Combustion and Flame 277 (2025) 114130

Contents lists available at ScienceDirect

Combustion and Flame

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

An improved machine learning method for thermochemistry tabulation, with application to LES-PDF simulations of piloted diffusion and swirl-bluff-body stabilised flames with NO_x formation

Tianjie Ding, W.P. Jones, Stelios Rigopoulos *

Department of Mechanical Engineering, Imperial College London, South Kensington Campus, London, SW7 2AZ, UK

ARTICLE INFO

Keywords:
Turbulent combustion
Machine learning
Artificial neural networks (ANNs)
Swirl-bluff-body stabilised flames
Large Eddy Simulation (LES)
NO_x formation

ABSTRACT

Many turbulent combustion modelling approaches require real-time computation of reaction source terms, which is time-consuming and represents a bottleneck of turbulent combustion simulations. In order to speed up the reaction computation process, an artificial neural network (ANN) thermochemistry tabulation methodology has been proposed and developed in our previous work (Ding et al., 2021). In the present work, this methodology is further developed and applied to thermochemistry that includes NO_x formation, which poses further challenges due to the small concentrations of the N-containing species. In particular, we build on the Multiple Multilayer Perceptrons (MMLP) concept, which aims to improve prediction accuracy by systematically combining multiple ANNs. A new method, MMLP-II, is proposed in this work, which trains different ANNs to predict states with different ranges of initial species concentration, in contrast to the previous MMLP-I method which trains several ANNs with different ranges of output magnitude. Both MMLP methods are applied to tabulate the complete GRI-3.0 mechanism and the resulting ANNs are tested on two different turbulent methane flames: Sandia flame D and Sydney flame SMA2. It is found that MMLP-II method can reduce the ANN error accumulation of minor species, and very accurate results are obtained in both turbulent flames. The successful application to two different turbulent combustion problems is indicative of the capacity for generalisation of the ANN tabulation approach. Finally, the reaction integration step is accelerated by a factor of about 15 with ANNs, thus rendering chemical kinetics no longer the bottleneck of the whole simulation.

Novelty and significance

A new ANN thermochemistry tabulation methodology is developed, aimed at providing the higher accuracy needed for predicting species with small concentrations, as encountered in NO_x chemistry. The methodology is applied to two different turbulent flames with NO_x formation: a piloted diffusion flame (Sandia D) and a swirl-bluff-body stabilised flame (Sydney SMA2). The results demonstrate that the method provides both high accuracy and capacity for generalisation. The LES-PDF method is employed here, but the method is applicable to any method that involves real-time calculation of thermochemistry.

1. Introduction

Stringent regulations are nowadays applied to emissions of pollutants such as nitrogen oxides (NO_x) from combustion processes. Even if hydrogen is used as fuel, the emission of NO_x is still inevitable. Therefore, it is necessary to optimise the design of combustion devices in order to reduce pollutant emissions and improve their efficiency. The turbulent combustion simulations via computational fluid dynamics (CFD) methods can greatly facilitate the design of combustors, as detailed information about the flow field and scalar fields is available from the simulation, which is difficult to be obtained via experimental measurements. For practical industrial combustion problems, direct numerical simulation (DNS) is computationally expensive because of the high grid resolution, while Reynolds-averaged Navier–Stokes (RANS) and large eddy simulation (LES) are more affordable techniques. However, RANS and LES require closure models, and the main challenge for simulations of turbulent combustion is posed by the reaction source term due to the turbulence-chemistry interaction. In recent years, machine learning has been finding an increasing number

* Corresponding author.
E-mail address: s.rigopoulos@imperial.ac.uk (S. Rigopoulos).

https://doi.org/10.1016/j.combustflame.2025.114130
Received 15 April 2024; Received in revised form 6 March 2025; Accepted 16 March 2025
Available online 15 April 2025

of uses in combustion modelling, to address the severe modelling and computational challenges in this field. A general review of the related literature, which is currently very large and rapidly growing can be found in Ihme et al. [1]. The following introduction will be limited to the problem of thermochemistry tabulation that is the objective of the present paper.

For certain turbulent combustion modelling methods, such as the steady flamelet and the Flamelet Generated Manifold (FGM), the thermochemical state is represented using only a few scalars such as mixture fraction and progress variable. While this is in itself a tabulation, these methods can be further accelerated with machine learning approaches; Ihme et al. [2] employed artificial neural networks (ANNs) to represent the flamelet/progress variable (FPV) tabulation table, with each ANN predicting a single filtered reactive scalar. Hansinger et al. [3] developed a single deep ANN to represent an FPV table, which outputs 14 combustion quantities. On the other hand, methods such as direct numerical simulation (DNS), transported probability density function (PDF) methods and multiple mapping closure (MMC) are regime-independent and solve transport equations for all reactive scalars, with the chemical source terms directly integrated. However, such methods face severe computational demands due to the requirement for transporting and integrating all species. Comprehensive mechanisms involve very large numbers of species and reactions and, furthermore, the integration of the ordinary differential equations (ODEs) of chemical kinetics is usually stiff, thus being the bottleneck in such simulations.

One solution to alleviate the problem is to use reduced mechanisms, which can be obtained via the traditional Quasi-Steady State Approximation (QSSA) and Partial Equilibrium (PE) methods [4], as well as via systematic methods such as Computational Singular Perturbation (CSP) [5–7], Level of Importance (LOI) [8] or Rate-Controlled Constrained Equilibrium (RCCE) [9,10]. However, in many cases where minor species such as NO_x or polycyclic aromatic hydrocarbons (PAHs) are of particular interest, the mechanism cannot be greatly reduced. Another solution to speed up the computation of chemical kinetics is to employ tabulation methods. The chemistry integration can be either pre-computed, as in the Look-up Table approach [11]), or computed on the fly, as in In-Situ Adaptive Tabulation (ISAT) [12], and interpolation from the resulting table is applied to compute the reaction in the turbulent combustion simulation. However, the Look-up Table approach is suitable only for simple mechanisms because the memory requirements grow exponentially with the number of species. On the other hand, the speed-up of on-the-fly tabulation strongly depends on whether the new states are similar to previously encountered ones. An alternative tabulation approach is to use ANNs, as these are machine learning models capable of nonlinear function approximation. For a given chemical mechanism, the thermochemical state after a reaction time step is a function of the initial state, and ANNs can be trained to approximate this highly nonlinear function. Compared with conventional tabulation methods, the memory requirements for ANNs are negligible, so the ANN tabulation method is not limited by the number of species and can accommodate complex mechanisms.

Early works on ANN tabulation of thermochemistry were conducted by Christo et al. [13,14] and Blasco et al. [15]. Although the chemical mechanisms used in these works were quite simple, the results obtained showed the potential of ANNs for chemistry tabulation. However, there are two major issues in further extending the ANN method to more complex turbulent combustion problems. The first one is related to accuracy: a single simple ANN is not capable of accurate predictions for a complex mechanism involving tens or even hundreds of species. The second issue is related to generalisation, as the combustion thermochemical space is multidimensional and features different chemical kinetics at different regions. If every time that the ANNs are to be applied to a new combustion problem, the same or a similar problem must be simulated beforehand with direct integration to generate the training data, then the computational benefits of ANNs would be negated.

To address the first issue, one solution is to systematically combine several simple ANNs in order to form a stronger predictor. Blasco et al. [16,17] introduced the idea of composition partitioning, which splits the chemical composition space into several subdomains based on mixture fraction and temperature, with each subdomain fitted by an individual ANN. Later, the same group introduced the use of a Self Organising Map (SOM) to automatically perform clustering of training data [18], with specialised ANNs trained for each cluster. This approach was employed by Chatzopoulos et al. [19], Franke et al. [20] and An et al. [21]. An alternative approach was proposed by Readshaw et al. [22], who employed individual ANNs for predicting every single species. It was found that this approach yields better generalisation due to the fact that each ANN 'sees' the entire composition space; by contrast, the clustering approach was prone to producing highly accurate but too specialised ANNs.

An additional problem related to the first issue is the wide range of the magnitudes of species concentrations or reaction rates. For example, for a methane flame, the concentration of CH_4 may vary from 10^{-3} kmol/kg to less than 10^{-10} kmol/kg. In order to ensure the ANN prediction accuracy for data with different magnitudes, Ding et al. [23,24] proposed the multiple multilayer perceptron (MMLP) method, which employs several ANNs for predicting a single species, each trained to predict different magnitudes of species concentration change. Another possibility is the use of deep neural networks (DNN), which have many hidden layers and increased approximation capability. In the work of Wan et al. [25], two DNNs were employed to predict the reaction rate of 11 species, with one focussing on reactions with slow burning rates, and the overall tabulation was coupled with DNS. In the work of Zhang and Mao et al. [26,27], a Box–Cox transformation (BCT) was applied to preprocess the combustion data in order to deal with the issue of multi-scale data, and then a DNN was trained. Their method was first successfully applied to a 9-species mechanism on a H_2 flame and later on more complex fuels and mechanisms [28]. In their recent work [29], they performed combustion simulations with DNN-accelerated chemistry integration on GPUs to achieve further speed-up.

The second issue, i.e. that of generalisation without training on data from the same problem, is probably the most challenging problem in ANN chemistry tabulation. In order to fully realise the speed advantage of ANNs, the training data generation methods should be simple enough to provide fast computation and the resulting ANNs should be applicable to a wide range of realistic turbulent combustion problems. Since the ANN cannot ensure accurate predictions if the given data is not covered by the training dataset, the ANN training data must be carefully chosen so that they can properly cover the composition space of interest. Many efforts have been made to develop such data generation methods. Sen et al. [30,31] generated the training data via stand-alone linear eddy mixing (LEM) simulations, with applications to LES of syngas/air flames. Chatzopoulos et al. [19] proposed a method of generating training data via dynamic laminar flamelet simulations. The SOM method was also employed in their work and the ANNs were successfully applied to the RANS-PDF simulation of DLR-A/B flames. This flamelet sampling method was further validated by Franke et al. [20], with application to an LES-PDF simulation of the Sydney L flame. Wan et al. [25] generated the training dataset via a canonical stochastic micromixing problem and the ANNs were applied to a DNS of a turbulent non-premixed syngas oxy-flame. In the work of An et al. [21], the training data was collected via a RANS simulation and then the ANNs were applied to an LES simulation of a rocket-based engine. Zhang et al. [26] proposed a multi-scale sampling method and applied it to a 9-species mechanism for H_2 flames. Ding et al. [23] proposed combining the flamelet sampling method with a random data generation method in order to enhance the generalisation

capacity of the ANNs. This approach was applied to a range of problems including turbulent non-premixed methane-air flames [22,23] and flames with blended fuels such as methane/hydrogen mixtures [24] and, more recently, to turbulent premixed methane flames [32] and dimethyl ether non-premixed flames [33]. A method for improving element conservation in ANN-tabulated thermochemistry has also been proposed by Readshaw et al. [34].

While several studies have demonstrated the potential of the ANN tabulation method in a range of problems, none of these (to the authors' knowledge) involved detailed reports of nitrogen chemistry. Not only is the prediction of NO_x important for emissions control, but it also poses further issues for ANN tabulation due to the low concentrations of most species involved in nitrogen chemistry. The question is therefore whether the existing ANN method is sufficient, and it is the objective of this paper to explore this question. As it will be shown, the challenge posed by N chemistry motivates the development of an improved method, termed MMLP-II. This method is developed and compared with the original MMLP-I method. The mechanism to be tabulated is GRI-3.0 [35], including its N chemistry component. The tabulated chemistry is tested on two flames. The first one is Sandia flame D, a typical piloted diffusion flame previously simulated by Ding et al. [23] with the GRI-1.2 mechanism [36] and without nitrogen chemistry. Furthermore, a swirl-bluff-body stabilised flame, the Sydney SMA2, is simulated, again with GRI-3.0 and N chemistry. This flame represents a more stringent test due to recirculation, which results in the ANNs being called many times in succession along the pathline of a fluid element and thus in potential for error accumulation. As with our previous work, it should be emphasised that the simulations of the turbulent flames have been accomplished with ANNs trained on data generated in an abstract way and not with data from the same problem, and the fact that they can yield excellent predictions on two different turbulent flames is a strong indication of the generalisation capacity of the approach. Furthermore, the methodology of the paper is applicable to a wider range of combustion problems involving N-chemistry, such as ammonia combustion.

It should be noted that, while the LES-PDF approach is employed in the present paper, the machine learning methods developed are equally applicable to any other method that involves the direct coupling of chemistry and flow, thus requiring real-time numerical integration of chemical kinetics. Such methods include direct numerical simulation (DNS), conditional moment closure (CMC), unsteady flamelet, multiple mapping closure (MMC), thickened flame model, linear eddy model (LEM), partially stirred reactor (PaSR) as in OpenFOAM and laminar flame computation.

The paper is structured as follows. Section 2 describes the methodology, with emphasis on the new MMLP-II method. Section 3 evaluates and compares the performance of both MMLP methods on one-dimensional problems. In Section 4, both methods are applied to the Sandia D and Sydney SMA2 turbulent flames and the results are discussed, followed by the conclusions.

