[60] Y. Shi, W.H. Green Jr., H.-W. Wong, O.O. Oluwole, Redesigning combustion modeling algorithms for the graphics processing unit (GPU): Chemical kinetic rate evaluation and ordinary differential equation integration, Combust. Flame 158 (5) (2011) 836–847.

[61] Y. Shi, W.H. Green, H.-W. Wong, O.O. Oluwole, Accelerating multi-dimensional combustion simulations using GPU and hybrid explicit/implicit ODE integration, Combust. Flame 159 (7) (2012) 2388–2397.

[62] F. Ghioldi, F. Piscaglia, Acceleration of supersonic/hypersonic reactive CFD simulations via heterogeneous CPU-GPU supercomputing, Comput. & Fluids 266 (2023) 106041.

[63] N. Peters, The turbulent burning velocity for large-scale and small-scale turbulence, J. Fluid Mech. 384 (1999) 107–132.