2. ANN tabulation methodology

2.1. Overview

For turbulent combustion simulations, the convection–diffusion-reaction governing equations of reactive scalars are solved via the operator splitting method, which allows the reactions to be computed separately. In the reaction computation step, the reactive scalars after a simulation time step δt will be calculated given their initial values. This computation can be achieved by numerically integrating the following ODEs which describe the time evolution of the reaction system:

$$\frac{d\psi_i}{dt} = F(\psi_i) \tag{1}$$

Here ψ_i is an array containing the species concentrations (y_i) and temperature (T), while the function F is determined by the chemical kinetics of the given mechanism. The species concentration after a time step δt can be regarded as a function of initial species concentration and temperature, which is implicitly determined by the direct integration of Eqs. (1):

$$y_i^{\delta t} = G(y_i, T) \tag{2}$$

Since the temperature is also a function of specific enthalpy (h) and species concentration, so the change of each species concentration after a time step δt can be seen as:

$$\delta y_i = H(y_i, h) \tag{3}$$

Note that the specific enthalpy includes the enthalpy of formation and therefore does not change due to chemical reaction.

The direct integration of the chemical kinetics ODEs (Eqs. (1)) is very expensive. The objective of the ANN tabulation approach is to replace these ODEs with a much faster, algebraic model. The ANN model employed in this work is the multilayer perceptron (MLP). The basic unit of an MLP is a neuron, which can linearly or non-linearly map a set of inputs to an output. The inputs to a neuron will be first linearly combined using the corresponding weights and bias, and then an activation function is used for output mapping. An MLP consists of several layers of neurons that endow it with the capability to fit the highly nonlinear function $H(y_i, h)$.

The basic ANN architecture underlying the methods in our work is the multilayer perceptron (MLP), a detailed description of which can be found in several books such as Hagan et al. [37]. The MMLP methods (both the previous one and the new one developed here) employ several MLPs in succession, and the structure of this process will be described in Section 2.3. Here, we describe the basic MLP unit that we employ for predicting a single species, depicted by a schematic in Fig. 1. The GRI-3.0 mechanism has 52 species if Ar is not considered, so the number of inputs totals 53 (52 species + enthalpy). Each MLP only predicts the concentration change of a single species, so there is only one output. As the chemistry of GRI-3.0 is more complex than that of GRI-1.2, three hidden layers, each containing 40 neurons, are employed in this work while previous works [23,24] only use two hidden layers for GRI-1.2.

To accurately fit the target functions, the ANN parameters must be tuned to the desired values, and this is accomplished via the training process. For supervised learning, data with input–output values (computed via direct integration) are presented to the ANN and a learning rule is employed to adjust the network parameters in the training process. Most learning rules for function approximation are based on the backpropagation algorithm. In this work, the Bayesian Regularisation method is applied, which can avoid overfitting, making the ANNs more robust than those trained using standard backpropagation methods. The backpropagation algorithm employed for Bayesian Regularisation is the Levenberg–Marquardt algorithm, which is based on the Gauss–Newton optimisation method. More details on these methods can be found in several books [37–39].

2.2. The hybrid flamelet/random data (HFRD) generation method

In the present work, the hybrid flamelet/random data (HFRD) method is applied to generate the training set. This method endows the ANNs with considerable generality and has been successfully applied to different combustion problems in previous works [22–24,32,33]. The basics of this method will be described here succinctly, and only the modifications related to NO_x formation will be described in detail.

The HFRD method comprises mainly two steps: sampling via flamelet simulations and random data generation. For the first step, an ensemble of 1-D laminar counterflow diffusion flamelet simulations is performed to collect data samples. This canonical combustion problem is governed by the following equation [40]:

$$\rho\frac{\partial y_i}{\partial t} = \rho\frac{\chi(z)}{2}\frac{\partial^2 y_i}{\partial z^2} + \dot{\omega}_i \tag{4}$$

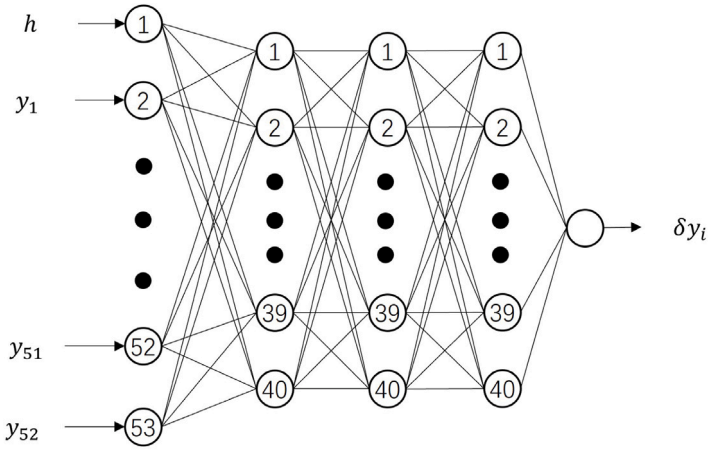

Fig. 1. MLP structure for GRI-3.0 tabulation.

Table 1
Initial conditions and sampling range of flamelet simulations.

	Fuel temperature	Oxidant temperature	Strain rate	Mixture fraction
Minimum	100 K	100 K	1/s	0.02
Maximum	400 K	400 K	600/s	0.12

Here z is the mixture fraction, y_i is the species concentration, $\dot{\omega}_i$ is the reaction rate and χ is the scalar dissipation rate given by the following equation:

$$\chi(z) = \frac{S}{\pi} \exp\left[-2(\mathrm{erfc}^{-1}(2z))^2\right] \tag{5}$$

where S is the strain rate.

For each individual flamelet simulation, the initial temperature and strain rate are randomly chosen, and a composition state with a random mixture fraction within the flammability range is collected at each time step. In the present work, a total of 200 flamelet simulations are performed, and the minimum and maximum values for initial temperatures, strain rate and mixture fraction sampling range are listed in Table 1.

In the previous works mentioned above, the flamelets were initialised by either setting a small part of the mixture fraction space to equilibrium in order to collect data representing ignition or setting the entire space to equilibrium, which can help collect more data close to the equilibrium state. In the present work, the constrained equilibrium state is used instead of unconstrained equilibrium, in order to collect states with proper NO_x concentration. In many cases of diffusion flames, including industrial combustion chambers, the residence time is small compared to the time scale needed for NO_x to reach equilibrium values, so NO_x are always far from its equilibrium state in these cases. If equilibrium values of NO_x are used for initial conditions of flamelet simulations, then a large number of states with high NO_x concentration will be collected, which will not be encountered in the turbulent diffusion flames. If constrained equilibrium is applied to the initial conditions, where the nitrogen-containing species are set to zero while other species are set to their equilibrium state, then the NO_x will evolve from zero to the final steady flamelet state during the simulation, producing better state samples with reasonable NO_x concentration.

After the flamelet simulations, the flamelet dataset is used as a basis to generate the random dataset. For each data state (enthalpy plus species concentration) of the flamelet dataset, a new random state can be generated accordingly. The equations (Eqs. (6) to (9)) for generating the random state have the same form as those in our previous works [24], except that the values of coefficient a and b are set to 20 ($a = b = 10$ in Ref. [24]) in order to produce a more suitable dataset for the GRI-3.0 mechanism. The variable c in Eq. (6) to (8) takes a random value uniformly distributed between −1 and 1. Compared with the GRI-1.2 mechanism in our previous work [24], which has 31 species, the GRI-3.0 mechanism has 52 species. The increased number of species requires that the coefficients a and b be increased to ensure that the generated random states can satisfy the constraints in Table 2, which will be mentioned below[1]. The same values for a and b are also used in the tabulation of a dimethyl ether mechanism [33], which has a similar number of species with GRI-3.0.

$$h_{rnd} = h + \frac{c}{a}(h_{max} - h_{min}) \tag{6}$$

$$y'_j = y_j^{(1+c/b)} \quad (j \neq j_{N2}) \tag{7}$$

$$y'_{N_2} = y_{N_2} + \frac{c}{a}(y_{N_2,max} - y_{N_2,min}) \tag{8}$$

$$y_{j,rnd} = \frac{y'_j}{\sum y'} \tag{9}$$

The new random state should also satisfy element ratio constraints and flammability constraints to ensure that they are physically realistic. The constraints (minimum and maximum values) employed in this work are shown in Table 2. For the detailed procedure for generating random data based on flamelet samples, the readers can refer to our previous works [22,23]. In the present work, a total of about 1,300,000 random data are generated. All direct integrations are performed with the VODE solver [41] and the time step employed is 10^{-6} s, which matches that of the LES-PDF simulations. If variable time step is desired, then the ANNs trained for the shortest time step can be called in succession; alternatively, it is possible to have more than one sets of trained ANNs with different time steps and combine them for further speed-up. A recent work [32] tested the robustness of the training with respect to the time step and found that no consistent trend in error accumulation rate with increasing time step can be discerned.

[1] For a random state, the element ratio is a function of species concentrations, which can be considered as random variables. With more species, the element ratio will be determined by more random variables (random species concentrations), which makes it more difficult to be within a certain small range. To ensure the element ratios are within the given range, we have to decrease the randomness of each random variable, so we need to use larger values for a and b.

Table 2
Constraints for random data generation.

	H/C ratio	O/N ratio	Mixture fracture	Temperature (K)
Minimum	3.9	0.255	0.02	500
Maximum	4.1	0.275	0.12	–

2.3. Multiple Multilayer Perceptrons (MMLP) and development of MMLP-II

In most ANN training cases, the ANNs are trained to minimise the absolute error of the prediction. However, since we are predicting the concentration change during the reaction step, what is also important is the *relative error of the prediction with respect to the magnitude of the concentration change*. For example, the same value of absolute error may be acceptable if the concentration change is large, but may incur a large relative error to a small concentration change at a thermochemical state near equilibrium. In most turbulent combustion cases, the species concentrations span several orders of magnitude. Therefore, ensuring acceptable prediction errors for data with varying magnitudes poses a big challenge for ANN training.

There are mainly two ways to deal with the issue caused by the multi-scale chemical data. This first solution is to apply proper data preprocessing methods for deep neural network training, such as the Box–Cox transformation (BCT) [26] and the logarithmic-normalised transformation [42]. After applying these nonlinear transformations, the errors caused by the originally low-magnitude data can contribute more to the updating of network parameters in the training process, so that the final ANNs can provide low relative errors for data with small magnitudes. However, deep ANNs are usually required to ensure high prediction accuracy in this case. Since the network is required to provide accurate prediction for data with all different magnitudes, which includes all chemical dynamics, the fitting capability of an ANN with a simple structure is not enough in these cases [24]. It should also be mentioned that a deep ANN would also be more expensive to compute; in the context of our approach, the ANNs employed have a simple structure in order to allow for fast computation, which is the objective in a machine learning approach that aims to accelerate the computation of thermochemistry.

Another solution is based on the idea of ensemble learning by training multiple simple ANNs focusing on data with different scales. In our previous works [23,24], the Multiple Multilayer Perceptrons (MMLP) approach was proposed, which trains different MLPs to predict states with different magnitude ranges of outputs (species concentration changes). In the present work, this MMLP approach is further developed. For the convenience of distinguishing different methods, the previous method is referred to as MMLP-I while the new method is termed MMLP-II. Both MMLP methods focus on improving the prediction accuracy for data with small output magnitudes by systematically combining multiple MLPs to form a stronger predictor. The MMLP-I method will be briefly introduced while MMLP-II will be described in detail.

While the main feature of the new MMLP method will be more clearly demonstrated in Sections 3 and 4, where it is applied to different combustion problems, we should mention here the reasons for its development. Although the MMLP method has been successfully applied to different turbulent combustion problems using the GRI-1.2 mechanism in our previous works, some errors for minor species, though very small, still pertain at certain locations where the data are likely to go through many reaction steps. In this work, where the aim is to tabulate NO_x chemistry, more minor species are involved, some with very small concentrations, and it is desired to predict them accurately as they are important output variables rather than intermediate steps. Furthermore, the recirculation in the swirl-bluff-body stabilised flame results in fluid elements undergoing more reaction steps in their pathline before being flushed out of the reaction zone. The objective of the MMLP-II method is to further improve the prediction accuracy for certain composition states, thus reducing the accumulated errors of minor species.

2.3.1. MMLP based on output range (Type I)

One way to improve the prediction accuracy for data with small output magnitudes is to train additional MLPs targeting these data only. This is the main idea of the MMLP-I method. The advantages and the main process of this approach are briefly demonstrated here using the training results of CH_4. A more detailed analysis can be found in our previous work [23].

First, an MLP is trained to predict the change of CH_4 concentration using the random data and then tested on the flamelet dataset. The results are shown in Fig. 2(a), where the MLP predictions are plotted against the target outputs obtained via direct integration. Overall, the MLP predictions are accurate. However, if zooming into the range $(-10^{-5}, 10^{-5})$, represented by the small black square in Fig. 2(a), the relative errors become significant, as shown in Fig. 2(b). Although the absolute errors of MLP predictions are still very small, of the order of 10^{-7}, the target outputs also have a small order of magnitude, thus leading to significant relative errors.

In order to improve the MLP prediction accuracy on data with small output magnitudes, a new MLP can be trained to focus on these data only. For the data shown in Fig. 2(b), the random data with corresponding small output magnitudes are used to train a new MLP. The flamelet data in Fig. 2(b) are then used to test a new MLP, and the results are shown in Fig. 2(c). It can be clearly seen that a significant improvement in prediction accuracy is obtained for the new MLP, compared with the original MLP in Fig. 2(b). Similarly, the data with smaller output magnitude (e.g. output within range $(-10^{-6}, 10^{-6})$ in Fig. 2(c) may also have large relative errors. If necessary, the prediction accuracy of these data can be further improved by training another new MLP focussing on these data only. However, this process cannot be repeated infinitely, because as the data output range narrows, the improvements in the prediction accuracy of the new MLP becomes more and more marginal and may even deteriorate.

For the training with the MMLP-I method, multiple MLPs are trained for each species targeting different ranges of output change. The training starts with the widest output range, and then the output range is narrowed to train the new MLP. The training process continues until the new MLP fails to yield further improvements in the prediction accuracy. When applying these MLPs, the one with the narrowest output range is first applied to calculate the output. If the result is within the MLP range, then it is accepted. Otherwise, the previous MLP with a wider range will be used to repeat the above process until the result falls within the predicting MLP range. More details about the training procedure and implementation process of the MMLP-I method can be found in Ding et al. [23].

2.3.2. MMLP based on input range (Type II)

The motivation for developing MMLP-II is best explained with an example. Consider the data for two cases shown in Table 3. Both composition states A and B have the same concentration change, δy_i). However, for case A, the initial concentration, y_i, is ten thousand times greater than δy_i. Now, suppose that both states are computed with the same MLP in the MMLP-I algorithm, and the result, δy_i^{MLP}, is the same in both cases. In case A, the relative error of the approximation is very small, but in case B it is very large. This illustrates the fact that the magnitudes of the concentration values (the *inputs* to the MLP) may be more important than those of the concentration changes (the *outputs*). Therefore, it is necessary to further improve the MLP prediction accuracy for data with small input values and ensure that the ensuing MLP errors are much smaller than the magnitude of the inputs (i.e., $|\delta y_i^{MLP} - \delta y_i| \ll y_i$), so that the influence of the prediction errors on the simulation results can be minimised.

It is also found that data with small input values take up a considerable portion of the entire dataset for most species. This is the case for random data, flamelet data and the data arising from the turbulent combustion problems that are considered in the present work. The histogram of CH_4 concentration from the flamelet dataset is shown in

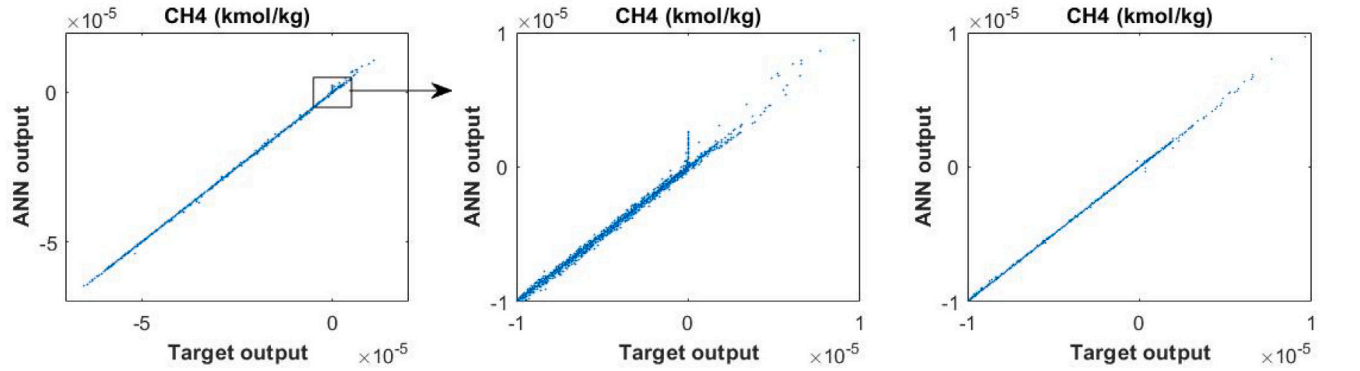

(a) Test on flamelet data (b) Partial enlargement of (a) (c) Improvement of results in (b) with the new MLP

Fig. 2. Comparison of MLP predictions for CH_4.

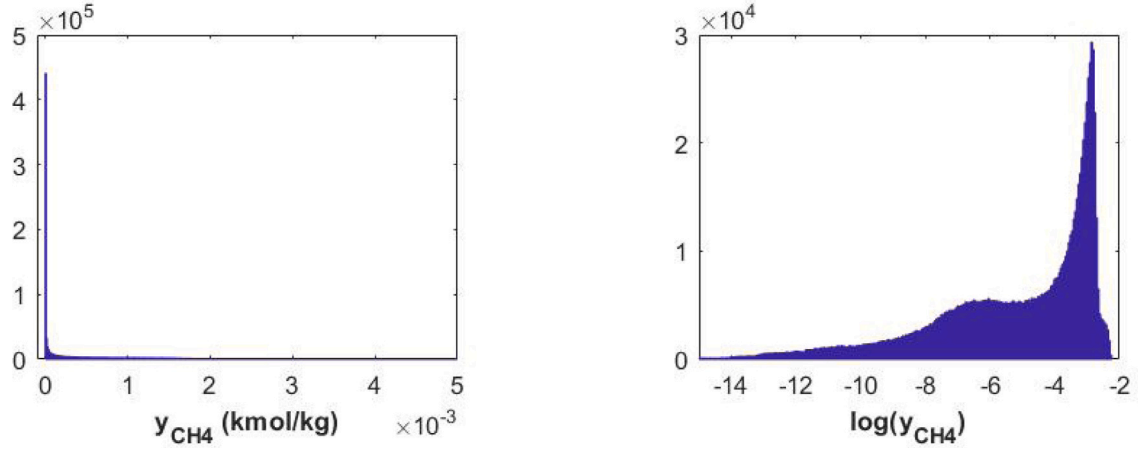

(a) Histogram of y_{CH_4} (b) Histogram of $\log(y_{CH_4})$

Fig. 3. Histograms of CH_4 concentration for the flamelet dataset.

Table 3
Examples of the influence of input magnitude on relative errors.

Case	A	B
Input y_i	10^{-2}	10^{-5}
Concentration change δy_i	10^{-6}	10^{-6}
Concentration after reaction $y_i(\delta t)$	$1.0001 \cdot 10^{-2}$	$1.1 \cdot 10^{-5}$
MLP prediction δy_i^{MLP}	10^{-5}	10^{-5}
MLP predicted concentration $y_i^{MLP}(\delta t)$	$1.001 \cdot 10^{-2}$	$2 \cdot 10^{-5}$
absolute value of relative error of $y_i^{MLP}(\delta t)$	0.09%	82%

Table 4
Distribution of several species for the flamelet dataset.

Input range (kmol/kg)	$< 10^{-8}$	$< 10^{-6}$	$< 10^{-4}$	$< 10^{-2}$
Data percentage of H	1.98%	5.73%	63.75%	100%
Data percentage of O	4.14%	31.02%	79.03%	100%
Data percentage of OH	2.37%	7.82%	60.34%	100%
Data percentage of CH_4	14.38%	32.98%	55.50%	100%

Fig. 3(a), where the data are classified into 500 bins. The maximum value of y_{CH_4} is of the order of 10^{-3} kmol/kg, and it can be seen that the vast majority of the data are of a smaller order of magnitude than 10^{-3}; as a result, the bars close to zero have extremely high counts. For clarity, the histogram is plotting the logarithmic value of the CH_4 concentration ($\log(y_{CH_4})$) in Fig. 3(b). The data within the range of $(10^{-3}, 10^{-2})$ only account for a very small percentage of all the data, while the majority of the data are below 10^{-3} kmol/kg. About 50% of the data are within the range of $(0, 10^{-4})$. The distribution of CH_4 and several other species from the flamelet dataset is shown in Table 4. For all of these four species, more than half of the data are below the value of 10^{-4} kmol/kg. For CH_4 and O, about 30% of the total data have input values below 10^{-6} kmol/kg. A similar distribution applies to other minor species, including those related to NO_x formation.

The above analysis suggests that the input magnitude may be a more important criterion for training separate MLPs, especially for species with small concentrations, and this is the basis of the MMLP-II. Similar to the MMLP-I method, the prediction accuracy of data with small input magnitudes (i.e. concentrations) can be improved by training new MLPs focussing on these data only.

Examples of CH_4 are shown here again to demonstrate the main process and the features of the MMLP-II approach. For both random and flamelet datasets, the maximum input value of CH_4 is of the order of 10^{-3} kmol/kg. Looking at the performance of an MLP trained with the whole dataset, the predictions initially seem reasonable Fig. 2(a). However, looking at data with input range $y_{CH_4} < 10^{-6}$ kmol/kg only, the errors are significant, as shown in Fig. 4(a). To improve the prediction accuracy for data with small input magnitudes, new MLPs can be trained using only the random data whose input for the target species (i.e. the species to be predicted by the particular MLP) are within the corresponding input range. Fig. 4(b) shows the training results of the new MLP. It can be clearly seen that the prediction accuracy is greatly improved.

At some point, it ceases to be useful to keep training new MLPs, because as the input range reduces, the prediction accuracy improvement brought by the new MLP becomes less significant. Moreover,

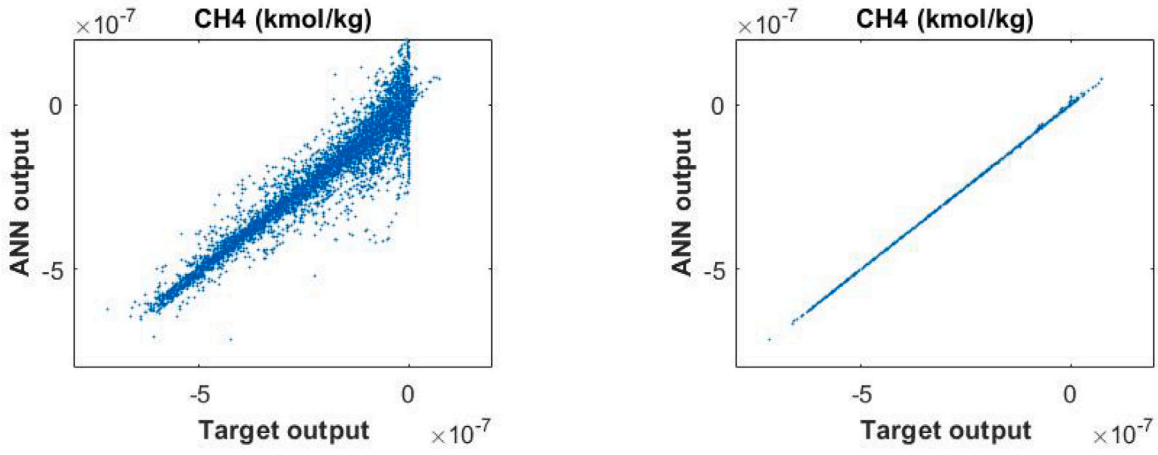

(a) data with $y_{CH_4} < 10^{-6}$ koml/kg (b) data with $y_{CH_4} < 10^{-6}$ koml/kg

Fig. 4. CH_4 prediction results of flamelet data with small input magnitude values when a single MLP is used for the entire dataset (a) and improved prediction with an MLP specially trained for this range (b).

the amount of data also decreases with decreasing input range, thus rendering new MLPs redundant at the point that they can only be applied to a very small portion of data.

The procedure for training with MMLP-II will now be described. First, an MLP is trained for the target species using data from the entire dataset. The test dataset (for which the flamelet dataset can be utilised) is then applied to evaluate the performance of the MLP. When the MLP fails to produce accurate results, a smaller input range is prescribed for the next MLP, denoted as (0, k). The value of k is usually chosen from 1/1000 to 1/10 of the previous maximum input value, based on the performance of the MLP. The training data whose input of the target species falls in (0, k) are selected to train the new MLP, and training results are evaluated using the test data within the given input range. If the prediction accuracy is improved, then the new MLP is accepted. The above process is repeated until the new MLP fails to further improve the prediction accuracy or there are not enough data within the given input range. However, for some major species, the concentrations of most of the data are of a similar order of magnitude (e.g.,N_2, O_2, CO_2 and H_2O), so a single MLP is employed for each of these species and the above procedure for MMLP-II is not required.

When applied in the context of a CFD simulation, the implementation of MMLP-II is quite straightforward. For each species at each composition state, the MLP with the correct input range for that species is selected and applied to compute the output (concentration change) for the same species. This is repeated for all species in the given composition state.

2.3.3. Comparison of the two MMLP methods

A comparison of the data distribution among different MLPs between the two MMLP methods is shown in Fig. 5 using the CH_4 input–output space, where the input concentration of CH_4 (y_{CH_4}) and its absolute change during the reaction ($|\delta y_{CH_4}|$) are plotted. The data shown in this figure is the flamelet dataset and the axes are in log scale. As explained in the previous section, the MMLP-I method trains multiple MLPs based on the output range while MMLP-II is based on the input range. For MMLP-I method, three MLPs with absolute output ranges $(0, 10^{-4})$, $(0, 10^{-5})$ and $(0, 10^{-6})$ are trained and these ranges are indicated by the horizontal red lines in Fig. 5. Three MLPs are also trained for the MMLP-II method with input ranges $(0, 10^{-2})$, $(0, 10^{-6})$ and $(0, 10^{-8})$, represented by the vertical black lines in Fig. 5. Based on the input and output range of the trained MLPs, the entire dataset can be split into 9 areas in Fig. 5. However, the output (the change) is usually smaller (in absolute value) than the input (the concentration), unless the latter value is extremely small, so the data are located in only 5 areas, marked in Fig. 5 using circled numbers.

The root mean square errors (RMSE) obtained by applying both MMLP methods on data located in different areas of the composition space are calculated and compared in Fig. 6. For data in area 1, both methods use the first MLP which covers the entire composition space in both methods, so the prediction accuracy is the same. For areas 2 and 3, new MLPs are trained in the case of MMLP-I, but MMLP-II still uses the first MLP, so the accuracy of MMLP-I is higher. A different picture emerges for data in area 4, where the MMLP-I method uses its third MLP that employs the random data in areas 3, 4 and 5 for training. By contrast, MMLP-II uses its second MMLP which is based only on the data in areas 4 and 5 for training and is, therefore, more specialised than the third MLP of MMLP-I. Therefore, the RMSEs of MMLP-I in areas 3, 4, and 5 are quite similar while the performance of the MMLP-II on data in area 4 is better than that of MMLP-I. Similarly, MMLP-I is also outperformed on data in area 5. These results affirm that the MMLP-II method performs better at low concentrations, which is consistent with the main concept in its design.

Although the performance of both MMLP methods is demonstrated here using the example of CH_4, the conclusions drawn are relevant for many other species. Except for several major species with large mass fractions (e.g. O_2, N_2, CO_2), most species, especially minor ones, share a similar distribution with the one shown in Fig. 5. Overall, the MMLP-II method can be used to improve the prediction accuracy of the minor species with low concentrations, which are more sensitive to prediction errors. This feature will be evidenced in the results presented in Sections 3 and 4, where both MMLP methods are applied and their performance on different combustion problems is compared.

For MMLP-I, 2 to 3 MLPs are trained for each species, resulting in a total of 118 MLPs. For MMLP-II, a total of 161 MLPs are trained, with each species represented by 1 to 5 MLPs. All training is performed using the MATLAB neural network toolbox. The training time of a single ANN is about 70 h using a single Xeon 6132 2.6 GHz core, and the complete training process can be completed quickly if the training is done in parallel in a multi-core system.

2.3.4. CPU time reduction

The CPU time required for the computation of the reaction step is significantly reduced in both MMLP methods, compared with the time required for direct integration. Fig. 7 shows the CPU time taken in the case of each method, normalised by the time required for direct integration. It can be seen that a speed-up ratio of about 15 is

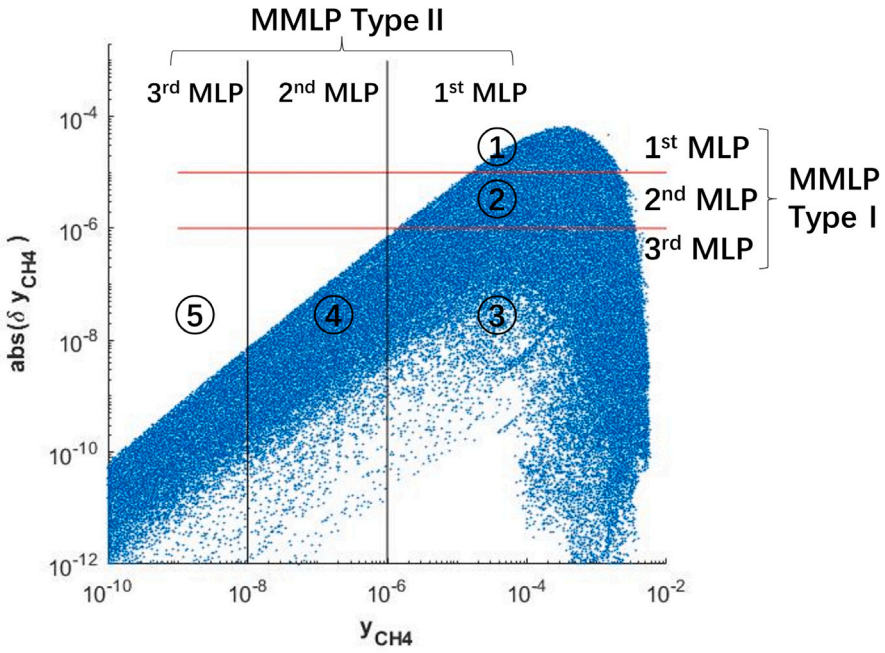

Fig. 5. Comparison of data division between MMLP-I and MMLP-II, shown on the abs(δy_{CH_4})-y_{CH_4} space of the flamelet dataset.

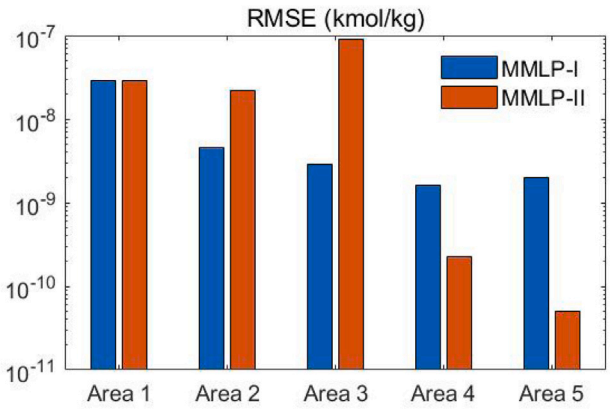

Fig. 6. RMSE comparison on different data areas between MMLP-I and MMLP-II.

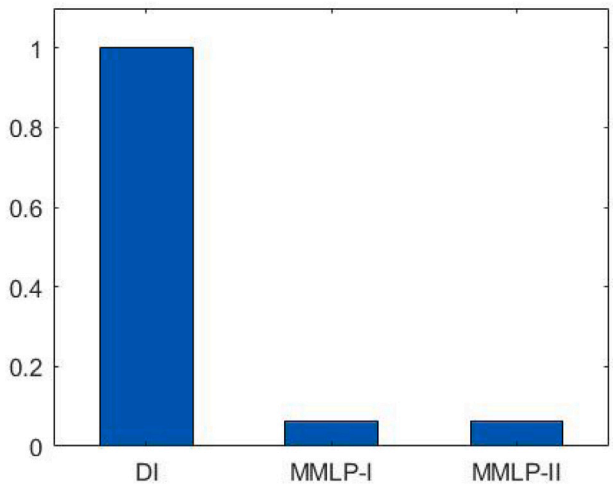

Fig. 7. The CPU time required for computing the reaction step with direct integration and the two MMLP methods, normalised by the time required for direct integration.

attained by both MMLP methods. If the reaction step takes about 60% to 90% of the total simulation time, a speed-up ratio of 15 brings a reduction of the total simulation time that ranges from 56% to 84%. With this level of reduction, the computation of chemical kinetics ceases to be the bottleneck of the entire simulation. Higher speed-up ratios are not necessary because they would bring a marginal decrease in the total simulation time, which is now largely due to the need for transporting a large number of species as well as computing the flow field. A similar situation is encountered with the parallelisation of codes, where a code that is 90% parallelisable is not suitable for computation with more than 10–20 cores; that would be wasteful. An example of a simulation where the chemical reaction step was particularly expensive can be found in Liu et al. [33], where dimethyl ether (DME) flames were simulated and the reaction step took 95.7% of the total time in the case of direct integration. Application of an ANN tabulation method similar to the one used in the present work (based on MMLP-I) brought up a speed-up ratio of 16.3 in the reaction step, corresponding to a speed-up ratio of about 10 in the total CPU time.

3. Application to one-dimensional laminar flamelets

The ANNs are first tested on one-dimensional laminar flamelet simulations. The governing equation of this combustion problem [40] shows that the controlling parameter is the strain rate. The simulation of flamelets with low strain rate requires more time steps to converge and the concentration changes are very small close to equilibrium, which results in further accumulation of errors, and is thus a good test problem to evaluate the performance of both MMLP methods.

Results show that both MMLP methods can produce excellent results for both major and minor species when the strain rate is $S \geq 10\ s^{-1}$, with the direct integration (DI) and ANN results practically coinciding, so these results will not be shown here; note that such results have been shown with MMLP-I in Ding et al. [23]. Instead, in order to show the difference of the two methods, we focus on a flamelet simulation with $S = 5\ s^{-1}$ and show the mass fraction profiles of minor species related to NO_x formation in Fig. 8. The mass fractions are plotted over

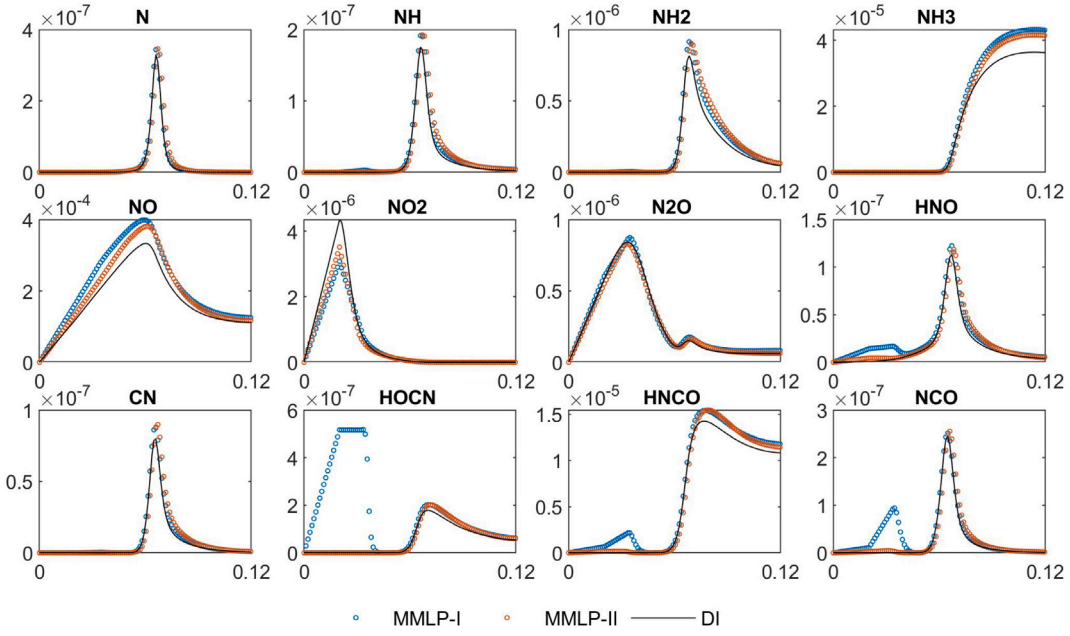

Fig. 8. Mass fractions profiles of the flamelet simulation with $S = 5\ s^{-1}$. Blue dots show results of MMLP-I, red dots are those of MMLP-II and black solid lines are those of DI. It can be seen that MMLP-I exhibits substantial errors particularly in HOCN, and NCO, which are remedied by MMLP-II.

mixture fractions around the flammability range. It can be seen that the results of MMLP-I on several species exhibit an overprediction at low mixture fraction locations while the simulation results of MMLP-II of these species are quite good. The worst case for MMLP-I is HOCN. These minor species have very low concentrations at low mixture fraction locations, and it is exactly this feature that MMLP-II aims to remedy. Once MMLP-II is applied, these errors are mostly fixed and HOCN, in particular, is very accurately predicted.

4. Application to LES-PDF simulations of turbulent combustion

In the present work, the Sandia piloted diffusion methane/air flame D and the Sydney methane/air swirl-bluff-body stabilised flame SMA2 are chosen as test cases for our ANN methodology. These two types of standard flames are well studied and the detailed experimental data are available on "International Workshop on Measurement and Computation of Turbulent Nonpremixed Flames" (TNF Workshop) [43]. In this work, the flames are simulated with the LES stochastic field method, and the implementation is identical to that described in the work of Jones [44] and our previous works [22,23]. Eight stochastic fields are applied, and equal diffusivities are assumed with a Schmidt number of $\sigma_{sgs} = 0.7$. The simulations are performed with our in-house CFD code BOFFIN [45].

For each flame, simulations were carried out with direct integration (DI) and with different MMLP methods for the reaction source term computations. The optically thin radiation model is employed in this work to account for the radiative heat loss. Only CH_4, CO, CO_2 and H_2O are considered for radiation and their radiative properties are given based on the RADCAL model [46]. A detailed description of the model can be found in Ref. [47]. The statistical quantities (mean and RMS) were collected for a period of more than 5 flowthrough times (based on the bulk velocity of the jet), and the collection started once the flame had become fully developed.

4.1. Turbulent piloted diffusion flame - Sandia D

The Sandia flames are piloted diffusion jet flames, studied experimentally by Barlow and Frank [48]. This flame series, especially flame D, has been widely used to validate different combustion simulation methods. However, a limited number of works have simulated NO formation with real-time chemical source term computation, including the works of Wang et al. [49], Raman et al. [50], Cao et al. [51,52] and Jaravel et al. [53]. This flame has also been used as a test case in previous work of the authors [23] for ANNs with MMLP-I and without the nitrogen chemistry (using the GRI-1.2 mechanism [36] rather than GRI-3.0). The burner of Sandia flames has a main jet diameter (D) of 7.2 mm and a pilot diameter of 18.2 mm. The jet composition is 25% CH_4 and 75% air by volume. The annular pilot burnt composition has the same specific enthalpy and equilibrium composition as CH_4/air mixture at equivalence ratio $\phi = 0.77$. For flame D, the bulk velocities of the jet, pilot and coflow are 49.6 m/s, 11.4 m/s and 0.9 m/s respectively. The simulation domain has dimensions $36 \times 7.5 \times 7.5$ cm and a Cartesian finite volume grid with 3M cells is employed.

The simulation results are shown in Figs. 9 and 10. The radial profiles of major species and temperature at 4 different axial locations are shown in Fig. 9. The ANN results are in excellent agreement with DI. At $Z = 7.5$ D, 15D and 30D, the mean results and RMS results of both MMLP methods and DI are almost visually indistinguishable. At $Z = 45$ D, the results are still very accurate, apart from an overprediction of the RMS value for CH_4 by MMLP-I and a small overprediction of the mean values of CO and H_2 by MMLP-II.

The results for several nitrogen-containing species are shown in Fig. 10. For both MMLP methods, most of these species are accurately predicted by ANNs. The mean and RMS results of NO are excellent at all 4 axial locations. For NO_2, the means values are highly accurate, but the peak RMS values at location $Z = 15$ D exhibit some errors. At location $Z = 45$ D, MMLP-II slightly overpredicts the mean values of some species (HCN, HCNO, HNCO) near the centreline, while MMLP-I can still produce accurate mean profiles.

Overall, both MMLP methods provide accurate predictions for the case of Sandia flame D. The only exception is HOCN, which shows inaccurate results at locations far from the centerline (low mixture fraction locations) when using MMLP-I method, especially at location Z = 45D. This species is found to be the most difficult to accurately predict for GRI-3.0 when using MMLP-I because of error accumulation at low mixture fraction conditions. However, these errors in HOCN are not significant enough to have any apparent negative effect on the ANN predictions of other species, so the rest of the predictions with

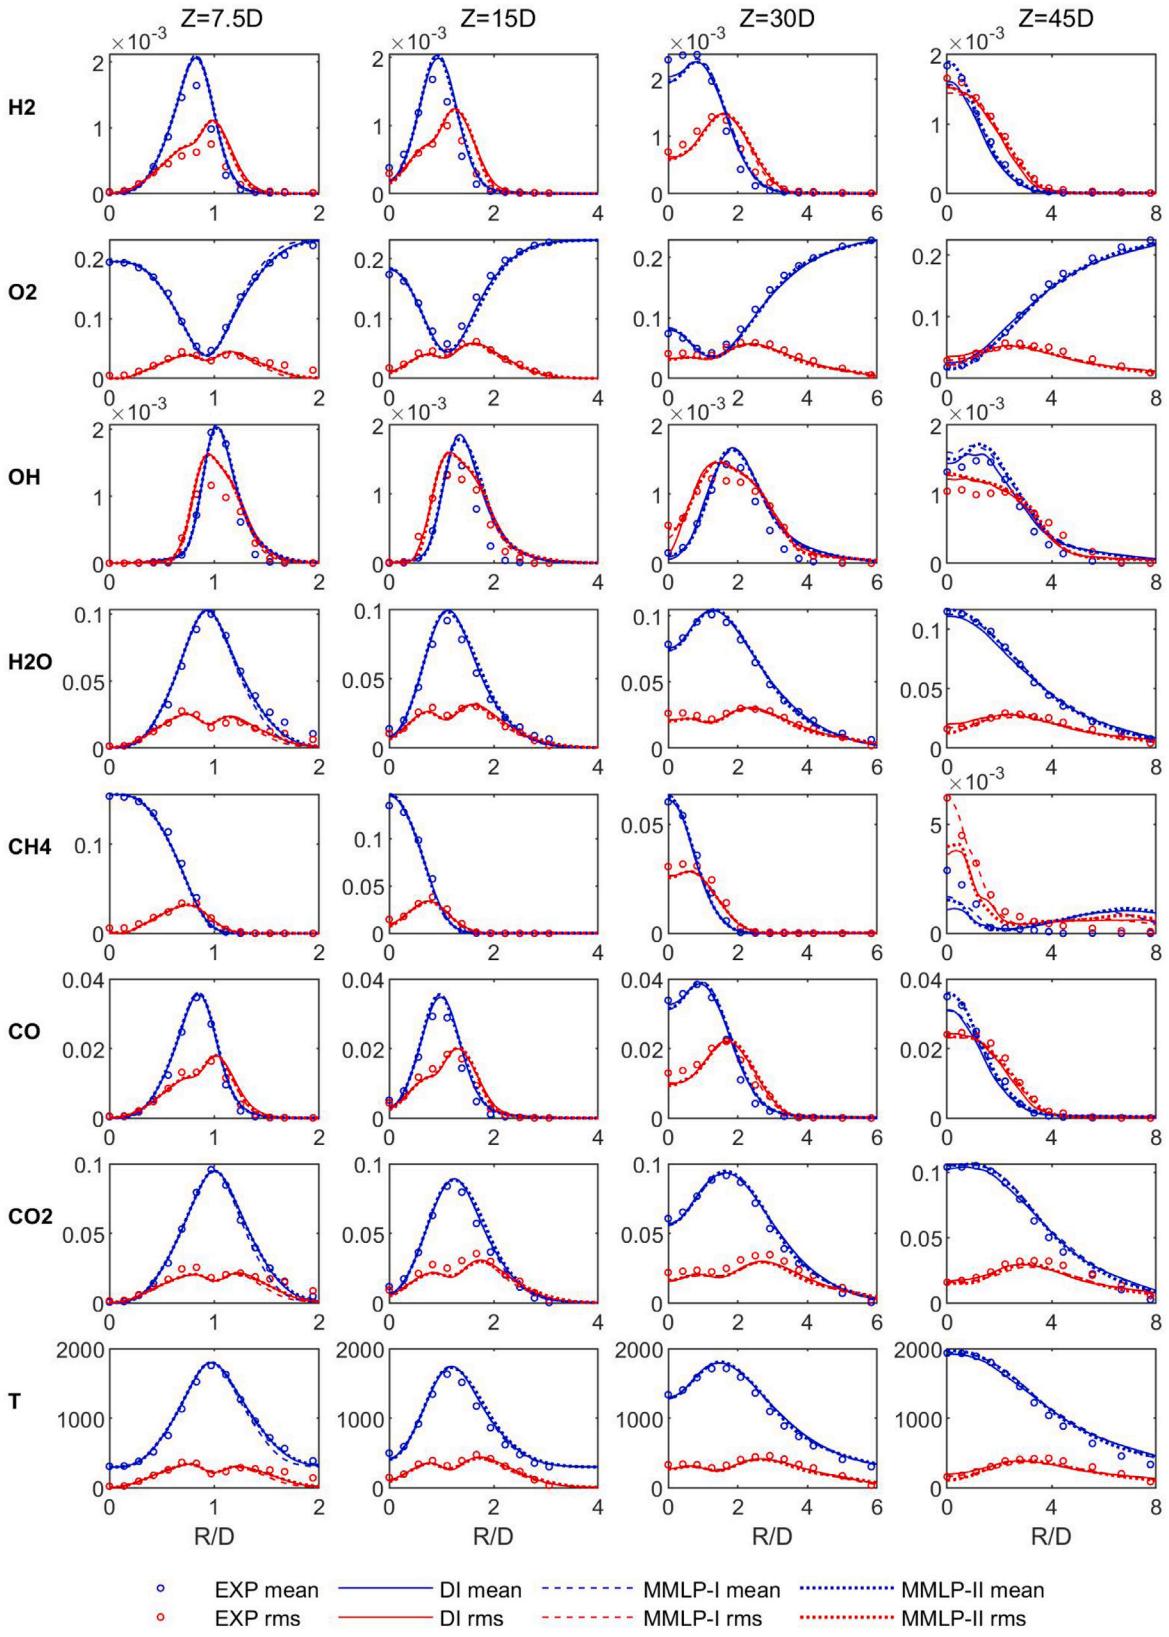

Fig. 9. Major species and temperature in flame D.

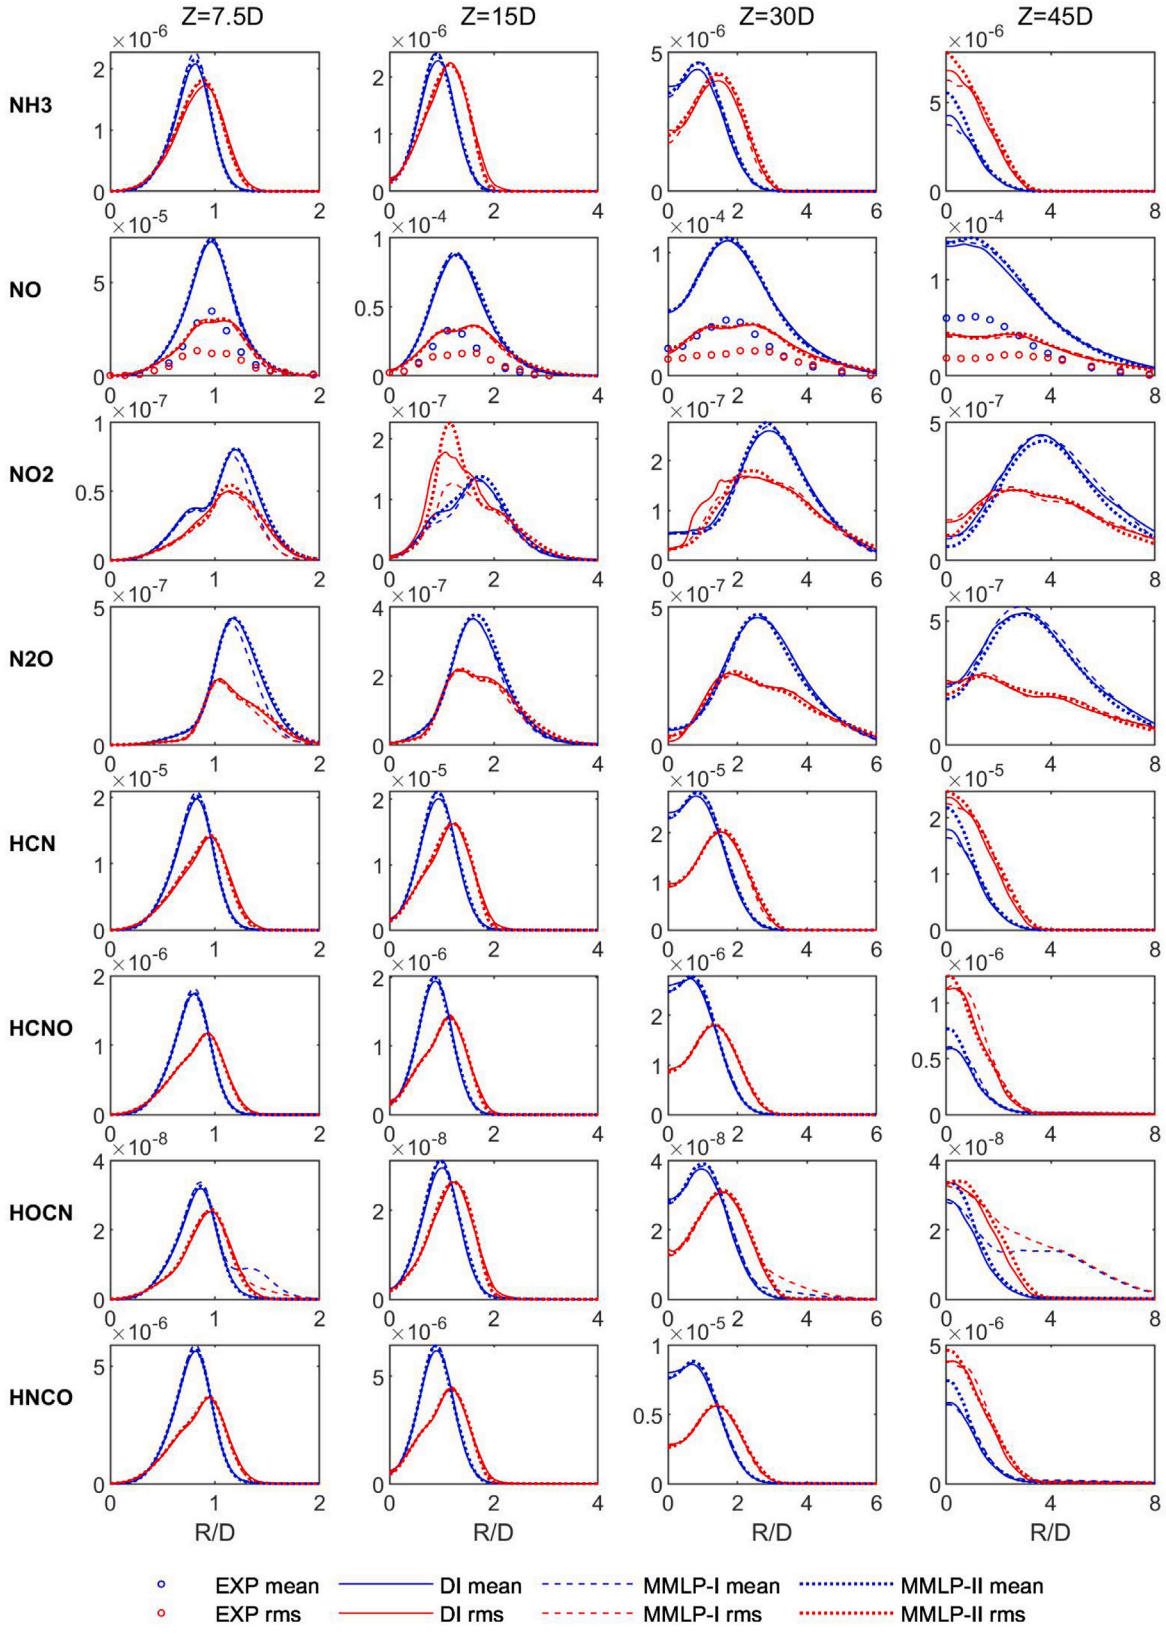

Fig. 10. Nitrogen-containing species of flame D.

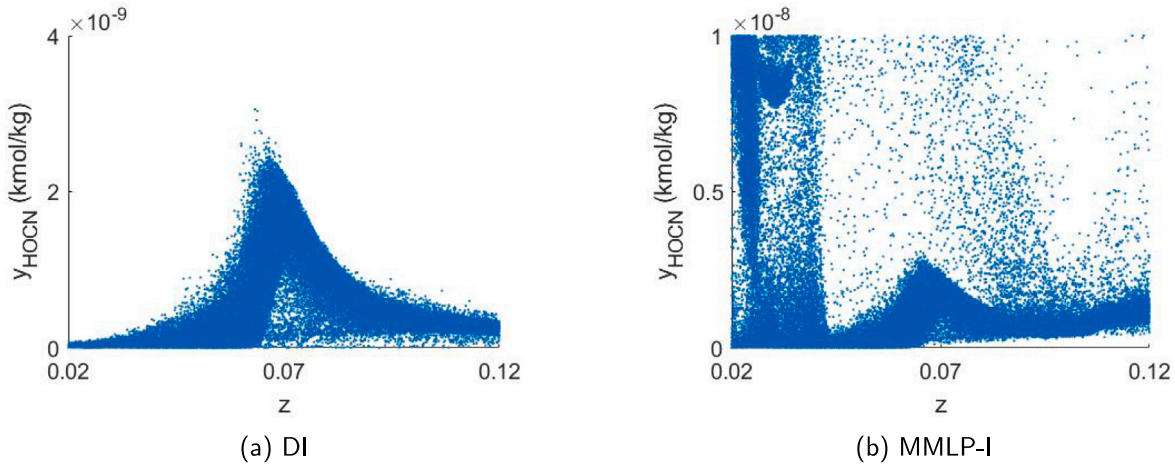

Fig. 11. Instantaneous results of HOCN vs mixture fraction after 10,000 reacting simulation steps.

MMLP-I are still excellent. The ANN error accumulation of HOCN can be alleviated using MMLP-II method, as shown in the corresponding plots in Fig. 10, where the HOCN profiles produced with MMLP-II show good agreement with those of DI except for the slight overprediction near the centerline at Z = 45D.

The performance of either MMLP method on the Sandia flame is consistent with the findings reported in Section 3, where MMLP-II improved the performance of minor species and HOCN in particular. Also, the small deviations of MMLP-II for the peak values of H_2 and CO are explained by the fact that, as discussed in Section 2.3.4, the MMLP-I outperforms MMLP-II on composition states with large input species concentration. Overall, the results of both methods are in good agreement with DI and the only major disagreement is that of HOCN with MMLP-I.

Regarding the comparison with experimental data where available, the simulation results for the major species and temperature are in good agreement while NO is overpredicted by a factor of about 2 (see Fig. 10). The overprediction of NO is mainly caused by the mechanism GRI-3.0, and has also been observed in other works [49–51]. In order to obtain accurate predictions of NO, the GRI-2.11 mechanism could be used instead. However, from the aspect of ANN chemistry tabulation, GRI-3.0 is more challenging to tabulate than GRI-2.11 because it involves more species and reactions, so it was a better choice for the present work. It should be emphasised that the experimental data are shown for reference only, as the performance of the ANN tabulation should be judged only based on the comparison between ANN and DI results.

4.2. Swirl-bluff-body stabilised flame - Sydney SMA2

Swirling flows and bluff-body stabilisers are widely employed in industrial applications. In order to facilitate the development of turbulent combustion simulations for swirling flames, a series of bluff-body swirl-stabilised flames, namely the Sydney swirl flames, with different inflow velocities, swirling velocities and fuel mixtures, have been developed [54–57] as benchmarks. The introduction of the bluff body and the swirling flow leads to more complicated structures such as the recirculation zone and neck zone. Therefore, these flames are good validation cases to test our ANN methodology on more complex combustion problems. In this work, the SMA2 flame is chosen to evaluate the performance of the two MMLP methods.

The burner of the Sydney swirl flames has a fuel jet with a diameter of 3.6 mm, centred in a bluff body with a diameter of 50 mm. The bluff body is surrounded by an annular swirling flow with an outer diameter of 60 mm. The entire burner is surrounded by an external ambient air coflow with a velocity of 20 m/s. For flame SMA2, The jet composition is 1/3 methane and 2/3 air by volume, with a bulk velocity of 66.3 m/s. The axial and tangential velocities of the annular swirling air are 16.3 m/s and 25.9 m/s, respectively. The simulation domain has cylindrical dimensions of 25 (axial) × 7.5 (radial) cm. The grid is uniform in the axial direction, with a cell size of 1 mm. In the radial direction, the cell size is about 0.6 mm in the centre and then expands smoothly towards the coflow region. The entire grid has a total of about 2M cells.

The MMLP-I method was first applied to the SMA2 simulation. However, it was found that the error in HOCN accumulates significantly in this case. By the time the simulation had completed 10,000 time steps after ignition, the error of HOCN had reached an unacceptable level. The corresponding instantaneous results are shown in Fig. 11, where the HOCN concentrations (y_{HOCN}) of all the reacting states are plotted against mixture fraction (z). When using DI, the peak HOCN concentration at this stage is about $3 \cdot 10^{-9}$ kmol/kg, located at $z \approx 0.7$, and the concentrations at low mixture fraction locations ($z < 0.04$) are quite small. However, when applying MMLP-I, a large amount of highly overpredicted data appear, as shown in Fig. 11(b). Most of the large errors occur at low mixture fraction locations, and the maximum HOCN concentration in this case is over three times higher than that of DI (note that the Y-axis of Fig. 11(b) covers a range that is one order of magnitude higher than that of Fig. 11(a)). While HOCN may not be of interest in itself, simulations showed that, in the case of the SMA2, the very large errors affected the prediction of other NO_x species as well as many other species. After the flame was fully developed, it was found that the maximum N_2O and NO_2 concentrations were about 10 times higher than those obtained with DI. The maximum NO concentration was also overpredicted by a factor of two. Several major species were also affected by the error accumulation in minor species. The mean valueCH_4 and H_2 showed errors of about 25% near the centreline at high axial locations. It was therefore concluded that the MMLP-I method is not appropriate for this flame.

It is worth commenting on the reasons for this error accumulation in Sydney SMA2, which was not the case in Sandia D. The bluff body of the SMA2 flame causes recirculation zones, such that the fluid elements will go through many simulation steps along their pathlines before being flushed out of the simulation domain. By contrast, in a piloted jet flame such as Sandia D, the pathlines will move upstream till the outlet boundary, and the accumulated errors of HOCN will be flushed out before becoming unacceptable.

The MMLP-II method was then applied to the Sydney SMA2. The instantaneous results of HOCN when the flame is fully developed are shown in Fig. 12. We can clearly see that the results of MMLP-II are in good agreement with those of DI. The overall distribution, the peak value and the peak locations are similar in both cases. Continuing

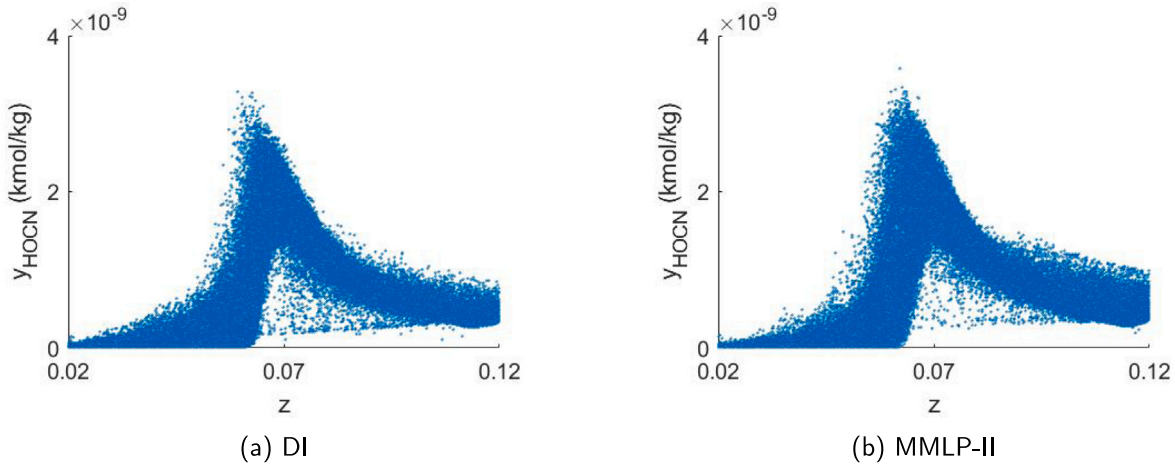

Fig. 12. Instantaneous results of HOCN vs mixture fraction after the flame is fully developed.

the simulation, the mean and RMS values of species and temperature are collected at 5 axial locations: Z = 10, 20, 50, 70, 100 mm. The comparison of the statistical quantities between MMLP-II and DI are shown in Figs. 13 to 14.

The profiles of major species' mass fractions and temperature are shown in Fig. 13, as computed by ANNs (MMLP-II) and DI, while the experimental data are also shown. The results of MMLP-II match the results of DI very well, except for some small discrepancies for the mean profiles at location Z = 10 mm and 20 mm. Results for minor species are shown in 14, together with experimental data in the case of NO. For most species, the ANN simulation results are excellent at Z = 50 mm and 70 mm. At Z = 100 mm, small errors can be seen for some species, but overall the ANN profiles are still very accurate. At Z = 10 mm and 20 mm, the peak values of several species are overpredicted by ANNs. The bluff body in this flame causes recirculation zones at low axial locations and, as mentioned before, this can result in error accumulation, which is most likely the reason for these discrepancies. Having said that, the ANN results are still very accurate for the major species at these locations and reasonably good for the remaining species.

While comparison with the experimental data is not relevant for assessing the ANN accuracy, it is worth commenting that the results are overall in reasonably good agreement, although there are considerable discrepancies for some major species, particularly CO and H_2, while NO is overpredicted by a factor of two. There are several factors that play a role on this, such as the grid resolution and the use of GRI3.0 for NO prediction (see comments in Section 4.1). Further improvements on these aspects are out of the scope of the present work.

5. Conclusions and future work

In the present work, the ANN tabulation methodology for thermochemistry kinetics was further developed and applied to piloted diffusion and swirl-bluff-body stabilised flames with NO_x formation. In particular, a method termed Multiple Multilayer Perceptron - II (MMLP-II) was developed to provide the higher accuracy needed for predicting species with small concentrations, as encountered in NO_x chemistry. Compared with the previous MMLP-I method [23], which trains multiple MLPs based on the output magnitude of a certain species, the MMLP-II method proposed here trains MLPs focussing on data with different input ranges for the species being predicted.

A first test of the MMLP-II method was performed on 1-D laminar flamelets, and results were compared with direct integration (DI) of chemical kinetics. While both MMLP methods produced excellent results in the vast majority of cases, discrepancies were noted with MMLP-I at very low strain rates and low values of mixture fraction for some N-containing species, particularly HOCN, due to the low concentration that they exhibited under these conditions. This issue was remedied by MMLP-II, reaffirming the capacity of this approach to address problems of this nature.

The two aforementioned turbulent flames were subsequnetly simulated with both MMLP methods, as well as with DI. In the case of Sandia D, both methods produced very good results, and the only significant discrepancy was for HOCN in the case of MMLP-I. This was remedied by MMLP-II, consistent with the observations in the case of the flamelets. However, the errors in HOCN did not corrupt the predictions of other species and, since this particular species is not likely to be of interest in itself, it can be concluded that MMLP-I is sufficient in this case. Application to the Sydney SMA2 flame showed a different picture, however. In this case, due to the recirculation zone, the errors in HOCN accumulated to an unacceptable level and affected NO and other species. Therefore, MMLP-II was applied and indeed achieved very good agreement with DI results for all species. It should be emphasised that the ANNs used to simulate the turbulent flames were trained with a general approach of data generation based on canonical problems and partial randomisation, rather than with data collected from the problems to be simulated. Therefore, the fact excellent predictions were obtained on two different turbulent flames indicates the strong capacity of the approach for generalisation.

A speed-up factor of about 15 for the reaction source term was accomplished, which can be regarded as sufficient for most simulations. If, for example, in a certain combustion LES implementation, chemical kinetics takes 90% of the total CPU time with DI, then 100 s of overall simulation time can be reduced to at best 10 s (if the reaction step takes zero time); the proposed ANNs methodology would already bring it down to 16 s and only marginal gains could be achieved with further speed-up. It should also be stressed that, while the turbulent combustion simulations in the present work were performed using the LES-PDF method, the ANN tabulation approach is equally applicable to any combustion modelling approach that requires real-time computation of the reaction source term.

Further advances in the MMLP method could include combining both MMLP methods so that the optimal MLP can be selected for the given composition states. For example, in Fig. 5, the data in areas 1, 2, and 3 could be predicted by MMLP-I while the data in areas 4 and 5 could be predicted by MMLP-II. In the present work, the MMLP-II method was sufficient on its own, but such a combination may prove beneficial for more complex mechanisms. Furthermore, the ability of MMLP-II to provide accurate prediction of minor species can be exploited in other problems of chemical kinetics, and particularly in soot, where it is important to predict accurately the precursors. Most important, the methodology presented in this paper is applicable to

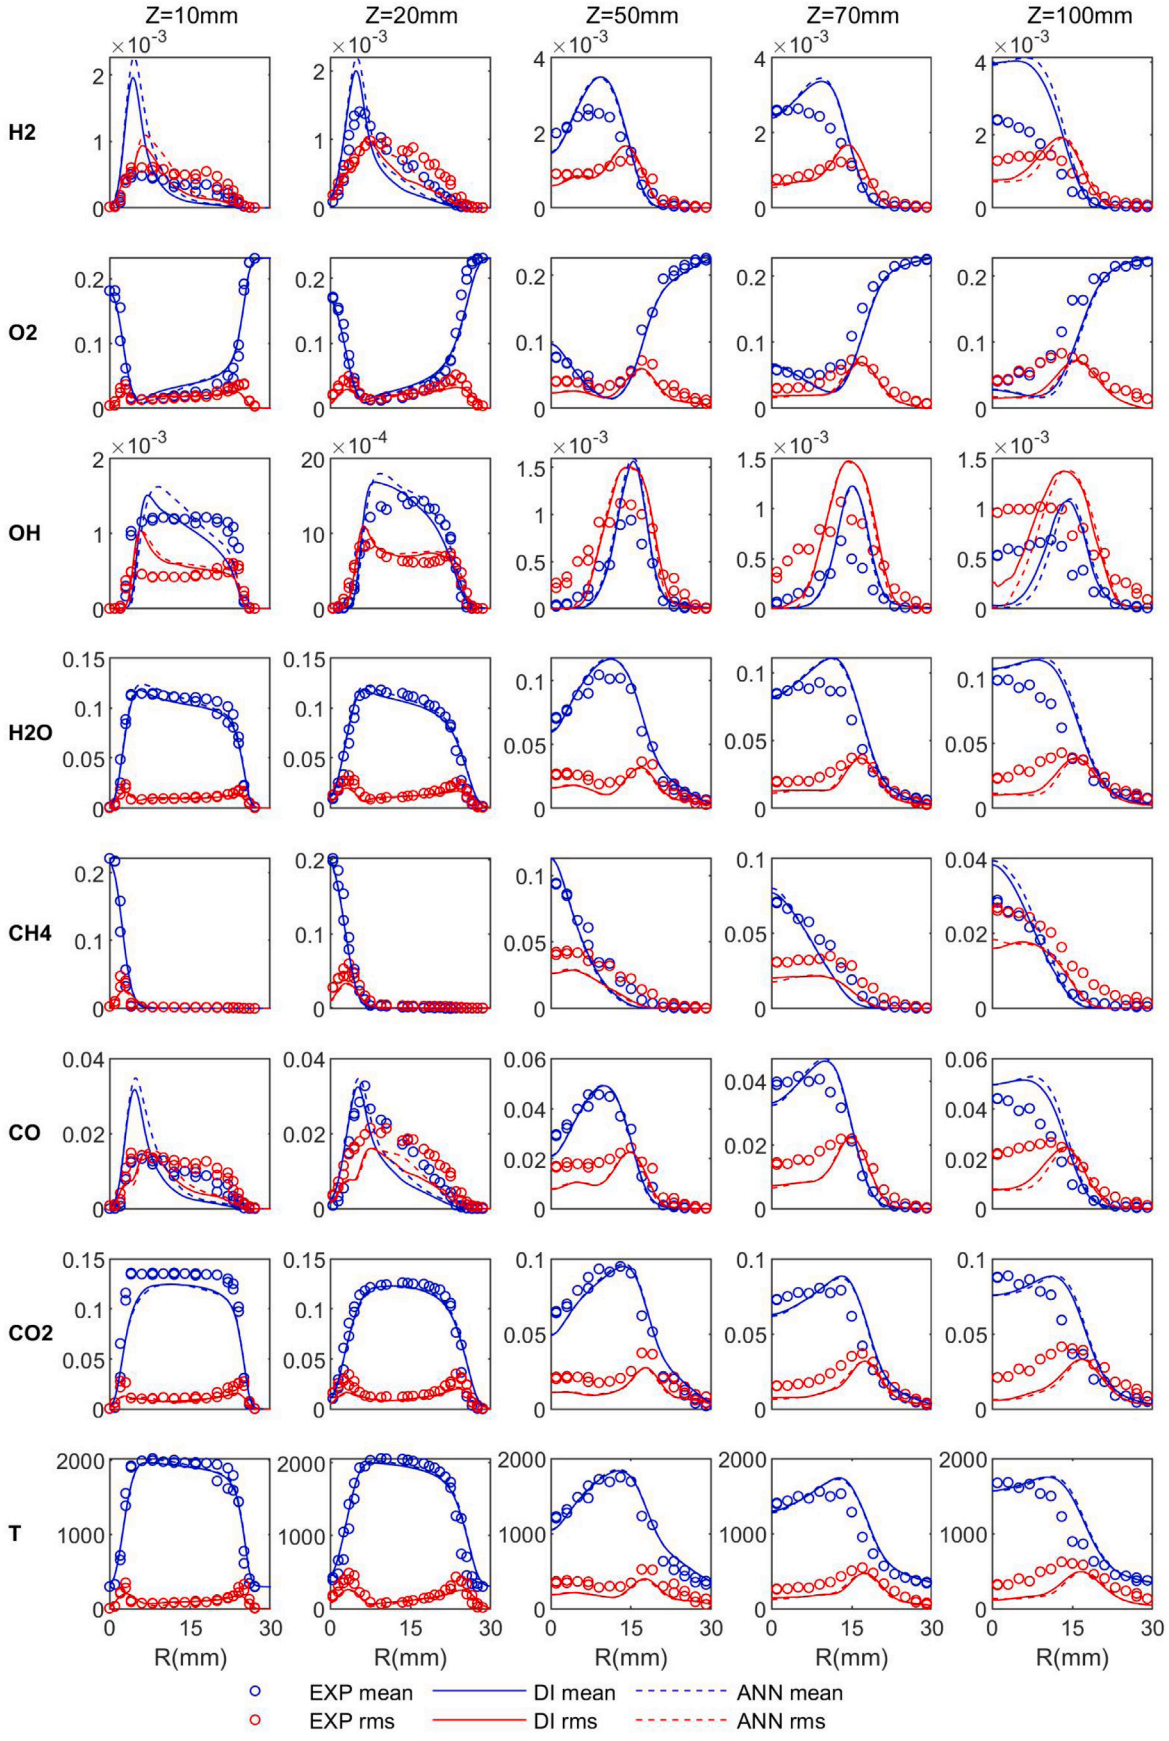

Fig. 13. Major species and temperature of flame SMA2. Comparison between experiment, DI and ANN (MMLP-II).

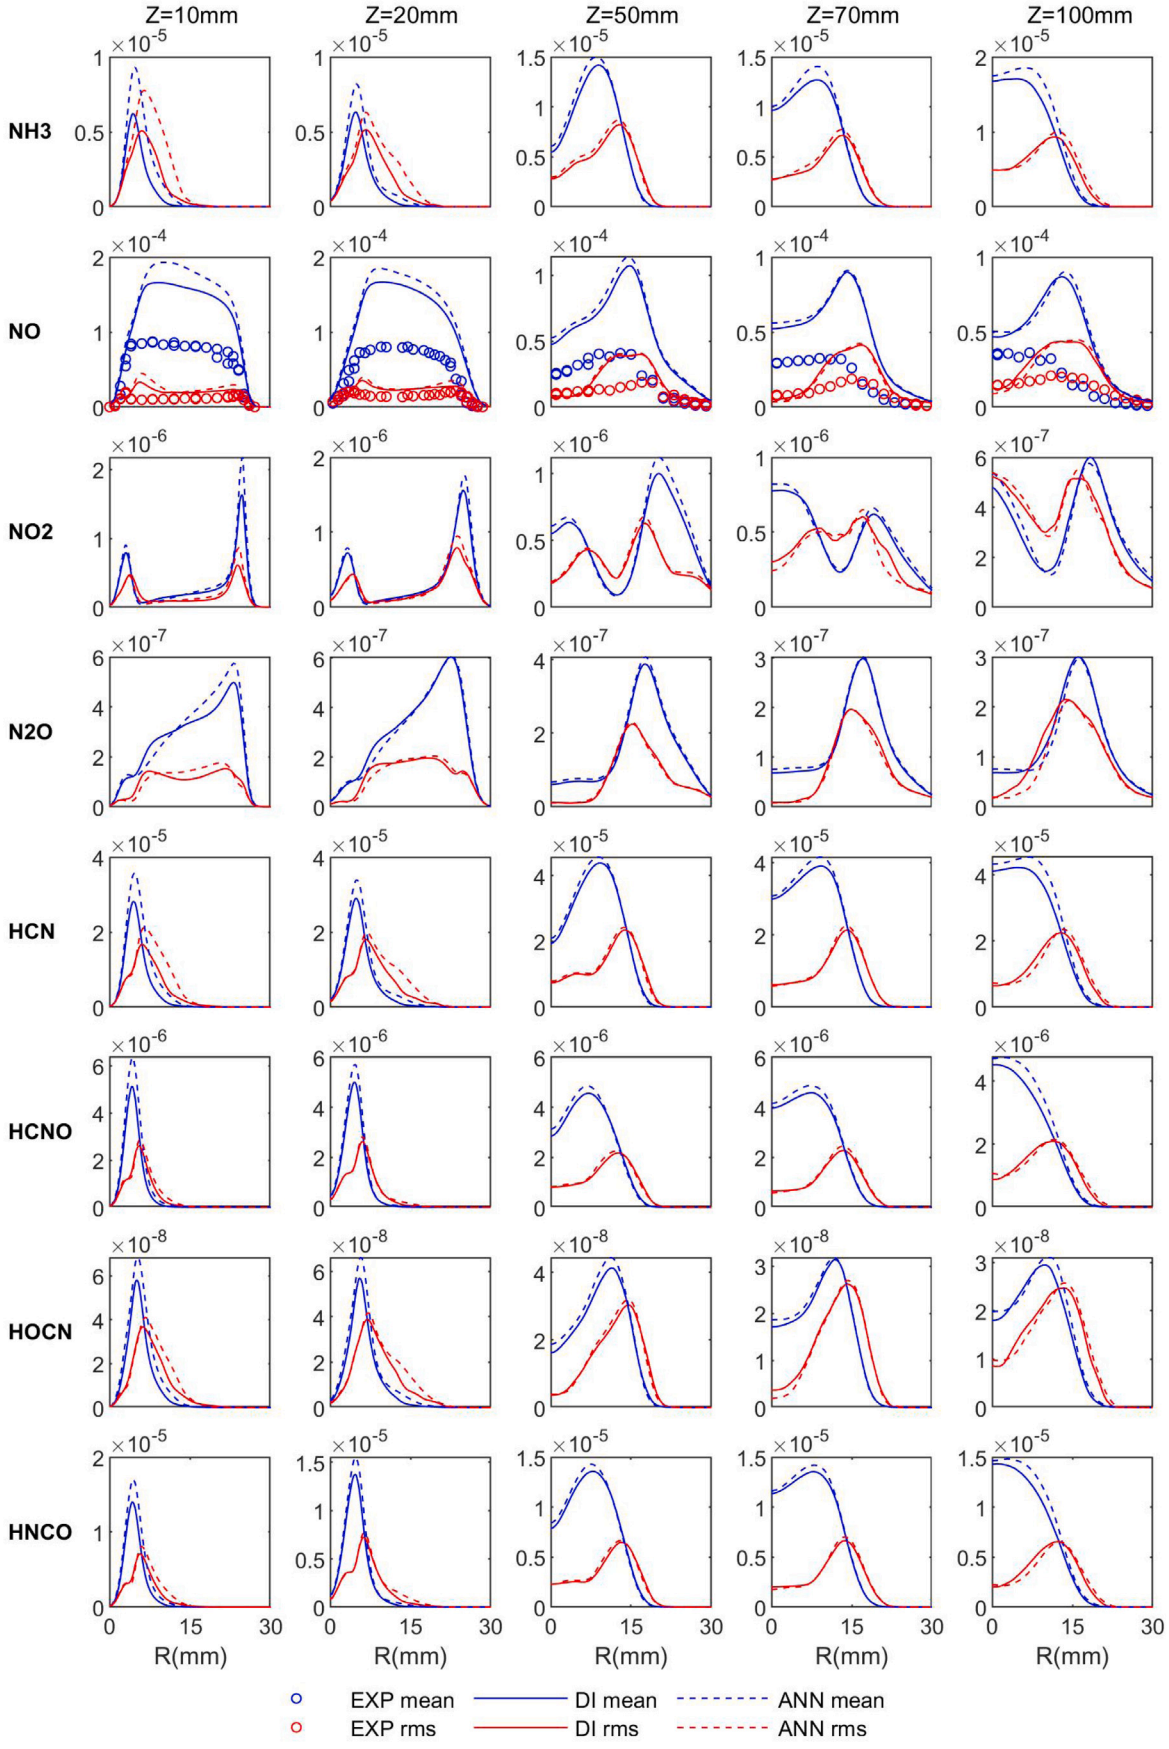

Fig. 14. Nitrogen-containing species of flame SMA2. Comparison between experiment, DI and ANN (MMLP-II).

a wider range of combustion problems involving N-chemistry, which is important especially in light of the recent interest in ammonia combustion.

CRediT authorship contribution statement

Tianjie Ding: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **W.P. Jones:** Writing – review & editing, Supervision, Software. **Stelios Rigopoulos:** Writing – review & editing, Supervision, Project administration.

Acknowledgements

Tianjie Ding gratefully acknowledges the financial support provided by the China Scholarship Council, China (CSC) and the Department of Mechanical Engineering, Imperial College London. For the computational resources, we wish to acknowledge the Imperial College Research Computing Service (https://doi.org/10.14469/hpc/2232), the UK Materials and Molecular Modelling Hub partially funded by EPSRC, United Kingdom (EP/P020194/1 and EP/T022213/1) and the ARCHER UK National Supercomputing Service (http://www.archer.ac.uk) via the UK Consortium on Turbulent Reacting Flows (www.ukctrf.com) supported by EPSRC, United Kingdom grant EP/R029369/1.

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.

References

[1] M. Ihme, W.T. Chung, A.A. Mishra, Combustion machine learning: Principles, progress and prospects, Prog. Energ. Combust. 91 (2022) 101010.
[2] M. Ihme, C. Schmitt, H. Pitsch, Optimal artificial neural networks and tabulation methods for chemistry representation in LES of a bluff-body swirl-stabilized flame, P. Combust. Inst. 32 I (2009) 1527–1535.
[3] 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. (2020) 1–27.
[4] F.A. Williams, Combustion Theory, second ed., Perseus Books, Reading, Massachusetts, 1985.
[5] S. Lam, D. Goussis, Understanding complex chemical kinetics with computational singular perturbation, Symp. (Int.) Combust. 22 (1989) 931–941.
[6] S.H. Lam, D.A. Goussis, The CSP method for simplifying kinetics, Int. J. Chem. Kinet. 26 (1994) 461–486.
[7] A. Massias, D. Diamantis, E. Mastorakos, D. Goussis, An algorithm for the construction of global reduced mechanisms with CSP data, Combust. Flame 117 (1999) 685–708.
[8] T.T. Lovas, D. Nilsson, F. Mauss, Automatic reduction procedure for chemical mechanisms applied to premixed methane/air flames, P. Comb. Inst. 28 (2000) 1809–1815.
[9] J.C. Keck, D. Gillespie, Rate-controlled partial-equilibrium method for treating reacting gas mixtures, Combust. Flame 17 (1971) 237–241.
[10] W. Jones, S. Rigopoulos, Rate-controlled constrained equilibrium: Formulation and application to nonpremixed laminar flames, Combust. Flame 142 (2005) 223–234.
[11] J.-Y. Chen, W. Kollmann, R. Dibble, Pdf modeling of turbulent nonpremixed methane jet flames, Combust. Sci. Technol. 64 (1989) 315–346.
[12] S. Pope, Computationally efficient implementation of combustion chemistry using in situ adaptive tabulation, Combust. Theor. Model. 1 (1997) 41–63.
[13] F. Christo, A. Masri, E. Nebot, T. Turanyi, Utilizing artificial neural network and repro-modelling in turbulent combustion, in: IEEE International Conference on Neural Networks - Conference Proceedings, vol. 2, 1995, pp. 911–916.
[14] F. Christo, A. Masri, E. Nebot, S. Pope, An integrated PDF/neural network approach for simulating turbulent reacting systems, Symp. (Int.) Combust. 26 (1996) 43–48.
[15] J. Blasco, N. Fueyo, C. Dopazo, J. Ballester, Modelling the temporal evolution of a reduced combustion chemical system with an artificial neural network, Combust. Flame 113 (1998) 38–52.
[16] J. Blasco, N. Fueyo, J. Larroya, C. Dopazo, Y.-J. Chen, A single-step time-integrator of a methane-air chemical system using artificial neural networks, Comput. Chem. Eng. 23 (1999) 1127–1133.
[17] J.-Y. Chen, J. Blasco, N. Fueyo, C. Dopazo, An economical strategy for storage of chemical kinetics: Fitting in situ adaptive tabulation with artificial neural networks, P. Combust. Inst. 28 (2000) 115–121.
[18] J. Blasco, N. Fueyo, C. Dopazo, J.-Y. Chen, A self-organizing-map approach to chemistry representation in combustion applications, Combust. Theor. Model. 4 (2000) 61–76.
[19] 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, P. Combust. Inst. 34 (2013) 1465–1473.
[20] L. Franke, A. 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.
[21] 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. Energ. 45 (2020) 29594–29605.
[22] 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 (2021) 035154.
[23] 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 231 (2021) 111493.
[24] T. Ding, S. Rigopoulos, W. Jones, Machine learning tabulation of thermochemistry of fuel blends, Appl. Energy Combust. Sci. (2022) 100086.
[25] 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.
[26] T. Zhang, Y. Yi, Y. Xu, Z.X. Chen, Y. Zhang, E. Weinan, Z. Xu, A multi-scale sampling method for accurate and robust deep neural network to predict combustion chemical kinetics, Combust. Flame 245 (2022) 112319.
[27] R. Mao, M. Lin, Y. Zhang, T. Zhang, Z.J. Xu, Z.X. Chen, DeepFlame: A deep learning empowered open-source platform for reacting flow simulations, Comput. Phys. Comm. 291 (2023) 108842.
[28] Z.-Q.J. Xu, J. Yao, Y. Yi, L. Hang, W. E, Y. Zhang, T. Zhang, Solving multiscale dynamical systems by deep learning, 2024, arXiv:2401.01220.
[29] R. Mao, M. Zhang, Y. Wang, H. Li, J. Xu, X. Dong, Y. Zhang, Z.X. Chen, An integrated framework for accelerating reactive flow simulation using GPU and machine learning models, P. Combust. Inst. 40 (1–4) (2024) 105512.
[30] B. Sen, S. Menon, Linear eddy mixing based tabulation and artificial neural networks for large eddy simulations of turbulent flames, Combust. Flame 157 (2010) 62–74.
[31] B. Sen, E. Hawkes, S. Menon, Large eddy simulation of extinction and reignition with artificial neural networks based chemical kinetics, Combust. Flame 157 (2010) 566–578.
[32] T. Readshaw, L.L.C. Franke, W.P. Jones, S. Rigopoulos, Simulation of turbulent premixed flames with machine learning - tabulated thermochemistry, Combust. Flame 258 (2023) 113058.
[33] A. Liu, T. Ding, R. Liu, S. Rigopoulos, K. Luo, Machine learning tabulation of thermochemistry for turbulent dimethyl ether (DME) flames, Fuel 360 (2024) 130338.
[34] T. Readshaw, W.P. Jones, S. Rigopoulos, On the incorporation of conservation laws in machine learning tabulation of kinetics for reacting flow simulation, Phys. Fluids 35 (2023) 047103.
[35] G. Smith, D.M. Golden, M. Frenklach, N. Moriarty, B. Eiteneer, M. Goldenberg, C. Bowman, R. Hanson, S. Song, W. Gardiner Jr., V. Lissianski, Z. Qin, GRI-mech home page, 2024, URL http://www.me.berkeley.edu/gri_mech/. (Accessed 17 March 2024).
[36] M. Frenklach, H. Wang, M. Goldenberg, G.P. Smith, D.M. Golden, C.T. Bowman, R.K. Hanson, W.C. Gardiner, V. Lissianski, GRI-MECH: An Optimized Detailed Chemical Reaction Mechanism for Methane Combustion, Tech. Rep., Gas Research Institute Topical Report No. GRI-95/0058, 1995.
[37] M.T. Hagan, H.B. Demuth, M.H. Beale, O. De Jesús, Neural Network Design, second ed., Martin Hagan, 2014.
[38] D. Livingstone, Artificial Neural Networks, vol. 458, Springer, 2009.
[39] Y. Goodfellow, A. Courville, Deep Learning, MIT Press, 2016.
[40] N. Peters, Turbulent Combustion, Cambridge University Press, 2000.
[41] P.B. Brown, G.D. Byrne, A.C. Hindmarsh, VODE: A variable-coefficient ODE solver, SIAM J. Sci. Stat. Comp. 10 (1989) 1038–1051.
[42] S. Zhang, C. Zhang, B. Wang, CRK-PINN: A physics-informed neural network for solving combustion reaction kinetics ordinary differential equations, Combust. Flame 269 (2024) 113647.
[43] International workshop on measurement and computation of turbulent flames, 2024, URL https://tnfworkshop.org/. (Accessed 17 March 2024).
[44] W. Jones, V. Prasad, Large Eddy Simulation of the Sandia Flame Series (D-F) using the Eulerian stochastic field method, Combust. Flame 157 (2010) 1621–1636.
[45] W.P. Jones, F. di Mare, A.J. Marquis, LES-BOFFIN: user's guide, 2002.
[46] W.L. Grosshandler, RADCAL: a narrow-band model for radiation, Calc. A Combust. Environ. NIST Tech. Note 1402 (1993).

[47] R. Barlow, A. Karpetis, J. Frank, J.-Y. Chen, Scalar profiles and NO formation in laminar opposed-flow partially premixed methane/air flames, Combust. Flame 127 (2001) 2102–2118.
[48] R. Barlow, J. Frank, Effects of turbulence on species mass fractions in methane/air jet flames, Symp. (Int.) Combust. 27 (1998) 1087–1095.
[49] H. Wang, Y. Chen, PDF modelling of turbulent non-premixed combustion with detailed chemistry, Chem. Eng. Sci. 59 (2004) 3477–3490.
[50] V. Raman, R. Fox, A. Harvey, Hybrid finite-volume/transported PDF simulations of a partially premixed methane-air flame, Combust. Flame 136 (2004) 327–350.
[51] R. Cao, S. Pope, The influence of chemical mechanisms on PDF calculations of nonpremixed piloted jet flames, Combust. Flame 143 (2005) 450–470.
[52] R. Cao, H. Wang, S. Pope, The effect of mixing models in PDF calculations of piloted jet flames, P. Combust. Inst. 31 I (2007) 1543–1550.
[53] T. Jaravel, E. Riber, B. Cuenot, P. Pepiot, Prediction of flame structure and pollutant formation of Sandia flame D using Large Eddy Simulation with direct integration of chemical kinetics, Combust. Flame 188 (2018) 180–198.
[54] P.A.M. Kalt, Y.M. Al-Abdell, A.R. Masri, R.S. Barlow, Swirling turbulent non-premixed flames of methane: flow field and compositional structure, P. Combust. Inst. 29 (2002) 1913–1919.
[55] Y.M. Al-Abdeli, A.R. Masri, Stability characteristics and flowfields of turbulent non-premixed swirling flames, Combust. Theor. Model. 7 (2003) 731.
[56] Y.M. Al-Abdeli, A.R. Masri, Precession and recirculation in turbulent swirling isothermal jets, Combust. Sci. Technol. 176 (2004) 645–665.
[57] A.R. Masri, P.A.M. Kalt, R.S. Barlow, The compositional structure of swirl-stabilised turbulent nonpremixed flames, Combust. Flame 137 (2004) 1–37.